From theta periods to the Riemann hypothesis

Our first aim is to reconstruct Levent Alpöge's argument for the Riemann hypothesis in ordinary mathematical language. We work toward a proof that a reader can check independently of Lean, using established theorems and explicit derivations. Our second aim is to identify genuinely new mathematics by comparing the argument's results and methods with what was known before the announcement. Alpöge's announcement and Sahibzada Allahyar's reconstruction introduce the argument [1, 3].

The proof map is a directed acyclic graph that can be read from the main implications down to individual arguments. At each level, preceding conclusions and stated premises must supply every input on the same data, with the required normalizations and quantifier order. A refinement must recover the same conclusion under the same hypotheses. AI agents use the map to select and refine a step, write its proof and update its inputs, outputs and links together. An independent agent checks the proof. An Astra reviewer assesses substantive candidates afresh from their precise statements and the earlier literature.

For results from the literature we check the exact theorem or passage in its original source and state the special case used here. We have sought to fill in as much of the surrounding argument as possible by direct derivation, consulting the pinned Lean release for difficult constructions, normalizations or arguments [2].

The inventory brings together nineteen AI-proposed candidates for new mathematics, grouped by subject. Human specialist assessment is pending.

The chain of roughly $\sqrt{x}$ theta lifts plays the role of multiplicity in a Stepanov-type argument [1, pp. 1, 4]. The analytic estimate concerns the centered sum

$$G_1(x) = \sum_{n\geq 1} \Lambda(n)(x/n-1)_+ - x\log x + (1+\gamma)x \qquad (x>0). \tag{1}$$

Here Λ is the von Mangoldt function, γ is Euler's constant and $r_+ = \max(r,0)$. For every sufficiently small fixed $\varepsilon > 0$, the assembled estimates must give arbitrarily large Y with

$$G_1(x) \leq x^{1/2+9\varepsilon} \qquad \text{for every real } x \in [Y, Y^{3/2}].$$

The centering cancels the polar part at 1 in the Mellin integral over $x \geq 1$. For a zero ρ with $1/2 < \Re\rho < 1$, choose $1/2 < \theta < \Re\rho$ and fix a sufficiently small $\varepsilon > 0$ with $1/2 + 9\varepsilon \leq \theta$. With the interval bound, the sign-change theorem excludes the simple pole at ρ [76]. The functional equation and classical boundary facts then give RH. Section 6 also gives the Dirichlet extension from the corresponding window bound for each primitive target.

Original argument by Levent Alpöge.
Reading edition by The Theta Periods to RH Reading Edition contributors, 22 November 2026.
Original text and figures in this edition are licensed under CC BY 4.0. Third-party materials retain their terms.
Project repository Edition v2026.11.22

Contents

Candidates for new mathematical results **3**

1 CM fields with character and Selmer bounds **9**
1.1 The auxiliary primes and character estimates . . . 9
1.2 The Selmer bound on the same scales . . . 14
1.3 Level primes, sampling sets and character estimates . . . 15

2 Theta lifts and a common period lattice **18**
2.1 The towers and the top pairing . . . 19
2.2 Changing the path of the local tests . . . 19
2.3 The principal-level row test . . . 23
2.4 Bessel periods and Poincaré series . . . 69
2.5 From top theta periods to bottom Bessel periods . . . 100
2.6 The coefficient-field map and finite-place tests . . . 103
2.7 The matrices in common row coordinates . . . 136

3 Relative traces and the span of the periods **137**
3.1 Unfolding the full Gram operators . . . 137
3.2 The single-word estimate . . . 139
3.3 Squaring the orbital sum at finite level . . . 141
3.4 The leading contribution to the mixed Gram blocks . . . 142
3.5 Trace moments after restricting the columns . . . 144
3.6 The span of restricted columns . . . 148
3.7 Projection to the common row dual . . . 148
3.8 Counting the distinguished rows . . . 148
3.9 From ranks to the mean slope . . . 152

4 Period lattices and centered prime-power sums **154**
4.1 From valuations to integral maps . . . 154
4.2 Valuation gains and Frobenius classes . . . 218
4.3 Degrees and slopes of the source lattices . . . 221
4.4 The CM period and the centering term . . . 417
4.5 The upper bound at each translate . . . 423

5 Fourier inversion and the pointwise bound **424**
5.1 Sampling degrees and weights at the tied indices . . . 424
5.2 Finite Fourier inversion . . . 425
5.3 Reduction to the large weights . . . 426
5.4 The translate average . . . 427
5.5 Selecting the target characters . . . 427
5.6 The square-root scale . . . 428

6 Bounds on intervals and the critical line **429**
6.1 One family for the lattice and Fourier estimates . . . 429
6.2 The centered zeta transform . . . 438
6.3 One-sided windows and the critical line . . . 439
6.4 Extension to Dirichlet L-functions . . . 440

References **442**

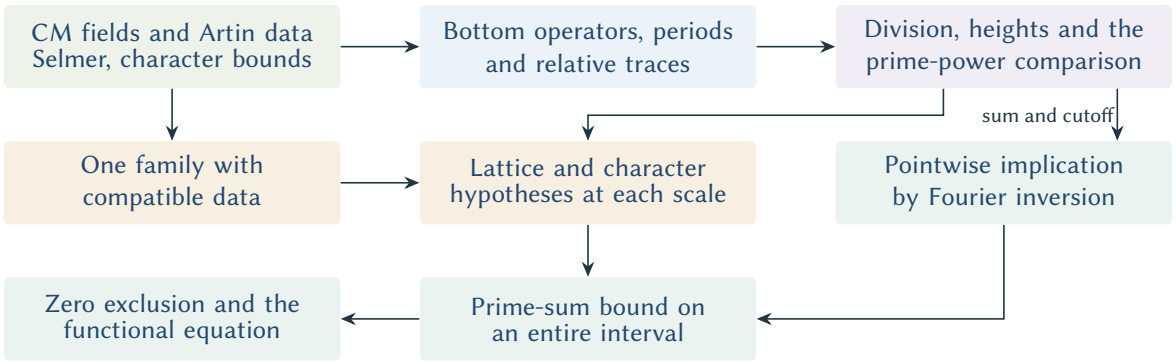

Figure 1. The period and trace arguments of Sections 2–3 share a box. Section 1 supplies Selmer and Artin data for Section 4 and character estimates for Section 5. Apply the lattice inequality once its period, rank and height hypotheses hold on the selected records. Then apply the Fourier implication with the same cutoff and centering term and the character estimates on the base field. The ten-stage proof overview links these implications to their statements and arguments in the paper.

Candidates for new mathematical results

We compare the nineteen AI-proposed candidates below with earlier results under the hypotheses, normalizations and uniformity needed here. A short deduction from earlier work can supply the required conclusion. A broader range or a reusable argument can make a candidate worth recording. The sections specify the scope of the independent proof checks.

Arithmetic windows and Galois cohomology

Uniform bounds for higher fine Selmer groups. Fix the primitive target ψ, $0 < \varepsilon < 1/52$ and the dyadic profile h from Section 1. The release gives $J \geq 1$ and X_0. For every $X \geq X_0$ and every pair q, d with the prescribed sizes and congruences at $Z = X^{1/J}$, there is $W \geq Z$ with $W^2 \leq X$. At every real $y \in [W, W^2]$, every CM field E of the prescribed type at $(h(y), q, d)$ satisfies

$$\sum_{\ell \in S} (\log \ell) \sum_{1 \leq r < \lfloor y^{401} \rfloor} s_\ell(E, r) \leq \exp(y^{\varepsilon}/2) \qquad \text{for every finite set } S \text{ of primes.}$$

Here $s_\ell(E, r)$ is the length of the localization kernel for the Tate twist $(-r)$. Duality bounds these lengths at each fixed field and twist in terms of $K_{2r}(\mathcal{O}_E)$ [13, 14]. For $\psi = 1$, prime-counting, zero-density estimates and the explicit formula give the required three quadratic-twist bounds, their logarithmic-derivative bounds, and the large-argument estimate for all primitive host associates [4, 5, 6, 7]. The windows at scale x give the bound at $\lfloor x^{400} \rfloor$ and hence at the realized cutoff t for the finite-place defect in division.

The **lifting of integral finite cohomology classes with separated Hodge jumps** occurs at inert places in Section 4. Let $K/\mathbf{Q}_\ell$ be unramified of degree n, $\ell \geq 11$, and let $T_1 \twoheadrightarrow T_2$ be a surjection of crystalline lattices over a coefficient ring $\mathcal{O}$ whose fraction field splits K. Write r for the rank of T_2. For each embedding of K, its decreasing-filtration jumps lie in a central interval and one signed interval. Large intervals can occur singly or in pairs in Frobenius order, with at least j small labels between successive groups. The large intervals have the prescribed comparable size, separated from the central interval as in Section 4. Write H for the common bound on the absolute endpoints of all the intervals and g for the sum of the three small width parameters. For the full inverse-image Bloch–Kato finite condition the bound is

$$\frac{\operatorname{length}_{\mathcal{O}} \operatorname{coker}[H^1_f(K, T_1) \longrightarrow H^1_f(K, T_2)]}{\operatorname{length}_{\mathcal{O}}(\mathcal{O}/\ell\mathcal{O})} \ll (n+1)r\left(\frac{H}{(\ell-1)^{j+1}} + (g+1)(j+1) + \log_\ell \max(1, H)\right).$$

The cyclic spacing suppresses the contribution from the largest jumps even when their average width exceeds the Fontaine–Laffaille range of Zhu's Theorem B. An integral description of H^1_f for an individual

lattice with one-sided weights is available from Čoupek, Gazaki and Marmora [67, 68]. In the proof of the lifting bound, the intervening small jumps control the loss in integral Frobenius coordinates. At two large jumps of opposite signs, compatibility of the two lifts must be proved in both cyclic orders. The local lifts then give one exponent annihilating the cokernel, independently of the number of pairs. This supplies the corresponding inert-place term in the Hecke lattice comparison.

The **finite condition for torsion extensions with prescribed crystalline lifts** concerns a fixed lattice. Let $K/\mathbf{Q}_\ell$ be unramified, $\ell \geq 5$, and let the coefficient DVR $\mathcal{O}$, with uniformizer λ, contain the integers of every conjugate of $K(\mu_\ell)$. Fix a crystalline lattice T with labelled filtration jumps U_i, and characters $\chi_c = \chi_1 \varepsilon_\ell^{1-c}$, $1 \leq c \leq t$, with labelled jumps $w_i + c - 1$. Section 4 gives the separation and residue conditions on these jumps at each embedding. For a finite $\mathcal{O}$-module X with trivial action, a class in $H^1(K, T(\chi_1^{-1}) \otimes X)$ gives extensions of $(\mathcal{O}/\lambda^a)(\chi_1)$ by $T \otimes X$ for $\lambda^a X = 0$. Test whether for every sufficiently large a the extension is a quotient of nested stable lattices in a finite direct sum of crystalline representations. Each summand must have the full labelled jumps and residual characteristic polynomials of $(T \otimes \operatorname{Frac} \mathcal{O}) \oplus \bigoplus_c \chi_c$. The classes passing the test form exactly the image of

$$H^1_f(K, T(\chi_1^{-1})) \otimes_{\mathcal{O}} X \longrightarrow H^1(K, T(\chi_1^{-1}) \otimes_{\mathcal{O}} X).$$

Here the integral finite group is the inverse image of the rational Bloch–Kato group. With a lattice in χ_c in place of T, the same test gives zero for $c \geq 2$ and the unramified subgroup for $c = 1$. In particular, an isomorphism of two crystalline lattices modulo λ^a identifies their finite images at that precision when the characters and labelled type agree. This conclusion reaches some profiles beyond the Fontaine–Laffaille range. For short crystalline representations over $\mathbf{Q}_\ell$, Atobe and coauthors lift a residue-field subquotient extension together with suitable endpoints [126, Lemma 7.6]. Here the endpoints are prescribed. The three local conclusions give zero local loss at the large primes satisfying the separation conditions.

At each of the two type primes a different **integral lifting bound** uses the Frobenius action on tame inertia characters. Let $T_1 \twoheadrightarrow T_2$ be a surjection of stable lattices in potentially crystalline representations of $G_{\mathbf{Q}_p}$, over the integers of a finite extension of $\mathbf{Q}_p$. Suppose their characteristic-zero inertial types have characters in a common Frobenius-stable set Σ of at most N tame characters, including 1, with $N \geq 2$. Their de Rham filtration jumps must lie in a common set of at most $N + 1$ integers in $[-\eta, H]$ containing zero. Write q for the maximum Frobenius orbit length of a ratio of two characters in Σ. If $1 \leq \eta < H$ and $p > 40(HN^2 + N^3)$, then

$$p^{7(Nq+\eta(N+q))} H^1_f(\mathbf{Q}_p, T_2) \subseteq \operatorname{im}\left[H^1_f(\mathbf{Q}_p, T_1) \longrightarrow H^1_f(\mathbf{Q}_p, T_2)\right].$$

Here the integral finite subgroup is the full inverse image of the rational Bloch–Kato subgroup. Once the prime satisfies the threshold, the exponent is independent of the positive endpoint H, the tame splitting degree and the multiplicities of the characters and jumps. The covariant Breuil–Kisin comparison gives left exactness and an unspecified power of (p, u) killing the final module cokernel [68, Proposition 3.4]. The quantitative argument first separates the characters on two ranges of Euler degrees and then corrects the monodromy of a lifted finite class. At g_0, g_1 the argument uses $q \leq \lfloor N^2/4 \rfloor$ for the graph lattices of the cohomology maps and their Tate duals. A length bound for the local Hecke defect also uses the rank of each target.

The **integral Selmer bound for tame anticyclotomic twists** concerns the finite list of cuspidal lattices T_s of odd rank $r \geq 100$ at one stage over the CM field E. Their conjugate Tate duals are $T_s \epsilon_\ell^t$, with t positive and odd. Here ϵ_ℓ is the cyclotomic character. The residual representations are absolutely irreducible. Use the prescribed Galois representations and finite local conditions of Section 4, including the Hodge gap, the conditions on the Tate duals and the common tame inertia element. Put $v_\ell = \min_{g \in G_E} v_\ell(\epsilon_\ell(g)^{rt} - 1)$, with $v_\ell(\ell) = 1$. Given $\delta_A, \delta_D > 0$ and the prescribed finite set to avoid, choose one finite ℓ-power anticyclotomic character ψ, trivial at every decomposition group over ℓ, with tame ramification at split conjugate pairs Q. For each a prime to ℓ, let $L_{a,s}$ be the normalized length of the

relaxed finite Selmer group of the twist by $\psi^{a\sigma_s}$, with the sign σ_s assigned to s. The desired conclusions are

$$\sum_s (L_{a,s} - \nu_\ell)_+ \le \delta_A \quad ((a,\ell)=1), \qquad t\sum_{v\in Q} v_\ell(\psi(\tau_v)-1) \le \delta_D,$$

where τ_v generates the inertia image. On the classical family $\nu_\ell = 0$ at odd primes and $\nu_2 = 1$. The distinction matters for lengths over ramified coefficient fields. With these hypotheses, Section 4 obtains a bounded-degree coefficient of valuation at most ν_ℓ in each integral Selmer determinant for one choice of split prime pairs. A simultaneous specialization then gives the common character and both bounds. To pass to the integral Hecke row, Section 4 compares the vacuum and the algebraic section at the same marked point of every cover. Their changes of frame cancel, leaving one phase and one harmonic scalar for the stage. Projecting the finite row vectors to the quotients by their complementary planes then identifies the specialized columns in the prescribed twisted lattice. Bounds for its congruence index also require the local Galois estimates. Earlier theorems control residual Selmer groups in suitable cyclic twists or give distributions in specified families of twistable modules [91, 92]. The bound here feeds the cuspidal contribution to the integral Hecke comparison.

Theta and Bessel periods

At a selected translate with the prescribed unitary towers, real tests, levels and transport data, the **theta–Bessel period comparison** uses one finite central element and row isometry sV for all words. For each w, one unitary U_w and one $a_w \neq 0$ give, for every twisting tuple $\boldsymbol{\mu}$, ramified column ϕ and projected row F_m,

$$b_{w,\boldsymbol{\mu}}(\phi, F_m) = \frac{d_{w,\boldsymbol{\mu}_0}}{d_{w,\boldsymbol{\mu}}}\, a_w \big\langle sV\iota_{\boldsymbol{\mu}} F_m, \mathcal{P}_w J_w U_w \phi \big\rangle.$$

Some positive power of the ratio $d_{w,\boldsymbol{\mu}_0}/d_{w,\boldsymbol{\mu}}$ is rational and supported on the twisting primes of $\boldsymbol{\mu}$. Seesaw compares matched kernels. The Bessel formula gives normalized squared periods of factorizable vectors in a fixed tempered representation with discrete Hermitian weak base change [21, 26]. The simultaneous comparison permits descent to a common row dual for the rank argument. The support property preserves valuations at the finite places where the tuple avoids the residue characteristic.

At a split finite place, the **transfer around a square** takes an admissible test at depth $e \ge 1$ with the prescribed ramified unitary determinant ratio of conductor exponent at most e. Partial Fourier transforms and a primitive-frame integral give the opposite test with integral support, the same depth and the required level and character actions. In these Fourier models the large zeta section equals the complementary small section with the product determinant character when $\mathrm{vol}(\mathrm{GL}_a(\mathcal{O})) = 1$. Related partial transforms construct split doubling sections [20]. The matching lets the calculation of integral row tests change paths in the grid of lifts.

Detection at the prescribed principal level is proved in Section 2 for the prescribed grids. Fix the precision and depth constant $C_F \ge 6510$. Beyond one threshold, take an admissible grid over E/F, its bottom Hermitian space Y, and $V = Y \perp \langle 1\rangle$. For every nonzero cuspidal multiplicity map t on $U(Y)$ of the prescribed real type at principal level $(g_0 g_1)$, and every nonzero vector w in its smooth real tensor space, the assertion gives a cusp form a on $U(V)$ of the prescribed discrete-series type, with a K_∞-finite testing vector, such that

$$\int_{[U(Y)]} a(\iota h)\, t(w)(h)\, dh \neq 0.$$

The finite level for a is the principal subgroup for the one ideal $\mathfrak{m}$ specified by the grid before t, w vary. Both level characters are trivial. The gain over qualitative period nonvanishing is uniform control of that level as the field, rank and real weights vary. For a Hermitian Arthur parameter on the product group, the global GGP theorem relates a nonzero central value to a nonzero period on some relevant unitary pair. Harris's restriction theorem starts with a cuspidal representation of the larger group with

sufficiently regular discrete series at infinity. Here the smaller row and the larger principal level are both prescribed [23, 25]. The proof bounds a translated real matrix coefficient relative to its reproducing scalar. At a point where the row is large, the prescribed depth makes the nonidentity terms in the Poincaré period smaller than the stabilizer term. Compact averaging then supplies a K_∞-finite testing vector. The spanning argument expresses the core tests in terms of descending lifts. The ascent argument transports the nonzero Bessel pairing along the prescribed core tower.

Relative traces and automorphic multiplicities

The **mixed Gram estimate** uses the levels, real types and rank conditions of the tied family. At each translate the normalized Poincaré operators $L_w = c_0(w)^{-1/2}\mathcal{P}_w$ act on full columns of dimension d_w and share a Petersson target. With $b = \exp(x^\epsilon)$ and reference dimension $D_\star$ the estimate is

$$\|L_w^* L_v\|_{\mathrm{HS}}^2 \le (1 + b^{-6})\, d_w d_v / D_\star \qquad (w \ne v).$$

One threshold works for all selected translates and distinct words. Relative trace unfolding and density formulas give the orbit reindexing and primitive counts [40, 48]. The signed scalar main term and nongeneric corank-one remainder, with absolute bounds on other errors, control the mixed terms of the second moment for the period span.

The **single-word estimate** uses the full columns from the mixed Gram estimate and the balance $d_w\Gamma^2 \asymp D_\star$ of (55), where $\Gamma = \exp(C_\Gamma m^5 f \log b)$, $f = [F : \mathbb{Q}]$ and k, m are the rank parameters. Beyond one threshold, every selected translate and every word satisfy

$$\frac{\|L_w^* L_w - I\|_{\mathrm{HS}}^2}{d_w} \le 2\exp(-kf\log k) + 2C_0/\Gamma.$$

Schur orthogonality gives the identity term. Bounds on the degenerate and nondegenerate orbital remainders control the mean-square deviation of the squared singular values from 1. This gives the diagonal trace moments and, with the ramified column count, a nonzero bottom reference pairing.

The **spectral multiplicity estimate** fixes a discrete series at each real place for the unitary group G of common signature (p, q) and rank $r = p+q$, with $q^2 \le r$. Use a finite projector e of level $g \le b^A$ with the prescribed split-prime support. The cross-sign Harish–Chandra gaps must be at least $Cr^{19}f^5\log(b\ \mathrm{rd}(E))$. The projector's rank D on the multiplicity space satisfies $|D - D_\star| \le \exp(-rf\log r/16)D_\star$. The main term $D_\star$ is the product of the formal degrees, $\mathrm{vol}([G])$, and $\mathfrak{Re}(1)$. Section 3 compares this with multiplicity theorems for a fixed group and field. The uniform error as r, f vary lets the moments measure the span against the actual target dimension.

Integral lattices and divisibility

For the **short window for perturbed determinant jets**, let v be a nonarchimedean valuation on a characteristic-zero field with $v(p) = 1$, $p \ge 5$ prime. Take $q = p^h$ for odd $h \ge 1$, $\rho = (p^2-1)^{-1}$, $e = 1 + (q-p)\rho$ and $1 \le r \le n$. Let X, Y be r-by-r matrices of independent variables and H a formal power series in those variables. With ordinary derivatives set

$$\Lambda_y = \det(\partial_X)^{qy}\det(\partial_Y)^y \exp\left(-p^{-e}\sum_{i,j} X_{ij}^q Y_{ij} + H(X,Y)\right)\Bigg|_{X=Y=0}.$$

Every nonzero term of H has positive Y-degree and coefficient valuation plus weight at least $1/(p-1)$. The weights per variable are $\rho+3/(8q)$ for X and $\rho+1/2-1/(8q)$ for Y. Every interval of $\lceil 64r^2(h+3+3\log_p(r+1))\rceil$ consecutive nonnegative integers contains an order y with $v(\Lambda_y) \le ry((q+1)p\rho - 1) + n(\log_p n + 1)$. Cayley's identity gives the factorial normalization [58]. The bound for the resulting Newton coefficients

makes the short interval accessible by finite differences. In the local normal form this supplies a test of controlled valuation for an odd Frobenius block.

The **realization of Frobenius gains by integral division** uses the prescribed Artin data at sufficiently large designated scales, with sampling degrees, stage weights and metrized period lattices at each tied translate. Local division realizes valuation gains on the source lattices enlarged by fractional ideals. The inertia-weighted sum of ramps (95) through $\lceil 2xb \rceil$ is at most the weighted valuation gain net of restriction plus $C_{\mathrm{loss}}|T|mbx^{e_L\varepsilon}$, with the same bound for the weighted index allowance in (97). Height and slope inequalities and the CM period normalization give the comparison in the prescribed metrics [51, 69, 70]. Combining those inputs with the realized gains gives the centered inequality $D(u) \le E_x$ of (167) at every tied translate u. The Fourier projection uses this inequality with the same sampling data and cutoff.

There is also a **weightwise integral refinement** of the comparison between Weyl and dual Weyl lattices. Let $\lambda = (a^t, \kappa)$ be a dominant polynomial weight, and let Δ, ∇ use the same primitive highest vector. At grade i, take the monic least common multiple $D_i(T)$ of the polynomials $\prod_{(b,c)\in\tau/\kappa}(T-(c-b))$ over the allowed skew shapes with nonzero value at a. Each shape has i boxes, at most t in any column, and τ has length at most that of the second block. If p^{e_L} clears the Levi weight lattices up to a given depth measured by a decreasing linear functional, then at every weight v of grade i up to that depth one has

$$p^{e_L+v_p(D_i(a))}\nabla(\lambda)_v \subseteq \Delta(\lambda)_v.$$

Taking the polynomial least common multiple before specializing retains distinct factors whose values coincide modulo p. The formula includes singular specializations and primes dividing $i!$. Nakahama computed the rational scalar on each isotypic component [114, Theorem 1.2]. The candidate is the resulting denominator on the integral weight lattice, with the original cutoff on its Levi coefficients. For the bound summed over primes in this paper, ordinary rational projectors with a common factorial denominator already suffice. The sharper statement has a separate application wherever a bound at one prime and one weight is required. Its historical priority is open.

A local estimate in the height calculation concerns **divisibility of a determinant of invariant derivatives**. Let p be odd and G a commutative formal group over the valuation ring of an algebraically closed valued field of characteristic zero, with $v(p) = 1$. For an even integer $h > 0$, suppose μ_{p^h-1} acts integrally with logarithmic types $\zeta \mapsto \zeta^{p^s}$, $s \in \mathbf{Z}/h\mathbf{Z}$. Assume that each odd type has tangent dimension at most j. Choose any $i \ge 1$ integral tangent vectors in a single even type whose matrix has a unit i-minor. Write $D_a^{(b)}$ for the invariant derivative along vector a in copy b of G^i. Let δ_c be the infimum, over all integral power series f on G^i, of $v((\det(D_a^{(b)})^c f)(0))$. For every integer $c \ge 1$ some integer $0 \le r \le \min(i, j(3+\log_p c)/2)$ satisfies

$$\delta_c \;\ge\; (i-r)c\,\frac{p}{p^2-1} - i^2(1+\log_p c).$$

The coefficient $p/(p^2-1)$ already occurs for a one-dimensional Lubin–Tate group for the unramified quadratic extension, with its primitive dual character [116, §V][117, Appendix, Theorem (c)]. The candidate is the bound on how many directions can lose this gain in an arbitrary saturated frame of mixed type. The proof chooses successive bases of p-typical curves and controls the dimensions lost at each step. Keeping the separate radius at every step gives a further consequence in this reconstruction,

$$\delta_c \;\ge\; ic\,\frac{p}{p^2-1} - jc\,\frac{p^3}{(p^2-1)^2} - i^2\log_p(c+1).$$

For fixed p, i, j the loss proportional to jc in this formula has a coefficient independent of c. On the arithmetic family a bounded number of steps already suffices at the small inert primes. This estimate concerns every integral germ. The separate bound for perturbed determinant jets selects a term in the jet.

The **Frobenius pairing on two prescribed tangent planes** occurs before the estimate for perturbed determinant jets is applied. Take a prime $p \geq 5$ in either alternating case of the local index calculation and its polarized Raynaud extension with unramified CM action. On each of two paired Frobenius cycles the tangent ranks alternate between $n + j$ and j'. In the self-conjugate case use two copies of the factor as in Section 4. Fix an odd $0 < i < D$, where D is the full cycle length. Prescribe integral rank-n direct summands in the large types a and $\overline{a + i}$ on opposite sides. Put $q = p^i$ and $E_i = 1 + (q - p)/(p^2 - 1)$. There are saturated subplanes of equal rank $s \geq \max(0, n - ij' - j)$ inside these summands such that the mixed coefficient of the polarized biextension, in formal logarithm coordinates on them, satisfies

$$p^{E_i}[Gh]_{q,1} \bmod \mathfrak{m} = -\sum_{b=1}^{s} x_b^q y_b$$

in suitable integral bases. Here h is its rigidified additive gauge, G is the unit scalar of the level, and $\mathfrak{m}$ is the maximal ideal in the valuation ring of the completed algebraic closure at p, with $v(p) = 1$. The coefficient thus gives a perfect pairing after Frobenius twist. The proof bounds the other coefficients on the same planes, which allows the determinant estimate to be applied. Classical display theory constructs the biextension from the polarization [131, §4.3]. The candidate concerns the uniform rank and the separation of this coefficient when both tangent planes are prescribed.

Heights, section lattices and Fourier–Jacobi coefficients

The **finite-degree bound for metrized adjoint section lattices** gives, for a birational moving system of degree d_Z on a flag component of dimension s, a degree lower bound with leading term $\binom{\ell-1}{s} d_Z\, \ell h_Z$, with normalized Chow height h_Z and the chosen degree $\ell \geq 3s + 5$. It uses controlled multiplication, short integral frames and rank and metric comparisons. The Chow determinant method has precedents in David–Philippon and Chen. The chosen degree may lie below a projective Hilbert-point threshold. After descent from the cover and truncation, the component bounds control the covolumes of cusp-section lattices in the indefinite tower. These covolumes enter the division estimates. The reference slopes on the definite column spaces require a separate calculation.

The **height of the metrized full flag at finite cutoff** enters the mean slope on a nonterminal unitary cover. Let the abelian CM field E have degree $2f$, $f \geq 2$, and take the CM type induced from the specified imaginary quadratic field. Each label has signature (p, q), with $q \geq 1$ and $N = p + q \geq 3q + 1$. Use the prescribed good level M' with auxiliary parameter $A = 1$ and at most $4fN^2$ exceptional prime divisors of $2|\operatorname{disc} E|M'$. The compatible integral models are chosen before the weight and cutoff. For a strictly decreasing weight u, require that, at each label, every one of the last q entries be at distance at least $g > 0$ from each other entry. Take an integer B with $B + 1 > 4fN^2(1 + \|u\|_\infty/g)$. The integral Hodge–Tate lattices are used at the exceptional primes and at all primes up to B.

Put $D = f\binom{N}{2}$. Let $H_B(u)$ be the sum of the normalized arithmetic self-intersections of the metrized weight line on the components, divided by $D + 1$ times the sum of their geometric degrees. With the prescribed quadratic character χ_{d_0} set $\Lambda_B = \log(B \operatorname{rd} E) + |\Re(L'/L)(1, \chi_{d_0})|$. Section 4 obtains the absolute lower bound

$$H_B(u) \;\geq\; -C\left(\Lambda_B \sum_{\tau,i} |u_{\tau i}| + q \sum_{\tau} (\max_i u_{\tau i} - \min_i u_{\tau i})\right).$$

This controls a height used to estimate the mean of the component Chow heights in the finite-degree adjoint bound. Geometric proportionality gives the ordinary flag degrees [100]. Arithmetic volume formulas for a Hodge line in signature $(n - 1, 1)$ over an imaginary quadratic field, or with only one noncompact real place, concern fixed integral models [101, 102]. The equivariant arithmetic Chern character formulas for abelian schemes in [103] apply away from the coefficient discriminant. Here each newly selected good prime changes the integral flag. Section 4 computes that change on the common finite model, including its toroidal boundary, and bounds the resulting prime tail. On each fixed cover the mixed intersections

converge to those for the prescribed Hodge metric. The proof tests weights on the walls of the dominant chamber with positive last entry at each real place. Polynomial continuation then includes the signed weights that occur in the estimate. Two determinant heights account for the terms that survive the wall tests. This gives the displayed component mean at the original cutoff. Applying it to a lattice of sections also uses the finite-place trace and the norms of sections on that cover.

The **detection of a prescribed Fourier–Jacobi coefficient at a CM boundary** concerns either principal-level cover at a nonterminal stage. Fix an isotropic line of the marked unitary lattice and a frequency ν with $0 < \sum_\tau \nu_\tau \le c_1$, where the sum is over the f embeddings in the CM type. If this coefficient of a complex cusp section is nonzero, there is a one-parameter degeneration with CM abelian part and torsion extension data on which a Laurent coefficient, followed by a monomial coefficient in the tangent frames, detects it. The resulting linear functional kills every section whose coefficient at this line and frequency vanishes. Let E be the CM field of degree $2f$, r the unitary rank, M the level, $d_E = |\operatorname{disc} E|$, N_A, N_c the two lattice comparison integers, and W the bound for the analytic weight spread. If $X \ge 1$ bounds f, r and the logarithms of M, d_E, N_A, N_c, c_1, W, the torsion denominator n_Q and the factor Λ_Q bounding the lengths of the CM tangent generators satisfy

$$\max(\log n_Q, \log \Lambda_Q) \le 200X^2.$$

The same functional has bounded denominators at every finite place and bounded operator norm at every complex embedding, in its metrized value line. It therefore gives a lower bound for the first minimum of the integral quotient by sections whose low coefficients all vanish. The proof begins at a point where the prescribed coefficient is large. Positivity of the rank-two Fourier indices bounds the reduced lattice minimum there, allowing the sign plane and the extension data to be varied together. Pasten uses Heegner points to bound norms of integral differentials on quaternionic Shimura curves [128, §§13–14]. Richter and Westerholt-Raum use specialization of Jacobi forms at torsion points for Siegel modular forms modulo a prime [129, proof of Proposition 2.4]. Here the uniform bounds on a test for the specified boundary coefficient enter both the rational rank and the integral height of the quotient.

1 CM fields with character and Selmer bounds

Fix a primitive Dirichlet character ψ of conductor M and a real exponent parameter $0 < \varepsilon < 1/52$. The zeta case takes $\psi = 1$ of conductor 1. For every nonprincipal Dirichlet character χ, $L(s, \chi)$ is holomorphic at $s = 1$ and $L(1, \chi) \ne 0$ [73, §25.15(i) and (25.15.9)]. For primitive nonprincipal χ define

$$G_\chi(y) = \sum_{n \ge 1} \Lambda(n)\chi(n)(y/n - 1)_+ + y\frac{L'(1, \chi)}{L(1, \chi)} \qquad (y > 0). \tag{2}$$

For $\chi = 1$, use the zeta sum G_1 of (1).

At each real scale x, the local division estimates require bounds on Selmer groups over a CM field E. The Fourier estimate requires prime-sum bounds for characters of $\mathrm{Gal}(E/\mathbf{Q})$. The two bounds will hold for the same choice of E at each point of suitable intervals.

1.1 The auxiliary primes and character estimates

At an auxiliary scale Z, put $w = Z^{\varepsilon/5}$. The primes q, d are distinct and coprime to M, with

$$q \equiv 3 \pmod 8, \qquad d \equiv 7 \pmod{16}, \qquad q \mid d + 1, \qquad p \mid d + 1 \quad (3 \le p \le w,\ p \text{ prime}),$$
$$\tfrac{1}{2}w < \log q \le w, \qquad \tfrac{1}{2}w(\log Z)^2 < \log d \le w(\log Z)^2.$$

Write $\eta_q = (\frac{\cdot}{q})$ and $\eta_d = (\frac{\cdot}{d})$ for the quadratic Dirichlet characters. For $h \ge 0$ let

$$E(h, q, d) = \mathbf{Q}(\zeta_{2^{h+2}})^+\, \mathbf{Q}(\psi)\, \mathbf{Q}(\sqrt{-q}, \sqrt{-d}),$$

where $\mathbf{Q}(\psi)$ is the subfield of $\mathbf{Q}(\zeta_M)$ fixed by the kernel of ψ. Since $E(h,q,d)$ is abelian and contains $\mathbf{Q}(\sqrt{-q})$, it is CM.

For an odd or 4-divisible modulus N, the cyclotomic Artin map identifies $(\mathbf{Z}/N\mathbf{Z})^\times$ with $\mathrm{Gal}(\mathbf{Q}(\zeta_N)/\mathbf{Q})$ by $a \mapsto (\zeta_N \mapsto \zeta_N^a)$. At any prime $\ell \nmid N$ the residue class of ℓ gives the arithmetic Frobenius. For a prime $p \equiv 3 \bmod 4$, the unique quadratic subfield of $\mathbf{Q}(\zeta_p)$ is $\mathbf{Q}(\sqrt{-p})$, and the restriction of the automorphism indexed by a mod p to it acts on $\sqrt{-p}$ by the sign $(\frac{a}{p})$. [12, Ch. V, Examples 3.1–3.2 and 3.11] Applying this to q, d gives the quadratic factors of conductors q, d. These are coprime to $\mathrm{lcm}(2^{h+2}, M)$ and to each other, so the Chinese remainder theorem gives $\Gamma = \mathrm{Gal}(E(h,q,d)/\mathbf{Q}) \simeq (\mathbf{Z}/2\mathbf{Z})^2 \times H$, with H the group of the compositum of the first two factors. For a coprime to $2Mqd$, let $h_{\mathrm{nat}}(a) \in H$ be the restriction of the cyclotomic automorphism $\zeta \mapsto \zeta^a$ to the compositum. Put $h_- = h_{\mathrm{nat}}(-1)$, so $2h_- = 0$. Write $\eta_d(a) = (-1)^{e_d(a)}$ and $\eta_q(a) = (-1)^{e_q(a)}$. The release's Artin map orders the signs as q, d and shifts the last component to $h_{\mathrm{nat}}(a) + e_q(a)h_-$. The Fourier coordinates interchange the two sign positions, giving

$$\mathrm{Art}(a) = \big(e_d(a), e_q(a), h_{\mathrm{nat}}(a) + e_q(a)h_-\big). \tag{3}$$

Both sign bits of -1 are 1, so complex conjugation has coordinates $c_\infty = (1,1,0)$. For $\psi = 1$ the compositum is totally real and $h_- = 0$.

For odd primes $\ell \nmid d$, splitting in $\mathbf{Q}(\sqrt{-d})$ is equivalent to $-d$ being a square modulo ℓ. For $\ell = 2$ it is equivalent to $-d \equiv 1 \pmod 8$, since $-d$ is an odd quadratic discriminant. Thus 2 and the specified odd primes dividing $d+1$ split. These include q and, once $w \geq 4M$, every odd prime dividing M.

The classical host and its Artin data. For $\psi = 1$, put $Q_S = 2^{h+2}$ and $Q = Q_S qd$. The field $E = E(h,q,d)$ is the fixed field in $\mathbf{Q}(\zeta_Q)$ of the units a satisfying $a \equiv \pm 1 \bmod Q_S$ and $\eta_q(a) = \eta_d(a) = 1$. Indeed, these are precisely the automorphisms fixing each of its three factors. This also identifies any field of the prescribed type in the Selmer assertion with E. In an ambient cyclotomic field of a multiple of Q, its defining subgroup is the inverse image of the same subgroup modulo Q. The group $H = (\mathbf{Z}/Q_S\mathbf{Z})^\times/\{\pm 1\}$ has order 2^h and is cyclic. The class of 5 generates it, since $v_2(5^{2^j} - 1) = j + 2$ by repeated squaring and no power of 5 is -1 mod 4. Here $H = H_0$. The fixed field $F = E^{\langle c_\infty \rangle}$ is totally real and $[F : \mathbf{Q}] = 2|H|$.

For $p \mid Q$, the p-power cyclotomic factor is totally ramified locally, by the Eisenstein polynomial of $\zeta_{p^a} - 1$ when $p^a \| Q$. The roots of unity of order prime to p lie in the maximal unramified extension, by lifting the roots in the residue field. Arithmetic Frobenius raises these roots to their p-th powers. Thus the inertia subgroup is the image of the p-factor of the group of units in the Chinese-remainder decomposition. In the coordinates of (3) this gives

$$I_2 = \{(0,0,s) : s \in H\}, \qquad I_q = \{0, (0,1,0)\}, \qquad I_d = \{0, (1,0,0)\}, \qquad I_p = \{0\} \quad (p \nmid Q). \tag{4}$$

None of these contains $c_\infty = (1,1,0)$, so E/F is unramified at all finite places. Put $c(e_d, e_q, s) = (-1)^{e_d}$ and $P = \ker c$. Both 2 and q split in $\mathbf{Q}(\sqrt{-d})$. Therefore every Frobenius lift at 2 or q, as well as its inertia subgroup, lies in P. For every lift F_d and every integer $j \geq 1$, the coset $jF_d + I_d$ has one element outside P, since c is nontrivial on I_d. At $p \nmid Q$ the Frobenius is $\mathrm{Art}(p)$. These assertions give the local inputs to (95) and its Fourier expansion. The character c induces the primitive character η_d.

For $h \geq 1$ there is just one nontrivial real character ν of H. On odd integers it is $(2/\cdot)$, since the class of 5 has value -1 and the map through units modulo 8 is even. Its primitive conductor is 8; for $h = 0$ there is no such ν. Consequently the real primitive host associates are precisely $\rho\eta_q^i\eta_d^j$, with $\rho = 1$ or ν when available and $i, j \in \{0,1\}$. Their conductors are $\mathrm{cond}(\rho)q^i d^j$. The same product of conductors holds if ρ is any primitive character induced by a character of H, since the three prime-power parts of Q are coprime.

For the two prime counts, use Thorner and Zaman's exceptional-zero convention. For $m \geq 2$ the product of Dirichlet L-functions modulo m has no zero in

$$\Re s \geq 1 - \frac{1}{13\log(m(|\Im s| + 3))}$$

apart from a possible simple real zero β_m of a unique real nonprincipal character χ_m mod m. If this zero exists, then $1-\beta_m \geq b_0 m^{-1/2}$ for an effective absolute $b_0 > 0$. [4, arXiv v2, preceding Theorem 1.1 and (1.2)]

For real $U \geq 8$ and $(a,m) = 1$, let $\lambda = 1$ if there is no exceptional zero. If there is one, let λ be the mean of $1 - \chi_m(a)t^{\beta_m - 1}$ on $(U/2, U]$, so that

$$\tfrac{1}{8}\min\{1, (1-\beta_m)\log U\} < \lambda < 2. \tag{5}$$

Put $\theta = 32/37$ if the exceptional zero exists and $7/12$ otherwise. If $\lambda U/(2\varphi(m)) \geq U^{\theta+1/100}$, the half-interval case of the relative prime number theorem is

$$\sum_{\substack{U/2<p\leq U \\ p\equiv a \bmod m}} \log p = \frac{\lambda U}{2\varphi(m)}\left(1 + O\left(\exp\left(-\frac{c\log U}{\log m + (\log U)^{2/5}(\log\log U)^{1/5}}\right)\right)\right).$$

The constant $c > 0$ and the implied constant are effective and uniform in U, m, a. [4, arXiv v2, Theorem 1.1 and (1.7)]

Use $(U/2, U] \subset (\sqrt{U}, U]$ for both candidate counts. For q take $m_0 = 8$, $a = 3$ and $U = e^w$. These candidates satisfy $q > w$ for large Z. With q fixed, the Chinese remainder theorem makes the congruences for d a reduced class modulo

$$m_0 = 16q \prod_{\substack{3\leq p\leq w \\ p \text{ prime}}} p, \qquad U = \exp\bigl(w(\log Z)^2\bigr).$$

For any real $v > 0$, Rosser and Schoenfeld give $\sum_{p\leq v}\log p < 1.01624v$ [72, Theorem 9, (3.32)]. At $v = w$ this gives $\log m_0 \leq 4\log U/(\log Z)^2$ for both candidate counts when Z is large. Both choices also have $\log U \geq Z^{\varepsilon/6}$. The bounds on β_{m_0} and λ give $\lambda \gg m_0^{-1/2}$. Thus the main term is $\gg U m_0^{-3/2} = U^{1-o(1)}$, meeting the condition $U^{\theta+1/100}$ since $\theta \leq 32/37$. The relative error is $o(1)$ uniformly in these moduli. Dividing by $\log U$ gives at least $U^{1-o(1)}$ candidates, uniformly in the first prime when choosing the second.

For each fixed $J \geq 1$ and all sufficiently large Z, we need one pair as above satisfying both assertions below. The classical proof for $M = 1$ follows. For fixed $M > 1$, Alpöge's theorem on simultaneous quadratic twists supplies such a pair.

1. For every character ξ modulo M and $(i,j) \in \{0,1\}^2 \setminus \{(0,0)\}$, let χ be the primitive character agreeing with $\xi\eta_q^i\eta_d^j$ on the positive integers coprime to $2Mqd$. For every real $y \in [Z, Z^J]$,

$$|G_\chi(y)| \leq y^{1/2+\varepsilon/2}.$$

2. For every real primitive character ρ of conductor dividing $8M$ and the same pairs (i,j), the primitive character agreeing with $\rho\eta_q^i\eta_d^j$ on those integers has no zero in

$$\Re s > 1 - 1/400, \qquad |\Im s| \leq 8M\exp\bigl(w + jw(\log Z)^2\bigr).$$

Choose q first and then test twists by η_d in the progression it determines. For any real $R \geq 3$, Thorner and Zaman's density theorem gives

$$\sum_{f\leq R}\sum_{\substack{\chi \bmod f \\ \chi \text{ primitive}}} \#\{\varrho : L(\varrho,\chi) = 0,\ \Re\varrho > 399/400,\ |\Im\varrho| \leq R\} \leq CR^{99/400}$$

with an absolute effective C. The count uses multiplicities and includes real exceptional zeros. [5, arXiv v4, (1.1)–(1.2) and Theorem 1.2] In item 2, the two stages use $R = 8MU$ for q and $R = 8Me^wU$ for d. A failing twist has conductor divisible by the candidate prime, with remaining factor at most $8M$ or $8Mq$,

respectively. That factor is smaller than every candidate at its stage. Thus distinct rejected primes give distinct primitive characters. Since $R = U^{1+o(1)}$, these exclusions remove $o(U^{9/10})$ candidates. For the general fixed target, the mean-square argument assumes every prime factor of the progression modulus is at most w or greater than Z^{400J}. The former holds for every factor of 8 and for the small primes in the second modulus. Its other factor q satisfies $q > e^{w/2} > Z^{400J}$ for sufficiently large Z. The release controls the maximum on each dyadic block by a second moment bound and uses it to choose one pair meeting both conditions.

For the characters in item 1, evaluation at Z also gives

$$\left|\frac{L'(1,\chi)}{L(1,\chi)}\right| \le 1 + \frac{1}{Z}\sum_{n\le Z}\Lambda(n)\lfloor Z/n\rfloor = 1 + \frac{\log(\lfloor Z\rfloor!)}{Z} \le 4\log y \qquad (y \ge Z).$$

The zeta case from classical prime estimates. Fix $J \ge 1$, put $\sigma = \frac{1}{2}+\varepsilon/8$ and $T_0 = Z^{2J}$. For $Q_0, T_0 \ge 2$ and $1/2 \le \sigma < 1$, Heath-Brown's real-character density bound is

$$\sum_{\substack{\chi \text{ real, primitive}\\ \operatorname{cond}(\chi)\le Q_0}} N(\sigma, T_0, \chi) \ll_\alpha (Q_0 T_0)^\alpha Q_0^{3(1-\sigma)/(2-\sigma)} T_0^{(3-2\sigma)/(2-\sigma)} \qquad (\alpha > 0),$$

where $N(\sigma, T_0, \chi)$ counts zeros, with multiplicity, in $\sigma \le \Re s \le 1$, $|\Im s| \le T_0$ [6, Corollary 1 and Theorem 3]. For the first prime, take $Q_0 = U = e^w$ and test η_q. For the second, take $U = \exp(w(\log Z)^2)$, $Q_0 = qU$ and test η_d and $\eta_q\eta_d$. The two single-prime characters are odd and their product is even. Their conductors are q, d, qd, respectively, so each test assigns distinct real primitive nonprincipal characters to distinct candidates at its stage. In each stage $Q_0 = U^{1+o(1)}$ and $T_0 = U^{o(1)}$. The conductor exponent is $1-\delta$ for $\delta = (2\sigma-1)/(2-\sigma) > 0$. Taking $0 < \alpha < \delta/4$ excludes $O(U^{1-\delta/2})$ candidates in each stage. Intersect with the tall-rectangle exclusions above at $M = 1$. They test $\rho\eta_q$ first, and $\rho\eta_d, \rho\eta_q\eta_d$ second, for all real primitive ρ of conductor dividing 8. Each failing character at a stage contains its candidate prime in its conductor. The bounds $O(U^{1-\delta/2})$ and $O(U^{99/400+o(1)})$ each exclude a proportion $o(1)$ of the $U^{1-o(1)}$ candidates already counted in the progression. We may choose q first and then d from the surviving primes, with the second-stage bounds uniform in q. This proves item 2 and leaves each of the three tested characters without a zero in the Heath-Brown rectangle.

For a primitive nonprincipal character χ of conductor f, let $D_\chi = L'(1,\chi)/L(1,\chi)$ and $A_\chi(y) = \sum_{n\le y}\Lambda(n)\chi(n)(y/n-1)$. The explicit formula over $\mathbf{Q}$ gives, for real $y \ge 2$,

$$A_\chi(y) + (y-1)D_\chi = \sum_\varrho \frac{y^\varrho - 1}{\varrho(1-\varrho)} + O(\log(2y)).$$

The sum is over nontrivial zeros with multiplicity. The respective archimedean terms for even and odd χ are $(y-1)F_1(y)/2$ and $(y-1)F_3(y)/2$, both $O(\log(2y))$ with an absolute constant [7, §2.1 and Theorem 1, (2.2.1)]. By the functional equation, the three selected real characters have $1-\sigma < \Re\varrho < \sigma$ for $|\Im\varrho| \le T_0$ [73, (25.15.5)]. The number $N_\chi(t)$ of zeros up to height $t \ge 1$ has main term $(t/\pi)\log(ft/(2\pi e))$ and error $O(\log(f(t+2)))$ for primitive χ of conductor $f > 1$ [8, Corollary 1.2]. Thus the number of zeros in $n \le |\Im\varrho| < n+1$ is $O(\log(f(n+3)))$. Use $N_\chi(n+1) - N_\chi(n-1)$ for $n \ge 2$, and $N_\chi(2)$ for $n = 0, 1$. For $|\Im\varrho| \le T_0$ the reciprocal denominator is $O_\varepsilon((1+|\Im\varrho|^2)^{-1})$. For $|\Im\varrho| > T_0$ use $|y^\varrho - 1| \le 2y$ and the reciprocal bound $|\Im\varrho|^{-2}$. Summing over the unit-height strips gives an absolutely convergent zero sum and the bound

$$|A_\chi(y) + (y-1)D_\chi| \ll_\varepsilon y^\sigma \log(3f) + \frac{y}{T_0}\log(3fT_0) + \log(2y).$$

The conductor and height satisfy $\log(3fT_0) \ll_J Z^{\varepsilon/5}(\log Z)^2$. For every $y \in [Z, Z^J]$ this gives

$$y^\sigma \log(3f) \ll_J y^{1/2+13\varepsilon/40}(\log y)^2, \qquad \frac{y}{T_0}\log(3fT_0) \ll_J Z^{-J+\varepsilon/5}(\log Z)^2.$$

At $y = Z$, use $|A_\chi(Z)| \le \log(\lfloor Z \rfloor!)$ and divide by $Z - 1$ to get $|D_\chi| \le 4 \log Z$ for large Z. Adding D_χ gives the centering in (2). Since $13/40 < 1/2$, this proves $|G_\chi(y)| \le y^{1/2+\varepsilon/2}$ for all real $y \in [Z, Z^J]$ with the same pair, proving item 1 for $M = 1$. For general M, the twisted characters in item 1 can be complex and the release's prime-pair argument applies to those characters.

Mean squares in the two progressions. Let $P = \{p : \sqrt{U} < p \le U,\ p \equiv a \bmod m_0\}$ be the prime set at either stage above. For $L = Z^A$, $A \ge J$ fixed, $|\kappa(n)| \le 1$ independent of p, and $y \in [Z, Z^J]$, set

$$S_L(y, p; \kappa) = \sum_{n \le L} \Lambda(n) \min(1, y/n) \kappa(n) \Big(\frac{n}{p}\Big).$$

Then $\sum_{p \in P} |S_L(y, p; \kappa)|^2 \ll |P| y \log y$ with an absolute constant for sufficiently large Z. The unrestricted prime-modulus inequality of Montgomery and Vaughan uses $U/\log(2U/L^2)$ times the energy of coefficients grouped by squarefree part [9, Lemma 9, (7)]. For the relative factor $|P|$, use the full-interval prime theorem modulo $m_0 s$ at $t \in [\sqrt{U}, U]$, for odd squarefree $1 < s \le L^2$ coprime to m_0. Here $\log(m_0 L^2) \ll \log U/(\log Z)^2$, so $t \ge (m_0 s)^{12}$ and the relative error is $O(\exp(-c(\log Z)^2))$ for some $c > 0$. Summing over the reduced classes above a with weights $(\frac{\cdot}{s})$ cancels the principal term. The secondary term survives only if the exceptional character restricts to $(\frac{\cdot}{s})$ on $(\mathbf{Z}/s\mathbf{Z})^\times$. The relative error is bounded using the sum of the nonnegative main terms over those classes, including a possible bias from a character induced from m_0 [4, arXiv v2, Corollary 1.4 and Remark 1.5]. In the matching case the secondary term sums to zero without the weights. Its signed contribution is negligible if $1 - \beta \ge c/\log(m_0 L^2)$, for fixed $c > 0$. For sufficiently small c, Landau–Page gives at most one primitive real character of conductor at most $m_0 L^2$ with a closer real zero [10]. Its conductor away from m_0 must be exactly s. Partial summation therefore gives $\sum_{p \in P} (\frac{p}{s}) = O_B(|P| Z^{-B})$ for every fixed $B > 0$, with at most one exceptional $s > 1$.

Chebyshev's bound and partial summation make the contribution to S_L from proper prime powers $O(\sqrt{y})$. Primes dividing m_0 contribute $O(\log m_0)$. For the remaining primes $r \le L$, the weights $b_r = (\log r) \min(1, y/r)$ satisfy $\sum_r b_r^2 \ll y \log y$, $\max_r b_r \le \log y$ and $\sum_r b_r \ll y(1 + \log L)$. Quadratic reciprocity on the fixed class modulo 4 reduces the distinct-prime correlations to the last character average with s the product of the two primes. At most one unordered pair is exceptional. Its weight is $O((\log y)^2)$, and the other correlations contribute $O_B(|P| Z^{-B} y^2 (1 + \log L)^2)$ to the unnormalized moment. This gives the stated bound by taking $B > 2J + 1$.

For fixed M and any $\tau > 0$, the moment also gives $|G_\chi(y)| \le y^{1/2+\tau}$ on $[Z, Z^J]$ for a proportion $1 - o(1)$ at each stage. Test the primitive associates of $\xi \eta_q$ at the first stage and of $\xi \eta_q^i \eta_d$, $i = 0, 1$, at the second, for all $\xi \bmod 8M$. Apply the tall-rectangle exclusions to those associates. For a surviving χ of conductor f, take $L = Z^{3200J}$ and $\delta_0 = 1/400$. We have $(\log f)^{10/\delta_0} \le L \le f^2$, and the zero-free rectangle contains $\Re s > 1 - \delta_0$, $|\Im s| \le L^{\delta_0}$. The $k = 1$ logarithmic-derivative approximation gives $-D_\chi = \sum_{n \le L} \Lambda(n) \chi(n)/n + O(L^{-\delta_0/4})$ [11, Proposition 3.2]. Thus $G_\chi(y) = -S_L(y, p; \kappa) + O(Z^{-J})$ with κ the fixed factor before the new quadratic twist. Almost everywhere in y, $y\, \partial_y S_L(y, p; \kappa) = S_L(y, p; \kappa \mathbf{1}_{n > y})$. On a block $[a_0, b_0] \subseteq [Z, Z^J]$, $b_0 \le 2a_0$, Cauchy–Schwarz bounds the squared maximum by twice the squared value at a_0 plus $2a_0$ times the integral of the squared derivative. The summed block maxima are $O(|P| a_0 \log(2a_0))$. Markov's inequality and a dyadic sum give a discarded fraction $O(Z^{-2\tau_0} \log(2Z))$ for $\tau_0 = \min(\tau, 1/4)$, in addition to the rectangle exclusions. The finite union over twists is uniform in q at the second stage. Evaluation at Z with exponent $1/2 + \tau_0$ gives $|D_\chi| \le 4 \log Z$.

1.2 The Selmer bound on the same scales

For a prime ℓ let S_ℓ be the places of E over ℓ, and let G_{E,S_ℓ} be the Galois group of the maximal extension unramified outside S_ℓ. Write D_v for the decomposition subgroup of a chosen prolongation of $v \mid \ell$ to an algebraic closure of E. Restriction to D_v uses its map to G_{E,S_ℓ}. For integers $r \geq 1$, put

$$s_\ell(E,r) = \operatorname{length}_{\mathbf{Z}_\ell} \ker\left(H^1(G_{E,S_\ell}, (\mathbf{Q}_\ell/\mathbf{Z}_\ell)(-r)) \longrightarrow \prod_{v\mid\ell} H^1(D_v, (\mathbf{Q}_\ell/\mathbf{Z}_\ell)(-r)) \right).$$

Here H^1 is continuous Galois cohomology and the inverse r-th power of the cyclotomic character acts on the coefficients. At finite places v, D_v identifies with the absolute Galois group of E_v.

Let F be a totally imaginary number field and let S be a finite set of finite places containing every place above ℓ. For a finite ℓ-primary $G_{F,S}$-module A, write $\mathrm{Sha}^i_S(F,A)$ for the kernel of localization from $H^i(G_{F,S},A)$ to $\prod_{v\in S} H^i(F_v,A)$. Poitou–Tate duality gives a perfect pairing of finite groups

$$\mathrm{Sha}^1_S(F,A) \times \mathrm{Sha}^2_S\big(F, \mathrm{Hom}(A, \mu_{\ell^\infty})\big) \longrightarrow \mathbf{Q}/\mathbf{Z}.$$

For this finite S, the global cohomology groups and the local targets with coefficients in A are finite. The duality theorem includes the infinite places, whose local cohomology in degrees 1 and 2 vanishes for such F. [13, Theorem I.4.10(a) and Corollary I.4.15] Apply this with $F = E$, $S = S_\ell$ and $A = (\mathbf{Z}/\ell^k\mathbf{Z})(-r)$, whose Tate dual is $(\mathbf{Z}/\ell^k\mathbf{Z})(r+1)$. Filtered direct limits commute with the localization kernels, and finiteness at each level makes the first derived inverse limits vanish. Passing to the direct and inverse limits identifies the kernel defining $s_\ell(E,r)$ with the Pontryagin dual of the ℓ-primary étale wild kernel

$$\ker\left(H^2(G_{E,S_\ell}, \mathbf{Z}_\ell(r+1)) \longrightarrow \prod_{v\mid\ell} H^2(E_v, \mathbf{Z}_\ell(r+1)) \right).$$

For a totally imaginary number field F, every prime ℓ and integer $r \geq 1$, the K-theory comparison identifies

$$H^2_{\text{ét}}(\mathcal{O}_F[1/\ell], \mathbf{Z}_\ell(r+1)) \simeq K_{2r}(\mathcal{O}_F)[\ell^\infty],$$

where $K_{2r}(\mathcal{O}_F)$ is finite [14, Theorems VI.8.2 and VI.8.4]. At finite coefficient levels, the étale cohomology of $\mathcal{O}_E[1/\ell]$ with locally constant ℓ-primary coefficients agrees with cohomology of its étale fundamental group G_{E,S_ℓ} because ℓ is invertible there [13, Proposition II.2.9]. It follows in the inverse limit that the wild kernel is contained in the ℓ-primary part of $K_{2r}(\mathcal{O}_E)$. This includes the case $\ell = 2$ since E is CM. For fixed E, t, the weighted sum over all primes is thus at most $\sum_{1\leq r<t} \log|K_{2r}(\mathcal{O}_E)|$. The required quantitative bound is

$$\sum_{\ell\in S} (\log \ell) \sum_{1\leq r<t} s_\ell(E,r) \leq B \qquad \text{for every finite set } S \text{ of primes.} \tag{6}$$

This inequality is uniform over the choices of algebraic closure and prolongations.

Fix a nonnegative integer $h(y)$ for each real $y \geq 1$ with $y^{2+8\varepsilon} \leq 2^{h(y)} \leq 2y^{2+8\varepsilon}$. The release derives the Selmer window assertion from local–global compatibility and base change. In the application, the needed compatibility is established from theta lifts. For the fixed ψ, ε, h it gives $J \geq 1$ and X_0. For every $X \geq X_0$ and every pair q, d with the size and congruence conditions at $Z = X^{1/J}$, there is $W \geq Z$ with $W^2 \leq X$. For every real $y \in [W, W^2]$, the field $E(h(y), q, d)$ satisfies (6) at $t = \lfloor y^{401} \rfloor$ with $B = \exp(y^\varepsilon/2)$. The proof uses deformation-theoretic counts and controls exceptional components along the dyadic tower.

For $\psi = 1$ fix, for example, $h(y) = \lceil (2+8\varepsilon)\log_2 y \rceil$ on $y \geq 1$. Apply the Selmer window assertion to this whole profile to obtain J, X_0. Take $X \geq X_0$ so large that $Z = X^{1/J}$ exceeds the prime-selection threshold for this fixed J, and choose the pair at Z. The assertion supplies $[W, W^2] \subseteq [Z, X]$ for this very pair. On $[W, W^{3/2}]$, use the field $E(h(x), q, d)$ at each real x. Since $Z \leq W \leq x \leq W^2 \leq X = Z^J$, the prime-pair bounds apply at x, and $x^{1/J} \leq Z \leq x$. The Selmer bound on that field holds at $\lfloor x^{401} \rfloor$ with

$B = \sqrt{b}$, where $b = \exp(x^\varepsilon)$. For every $Y_0 \geq 2$ and any fixed lower threshold, X can be taken so large that Z, hence W, exceeds both. The pair is fixed on each such interval, and the layer $h(x)$ may vary.

For $M > 1$, use the J of the window assertion when choosing the pair q, d at $Z = X^{1/J}$. Take X large enough for the prime estimates at that scale. The release's prime-pair theorem supplies it for a general fixed target. The resulting Selmer interval lies in $[Z, X] = [Z, Z^J]$, on which the prime-pair bounds hold. The subgroup

$$H_0 = \mathrm{Gal}\big(\mathbf{Q}(\zeta_{2^{h+2}})^+\mathbf{Q}(\psi)/\mathbf{Q}(\psi)\big) \leq H$$

is cyclic, and its index $o = [H : H_0]$ is the order of ψ. The intersection of the real dyadic field with $\mathbf{Q}(\psi)$ stabilizes for large h, since $\mathbf{Q}(\psi)$ has finite degree. Consequently $n = |H| = c_\psi 2^h$ for a constant $c_\psi \geq 1$. Put $k_0 = c_\psi^{1/(2+8\varepsilon)}$ and $Y = k_0 W$. If $W \geq k_0^3$, then $Y^{3/2} \leq W^2$, and both x and $y = x/k_0$ lie in $[W, W^2]$ for every $x \in [Y, Y^{3/2}]$. In particular $x^{1/J} \leq Z \leq x$, as required for the host construction. For sufficiently large X all the layers $h(y)$ lie in the stabilized range. Choose the dyadic layer $h(y)$ at x. Then

$$x^{2+8\varepsilon} \leq n \leq 2x^{2+8\varepsilon}.$$

Taking $W \geq k_0^{400}$ gives $y^{401} \geq x^{400}$, and $y \leq x$. The Selmer bound therefore holds at cutoff $\lfloor x^{400} \rfloor$ with $B = \sqrt{b}$, where $b = \exp(x^\varepsilon)$. Taking X arbitrarily large gives arbitrarily distant intervals with q, d fixed throughout each interval. The dyadic layer can vary with x. For $\psi = 1$ we have $o = c_\psi = 1$ and $H_0 = H$.

1.3 Level primes, sampling sets and character estimates

Sizes and cutoff for the classical host. At every sufficiently large x in the classical intervals, the field above satisfies

$$n = |H| = 2^{h(x)}, \qquad Q_S = 4n, \qquad x^{2+8\varepsilon} \leq n \leq 2x^{2+8\varepsilon}, \qquad f = [F : \mathbf{Q}] = 2n. \tag{7}$$

The prime-pair bounds and $x^{1/J} \leq Z \leq x$ give

$$\tfrac{1}{2}x^{\varepsilon/(5J)} < \log q \leq x^{\varepsilon/5}, \qquad \log d \leq x^{\varepsilon/5}(\log x)^2.$$

Also $\log d > \log q$ once $(\log Z)^2 > 2$. Hence $d > q > x^A$ eventually for every fixed A. As $\log Q_S = O(\log x)$, these upper bounds give $\log(Q_S q d) \leq 2x^{\varepsilon/5}(\log x)^2$ for large x. In particular every fixed power of $Q_S q d$ is $\exp(o(x^\varepsilon))$. For the root discriminant, put $K = \mathbf{Q}(\zeta_{Q_S q d})$ and write $X^{Q_S q d} - 1 = \Phi_{Q_S q d}(X)R(X)$ with R integral. At $\zeta_{Q_S q d}$ the product of the derivative of $\Phi_{Q_S q d}$ and R equals $Q_S q d$ times a root of unity. The nonzero norm of $R(\zeta_{Q_S q d})$ is an integer, and the norm of the derivative is the discriminant of the power basis in absolute value. The field discriminant divides it, so $\mathrm{rd}(K) \leq Q_S q d$. The discriminant formula for the tower $K/E/\mathbf{Q}$ gives $\mathrm{rd}(E) \leq \mathrm{rd}(K)$. Thus $\log \mathrm{rd}(E) = o(x^\varepsilon)$.

Use the ranks at this field from the common-family construction in Section 6.1. The prescribed rank relations are

$$x^{1/2-\varepsilon/16} \leq m \leq 2x^{1/2-\varepsilon/16}, \qquad x^{9/2+10\varepsilon} \leq k \leq 2x^{9/2+10\varepsilon}, \qquad r = k + m - 1, \quad m + 3 \leq k.$$

The integer t used as twist cutoff is specified by

$$\begin{aligned} r + t &= 8fr + 4 + \lceil x^{377} \rceil A_x, \\ A_x &= 8\left\lfloor \frac{10r}{8}\left(f + \frac{f^2(k-1)^3 x^{2\varepsilon}}{m}\right)\right\rfloor. \end{aligned} \tag{8}$$

These are the rank prescriptions. For $x \geq 1$ and $0 < \varepsilon < 1/52$ we have $f \leq 4x^{2.16}$, $k \leq 2x^{4.70}$, $r \leq 4x^{4.70}$, $m \geq x^{0.49}$, and $x^{2\varepsilon} \leq x^{0.04}$. Consequently $f^2(k-1)^3 x^{2\varepsilon}/m \leq 128x^{17.97}$, so $A_x \leq 5288x^{22.67}$. Using $\lceil x^{377} \rceil \leq 2x^{377}$ in (8) gives $t \leq 10708x^{399.67}$. For $x \geq 2^{100}$ the last bound is at most x^{400}. Since t is an integer

with $t \le \lfloor x^{400} \rfloor \le \lfloor x^{401} \rfloor$, the same field satisfies (6) at t with $B = \sqrt{b}$, by nonnegativity of the summands. It follows also that $4tr + 1 < q, d$ at sufficiently large x.

At large x the rank interval for m gives $x^{1/2-\varepsilon} \le m \le x^{1/2+\varepsilon/2}$, and the lower bound for k gives $xm \le k$. The upper bound for k, together with $n \ge x^{2+8\varepsilon}$, gives $k \le n^2 x^{1/2-\varepsilon}$, since $2 \le x^{5\varepsilon}$ eventually. These are the numerical rank bounds in Sections 4 and 5. For the trace comparisons, $m \ge 4$ eventually and $m^4 \le 16x^{2-\varepsilon/4} \le f \le k \le r \le x^5$ for large x. Moreover,

$$m^5 f \log b \le 128 x^{9/2+(139/16)\varepsilon} \le 128 k x^{-(21/16)\varepsilon}.$$

Thus every fixed multiple of $m^5 f \log b$ is smaller than k beyond a threshold depending only on that multiple and ε.

Level and sampling conditions at the classical host. Use the level primes and the sets T, T_2 on the same selected records of Section 6.1. Write $K_d = \mathbf{Q}(\sqrt{-d})$ and $Q = Q_S q d$. The selected level conditions specify two distinct primes g_0, g_1 in $[(kb)^8, 2(kb)^8]$ with

$$g_i = \pi_i \overline{\pi_i}, \qquad \pi_i \in \mathcal{O}_{K_d}, \qquad \pi_i \equiv 1 \bmod Q\mathcal{O}_{K_d}.$$

They also specify distinct primes $p, p' \le Q^{C_8}$ for fixed C_8, with $p, p' \equiv 1 \bmod \operatorname{lcm}(15, Q)$. These conditions are the level data of the chosen host. The auxiliary twisting prime q_0 on the actual record is the least prime congruent to 1 mod $\ell_0 Q$, where ℓ_0 is the least prime not dividing $Q(p-1)(p'-1)$. The host has unit ideal $\mathfrak{c}_K$. A prime divisor of $Q(p-1)(p'-1)+1$ bounds ℓ_0 by $Q^{1+2C_8}+1$. If $m_0 = \ell_0 Q$, the full-interval prime theorem below at m_0^A, for sufficiently large fixed $A \ge 12$, has relative error of modulus less than $1/2$. The full-interval version of (5) gives a positive main term in the class 1 mod m_0. Hence $q_0 \le m_0^A \le Q^{A(3+2C_8)}$. The congruences distinguish q_0 from p, p', since ℓ_0 divides neither $p-1$ nor $p'-1$. Since $g_i \equiv 1 \bmod Q$ by the norm identity, all five primes split completely in the cyclotomic field of modulus Q and hence in E and F. The logarithms of p, p', q_0 are $O(x^{\varepsilon/5}(\log x)^2)$, which is eventually smaller than $x^{\varepsilon/3}$. Thus these primes are below both $\exp(x^{\varepsilon/3})$ and $(kb)^8$. Also $\log g_i \le 8x^{\varepsilon} + 8\log k + \log 2 \le 9x^{\varepsilon}$ for large x.

For each g let

$$B_r(g) = \{g\} \cup \{\ell \text{ prime} : \ell^{4v_\ell(g^a-1)} > g^a \text{ for some } 1 \le a \le r\}. \tag{9}$$

The disjointness $B_r(g_0) \cap B_r(g_1) = \emptyset$ is part of the selected level conditions. Thus for each residue characteristic ℓ, one of the two installed g_i differs from ℓ and satisfies $4v_\ell(g_i^a - 1)\log \ell \le a \log g_i$ for every $1 \le a \le r$.

The sampling conditions give $T \subset H$, $T_2 \subset H_0 = H$, and, writing $N_T = |T|$ and $N_2 = |T_2|$, the bounds

$$\tfrac{1}{2} n x^{-\varepsilon/8} \le N_T \le 2n x^{-\varepsilon/8}, \qquad N_2 \ge \frac{n}{L_w} \exp\left(-C_{T_2}\sqrt{\log n \, \log L_w}\right).$$

Here C_{T_2} is fixed and $L_w = 10\lceil 405 \log x \rceil + 2\lceil 30/\varepsilon \rceil + 22$. Since $\log n = O(\log x)$ and $\log L_w = O_\varepsilon(\log\log x)$, the logarithmic loss in the bound for N_2 is $o(\log x)$. For sufficiently large x these same sets therefore satisfy

$$0 < \tfrac{1}{2} n x^{-\varepsilon/8} \le N_T \le 2n x^{-\varepsilon/8}, \qquad n x^{-\varepsilon} \le N_2 \le n. \tag{10}$$

Every element of T lies in the sole coset of H/H_0, so its share in that coset is 1.

For $M > 1$, at each sufficiently large x in the intervals Alpöge chooses ranks, level primes, an Artin map and sampling sets $T \subset H$ and $T_2 \subset H_0$ for E. The choices give $|T_2| \ge x^{-\varepsilon} n$ and the distribution of T among the cosets of H_0 used in Section 5. The rank bound gives the actual twist cutoff $t \le \lfloor x^{400} \rfloor$, so (6) holds at t with bound $\sqrt{b}$. Write Q_S for the modulus of the compositum of the dyadic and target-character fields. Since the primitive character ψ factors through Q_S, its conductor M divides Q_S. Also $Q_S \mid \operatorname{lcm}(2^{h+2}, M)$, so every prime factor of Q_S divides $2M$. For the classical host, the only nontrivial characters left after averaging over $H_0 = H$ are $\eta_q, \eta_d, \eta_q\eta_d$. Each satisfies $|G_\chi(x)| \le x^{1/2+\varepsilon/2}$ and $|L'(1,\chi)/L(1,\chi)| \le 4\log x$

for the same q, d at x. In particular the centering for c satisfies $|\ell_c| \le 4\log x \le x^{\varepsilon}$ for large x. Together with the discriminant bound this gives $\log \mathrm{rd}(E) + |\ell_c| \le 2x^{\varepsilon}$. For $M > 1$ the primitive characters discarded after averaging induce the nontrivial quadratic twists above. For them we use $|G_\chi(x)| \le x^{1/2+\varepsilon}$, by weakening the prime-pair estimate, and $|L'(1,\chi)/L(1,\chi)| \le 4\log x$.

The modulus satisfies

$$\log(Q_S q d) \le 2x^{\varepsilon/5}(\log x)^2.$$

Pullback by the Artin map gives a Dirichlet character modulo $Q_S q d$ for each character of $\mathrm{Gal}(E/\mathbf{Q})$. A primitive real nonprincipal associate has order two, and its kernel in the cyclotomic Galois group fixes a quadratic field. For a quadratic field $\mathbf{Q}(\sqrt{D})$ with squarefree D, the primitive Dirichlet character has conductor $|D|$ when $D \equiv 1 \bmod 4$ and $4|D|$ otherwise. Its dyadic conductor therefore divides 8 [12, Ch. V, Example 3.11]. By the coprime product decomposition of the residue group, a real primitive associate is $\rho\eta_q^i\eta_d^j$, with real ρ of conductor dividing $8M$. If q or d divides its conductor, then $(i,j) \ne (0,0)$. Item 2 excludes real zeros $\beta > 1 - 1/400$ of that associate by taking the ordinate zero. The other real nonprincipal associates have conductor dividing $8M$. Nonvanishing at 1 for this fixed finite set gives a common real zero-free interval ending at 1. For the classical host the real associates with a q or d factor are exactly the three auxiliaries and, when $h \ge 1$, their three twists by ν of conductor 8. The only other real nonprincipal associate is ν, when available. Thus one $\delta > 0$ works for all real nonprincipal associates of all these hosts, including those nontrivial on H.

For integers $m \ge 2$, $(a,m) = 1$, and real $u \ge m^{12}$, put $\mathcal{E}_m(u,a) = \chi_m(a)u^{\beta_m}/\beta_m$ in the exceptional case above and $\mathcal{E}_m(u,a) = 0$ otherwise. The full-interval multiplier $1 - \mathcal{E}_m(u,a)/u$ lies between 0 and 2. Thorner and Zaman's full-interval prime number theorem gives

$$\sum_{\substack{p \le u \\ p \equiv a \bmod m}} \log p = \frac{u - \mathcal{E}_m(u,a)}{\varphi(m)} + O\left(\frac{u}{\varphi(m)}\,\mathcal{R}(u,m)\right),$$

where

$$\mathcal{R}(u,m) = \exp\left(-c\frac{\log u}{\log m}\right) + \exp\left(-c\frac{(\log u)^{3/5}}{(\log\log u)^{1/5}}\right).$$

The constant $c > 0$ and the implied constant are absolute and effective. The bound on the multiplier gives the displayed error uniformly in the reduced class a. [4, arXiv v2, Corollary 1.4, Remark 1.5 and (1.7)]

Put $Q = Q_S q d$ and $\Psi_\chi(u) = \sum_{n \le u} \Lambda(n)\chi(n)$. For $u \ge y \ge \exp(x^{\varepsilon/3})/3$ and large x, we have $u \ge Q^{12}$ by the bound on $\log Q$. Weight the progression estimate by the character induced by χ modulo Q and sum over the $\varphi(Q)$ reduced classes. Orthogonality gives main term $u\mathbf{1}_{\chi=1}$. Since χ_Q is real, the exceptional term sums to zero except when the induced character equals χ_Q; in that case the term is $-u^{\beta_Q}/\beta_Q$. In particular this term is absent for a principal or complex associate. For any character ξ modulo Q with primitive associate ξ^*, the two L-functions satisfy

$$L(s,\xi) = L(s,\xi^*)\prod_{p \mid Q}(1 - \xi^*(p)p^{-s}).$$

The product is nonzero for $\Re s > 0$ [73, (25.15.4)]. Thus in the surviving case β_Q is a real zero of $L(s,\chi)$ and $1 - \beta_Q \ge \delta > 0$ for a fixed δ. The error from summing the residue classes is $O(u\mathcal{R}(u,Q))$. For each $p \le \sqrt{u}$ the proper powers of p contribute at most $\log u$, and the omitted prime terms with $p \mid Q$ contribute at most $\log Q$. Since $\beta_Q > 1/2$ when it exists, the exceptional term divided by u has modulus at most $2u^{-\delta}$. Writing $v = \log u$ now gives

$$\Psi_\chi(u)/u = \mathbf{1}_{\chi=1} + O(e^{-cv/\log Q} + e^{-c\sqrt{v}})$$

with $c > 0$ uniform for the fixed target and ε.

Put $\delta_\chi = \mathbf{1}_{\chi=1}$ and $R_\chi(u) = \Psi_\chi(u) - \delta_\chi u$. For $y \geq 1$, integration of the step function gives $A_\chi(y) = y\int_1^y \Psi_\chi(u)\,du/u^2$, also for $\chi = 1$ with A_1 the uncentered zeta sum. The logarithmic derivative of the Euler product and partial summation give

$$-\frac{L'(s,\chi)}{L(s,\chi)} - \delta_\chi \frac{s}{s-1} = s\int_1^\infty R_\chi(u)u^{-s-1}\,du, \qquad \Re s > 1.$$

[73, (25.2.11), (25.15.2)] By the preceding bound the integral converges absolutely at $s = 1$. The limit as real $s \downarrow 1$ equals $C_\chi = \int_1^\infty R_\chi(u)\,du/u^2$. For $\chi \neq 1$, $C_\chi = -L'(1,\chi)/L(1,\chi)$. For $\chi = 1$, the Laurent expansion $\zeta(s) = (s-1)^{-1} + \gamma + O(s-1)$ gives $-\zeta'(s)/\zeta(s) = (s-1)^{-1} - \gamma + O(s-1)$, so $C_1 = -1 - \gamma$ [73, (25.2.4)]. It follows with these exact centerings that

$$G_\chi(y) = A_\chi(y) - \delta_\chi y\log y - yC_\chi = -y\int_y^\infty R_\chi(u)\frac{du}{u^2}.$$

The first error integrates to $O((\log Q)e^{-c\log y/\log Q})$, with $\log y/\log Q \gg x^{2\varepsilon/15}/(\log x)^2$. The second is $O(e^{-c'\sqrt{\log y}})$, with $\sqrt{\log y} \gg x^{\varepsilon/6}$. Thus

$$|G_\chi(y)| \leq y\exp(-c_6 x^{\varepsilon/8}) \qquad \big(y \geq \exp(x^{\varepsilon/3})/3\big) \tag{11}$$

for all primitive associates of the host characters, including $\chi = 1$. The constant $c_6 > 0$ and the threshold in x are uniform for the fixed target and ε. For the target of conductor one, the square-root estimates apply to η_q, η_d, $\eta_q\eta_d$, and the large-argument estimate also applies to characters that are nontrivial on H.

AI-proposed candidate (uniform higher fine Selmer bounds on windows). Fix the target ψ, $0 < \varepsilon < 1/52$, and an admissible dyadic profile h. The release's Selmer window assertion chooses $J \geq 1$ and X_0 before X and the primes. For every $X \geq X_0$ and every pair with the prescribed sizes and congruences at $Z = X^{1/J}$, it gives $[W, W^2] \subseteq [Z, X]$. At every real $y \in [W, W^2]$, the field $E(h(y), q, d)$ satisfies (6) with $t = \lfloor y^{401}\rfloor$ and $B = \exp(y^\varepsilon/2)$, uniformly over all finite prime sets and all choices of decomposition data.

Poitou–Tate duality and K-theory give finiteness at fixed field and twist. Lim's $\mu\ell^n + \lambda n + O(1)$ growth formula for the ℓ-part fixes ℓ, the twist and the $\mathbf{Z}_\ell$-extension [15, Proposition 5.1]. In the bound of Bayer-Fluckiger, Emery and Houriet for the logarithmic order of $K_n(\mathcal{O}_E)_{\mathrm{tors}}$ modulo the subgroup generated by elements of order at most $\max([E:\mathbf{Q}]+1, 2n+2)$, the factor $[E:\mathbf{Q}]^{293n^5[E:\mathbf{Q}]^5}$ alone exceeds the window budget at $n \asymp y^{400}$ and $[E:\mathbf{Q}] \asymp y^{2+8\varepsilon}$ [16, Theorem 1.1]. The abelian equivariant Tamagawa number formula relates leading terms to integral cohomology and a regulator in positive rank. Bounding the order of the finite cohomology from such an identity by estimating its factors separately would require lower bounds for the regulator on the same fields and twists [17, Corollary 1.2] [18, §5.4, (26)].

For $\psi = 1$, assume the Selmer window assertion and choose q, d by the classical prime estimates at $Z = X^{1/J}$. These primes work on $[W, W^{3/2}]$, with W arbitrarily large. At each point the Selmer bound at cutoff $\lfloor x^{400}\rfloor$ feeds the local division of Section 4. The quadratic-twist estimates and the large-argument estimate (11) on the same fields feed Section 5. Section 6.1 joins these inputs with the level and period comparisons.

2 Theta lifts and a common period lattice

Fix one translate of the family chosen in Section 6.1, with CM extension E/F. A word $w \in W$ is a sequence of characters with associated column types and levels. For each word we compare its top theta periods with Bessel periods at the bottom of the towers. All words at this translate must use the same bottom row coordinates to compute the dimension of the sum of the period images.

2.1 The towers and the top pairing

Take even k, m with $k \geq m+3$ and $m \geq 2$, and odd $t \geq r+5$, where $r = k+m-1$, as in the selected rank data. The spaces X_j are definite of rank $k + j(t+1)$. The spaces Y_j have rank $r + jt$, and negative index $\lfloor (m-j)/2 \rfloor$ at the real places. For $j < m$, the construction embeds X_j in Y_j. At the top Y_m is isometric to a hyperplane $X' \subset X_m$ perpendicular to a prescribed anisotropic line. At a real place the entry $-e_i$ of the character used at stage $i < j$ contributes the string $-e_i + (t-1)/2 - a$, $0 \leq a < t$, to the row Harish–Chandra parameter at j. The corresponding column string is $-e_i + t/2 - a$, $0 \leq a \leq t$. The row parameter also contains the signed entries from stage zero, with the labels prescribed at stage j. A common shift by the half-integral splitting parameter is added to the row entries when j is odd. These are the parameters of the two towers.

Write H_w^{full} for the full space at the bottom column level. For each designated residue-unit group, averaging its action gives the invariant projector. The ramified column space $\mathcal{M}_w$ is the intersection of the kernels of these projectors. For a tuple $\boldsymbol{\mu}$ of row characters, write $\Theta_{w,\boldsymbol{\mu}}^{[m]}$ for the composite of the m upward lifts on the X-tower. For the unitary pair on a Hermitian space A and a skew-Hermitian space B over E, fix the additive character used in its Weil representation. Kudla's construction splits the pullback of the local metaplectic extension over $U(A) \times U(B)$ using χ_B on $U(A)$ and χ_A on $U(B)$. The restrictions of χ_A, χ_B to the ground field's multiplicative group are $\eta_{E/F}^{\dim A}, \eta_{E/F}^{\dim B}$, with the local quadratic character in each local construction. With Hecke characters having these restrictions on $\mathbb{A}_F^\times$, the local splittings assemble to a splitting over the adelic product. A Hermitian space can be given the skew-Hermitian convention by scaling its form by a nonzero trace-zero element of E. [29, §§3.1.1 and 4.1.1] Each step uses the prescribed additive and splitting characters and finite and real tests. In the definite class model the step $j \to j+1$ is the sum

$$(T_{j,w,\boldsymbol{\mu}}\phi)(c') = \sum_{c} \frac{K_{j,w,\boldsymbol{\mu}}(c,c')\phi(c)}{|\operatorname{Stab}_j(c)|}.$$

Here $K_{j,w,\boldsymbol{\mu}}(c,c')$ is a map of coefficient spaces. It incorporates the determinant characters, the unfolded Schwartz test, and the determinant factors in the input and output identifications of the coefficient spaces. The determinant characters and Schwartz tests both depend on the tuple. The identification at stage zero realizes the same input $\phi \in \mathcal{M}_w$ in the column model for that tuple.

Write $[G] = G(F)\backslash G(\mathbb{A}_F)$. Let ι_m denote the top inclusion of unitary groups and res_m the coefficient restriction to X'. If f_m is the automorphic form representing the top row F_m, the top period in bilinear conventions is

$$b_{w,\boldsymbol{\mu}}(\phi, F_m) = m_Y(\boldsymbol{\mu}, m) \int_{[U(Y_m)]} (\mathrm{res}_m \Theta_{w,\boldsymbol{\mu}}^{[m]}\phi)(h) f_m(h)\, dh. \tag{12}$$

The product denotes contraction in the coefficient model, with the lift restricted along $h \mapsto \iota_m(h)$. For a Petersson pairing conjugate-linear in its first variable, the first input is the conjugate of the lifted column. The factor $m_Y(\boldsymbol{\mu}, j)$ is the product of the prescribed local parahoric indices at stage j. It equals 1 at all stages for the base tuple $\boldsymbol{\mu}_0$ and at $j = 0$ for every tuple.

In the opposite direction, pairing the downward theta integral against the row multiplicity space at j gives the row map from $j+1$ to j. This integral uses the factor $m_Y(\boldsymbol{\mu}, j+1)$. Composing the row maps gives $\iota_{\boldsymbol{\mu}}$ into the common bottom row space E_{row}. We use its restriction to the image $\mathcal{V}_{\boldsymbol{\mu}}^0$ of the chosen top row projector. The definitions of the row measures and pairings specify the same coefficient and measure conventions in both directions.

2.2 Changing the path of the local tests

For two dual pairs (G, H') and (H, G') in the same symplectic group, suppose $H \subset G$, $H' \subset G'$, and the restricted Weil kernels agree on $H \times H'$ via the chosen splittings. In Kudla's seesaw identity, for cusp

forms f on H and f' on H', with compatible quotient measures and an absolutely integrable product of the common kernel with $f(h)f'(h')$, one has

$$\int_{[H']} (\Theta_{H\to G'} f)|_{H'}(h')\, f'(h')\, dh' = \int_{[H]} f(h)\, (\Theta_{H'\to G} f')|_{H}(h)\, dh.$$

Here each Θ integrates the corresponding kernel against the input cusp form. The bilinear convention places the Petersson conjugation in the choice of dual inputs. Both sides integrate the same function on $[H]\times[H']$. [21, §1, Definition 1.3 and (1.14)] The two paths below use compatible restrictions of the Weil splittings.

The integral row tests use an auxiliary grid with rank $r + jt + i(r + 6)$ at (i, j). A rightward step increases this rank by $r + 6$, and a downward step decreases it by t. The definite column lifts defining the top pairing follow the X_j tower of rank $k + j(t + 1)$. The local grid realization uses the horizontal increment specified in the auxiliary rank data. At a fixed tuple the square with lower left corner (i, j) compares the routes

$$\begin{aligned} &(i, j+1) \longrightarrow (i, j) \longrightarrow (i+1, j) && \text{(down, then right)},\\ &(i, j+1) \longrightarrow (i+1, j+1) \longrightarrow (i+1, j) && \text{(right, then down)}. \end{aligned}$$

AI-proposed candidate (level-preserving transfer of the split finite tests). At a prescribed split finite place v, work over F_v with uniformizer ϖ and the chosen additive character of conductor $\mathcal{O}_v$. Take the prescribed depth e, the maximum of one and the coset exponent at v. The two determinant characters α, β of the sections have the prescribed ramified unitary ratio $\nu = \alpha/\beta$, with conductor exponent at most e. The vertex levels are open subgroups of the integral general linear groups. Take an admissible test on the old route, supported on integral matrices and invariant modulo ϖ^e there, satisfying the old level and character actions. Hold the variables outside the square fixed. Inside it concatenate the two test matrices in column blocks, with the prescribed transpose on one matrix. The numbers a, a' of rows on the new and old routes are the ranks at $(i + 1, j + 1)$ and (i, j). The number n of columns is the sum of the other two ranks, so $n = a + a'$.

The local transfer applies the finite Fourier sum in the first column block P, with the factor $q_v^{-ea'|P|}$ for residue cardinality q_v. It gives a function Ψ on a'-by-n matrices. The transfer then uses the primitive-frame transform below and the inverse Fourier sum in the complementary block. Let Λ' be the row lattice after the first Fourier sum, with $\varpi^{-e}\mathcal{O}_v$ on the transformed coordinates and $\mathcal{O}_v$ on the others. For matrices with ℓ rows write

$$I_\ell(\Phi, X) = \int_{\mathrm{GL}_\ell(F_v)} \Phi(gX)\, \nu(\det g)|\det g|^{n/2}\, dg.$$

The measure for $\ell = a'$ gives the inverse-transpose of the prescribed lower level volume one. For $\ell = a$ the subgroup $\mathrm{GL}_a(\mathcal{O}_v)$ has volume one. Write $X_{\mathrm{st}} = (I_a\ 0)$, $Y_{\mathrm{st}} = (0\ I_{a'})$ for the complementary coordinate frames. For $h \in \mathrm{GL}_n(F_v)$, the full row rank of $Y_{\mathrm{st}}h$ confines the support of $\Psi(gY_{\mathrm{st}}h)$ to a compact set of matrices g. Relative to additive matrix measure, the absolute determinant weight in $I_{a'}$ is proportional to $|\det g|^{(a-a')/2}$, which is bounded on that set since $a - a' = t + r + 6 > 0$. Thus $I_{a'}(\Psi, Y_{\mathrm{st}}h)$ converges absolutely. For $X = X_{\mathrm{st}}h^{-\mathsf{T}}$ primitive in $\varpi^{-e}(\Lambda')^\vee$, the transform sets $\phi_A(X) = \nu(\det h)|\det h|^{n/2} I_{a'}(\Psi, Y_{\mathrm{st}}h)$. It is zero off the primitive frames. Changing h with X fixed changes $Y_{\mathrm{st}}h$ by a left basis change. The inverse determinant factor from the integral cancels the change in the external weight $\nu(\det h)|\det h|^{n/2}$.

The resulting zeta sections satisfy, for every h,

$$\begin{aligned} &\beta(\det h)|\det h|^{-a/2} I_a(\phi_A, X_{\mathrm{st}}h^{-\mathsf{T}})\\ &\qquad = (\nu\beta)(\det h)|\det h|^{a'/2} I_{a'}(\Psi, Y_{\mathrm{st}}h). \end{aligned} \tag{13}$$

For the first integral one can sum over Cartan shells. For fixed h, the basis changes making $X_{\mathrm{st}}h^{-\mathsf{T}}$ a primitive frame form one coset of $\mathrm{GL}_a(\mathcal{O}_v)$. Its volume is one and it lies in a single shell. Replacing h by

$\operatorname{diag}(g^{-\mathsf{T}}, I)h$ shows that the determinant weights for g cancel on the coset. This gives the complementary-frame identity. To check depth, lift a change of a primitive frame modulo ϖ^e to a principal-congruence column change. The conductor bound makes its determinant character trivial. This calculation and finite Fourier inversion give the prescribed support and depth on the new route. Here depth e means constancy modulo ϖ^e times the prescribed matrix lattice. Use $\mathrm{GL}_a(\mathcal{O}_v)$ at the new intermediate vertex and the original levels at the other vertices. This is the prescribed level at the selected interior vertices. The comparison of the level actions preserves admissibility with the required character on the new path. Eischen–Harris–Li–Skinner use a partial Fourier transform in matrix blocks to construct split local doubling sections with specified support and congruence characters. [20, §4.3, (58) and Lemma 4.3.2] In the square the transform acts on the prescribed depth space of admissible tests, and the section identity relates the two complementary ranks in the indicated measures.

2.2.1 The adelic path comparison

Let $\mathcal{L}_\gamma(\Xi)$ be the iterated normalized theta integral on a path γ with test Ξ. The arithmetic calculation uses the following adelic matching on its prescribed tests. At each square, choose real Fock tests on the new route in the same local quotient as the old ones, and use the specified nonsplit tests and the split transfer above. The comparison required after integrating the kernels at the two intermediate vertices is a single nonzero scalar $c_{ij,\boldsymbol{\mu}}$ in the prescribed quotient measures. At each prefix of the chosen sequence of squares the identity on the same cusp forms, with the same exterior factors, is

$$\mathcal{L}_{\mathrm{right,down}}(\Xi') = c_{ij,\boldsymbol{\mu}}\mathcal{L}_{\mathrm{down,right}}(\Xi). \tag{14}$$

The scalar is independent of the cusp form and of the chosen tensor expansion of the test for that square. The adelic comparison follows by matching the real and finite models of the integrated kernels, applying the regularized theta identity and producing a nonzero leading integral for each square. Integrated kernel comparison · Nonvanishing · Application on a path. The split-place equality (13) gives the local transferred section. The real and inert normalized quotient maps, the placed nonzero test, and the regularized comparison are proved in the square calculation below. Its use in (14) is on the admissible placed path tests. The convergence argument below verifies the incoming cusp type and the replacement tests at every prefix.

Given (14), expand the tested Schwartz tensor into a finite sum of pure tensors and apply the square equality with the exterior arrows fixed in each summand. The two quotient integrations can be interchanged when the joined kernel times the cusp forms is absolutely integrable, as in the seesaw above. Swapping each downward arrow past each rightward arrow takes m^2 swaps, so induction gives the inverse factor

$$(c^{\square}_{\boldsymbol{\mu}})^{-1}, \qquad c^{\square}_{\boldsymbol{\mu}} = \prod_{0\le i,j<m} c_{ij,\boldsymbol{\mu}}.$$

If the original expression is divided by a character value $\gamma_{\boldsymbol{\mu}}$, the opposite path thus has denominator $\gamma_{\boldsymbol{\mu}} c^{\square}_{\boldsymbol{\mu}}$.

2.2.2 Convergence through the chosen squares

Fix an installed tuple, its permitted finite row test and the real polynomial tensor on the initial path. First consider the finite factors. On the boundary $i = 0$ they are the placed rectangular type tests, the tame parahoric tests at the tuple primes, and the lattice vacuum elsewhere among the selected split places. The last boundary edge uses the frame of X'. Each symbol is supported on the integral matrix lattice and constant modulo the residue prime. For a type test this follows because the condition on its triangular blocks and the tame determinant characters depend only on the reductions. For the parahoric and lattice tests the same holds by their support conditions. At the other boundary $j = 0$, the coset tests are indicators of an integral framed lattice coset at the prescribed depth. Their translations are integral because the

marked rational maps take the source lattice to the target lattice. The local path function, the product of these edge tests, is therefore supported on the integral matrices and constant modulo the installed depth on each of its edges.

In the finite matrix action at an interior vertex the absolute values of integral unit determinants are one. The extra parity character is unramified and is also one there. Along $i = 0$ the two consecutive edge laws contribute the intervening type character and its inverse. At $(0, 0)$ the endpoint character on the column side is trivial. Along $j = 0$ the vertex group preserves the indicated coset and the edge determinant characters are trivial on its congruence subgroup. Consequently the initial path is admissible at the prescribed levels. Support and depth, the horizontal and downward level laws and their product on the chosen path.

The ratio at a cell is the character $\mu_{j+1}\xi_H$ divided by the horizontal character indexed by $m - i$. At the deep auxiliary place the numerator is unramified and the denominator is ramified. At the other fixed split grid places the denominator is nontrivial on the first principal congruence subgroup, where the two numerator characters are trivial. At the extra tuple places the horizontal and the reference vertical characters are unramified, while ξ_H is ramified. The level data give the stated upper conductor bound in each case. The parity character in the two splitting characters is unramified. The ordinary additive character has conductor $\mathcal{O}_v$ at these places. This verifies the hypotheses of the split transfer for each cell at the same chosen split half and depth. The opposite middle vertex is new to a monotone path through the cell. It receives the integral level in the transfer, and the exterior characters and levels remain those already used for the path. Induction on the ordered squares now proves admissibility at every prefix. The three conductor cases · and the prescribed order and depths.

The same frames give the finite factors of the actual boundary lifts. Along the downward boundary, the shallow residue symbol is the specified local lift vector in its two framed polarizations. Outside the selected split halves it is the normalized vacuum. The local changes of position leave only values of the relative finite character. In the last term this uses the determinant of the fixed frame on X' and its conjugate half. Along the rightward boundary the factors are the prescribed coset vectors. Multiplying the inverses of the downward character factors gives the initial scalar $\gamma_{\boldsymbol{\mu}}$ in the representation of the path tensor. If the downward factor at step j is c_j times the section in the split chart, this scalar is $\prod_j c_j^{-1}$. The boundary polynomials use the very same real vectors as the row chain. Replacing a fundamental domain for a rational quotient has no effect since the integrand descends and the Haar normalization is fixed by the vertex level. The placed boundary vectors and scalar · and the initialized analytic path.

For the cuspidality test at any intermediate vertex, work at one real place and suppress that index. Use $R_{ij} = r + jt + i(r+6)$, let $s = r+6$, and put $\delta_{ij} = -1/2$ for odd $i+j$ and 0 otherwise. Before taking the dual, the signed parameter at (i, j) is the following union translated by δ_{ij}: the entries $v_2, \dots, v_k$ with positive sign, the entries $-b_\ell$ with positive sign when $\ell \le j+1$ or $\ell - j - 1$ is even and negative sign otherwise, the positive t-strings centered at $-e_1, \dots, -e_j$, and the positive s-strings centered at $c_{m-i+1}, \dots, c_m$. A string of length d centered at c has entries $c + (d-1)/2, \dots, c - (d-1)/2$. Here each c_ℓ is the prescribed integral midpoint between the successive upper nodes. Both t and s are odd. Thus R_{ij} is odd exactly when $i + j$ is even. The splitting exponent for the parameter of the smaller member, after the determinant-cover correction, is the unshifted center plus δ_{ij}. The exponent for the larger member is the same center plus the shift of that member, by the opposite rank parity.

For a rightward step take $d = s$ and center c_{m-i}. The defining upper gap exceeds $2N_{\mathrm{amb}}^3 + 2s + 2$, where N_{amb} bounds the grid ranks. The centers are in descending order, with rounding error at most one. If $i > 0$, the largest old upper string ends at $c_{m-i+1} + (s-1)/2$. The cloud and the lower strings are below the old upper strings and below the first upper center. Every old entry is therefore at distance at least $(s+1)/2$ below the new center. In the negated larger Hermitian form the old labels reverse and the new string has

negative sign, as in the real theta rule. For a downward step compare the pair in increasing rank, with $d = t$ and center $-e_{j+1}$. The two inequalities

$$e_\ell - b_\ell \;\geq\; (t+3)/2, \qquad b_{\ell+1} - e_\ell \;\geq\; (t+3)/2$$

show that the old entries below this center are exactly the $-b_\ell$ with $\ell > j+1$. All old entries are at distance at least $(t+3)/2$. Increasing j reverses exactly these labels and inserts the positive t-string. These assertions include the boundary cases $i = 0$, $j = 0$ and $j+1 = m$, where the descriptions use the transported X' parameter. Paul's discrete-series theta rule gives the specified partners with these signed lists, including their duals in the bilinear theta integral. [95, Theorem 2.7(b) and Proposition 3.4(ii)] The signed lists, the s-step and the t-step.

Here is the lower-tower check in both directions. With the splitting center removed, the smaller parameter is integral with no zero, so for the odd rank difference d the integer k_π of Ichino's criterion is -1. The signature of the larger partner has $l = 0$ and $t_{\mathrm{Ichino}} = (d-1)/2$, after exchanging the two signatures when the inserted string is negative. The two nearby signed sets in the nonvanishing criterion are empty by the preceding distances. A lower member in that Witt tower has negative l and the lift is zero. In the downward direction the larger input has the complete centered d-string and no longer one. The signs of its component characters alternate along the string, because its signature labels there are constant. Hence $k_\pi = d$. The small partner has rank $n - d$, where $n = R_{i,j+1}$ is the input rank. This is the first permissible rank. Indeed the criterion for a discrete-series input with $k_\pi = d$ requires $l \geq 0$ in the balanced signature case and $l \geq d$ otherwise. The only exception at $l = -1$ requires an even-multiplicity weight, absent in a regular discrete series. Thus no member of smaller rank in this tower can occur. [39, §6.1, Theorem 6.1 and Remark 6.2]

For any of the finite tensors just described, the theta kernel and its derivatives have polynomial height growth on the quotient in the integration variable. The Heisenberg–Poisson bound used below for the triangle tests gives this estimate before integration. Rapid decay of a smooth cuspidal input therefore gives absolute integrability, locally uniform in the target and polynomially bounded in its height. Real smooth Howe duality on the specified pair gives the stated real type in the image. The tower constant-term formula for a theta lift of a cusp form expresses the proper-parabolic terms through lower theta lifts in that target tower. The real vanishing just checked kills each such term. Thus the next function is again a smooth cusp form in the specified type and satisfies the same rapid decay hypothesis for the next integration. [88, Theorems A and 2.23–2.24] [85, Proposition 3.1] The type on every edge and the path induction.

At an active split place, an admissible path function is on the finite product of its residue matrix spaces. Expanding in products of coset indicators and applying the finite vertex averages expresses it as a finite sum of allowed pure Schwartz tensors. Exterior factors can be held fixed while partitioning the two edges of a cell. Real polynomial tensors admit a finite expansion for the same reason of finite dimensionality in each degree, and outside the active finite places the vacua are fixed. Apply the preceding cusp argument to each summand. The absolute bounds for the two integrated kernels of the square, multiplied by the rapidly decaying incoming cusp form, justify Fubini with the fixed exterior variables. This compares the two paths in the same vertex measures. Together with the finite admissibility induction it proves convergence and the stated equality after every flip, for each installed tuple and every allowed test on the original cusp input space. Finite decomposition with exterior factors and convergence for the resulting sums.

2.3 The principal-level row test

Work in the prescribed principal-level model before imposing the auxiliary bottom level character. Let $j_{\boldsymbol{\mu}}^{\mathrm{top}} x$ be the automorphic realization of a rational top coordinate x in the good isotypic part, and let $\Theta_{\boldsymbol{\mu}}^{\downarrow}$ be the downward row chain in this model, with its fixed real vectors. The functionals ψ_a on the bottom

multiplicity are indexed by the prescribed grid data Φ_a and column tests Ψ_a. Each value is the Bessel integral of the column test against the further grid lift of the bottom row, in the prescribed measures. Write $p_{\boldsymbol{\mu},a}(x) = \psi_a(\Theta^{\downarrow}_{\boldsymbol{\mu}} j^{\mathrm{top}}_{\boldsymbol{\mu}} x)$.

2.3.1 Changing the bottom level

The level-change group of the row chain is $\prod_{v\in P_g} \mathrm{GL}_r(k_v)$, where P_g is the set of chosen split components at the two type primes and k_v their residue fields. In the frames at those places the flag blocks have sizes $r, t, \dots, t$, with one extra t-block on the larger vertex. To represent a residue element g_0, use a block diagonal lift $\mathrm{diag}(\widetilde{g}_0, I, \dots, I)$ in both vertices and the identity away from P_g. The split unitary group identifies with the general linear group at its chosen component. An integral lift there gives the inverse dual action at its paired component, so it preserves the self-dual framed lattice. The lifts exist because a matrix with invertible reduction has unit determinant. Both have flag character one. Two lifts with the same reduction differ by a principal-level element, whose character is canceled in the normalized right-translation action.

At $v \in P_g$ the row-step test in the Schrödinger model is supported on integral matrices whose reduction is $\binom{A}{0}$, with A block upper triangular, invertible and with initial r-block I_r. For the step $j+1 \to j$ its coefficient there is

$$\prod_{h=1}^{j} \chi_{h,v}(\det A_h)\, \chi_{j+1,v}(\det A)^{-1},$$

where A_h are its t-blocks. The same formula applies at the last step with X' as the larger vertex. For lifts g, h on the small and large vertex use the matrix substitution $Y \mapsto h^{-1} Y g$. The preferred lifts preserve this support, all the block determinants and the determinant of A. The split Weil factor evaluates the two splitting characters on $(\det g)^{-1}$ and $\det h$, with unit absolute determinants. The two characters agree with $\chi_{j+1,v}^{-1}$ on the units, including in the last step where the extra unramified character is trivial on units. The two values cancel. This proves simultaneous invariance of each step test under the preferred lifts. There is no action at the nonsplit places, at the twist places, or on the real Fock factor for these lifts. The split row-step test · The last-step model.

In the theta integral against a right translate on the large vertex, change variables by that translate on its rational quotient. Simultaneous invariance of the Schwartz factor changes the translated kernel to a right translate on the small vertex. This proves the intertwining identity for each step on the convergent automorphic integrals, and hence for their composition. Apply the group-algebra idempotents on both sides by linearity. The bottom good projector fixes the lift of a good top coordinate. The subgroup character projector, when imposed at both ends, restricts the same downward map to that level. If the uncut image is the full permitted bottom space, its image after the cut is the corresponding bottom projector image: apply the top projector to any preimage of a vector in this image.

Denote by $Q_{\boldsymbol{\mu},a}(x)$ the corresponding Bessel integral on the opposite path after the swaps. The remaining matching required for this same Φ_a, Ψ_a, x is the normalized expansion

$$\begin{aligned} Q_{\boldsymbol{\mu},a}(x) &= (\gamma'_{\boldsymbol{\mu}})^{-1} s_{\mathrm{AA}} \sum_{\beta} \left(\prod_{v} d_{\beta v} \right) A_{\boldsymbol{\mu},a,\beta}(x), \\ A_{\boldsymbol{\mu},a,\beta}(x) &= \kappa J_{\boldsymbol{\mu},a,\beta}(x). \end{aligned} \tag{15}$$

Here β runs over the supported local coset families of Φ_a, and $d_{\beta v}$ is the rational measure weight of its cell at v. The factor s_{AA} is the product of the prescribed vacuum factors at the other finite places. The term $A_{\boldsymbol{\mu},a,\beta}(x)$ is the sum of definite periods of the two ascending lifts with the real Fock tensors from unfolding and the finite tests from an expansion of that cell. The arithmetic term $J_{\boldsymbol{\mu},a,\beta}(x)$ is the corresponding class sum of the normalized jet pairings. Its summands at each class have the reciprocal of the product of the row and column stabilizer orders as weights. The measures in $A_{\boldsymbol{\mu},a,\beta}$ are the measures for these particular term

levels. The second identity requires a single κ for the real Fock tensor used in the unfolding, independent of $\boldsymbol{\mu}, a, \beta, x$ and the choice of finite expansion. The two character values $\gamma_{\boldsymbol{\mu}}, \gamma'_{\boldsymbol{\mu}}$ belong to the subgroup generated by the local values of the prescribed finite-order character $\xi_{H,\boldsymbol{\mu}}$.

2.3.2 Finite row orbits and their measures

Only the vertices $I = \{1, \dots, m-2\}$ of the last downward column are integrated over a row-group quotient. Write R_j for the alternating Bessel and Fourier–Jacobi row groups, G_j for their ambient groups, and K_j for the full finite lattice stabilizers. Keep the fixed initial inverse staircase and the neutral endpoints. At an interior vertex the base finite point is the finite part of the fixed flag tilt δ_j. Its translates used here are

$$b_{\beta,j} = \delta_{j,f} \prod_{v \in S} s_{v,j} g_{\beta,v,j}.$$

At a tested split place v in S, identify $G_j(F_v)$ with $\mathrm{GL}_{n_j+c_j}(F_v)$, where n_j is the definite column rank and $c_j = m-1-j$. The sign $s_{v,j}$ is diagonal and integral. It acts trivially on the middle block and changes only the designated trailing flag entries at odd j. The matrices $g_{\beta,v,j}$ represent the local double cosets for $R_j(F_v)$ and $K_{j,v} = \mathrm{GL}_{n_j+c_j}(\mathcal{O}_v)$. The local row group is the product of the middle $\mathrm{GL}_{n_j}(F_v)$ and the unipotent group with entries in positions of strictly increasing flag weight. The middle block has weight $c_j + 1$; the successive one-dimensional flag weights are $2, 4, \dots, 2c_j$. Conjugation by the tilt and by the signs preserves this unipotent group and acts trivially on the middle Levi. Thus the labels are the double cosets in the stated frames, with no further identification. To see surjectivity and injectivity of this description on the ordinary-place support, lift each local left factor on one split half by its Hermitian inverse transpose on the other half, and take the identity elsewhere. Lift the right integral factors in the same way. Self-duality gives membership in K_j on both halves. The restricted-product factorization then gives the required global left and right factors. Conversely each global equality of double cosets gives its local equality by projection. The real component, which is taken separately in the quotient integral, does not affect a label. The fixed tilt, base points and representatives · The split row subgroup · Local representatives in the finite row quotient.

Here the ordinary-place support means the single orbit of the base lattice at every finite place outside S. In the actual unfolded product, nonvanishing forces each two-arrow contraction to be nonzero. Its two flat factors are then both nonzero. The ordinary-place support calculation can be made one arrow at a time. With the base moves already applied, use the self-dual lattice vacuum at a place outside S. The Bessel-to-Fourier–Jacobi condition says that a nonzero flat factor with source in its base orbit forces the next lattice point into the Fourier–Jacobi base orbit. The Fourier–Jacobi-to-Bessel condition forces the following point into the Bessel base orbit. Starting from the fixed initial point, induction gives the ordinary-place support at every interior vertex. At the fixed auxiliary split places within S, only the reference staircase is allowed. There the local input is the one-staircase support of the transferred tensor at the exact test depth. It is applied to the *flattened* tensor, rather than to each entry of the original path test. These two support conditions refer to the placed flat contractions described below. Induction with the two ordinary-place support tests · The fixed-place staircase at the actual depth.

For the cell of $b = b_{\beta,j}$, its mass in the quotient $R_j(\mathbb{A}_{F,f}) \backslash G_j(\mathbb{A}_{F,f})$ is

$$M_{j,\beta} = \frac{\mathrm{vol}(K_j)}{\mathrm{vol}_{R_j}(R_j(\mathbb{A}_{F,f}) \cap b K_j b^{-1})}.$$

Indeed its parametrization by the right compact group is the compact quotient of K_j by $b^{-1} R_j b \cap K_j$. Haar integration over the stabilizer gives the displayed denominator. Write N_j for the unipotent radical of R_j and $J_{\beta,j}$ for the projection of the stabilizer to the middle Levi. The latter is the term level of the definite column at this vertex. Locally, this is the condition that a middle-block element have a lift over the given flag whose action preserves $g_{\beta,v,j} \mathcal{O}_v^{n_j+c_j}$ in the stated frame. The tilt and sign commute with the middle

block and normalize its flag radical. At ordinary places the orthogonal self-dual decomposition gives the full middle lattice stabilizer. This identifies the projection with the product of the prescribed term levels. Projection to the definite column level.

The Haar measure on R_j is the product of the unipotent and middle-block measures. In the split model the two sets of rectangular coordinates on which the Levi has nontrivial determinant are dual, so the middle action has modulus one. Fubini on the compact stabilizer gives a denominator proportional to $\mathrm{vol}_{N_j}(N_j \cap bK_jb^{-1})\,\mathrm{vol}(J_{\beta,j})$, with the proportionality independent of the cell. When a Haar measure normalized to give the reference column level $J_{0,j}$ volume one is replaced by the measure giving $J_{\beta,j}$ volume one, the conversion factor multiplying $M_{j,\beta}$ is $\mathrm{vol}(J_{\beta,j})/\mathrm{vol}(J_{0,j})$. There is no column-level change at the endpoints and no change on the other definite path. Thus the variable Levi volume cancels exactly. For the remaining factor at v set

$$\begin{aligned} U^0_{v,j} &= N_j(F_v) \cap K_{j,v}, \\ U^\beta_{v,j} &= N_j(F_v) \cap g_{\beta,v,j} K_{j,v} g^{-1}_{\beta,v,j}, \\ d_{\beta v,j} &= \frac{[U^0_{v,j} \,:\, U^0_{v,j} \cap U^\beta_{v,j}]}{[U^\beta_{v,j} \,:\, U^0_{v,j} \cap U^\beta_{v,j}]}. \end{aligned}$$

The fixed tilt multiplies both volumes by the same determinant. The integral diagonal sign preserves $U^0_{v,j}$ and its volume. Hence the remaining factor is the fixed base measure factor times $\prod_v d_{\beta v,j}$. This gives $d_{\beta v} = \prod_{j\in I} d_{\beta v,j}$ with the inverse-volume orientation in (15). The fixed base factors and the ordinary-place vacuum values can be included in the common normalization of the real Fock expansion. The measure conversion defining each term is performed before evaluating its definite periods. Relative indices in the actual frames · Product of term-level measure conversions.

2.3.3 The phase of the finite expansion

Use determinant half-roots d_j of the unitary base points $\delta_{j,f}$ and the initial inverse staircase. A half-root of a determinant u satisfies $d_j/\overline{d}_j = u$. Its choice drops out of $\xi_{H,\boldsymbol{\mu}}(d_j)$, since the character is trivial on the fixed-field adeles. The root of a real translate may be taken with finite part one. In a paired downward contraction at $j, j+1$, changing the twist by ν changes the two operators by

$$\nu(d_j)^{-1}\nu(d_{j+1})\,\nu(d_{j+1})^{-1}\nu(d_{j+2}) = \nu(d_j)^{-1}\nu(d_{j+2}).$$

The intermediate determinant cancels with this sign under the linear flat contraction. The product on even j telescopes, and the final root can be taken to be one. In the reference-vacuum comparison the factor used is

$$Z_{\mathrm{base}}(\boldsymbol{\mu}) = \prod_{j \text{ even},\, j<m} \frac{\xi_{H,\boldsymbol{\mu}}(d_{j+2})}{\xi_{H,\boldsymbol{\mu}}(d_j)}.$$

Paired determinant twist · Normalization at the base points.

On the opposite path at a selected v, for a matrix assignment g write

$$X_{v,\boldsymbol{\mu}}(g) = \prod_{b \text{ downward}} \frac{\xi_{H,\boldsymbol{\mu},v}(\det g_{b,\mathrm{source}})}{\xi_{H,\boldsymbol{\mu},v}(\det g_{b,\mathrm{target}})}.$$

This is the ratio of the local determinant factors on switching from the untwisted to the twisted model. The rightward arrows use the other splitting character, so they contribute one. The absolute-value factors of the integral sign assignment s_v are one. At the retained-label places the untwisted determinant factors at s_v are also one, because the downward factors are unramified on those units and the rightward vertices are unchanged. At the fixed auxiliary split places one moves from the retained reference staircase g^0_v to

the base vacuum. The determinant factors of the untwisted model there are fixed. The remaining ratio is therefore

$$Z_{\mathrm{loc}}(\boldsymbol{\mu}) = \prod_{v\in S} \bigl(X_{v,\boldsymbol{\mu}}(s_v)X_{v,\boldsymbol{\mu}}(g_v^0)\bigr)^{-1},$$

using $g_v^0 = 1$ at the places where the cell label is retained. The common scalar of the reference model can be placed in the first even factor. When expressing the moved definite column vacua in the finite dictionary there is one further product $Z_{\mathrm{move}}(\boldsymbol{\mu})$ of determinant twist values. Write x_j for the determinant half-root of the prescribed finite moving element on the definite vertex j. The two splitting characters of arrow j change by the same twist. With the source and target conventions in the moved column action its multiplier is $\xi_{H,\boldsymbol{\mu}}(x_j)^{-1}\xi_{H,\boldsymbol{\mu}}(x_{j+1})$. The reference model is the same moved vacuum at the base tuple. These character values factor out before summing the vacuum expansion. Their product over the column arrows is Z_{move}. Local determinants of signs and base moves · Moved definite column vacua.

The scalar removed from the finite expression is consequently

$$\gamma'_{\boldsymbol{\mu}} = Z_{\mathrm{loc}}(\boldsymbol{\mu})Z_{\mathrm{move}}(\boldsymbol{\mu})/Z_{\mathrm{base}}(\boldsymbol{\mu}).$$

In particular the ratio uses the same determinants for every term of a cell. These are values of the given finite-order twist and their inverses, so the simultaneous test field contains the phase. The finite character $\gamma_{\boldsymbol{\mu}}$ from the preceding swaps is applied separately. Phase in the flat expansion · Assembly with the moved column factor.

2.3.4 The placed finite Fourier contractions

The local models at the points of a cell. Fix the tuple, one split place $v \in S$, and one row cell. Write B_j for the columns of the row frame at the chosen split component and P_j for the frame of its definite core. These are the two frame assignments in the simultaneous choice. Write $\epsilon_j = s_{v,j}$, $g_j = g_{\beta,v,j}$, S_0 for the initial staircase, and $S_j = 1$ for $j > 0$. The finite one-place point r has coordinates $1, \epsilon_j g_j, 1, 1$ at the initial, interior, and last two vertices, respectively. The reference point r_0 has coordinates $S_0^{-1}, 1, \dots, 1$. In particular the matrix of $\rho_j = r_{0,j}^{-1} r_j$ in B_j is $\epsilon_j g_j$ at every vertex, with $g_0 = S_0$ and $g_{m-1} = g_m = 1$. The two point assignments are identity off v. On the other split component use their inverse Hermitian transposes. Let h_j be the common factor, which has the local component of the fixed tilt at v, and put $a_j = h_j r_{0,j}$. At the last two vertices all three points can be taken equal to the identity. Thus the required relative product is $a_j\rho_j = h_j r_j$. Placed and reference points · Their local matrices.

Use the matrix coordinate X of a map from vertex j to vertex $j+1$ on this split component. If ϕ_b denotes the given split test at the downward grid arrow b, the coordinate in ϕ_b is X^{T}. On the chosen component the global Hom lattice at these vertices restricts to all integral matrices in B_j, B_{j+1}. Indeed the inverse frames on the two sides send the lattice-inclusion condition to entrywise integrality. The supported split half is zero at every other component, and the other half of the trace-dual Hom lattice is its adjoint. Consequently the one-place model $\mathcal{M}_b$ with these frames sends the integral indicator O_b to the reference vector v_b and $\phi_b(X^{\mathsf{T}})$ to the given test vector f_b. This is the local realization of the specified modification of the global vacuum. Equivalently, one obtains the one-place model by applying Heisenberg translations and character projectors at v to v_b; finite Fourier inversion on the quotients of the integral matrix lattice gives $\mathcal{M}_b(\phi_b(\cdot^{\mathsf{T}})) = f_b$. Split Hom lattice in the chosen frames · Transpose and realization of the given local vector.

The chosen oscillator action at a supported point of local matrices (k_j, k_{j+1}) is multiplication by $D_b(k)$ and pullback by $X \mapsto k_{j+1}^{-1} X k_j$. In terms of the downward grid arrow b with $\mathrm{src}(b) = j+1$ and $\mathrm{tgt}(b) = j$, the factor is

$$D_b(k) = \frac{\kappa_{b,j+1}(\det k_{j+1})}{\kappa_{b,j}(\det k_j)} |\det k_j|_v^{\mathrm{rk}(B_{j+1})/2} |\det k_{j+1}|_v^{-\mathrm{rk}(B_j)/2}.$$

Here $\kappa_{b,i}$ is the block character at b divided by the stipulated parity power of η_v at vertex i. The two splitting characters evaluate on single-component ideles with precisely these values, since projection of

the idempotent idele at the chosen component gives $\det k_i$, and the modulus uses the inverse transpose on the grid arguments. The powers of the absolute value are the square roots of the Jacobians on the matrix space for this component. Edge factor with the inverse transpose · Identification of splitting characters · Weil action on both vectors.

In the moved frames $B'_i = a_i B_i \epsilon_i$, with core indices first and flag indices second, set

$$H(X) = \epsilon_{j+1}^{-1} X \epsilon_j, \qquad \mathcal{M}'_b(F) = \omega_b(a_j, a_{j+1}) \mathcal{M}_b(F \circ H).$$

Indeed the same polarized vector in the two frames satisfies

$$a_{j+1}(B_{j+1} X B_j^{-1}) a_j^{-1} = B'_{j+1} H(X) (B'_j)^{-1}.$$

The opposite-half operators are transported by the adjoint identity for this equation. Thus $\mathcal{M}'_b$ has the translation and multiplication laws of the Schrödinger model in B'_j, B'_{j+1}, with no choice of scalar in this transport. Since the two ϵ_i are integral invertible matrices, $O_b \circ H = O_b$. Substituting the relative point in the action gives

$$\begin{aligned} \mathcal{M}'_b(O_b) &= \omega_b(h_j r_{0,j}, h_{j+1} r_{0,j+1}) v_b, \\ \mathcal{M}'_b(\Phi_b) &= \omega_b(h_j r_j, h_{j+1} r_{j+1}) f_b, \\ \Phi_b(Z) &= D_b(\epsilon g)\, \phi_b\left(g_j^{\mathsf{T}} Z^{\mathsf{T}} g_{j+1}^{-\mathsf{T}}\right). \end{aligned}$$

In this equation Z is reindexed in the order of B'_j and B'_{j+1}. The identity

$$(\epsilon_{j+1} g_{j+1})^{-1} (\epsilon_{j+1} Z \epsilon_j^{-1}) (\epsilon_j g_j) = g_{j+1}^{-1} Z g_j$$

accounts for the absence of ϵ from the argument of ϕ_b. Transport of the polarization · Placed inputs in the same local model.

At $v \in S$ write $k = F_v$, normalize additive Haar measure by $\mathrm{vol}(\mathcal{O}_v) = 1$, and use the given character ψ_v with annihilator $\mathcal{O}_v$. For an arrow $j < m-1$ let $c = m-2-j$. A downward matrix has block sizes (n_{j+1}, c) by $(n_j, c+1)$. Start with $\left(\begin{smallmatrix} x & 0 \\ 0 & (1_c \ 0) \end{smallmatrix}\right)$ in these coordinates. For $p = \lfloor c/2 \rfloor + 1$ allow free entries in the lower-left rows $i < p$, upper-right columns $i > p$, and strictly upper-triangular positions of the flag block. In the upper-right column p retain y when j is even. In the lower-left row p retain u when j is odd. These retained entries are outside the free positions. Denote the resulting affine matrices by $C_j(x, y; e)$ and $R_j(x, u; e)$, where e runs through the free entries, and let $s(A)$ be the sum of the first superdiagonal of the flag block. For the two moved local Schrödinger inputs Φ_1, Φ_2 the two flat symbols in an interior pair $j, j+1$ are

$$\begin{aligned} F_1(x, y) &= \int \Phi_1(C_j(x, y; e))\, \psi_v(-s(C_j(x, y; e)))\, de, \\ F_2(x', u) &= \int \Phi_2(R_{j+1}(x', u; e'))\, \psi_v(-s(R_{j+1}(x', u; e')))\, de'. \end{aligned}$$

Here Φ_i includes the shifted frame action, its determinant factor and the prescribed coordinates of the split transfer. The sign of the row-column contraction in these frames is

$$\mathfrak{C}_v(\Phi_1, \Phi_2)(x, x') = \int F_2(x', u) F_1(x, y) \psi_v(-uy)\, dy\, du. \tag{*}$$

In the terminal pair $j = m-2$ there are no free entries in the first matrix and no row integration after it. At the last vertex write $x' = \binom{Y}{u}$ with the extra line in the prescribed position. The terminal formula is

$$\mathfrak{C}_v^{\mathrm{last}}(\Phi_1, \Phi_2)(x, x') = \Phi_2(Y) \int \Phi_1((x \ \ y)) \psi_v(-uy)\, dy.$$

The endpoint decomposition has z_0 as marked line vector; the projection of its image into the preceding line has coordinate $-\xi_v a\, \zeta_v$, with $a = 1$ at this last pair. This is the minus sign in the terminal Fourier

argument. The zero Witt index on the remaining source removes the auxiliary coordinates by evaluation at zero. The final position of the marked row is that of x' in the normalized top frame. The free entries and the paired transform · Choice of free entries and their character · The placed augmentation at the last vertex.

Witt columns and the two contractions. The frame at an even Bessel vertex of Witt index ρ is obtained by inserting P_j on the definite core, putting $s\xi_v^{\rho-i}e_i$ in flag position i, putting $s\zeta$ in the marked position ρ, and putting $s(-h/\xi_v)(-1/\xi_v)^{\rho-1-i}e_i'$ in the reflected position $2\rho-i$, where $0 \le i < \rho$, $s = \xi_v^{-\rho}$, $h = -(-\xi_v)^\rho \xi_v^\rho$, and the marked line vector ζ is included in the core of the Witt decomposition. The two Witt vectors in the original orthogonal flag coordinates are $\xi_v^i w_i$ and $(-\xi_v)^{-i} w_{2\rho-i}$. Inserting these gives w_i and $-w_{2\rho-i}$, and ζ has coordinate ξ_v^ρ in the middle line, giving w_ρ. These agree with the columns B_j'. At the initial vertex the powers of the staircase in B_0 cancel against S_0^{-1} on the flags and have exponent zero on the core.

At the intervening Fourier–Jacobi vertex with index ρ the corresponding columns are P_{j+1} on the core, the isotropic columns $s\xi_v^{\rho-i}e_i$, and the reflected columns

$$s(h/\xi_v)(-1/\xi_v)^{\rho-1-i}e_i' \qquad \text{in positions } 2\rho-1-i.$$

The two multipliers simplify to ξ_v^{-i} and $(-\xi_v)^i$. These are the tilt multipliers on e_i, e_i'. On the reflected flag the frame already contains -1, canceled by ϵ_{j+1}. For the following Bessel vertex the normalization is $s' = \xi_v s$ and its last norm satisfies $-\xi_v^2 h' = h$. The same middle frame therefore serves both arrows of the pair, with the same marked vectors and dual pairings. At Witt index zero the flag formulas have no isotropic columns. Mixed-model frame columns · Bessel columns in the chosen frame · Fourier–Jacobi columns in the same frame.

Consider now the actual mixed-coordinate operators on $\mathcal{M}_b'(\Phi_b)$. Heisenberg modulation at an unread entry Z_{ia} multiplies by $\psi_v(t(Z_{ia}-Z_{ia}^0))$, where Z^0 is the matrix with the marked flag block $(1_c \;\; 0)$ and zero off the core and this block. Its image under the actual flat map is unchanged. This follows by pairing the respective column or row with its dual: in the split component the Heisenberg pairing is the trace pairing, and the above inverse scales on the dual columns cancel the scales on the evaluated columns. The marked columns pair with coefficient 1 at their diagonal positions, and the orthogonal-to-core rows have pairing zero with the core. For a free elementary entry E, the translation $\Phi_b(Z+E)$ has flat image $\psi_v(s(E))$ times the original image. In the upper flag block this follows from $s(UZ) = s(Z) + \sum U_{i,i+1}$ on the affine slice when U is upper unipotent with diagonal one. Multiplying by a simple root on the target uses its inverse, so its factor is $\psi_v(-\sum U_{i,i+1})$. Left multiplication by the simple roots accounts for the free entries of the earlier flag columns, using the diagonal 1 in each such column as the pivot. In the remaining free rectangular blocks the corresponding matrix unit pairs orthogonally with the marked blocks and gives the factor 1. A free transvection meeting the marked column pairs only with its neighboring flag vector and gives exactly that superdiagonal entry. These identities come from moving translations through the defining mixed Fourier integrals. Their coordinate changes are additive translations or triangular shears with unit diagonal. Free translations with the Bessel marked column · The reverse column evaluations · Flag translations for the reverse operator.

Consequently the quotient by these relations is the indicated partial Fourier integral F_1 or F_2. One can see the whole quotient by first imposing the evaluations at the early nonfree flag entries and then using their diagonal pivots to perform the remaining simple-root translations. The other nonfree entries are evaluated by modulation. Translation in the remaining free entries integrates against $\psi_v(-s)$. On any finite partition by compact open rectangles the evaluation eliminates the rectangles missing the marked values, and the integration identifies each translated indicator with its prescribed character times the original indicator. Taking all such partitions proves the factorization for every Schwartz–Bruhat input. The surviving translations and multiplications on F_1, F_2 are the Heisenberg operators in the indicated core-plus-line frames, by applying the same trace pairing to the retained row and column. The linking map

between the lines pairs their coordinates by $-uy$, using $s' = \xi_v s$, $-\xi_v^2 h' = h$, and the fact that the link carries the second marked vector to the negative of the first. Thus the actual two-arrow contraction factors through $\mathfrak{C}_v$. Applying this factorization to $O_b, O_{b'}$ gives the reference contraction with the product of the integral indicators. If the symbol is expanded as a sum of products, linear descent in each core coordinate sends that expansion to the same sum of the Heisenberg operators applied to this reference contraction. In particular, the descended map kills every relation between two such expansions. Evaluation and integration for the column descent · Evaluation and integration for the row descent · The placed interior contraction.

Pullback at the final line. For $j = m - 2$, the first arrow retains the column y and has no free entries. Its two inputs after the Bessel flat map are therefore read in the core-plus-line frame (P_j, ζ) and the core frame P_{j+1}. On the next arrow the source has Witt index zero. The map called extension by zero here inserts the source core into the full source, which at index zero is its entire space. Its model is $Y \mapsto \Phi_2(Y)$, with empty complementary coordinates, and the reference model is $1_{\mathrm{int}}(Y)$.

Write the final column space as $T \perp L''$, where T is the last definite row space. Its frame is $(i_T B_m, i_{L''} z_0)$ with $z_0 = 1$. The map to the preceding marked line is the skew adjoint of τw, where $\tau z_0 = -\xi_v \zeta$. In particular it exchanges the two polarized halves. Its Hermitian normalization is $\tau^* H_L \tau = -\xi_0 \overline{\xi_0}\, H_{L''}$; in the two line frames $H_L = -1$ and $H_{L''} = 1$. For a column y on the first line and a row u on the last line, substitution in the trace pairing on the chosen component cancels ξ_v against the character normalization and gives $-uy$. Therefore the pullback of a line test c under the skew adjoint, in the Schrödinger model on the last line, is represented by $\widehat{c}^{-}(u) = \int c(y)\psi_v(-uy)\,dy$. Indeed translation of the new coordinate becomes multiplication by the old dual character and conversely, and Fourier inversion uses the same Haar measure of volume one on integral vectors, which is self-dual for the indicated conductor $\mathcal{O}_v$. The two polarizations therefore give the terminal symbol $\Phi_2(Y)\widehat{\Phi_1(x,\cdot)}^{-}(u)$. The fixed reindexing sends the T-rows to their original indices and the L''-row to the final index of P_m; it preserves the integral indicator and the coordinate pairing. This is exactly the augmented definite matrix argument at the last vertex. The same computation for the two reference indicators gives the product of the two integral indicators in the definite matrix coordinates. The first input at Witt index zero · The second input in the row summand · The skew adjoint and the negative Fourier transform · The augmented frame and its reindexing · The placed terminal contraction.

For integral indicators on the large matrices, the flat operators give the integral indicators in their smaller matrix spaces. Indeed all their fixed flag entries are zero or one, and integration in each free entry has volume one and trivial character there. The conductor formula $\widehat{1_{\mathcal{O}_v^n}} = 1_{\mathcal{O}_v^n}$ shows that the contracted interior symbol is $1_{M_{n_{j+1},n_j}(\mathcal{O}_v)}(x)1_{M_{n_{j+2},n_{j+1}}(\mathcal{O}_v)}(x')$, and the same holds for the augmented terminal symbol. These calculations fix the reference normalization with no residual Haar factor at the two-arrow contraction. For a moved symbol partition the support into rectangular compact open cosets. Fourier transformation of one such indicator is

$$\widehat{1_{b+L}}(u) = \mathrm{vol}(L)\psi_v(-ub)1_{L^\perp}(u).$$

Consequently the retained symbols are finite sums of products of functions on the two matrix spaces. Translation and modulation in either retained matrix factor act on the standard lattice vacuum by the corresponding Heisenberg operators in the same frame. This gives the finite dictionary by applying these products of translations and Fourier projectors to the reference product of indicators. A relation between two finite expansions is killed by the descended linear operator. Thus no independence assumption on the chosen indicators or translates is required. Covariances for the column descent · Covariances for the row descent · The paired Heisenberg action · Product expansion relative to the integral vacuum.

The one-staircase claim at the fixed auxiliary places uses the conductor and periodicity on these same symbols. Write e_0 for the test depth and ϖ for a uniformizer. Index flag vectors starting at 1, and put $d_i = \varpi^{e_0 i}$. At the initial step the source model of the support calculation uses the staircase $\mathrm{diag}(1_{n_0}, d_1, \dots, d_{m-1})$.

This is the matrix in the test argument; the acting base point used to untwist the vacuum is its inverse. Suppose at one step the source frame uses the staircase and write h for the target finite representative in the unshifted local symbol. Its lattice is $L = h\mathcal{O}_v^{n_{j+1}+c}$. A nonzero value of the joint contraction implies a point of the affine support where the moved full matrix is integral. Its i-th flag column v_i has flag coordinate 1 at i, zero below i, and zero middle block if $i \le \lfloor (c+1)/2 \rfloor$. Integrality says $d_i v_i \in L$.

For the reverse divisibility take $v \in L$ with zero flag coordinates below i and zero middle block when $i \le \lfloor c/2 \rfloor$. In the next flag column $i+1$ of the test argument translate by tv/d_i for $t \in \mathcal{O}_v$. It changes only free entries. Conjugating this change to the moved coordinates multiplies on the right by d_{i+1} and on the left by h^{-1}. The resulting change is integral at depth e_0, so the test value is invariant. The only change in the Fourier character is the factor $\psi_v(-tv^{(i)}/d_i)$, where $v^{(i)}$ denotes the i-th flag coordinate. The pair term $-uy$ is unchanged because both its entries are retained. Translation of the integral gives $\psi_v(-tv^{(i)}/d_i) = 1$ for all $t \in \mathcal{O}_v$. The conductor condition gives $v^{(i)} \in d_i\mathcal{O}_v$.

Let U have columns v_i on the flag and identity on the middle block. It belongs to the flag radical with the middle block at position $\lfloor (c+1)/2 \rfloor$. Then $U^{-1}L$ contains $d_i e_i$ on each flag line. Choose an integral basis of $U^{-1}L$ adapted to the block flag with middle block of position $\lfloor c/2 \rfloor + 1$. Such a basis is obtained by successive lifts of bases of the filtration quotients of the lattice, which are free over the discrete valuation ring. The two divisibilities just proved give $d_i\mathcal{O}_v$ as the projected lattice on the i-th successive flag line. Replacing its lift by $d_i e_i$ keeps a basis, because it changes the i-th pivot by a unit and adds only earlier coordinates. The middle basis can have entries only in the middle block and the earlier flag coordinates. The resulting basis matrix is a middle Levi element times an element of the flag radical times $\mathrm{diag}(1, d_1, \dots, d_c)$. Both radicals used in the two orders lie in R_{j+1}. Therefore h belongs to the reference staircase double coset. Applying this argument successively, with the chosen reference representative of that coset at each step, proves the single allowed staircase at these places. Here the depth is exactly the staircase depth e_0 of the transferred tensor, and the argument applies to the joint integral even if this tensor is a finite sum. Conductor detection by depth-preserving translations · Adapted basis of the target staircase lattice.

2.3.5 The ordinary-place support calculation

The ordinary-place calculation uses the moved self-dual row lattices, with the prescribed core lattice and reflection-paired flag at each vertex. To see the ideal normalizations, write $\mathfrak{a}, \overline{\mathfrak{a}}$ for the two host ideals and index the flag by $i = 0, \dots, c_j - 1$. Before the tilt its line ideals are

$$I_{j,i} = \mathfrak{a}^{c_j - i}\overline{\mathfrak{a}}^{-(i+1)}.$$

In an even row the normalized first Witt vectors are ξ_0^i times the flag vectors. In an odd row the finite tilt multiplies the first-half ideals by ξ_0^{-i}. Both conventions therefore give isotropic coefficient ideal $\xi_0^{-i} I_{j,i}$. On the reflected Witt line the coefficient ideal is its conjugate inverse, since $I_{j,i}\overline{I}_{j,c_j-1-i} = \mathcal{O}_E$ and the finite tilt on the reflected lines uses inverse conjugate multipliers. The even middle line, with Gram -1, has ideal $(\mathfrak{a}/\overline{\mathfrak{a}})^{\rho+1}$ if $c_j = 2\rho + 1$. At either parity the normalized neighboring isotropic ideals have quotient $(\xi_0)\mathfrak{a}\overline{\mathfrak{a}}$, the annihilator of the integral scalars for the trace character on the completion. This is the primitivity of the adjacent entries. The relation $\mathfrak{a} I_{j+1,i} = I_{j,i}$ gives the equality of coefficient ideals under the marked column map after including the Hom-lattice scale $\mathfrak{a}$. For Bessel to Fourier–Jacobi the column map uses unit first flag vectors at the target. In the other direction it uses the vectors scaled by ξ_0^i, ending with the marked vector ξ_0^ρ in the middle line. The flag ideals · The moved flag and core lattices · Transport of the first flag ideals · Transport including the marked vector.

The same ideal check gives primitivity of the last Bessel-to-Fourier–Jacobi covector. Choose a local generator g of the middle line ideal. The corresponding norm in the marked coordinate is $-\overline{\xi_0}^{\rho}\xi_0^{\rho}$. The two fractional factors in the pairing with the last reflected flag are $g\xi_0^{-\rho}$ and $g^{-1}\overline{\xi_0}^{-(\rho-1)}$. Their product with this norm is ξ_0, using $\overline{\xi_0} = -\xi_0$. Applying the Hom-lattice scale therefore recovers the trace dual just

computed. Thus the last covector has integral scalar stabilizer. The split calculation uses both summands of $E \otimes_F F_v$; at inert places the trace pairing is on the unramified quadratic extension. Terminal covector with its line and Hom-lattice factors.

Here is how the support implications follow from the local mixed-coordinate laws. Iwasawa for the self-dual Witt lattice writes the target point as a row unipotent, a diagonal Witt element, a core element and a compact factor. This decomposition follows by extending primitive isotropic frames to Witt bases of the lattice, splitting off the paired vectors successively. Strip off the compact factor by the unramified vacuum property and the unipotent by the equivariant character. For the first arrow the core factor is allowed in the row Levi, so only the diagonal $t = \mathrm{diag}(t_i)$ remains. Multiply the vacuum by the character in the dual coordinate to the i-th marked column. For a lattice translate the vacuum is fixed. After conjugating by the torus the diagonal pairing is multiplied by t_i^{-1}, because the marked columns are $X_0 U$ for upper unipotent U and $U_{ii} = 1$. Nonvanishing gives t_i^{-1} in the integral scalars by trace duality and the column ideal equality. Applying instead the elementary transvection in position $(i, i+1)$ of the Witt flag gives the pairing $a\, t_i/t_{i+1}$ with a in the above annihilator ideal. Its character is one, since the conjugated transvection fixes the vacuum. Hence t_i/t_{i+1} is integral. The last covector calculation in the previous paragraph gives t_ρ integral by testing the core-to-Witt Heisenberg translations. These conditions give all t_i and t_i^{-1} integral. This places the diagonal in the original stabilizer. Dual diagonal tests · Simple-root tests · Last covector test.

For the second arrow $g^{-1}X$ has diagonal Witt components t_i^{-1} and last component $\gamma^{-1}\zeta$ in the core, where γ is the remaining core element. Dual coordinate tests on the individual Witt components give t_i^{-1} integral. The simple-root translation of column $i+1$ by $t_i a$ times column i is a lattice translate when a belongs to the same annihilator ideal. Nonvanishing forces its character to be one, hence t_i integral. The last dual test, now projected to the core, gives $\gamma^{-1}\zeta$ in the normalized marked-line lattice. More precisely it gives $\gamma^{-1}(a\zeta)$ integral for every a for which $a\zeta$ was integral. Take a generator of this ideal. The two integral marked vectors have the same unit norm; their orthogonal complements in the self-dual core lattice are again self-dual. Over an unramified quadratic extension of odd residue characteristic these complements are isometric, by diagonalization and surjectivity of the norm on units. In the split case the vector and corresponding covector form a direct summand with unit pairing, so an integral change of bases on the two halves gives the same orbit. This proves that γ is a core element fixing the marked line times an integral stabilizer. Nonvanishing at the second arrow thus forces the Bessel row orbit. At an inert place the lattice Heisenberg tests use 2 a unit. At a split place use each isotropic half separately; their characters multiply to the full trace character. These tests, and the twisting away from the places where the test tensor has been changed, give the two conditions used in the ordinary-place induction. Shift and last-line tests on the Fourier–Jacobi-to-Bessel arrow · Integral orbit of the marked vector.

Combining the finite places. Outside S the flip and the modification of the global reference vector act through the split idempotents supported in S. They commute with Heisenberg translations at the other finite places, so the local factors there are the transported self-dual lattice vacua. For a lattice isometry between the old and moved Hom lattices the identity $\omega(g)H(z)\omega(g)^{-1} = H(gz)$ carries the defining translation and opposite-half multiplication invariances to the moved lattice. In particular the Heisenberg tests in the ordinary-place support calculation apply to the actual two flat maps. Their dual diagonal factors are obtained by multiplication inside the first mixed Fourier integral, their isotropic shifts inside the second mixed integral, and their unipotent covariance by translation of the flag integral in these same moved lattices. At nonsplit places these formulas use the Hermitian trace pairing of the lattice and its trace dual. At the ordinary split places they use the two inverse dual components. No variation of the finite cell label enters these identities.

The local factors combine by multilinearity because their Heisenberg operators at different places commute. More explicitly, a finite tensor relation of functions is detected by evaluations at tuples of points. For each finite family of functions, finitely many evaluations span the dual of its linear span,

by successive selection of a point where a nonzero function in the remaining kernel is nonzero. Hence applying the descended local maps to the factors preserves the relation. This applies at every paired position including the last one, and then to the finite tensor expansion over the places. At $v \in S$ the factors $D_b(\epsilon g)$ are included in the placed vectors above. In comparison with the symbols for the cells, the ratio is $\prod_b D_b(\epsilon)$, times the removed determinant character at $g = g_v^0$ at a fixed staircase place. The two sources of this ratio are the sign assignment and whether the cell test retains the character of its representative. All absolute-value factors at the sign assignment are one. This recovers Z_{loc}^{-1} after taking out the fixed untwisted product, so the normalization is the one used in the phase calculation. The common reference bi-vacuum, the moved definite vectors, and the Haar changes have the factors Z_{base}, Z_{move}, and $\prod_v d_{\beta v}$ computed there. The ordinary exceptional-place factors give the prescribed s_{AA}; the other ordinary places use the unramified vacuum. Substitution in the convergent unfolded integral and application of Fubini give the finite factors of the first line of (15), with its real tensor determined in the normalized class pairing below. The actual translated vacua off the selected places · Assembly over pairs at one place · Assembly with the reference vacuum.

The exceptional-place factor. At an exceptional prime w of E, write q_w for its residue cardinality and a, a' for the valuations of the two host ideals $\mathfrak{a}, \overline{\mathfrak{a}}$. Put $e_j = 0$ on the core coordinates and $e_j(i) = a(c_j - i) - a'(i+1)$ on the flags indexed from zero. These are the exponents of the diagonal representative of the original lattice, relative to the integral matrix frame. The undisturbed finite test on each arrow is the indicator of matrices in $\varpi^a M(\mathcal{O}_w)$. In the moved box at the origin, a coordinate in row h and column k of the downward matrix at j is in $\varpi^{b_{j,hk}}\mathcal{O}_w$, where

$$b_{j,hk} = a - e_j(k) + e_{j+1}(h).$$

The adjacent flag exponents differ by $a + a'$ on a tested superdiagonal position. The sum of the exponents of each paired retained row and column coordinate is also $a + a'$. The fixed terminal coordinate is zero at the origin. Therefore the Fourier phase is trivial on the box, by the $\varpi^{a+a'}\mathcal{O}_w$ conductor. Write B for the sum of $b_{j,hk}$ over the integrated free and retained coordinates. The box contributes volume q_w^{-B}.

Let ζ_X be the sum of the diagonal exponents at a last-column vertex X and 0 elsewhere. Put $r_X = \dim X$. For an arrow $b : s \to t$ on the opposite path the two inverse-transpose half-densities give $q_w^{(\zeta_s r_t - \zeta_t r_s)/2}$. The cell weight is $q_w^{M_w}$, where M_w is the sum of $e_j(h) - e_j(k)$ over entries (h, k) of the integrated unipotent radicals. Indeed conjugation multiplies each such entry by $\varpi^{e_j(h)-e_j(k)}$, and the weight is the inverse volume. Writing the vacuum in the reference trace-dual frame adds $\prod_{b\,:\,s\to t} q_w^{(a-a')r_s r_t/4}$. Thus the exceptional factor in the unfolding is

$$s_w = q_w^{M_w - B + \sum_{b\,:\,s\to t}(\zeta_s r_t - \zeta_t r_s)/2 + (a-a')\sum_{b\,:\,s\to t} r_s r_t/4}.$$

Successive ranks have opposite parity, so this is positive and is an integral power of $\sqrt{q_w}$. Their product is s_{AA}, which is nonzero and belongs to the field with these square roots. The exceptional representatives and the trace-dual correction · The box with trivial Fourier phase · The inverse unipotent volume · The positive normalized scalar.

For the test, put

$$V_{\boldsymbol{\mu},a}(x) = s_{\mathrm{AA}} \sum_{\beta} \left(\prod_{v} d_{\beta v} \right) J_{\boldsymbol{\mu},a,\beta}(x).$$

For the coefficient field K with chosen embedding ι_K, the field condition on the test includes the values of $\xi_{H,\boldsymbol{\mu}}$. They have common finite order for each tuple, so $\gamma_{\boldsymbol{\mu}}\gamma'_{\boldsymbol{\mu}}$ is $\iota_K(e_{\boldsymbol{\mu}})$ for a root of unity $e_{\boldsymbol{\mu}} \in K^\times$. Substituting the expansion and jet identity into the path formula gives

$$\iota_K(e_{\boldsymbol{\mu}}) c_{\boldsymbol{\mu}}^{\square}\, \psi_a(\Theta_{\boldsymbol{\mu}}^{\downarrow} j_{\boldsymbol{\mu}}^{\mathrm{top}} x) = \kappa V_{\boldsymbol{\mu},a}(x). \tag{16}$$

The product of the square constants appears separately from the root of unity. The scalar κ is nonzero as soon as one test has nonzero left side.

2.3.6 Rationality of the finite sums

In the coefficient model the column and row inputs at a class representative are polynomials over K. Transport by a rational unitary matrix through the specified field embeddings is polynomial substitution and multiplication by a determinant power. Place type factors are character values. A chain of normalized coefficient matrices over K preserves this polynomial model. At the top, the coefficient restriction in codimension one is the algebraic projection onto the required branching summand. One can take its adjoint projection in the polynomial model using the monomial pairing, with factorial diagonal entries. Normalizing by the product of the corresponding extremal-vector pairings divides by a nonzero element of K. Thus the resulting bilinear contraction is over K.

For the finite test coefficients at a split place, taking the Fourier sums preserves the field containing the occurring phases. The complementary integral also preserves it when it contains $\sqrt{q_v}$ and the occurring determinant character values. Here is a direct calculation. With additive matrix measure the weight on $\mathrm{GL}_\ell(F_v)$ is a rational measure normalization times $\nu(\det g)|\det g|^{n/2-\ell}$. For the old rank $\ell = a'$ the exponent $s = n/2 - \ell$ is nonnegative. Composition of the test with the injective full-rank frame map gives a compactly supported locally constant function of the matrix g. After rescaling, it is a finite linear combination of integral matrix coset indicators. In a coset $D + \varpi^e M_\ell(\mathcal{O}_v)$, integral elementary divisors allow D diagonal, with valuations less than e in its nonzero positions. The block over those positions is invertible of constant determinant valuation. The other block in the determinant is its Schur complement, ϖ^e times an unrestricted integral matrix translated by an integral matrix. The character of the first block is computed on a finite congruence quotient. For the unrestricted j-by-j matrix integral, the ramified character case with $j > 0$ gives zero by multiplying a column by a unit. For an unramified character, decompose the first column by the common valuation k of its entries and take its primitive part to the first basis vector by an integral change of rows. Expanding the determinant along that column gives the recursion

$$I_j = \frac{1 - q_v^{-j}}{1 - \nu(\varpi) q_v^{-(s+j)}} I_{j-1}, \qquad I_0 = 1.$$

All factors are in the field. This argument applies at every split swap; the final Fourier inversion again involves a finite sum.

In the two definite paths of this test, the field calculation uses a further normalization of the kernel entries. In the dictionary for an arrow p, the raw derivative table at a matched pair of points has the form

$$J_p^{\mathrm{raw}}(\alpha) = c_{0,p} \prod_\tau \omega_{p,\tau}^{|\alpha_\tau|} J_p^{\mathrm{alg}}(\alpha).$$

The table J_p^{alg} is obtained by differentiating the finite sum of reduced torsion translates of the moved vacuum. For the reversed-arrow convention the complex factors are conjugated. The $\omega_{p,\tau}$ are the chosen nonzero CM periods, with that convention included. The symbol $|\alpha_\tau|$ is the derivative degree at τ. In the coefficient operator the source determinant powers and the substitution of frame variables give a single degree $d_{p,\tau}$ at each τ, as in the polynomial calculation below. Hence the scalar $c_{0,p} \prod_\tau \omega_{p,\tau}^{d_{p,\tau}}$ passes through the entire operator. Its reciprocal is the normalization of the kernel matrix, so the normalized entries are the entries computed from J_p^{alg}.

The field for the actual finite dictionary. Use each occurring pair of finite twist character and marked ramified set when choosing the coefficient field. For such a pair the split tensors, residue levels and depths are fixed before choosing the auxiliary class realization. For each test and tuple the support of its finite unfolding consists of finitely many row cells. At the fixed staircase places there is just the prescribed cell.

At the other split places one can see finiteness directly from the free-coordinate translations. Suppose the previous representative and its inverse have bounded entries. At a point where the cell integrand is nonzero, an affine matrix A in the two flat symbols has $h^{-1}A$ bounded, where h is the next row representative. This follows from integral support after multiplication by the previous representative. Write $h = bk$ by block Iwasawa, where k is integral invertible and b is upper triangular in flag order with a middle block. Put the core before the central flag when these have the same row-group weight.

For each flag index i, multiply an elementary free increment $zE_{i,i+1}$ in the flag block by b. Its nonzero column has entries only in the free positions by triangularity, and h^{-1} carries it to $k^{-1}zE_{i,i+1}$. The fixed test period and the bound on the previous representative give one nonzero ideal of values of z for which this preserves the test. The phase change is $\psi_v(zb_{ii})$, since the paired retained entries stay fixed. Nonvanishing forces this phase to be one for every z in that ideal. The conductor hence bounds the flag diagonals b_{ii}. In the row Levi component, the inverse flag columns are bounded by those of $b^{-1}A$, because the affine matrix has flag pivots 1 and zero entries below those pivots. This holds also for the central block when a flag shares the core weight. In that case the Levi block has the form $\left(\begin{smallmatrix} C & d \\ 0 & a \end{smallmatrix}\right)$. Its inverse flag column bounds a^{-1} and $C^{-1}d\,a^{-1}$. After removing C by the row Levi action, the bound on a bounds $C^{-1}d$. Removing the flag radical now leaves a matrix and its inverse in bounded sets. The set of matrices with both bounds is compact. An open right level has finitely many cosets intersecting this compact set. Induction from the fixed initial staircase proves the required finite support, at every selected place. Support and depth after the split changes of path · Bounding the flag and the middle block · Depth-preserving increments in the actual affine slice · Finite induction on the representatives.

For a supported cell take the prescribed expansion by cosets at its least positive depth of expansion. At fixed depth the products of the coset indicators on the split halves are disjoint. The compact support of the symbol therefore bounds the possible nonzero products by a finite quotient on each of those halves. The constant coefficient of one product is the value of the local symbol at its representative. These coefficients are over K by the Fourier and matrix integrals above, with the specified characters and phases. The class spaces at the term levels are finite because the groups there are definite. The terminal column representative is the image of the last row representative, so it needs no new choice. For these finitely many classes, pure summands and coefficient indices, take the normalized matrix entries on both paths. The derivatives after pulling back the self-dual Hom lattices are those of the finite sums of torsion translates used below in §2.6.11. The Gaussian normalization removes the same fixed CM period powers on every monomial of the indicated homogeneous weight. The fourfold theta transport there gives algebraicity of each normalized entry, with the transpose and conjugation on the row path. Adjoining these finite lists, the character values and the required square roots of residue cardinalities yields the specified coefficient field on the representative family. Finiteness at the prescribed depth · The normalized class entries · Adjoining the finite lists on the representative family.

In passing to the selected grid the corresponding tuples have the same twist and marked ramification. Carry the definite spaces, lattices, local frames, rational coefficient frames and finite moving elements by their specified identifications. The initial row corner is X', with the distinguished lattice in its identity frame on both sides. The terminal embedding is the same orthogonal-summand embedding. This carry also uses the same normalized vacua, constants $c_{0,p}$ and period parameter. With these choices the split transfer gives the same tensor in the same local matrix arguments. Hence taking its flat Fourier symbols gives identical functions, with the same supported cells, levels and least depths. In particular the finite coefficient expansions can be compared in the same quotients. At the class points the two derivative tables then evaluate the same Heisenberg translates of the same normalized vacuum, with the same two finite moves and real monomials. The arithmetic entries in the selected grid are thus the previously adjoined numbers. This gives the values of $V_{\boldsymbol{\mu},a}(x)$ in $\iota_K(K)$ in (16). Transport of frames and finite moves · Transport of the normalized vacuum · Transport of the support and both series of entries · Field membership on the selected tests.

2.3.7 The normalized definite class pairing

The coefficients at the initial points. Write C_j and R_i for the definite vertices of the ascending column and row, respectively. In this paragraph a finite point X_j is on C_j, and Y_i is on R_i. At the terminal vertex put $X_m = \iota(Y_m)$, using the stabilizer embedding adjoining the identity on the line. The two initial coefficient functions are

$$u_0(X_0) = \rho_{C_0}^{\mathrm{fr}\,-1}\Psi_a^{\mathbb{C}}(X_0),$$
$$v_0(Y_0) = \rho_{R_0}^{\mathrm{fr}\,-1}\,\mathrm{read}\big(A_{\boldsymbol{\mu}}(j_{\boldsymbol{\mu}}^{\mathrm{top}}x, Y_0)\big).$$

Here ρ^{fr} is the prescribed coefficient change to the orthonormal real frame, and read takes the listed row weight to its dual coefficient convention. The column coefficient is the image of the given polynomial Ψ_a at the finite point. The row coefficient $A_{\boldsymbol{\mu}}(y, Y_0)$ is the evaluation of the coefficient section y of X' at the image of Y_0 under the fixed identification $R_0 \simeq X'$, followed by the coefficient comparison. This comparison extends the coefficient embedding by polynomial substitution and renames the indices under the rank identification. Thus $A_{\boldsymbol{\mu}}$ is the coefficient used in the arithmetic sum.

In the automorphic realization of these two functions the archimedean components enter by the inverse real group actions. In fact after writing $g = g_f g_\infty$ the two automorphic values are

$$\Lambda_{C,0}\big(\rho_C(g_\infty^{-1})u_0(g_f)\big), \qquad \Lambda_{R,0}\big(\rho_R(g_\infty^{-1})v_0(g_f)\big).$$

For the column this follows by evaluating the fixed real coefficient vector against $\Psi_a^{\mathbb{C}}$. For the row, in the polynomial realization the fixed real vector is paired with the inverse of the same coefficient comparison, then with the inverse real action at g_∞. The comparison is injective because it is scalar extension of the coefficient space followed by an injective renaming of the matrix variables. Its inverse on the image therefore yields exactly the initial row value, including for the original input $j_{\boldsymbol{\mu}}^{\mathrm{top}}x$. Finite and real parts commute, so no finite determinant character is introduced by this frame change. The two functionals $\Lambda_{C,0}$ and $\Lambda_{R,0}$ are fixed for the family. Initial column and row coefficient evaluations · The column coefficient realization · The row coefficient realization · Initial real functionals in the two frames.

Normalized recursion at a pair of finite points. For a finite product summand d let $J_{p,d}^{\mathrm{raw}}(X, Y)$ be the Fock derivative table for definite arrow p in its coefficient frames. If a_X, a_Y are the two fixed finite moves from the standard self-dual lattices, the table is read at (Xa_X, Ya_Y) on the original vacuum with the finite Heisenberg operators of d applied to it. Indeed those moves are on the right of the group arguments of the kernel, and they commute with the real creation operators. At the opposite-sign row arrow the matrix indices in the Fock monomial are interchanged and the CM table is conjugated. These are the derivative table and the sign convention in the prescribed coefficient operator.

For clarity, let a, b be the two ranks, and at each complex embedding let n_τ be the a-th entry of the polynomial part of the target highest weight, with indices starting at one. Given a source coefficient polynomial $u(H)$, the coefficient operator first multiplies it by $\prod_\tau \det(H_\tau)^{n_\tau}$. It then substitutes $H_\tau = Z_\tau^{\mathsf{T}} T_{\tau,[a]}$, where Z_τ is b by a and $T_{\tau,[a]}$ consists of the first a columns of the target polynomial variables. Replace each Z^α by its entry in the actual derivative table. This is the jet operator in the fixed polynomial coefficient realization. The first a entries of the target polynomial weight are the source polynomial weight plus n_τ, and the remaining entries are zero. Thus the result is already in the required target coefficient space, and the total Z_τ degree is the sum $d_{p,\tau}$ of the target polynomial weight. This proves directly that division by $b_p = c_{0,p}\prod_\tau \omega_{p,\tau}^{d_{p,\tau}}$ yields the prescribed normalized operator $\mathcal{J}_{p,d}(X, Y)$. Both the real period and its finite column or row reading use the same polynomial and the same points. The nonzero b_p is independent of the finite points, and on every column arrow it is independent of the tuple because the normalized reference vacuum and its periods are fixed before twisting. The row convention · The prespecified normalized matrix · The polynomial substitution · The actual transported column points · The actual transported row points.

One can see the real integration giving this same operator before the finite class sum. At each embedding write $k \in U(a)$ and $h_\infty \in U(b)$ in the two coefficient frames. The oscillator on the coefficient variables substitutes $Z \mapsto h_\infty^{\mathsf{T}} Z k^{-\mathsf{T}}$. If ℓ_s, ℓ_t are the last entries of the source and target weights in that convention, its scalar is $\det(k)^{n_\tau - \ell_s} \det(h_\infty)^{\ell_t}$. In the column frame let c be the center of the block inserted below the common entries. The infinitesimal character is read as highest weight by subtracting the half-sum of the positive roots. The last target entry is $c + a/2$, and each common entry is decreased by $(b-a)/2$. Hence $n_\tau - \ell_s = -c - b/2$. These are the two splitting exponents with the positive oscillator half-ranks, in their source and target order. In the reversed row frame let c instead be the center of the inserted block before reversing its sign and let $\varepsilon(n)$ be the parity of n. The two row exponents are

$$c - \varepsilon(b)/2 - b/2, \qquad -c + \varepsilon(a)/2 + a/2.$$

The latter is the last target entry. Subtracting the parity shifts and the difference of the half-ranks in the common entries gives the former as $n_\tau - \ell_s$. The source coefficient action in the integral supplies $\det(k)^{\ell_s}$ and substitutes $H \mapsto kH$. The two determinants on k therefore give exactly the power of $\det H$ in the jet operator.

For fixed initial real functional Λ, express it as a linear combination of finitely many polynomial evaluations. These evaluations span its dual since a polynomial vanishing at all invertible matrices is zero. For one such evaluation at $H = H_0$, absorb $\det(k)^{n_\tau}$ by using $f(H) = u(H) \prod_\tau \det(H_\tau)^{n_\tau}$, and divide at H_0 by the nonzero $\prod_\tau \det(H_{0,\tau})^{n_\tau}$. The average of $P(Zk^{-\mathsf{T}}) f(kH_0)$, taken coefficientwise in Z, is an intertwiner from the source polynomial representation to the rectangular polynomial space. This follows by changing variables in Haar, and extends from the compact group to its complexification. By the Cauchy decomposition, the space of these intertwiners is an irreducible representation of the target general linear group at each embedding. It is spanned by the substitutions $f(H) \mapsto f(Z^{\mathsf{T}} T_{[a]})$ at invertible T. Indeed one such substitution is nonzero by specializing Z on its first a rows, and its target translates give exactly these substitutions as T varies. The same argument for the product of the compact groups gives the assertion at all embeddings. Applying the finite derivative functional to this linear combination recovers the jet operator evaluated by a new real functional. The target determinant and substitution are its inverse real action. Multiplying the new functional by b_p passes from the raw to the normalized finite operator. The compact average of the rectangular substitution · The column exponents · The column averaging · The row exponents.

Thus, with probability Haar on the real source group, at each arrow the exact recursion for a pure summand and a fixed real polynomial pair P is

$$\int_{G_{p,s}(F_\infty)} K_{p,d,P}(Xk, h)\, \Lambda_{p,s}\big(\rho_{p,s}(k^{-1})u\big)\, dk = \Lambda_{p,t}\big(\rho_{p,t}(h_\infty^{-1}) \mathcal{J}_{p,d}(X, h_f)u\big).$$

The new functional depends on the real polynomial, the preceding functional and the fixed normalized vacuum, uniformly in the finite input. Iterate along both paths, starting with $u_0(X_0)$ and $v_0(Y_0)$. The terminal average over $G_{R_m}(F_\infty)$, where $X_m = \iota(Y_m)$, is the bilinear invariant pairing of the last two representations. The dual type of R_m occurs with multiplicity one in the restriction of the C_m-representation. Averaging the two terminal functionals with probability Haar therefore gives $\kappa(P)\mathcal{B}$, where $\mathcal{B}$ is the specified algebraic branching contraction. Here $\kappa(P)$ depends only on P and the fixed real data. The column and row recursions at matched points · The last compact average.

The prescribed term levels and their class weights. First refine all the compact open levels used in a cell to levels fixing each factor in its finite expansion. Such levels exist by local constancy of the input coefficients and Schwartz vectors, since the expansion is finite. On this refinement the real-average recursion gives a finite sum with weight $\mathrm{vol}(K')/|G(F) \cap XK'X^{-1}|$ for each class X at each integrated

vertex. Probability Haar in the real variable is the one just used in the recursion. The fine-level calculation also proves invariance under left rational changes of representative, by moving the real component of such an element through the compact average.

To descend to the prescribed levels use the entire finite expansion of the cell. At an interior column vertex a prescribed level element is the middle projection of a stabilizer of the chosen row double coset. Lift it to that stabilizer and change the row variable by its unipotent component. The flag characters from the two adjoining flat kernels cancel by the paired translations of the retained row and column. To check the determinant in this substitution, write $d = \det k_v$ on the split component. The lift and the projected core matrix have the same determinant d, because the flag lift is unipotent and conjugating by the cell representative preserves the determinant. It is a unit. The two half-densities are one and the parity factors from the same vertex occur in opposite powers. Thus the path factor at an interior column vertex j is $\mu_{j,v}(d)/\mu_{j+1,v}(d)$. The same factor is the product of the scalars on the two translated core vacua. At the initial column vertex only the factor $\mu_{1,v}(d)^{-1}$ occurs, and K^+ fixes Ψ_a. At an interior row vertex i the two edges give $\chi_{m-i,v}(d)/\chi_{m-i+1,v}(d)$, again the product of the two vacuum factors. The typed path gives invariance with precisely these path factors, which persist through the square changes of variables. At the last row vertex make the same change also in the column vertex by ι. The orthogonal complement is the fixed line, so this preserves the terminal Fourier argument and branching contraction. The last rightward and downward factors there give $\mu_{m,v}(d)/\chi_{1,v}(d)$. In each case the product of vacuum factors replaces the determinant normally outside the substituted test arguments. Hence the unchanged full tensor equality for the placed path becomes the unchanged sum of composed maps. This uses linearity of the Schrödinger realization and of the normalized operator in its derivative table. Away from the selected split places the level stabilizes the reference lattice and the vacuum factors on units are trivial. Core lift and the interior column determinant factors · Interior projected column level · Initial column level · Interior row level · The two terminal determinant factors · Joint terminal level.

At the first row vertex the comparison has a character. For k in the prescribed first row level let k_v be its split coordinate. This level lies in the actual X' level, in the frame used for the initial row coefficient. Its character on that coefficient is $\chi_{\boldsymbol{\mu}}(k) = \prod_v A_v(k_v)$, where

$$A_v(k_v) = \begin{cases} \displaystyle\prod_{c=1}^{m} \chi_{c,v}(\det \bar{k}_{v,c}), & v \text{ a type component,} \\ (\xi^R_{\boldsymbol{\mu},v})^{-1}(\det \bar{k}_{v,\mathrm{tail}}), & v \text{ a tuple component,} \\ 1, & \text{otherwise.} \end{cases}$$

Here the type product omits the initial r-block, and the tail block at a tuple place has size mt. These values are the diagonal-block determinants in the same residue matrices, since the first row frame is the prescribed X'-frame. The original last downward step test at this vertex transforms under the local path action with eigenvalue $A_v(k_v)^{-1}$. This follows by multiplying the block diagonal characters and canceling the determinant weights in the split Weil action on units. The vertex is exterior to every square being changed. The same local eigenvalue therefore persists in the opposite tensor by performing every square change with the exterior substitution held fixed.

More explicitly, let $D_v^{\mathrm{first}}(k_v)$ be the determinant factor of the first rightward edge in its split frames. Its half-densities are one on this level. Replacing Y_0 by $Y_0 k$ in the raw function arguments removes $\prod_v D_v^{\mathrm{first}}(k_v)$ from the path action. In the standard oscillator model, this same change is by $a_{Y_0}^{-1} k a_{Y_0}$ on the vacuum and its Heisenberg translates, where a_{Y_0} is the fixed first row moving element. The scalar on the vacuum is $\prod_v D_v^{\mathrm{first}}(k_v)$: its determinant in the moved frame is still $\det k_v$. Away from the selected split components the unit determinant factors are trivial. The trace-dual lattice condition guarantees preservation of the vacuum line. Applying the same argument substitution to all factors in the expansion now gives the multiplier $\prod_v A_v(k_v)^{-1}$ for the full normalized kernel at $Y_0 k$. The frames and degree normalizations are unaffected by k. Multiplying by the initial row coefficient, which has multiplier $\chi_{\boldsymbol{\mu}}(k)$, gives

invariance of the paired summand on the prescribed Y_0-class. Original and transferred first row level character · Block determinant and vacuum scalar comparison · The moved first row vacuum · Cancellation on the paired class value.

A coarse class above X, on passage to $K' \subseteq K$, has fine classes indexed by the orbits of $G(F) \cap XKX^{-1}$ acting on K/K'. The sum of the reciprocal stabilizer orders is $[K : K']/|G(F) \cap XKX^{-1}|$ by orbit–stabilizer. Multiplying by $\mathrm{vol}(K')$ and using $\mathrm{vol}(K) = 1$ gives the desired weight at the actual term level. Thus the full class sum is

$$\kappa(P) \sum_{X,Y} \frac{\sum_d c_d \, \mathcal{B}\big(\mathcal{J}_{C,d} u_0(X_0),\ \mathcal{J}_{R,d} v_0(Y_0)\big)}{\prod_{j=0}^{m-1} |\Gamma_{X_j}| \prod_{i=0}^{m} |\Gamma_{Y_i}|}.$$

Here $\mathcal{J}_{C,d}$ and $\mathcal{J}_{R,d}$ are the composed normalized arrow operators at the indicated points. The terminal column point is $\iota(Y_m)$, so it contributes no independent class mass. The two chains, initial coefficients and stabilizer products in this expression are exactly those in $J_{\boldsymbol{\mu},a,\beta}(x)$. The finite expansion enters only through its total tensor: the same evaluation argument as for the local contractions passes linear relations through the composed operators. It therefore gives the value for the prescribed arithmetic expansion as well as for any refinement used in the integrals. Term class representatives and weights · Regrouping the fine classes · The paired class expansion.

The preceding class calculation shows that every term as a functional of P is proportional to $\kappa(P)$. Use this as the common functional ℓ when it is nonzero. If it vanishes identically, choose any nonzero ℓ, since in that case the multiple is zero in every term. Choose a polynomial tensor σ_0 with $\ell(\sigma_0) = 1$. At a real quotient point and for a real input tensor P, the placed comparison gives a finite definite Fock combination $\sigma_P(p)$. Because the reference bi-vacuum and both initial real types are fixed, its functional value $\ell(\sigma_P(p))$ is independent of the finite inputs. The compact Fock reduction projects to the central weights of the two fixed real types. Their scalar circle weights fix the degree at each embedding, so only finitely many real Fock coefficients enter that functional. Its evaluation commutes with the finite Heisenberg operators. This explains the same linear functional for all terms, with the finite determinants already removed. If some reference definite period is nonzero, divide the placed identity for that term by its fixed value. This writes the scalar as a measurable, integrable function of the real quotient point by absolute convergence of the unfolded integral. If all the reference definite periods vanish, every placed term vanishes and one can take the integrated coefficient to be zero. Denote the resulting real quotient integral by $c(P)$. If $\sum_P b_P P$ is the original fixed real tensor, the representative of the real output is therefore

$$\sigma = \frac{\sum_P b_P c(P)}{s_{\mathrm{AA}}} \, \sigma_0.$$

The denominator is the nonzero exceptional-place product computed above. This is the real combination when s_{AA} is written explicitly in the first line of (15). Using the same σ in the class calculation replaces $\kappa(P)$ by $\kappa = \sum_P \sigma(P)\kappa(P)$. Thus the scalar from the real data is common across the tuples, tests, cells and initial coefficients. This proves $A_{\boldsymbol{\mu},a,\beta}(x) = \kappa J_{\boldsymbol{\mu},a,\beta}(x)$ in the original term measures. Real central characters in the Fock reduction · The common real combination with the exceptional vacuum factor · Linearity in the real tensor.

2.3.8 Ratios of square constants

To compare principal-level tests across tuples one needs $c_{ij,\boldsymbol{\mu}}/c_{ij,\boldsymbol{\mu}_0} \in \iota_K(K^\times)$ for each square. Write $\lambda_{\boldsymbol{\mu}}$ for the comparison factor on the standard definite triangle whose two normalized class sums are defined over K. Section 2.3.22 proves $\lambda_{\boldsymbol{\mu}}/\lambda_{\boldsymbol{\mu}_0} \in \iota_K(K^\times)$ by evaluating a fixed complementary pair of minor products on that triangle. Equation (19) gives the ratio on any of the prescribed bottom triangles in the cell.

In the comparison of the square with this triangle, put $\mathcal{R}_{\boldsymbol{\mu}}(s) = \Lambda_{\boldsymbol{\mu}}(s)/\Lambda_{\boldsymbol{\mu}_0}(s)$ for the specified Hecke L-values of that square. Put $p_j = \lfloor (m-j-1)/2 \rfloor$, $q_j = \lfloor (m-j)/2 \rfloor$, and $n_0 = \lfloor (t+r+6)/2 \rfloor$ from the difference of its two intermediate ranks. The required identity in the actual normalizations is

$$\frac{c_{ij,\boldsymbol{\mu}}}{c_{ij,\boldsymbol{\mu}_0}} \prod_{k=0}^{q_j-1} \mathcal{R}_{\boldsymbol{\mu}}(n_0+k+1) = a_{ij,\boldsymbol{\mu}} \frac{\lambda_{\boldsymbol{\mu}}}{\lambda_{\boldsymbol{\mu}_0}} \prod_{k=0}^{p_j-1} \mathcal{R}_{\boldsymbol{\mu}}(n_0+q_j-k). \tag{17}$$

where $a_{ij,\boldsymbol{\mu}} \in \iota_K(K^\times)$ is the product of the measure and local step factors. For nonzero L-values the common factors cancel. When $m-j$ is odd, $p_j = q_j$ and the two products have the same indices. When $m-j$ is even, $p_j = q_j - 1$ and the only uncanceled factor is $\mathcal{R}_{\boldsymbol{\mu}}(n_0+1)$ on the left. Consequently only this latter ratio in even cells needs a field calculation. Triangle and critical-value factors · Application at the fitted points.

2.3.9 The square and the descending triangles

Keep the square, the tuple family and the coefficient embedding fixed. For a triangle at stage (p,q), the two source ranks are A_0+2p and B_0+2q, where $A_0 = R_{i+1,j+1} - 2p_j$ and $B_0 = R_{ij} - 2q_j$. The negative indices of the two sources and the target are $p, q, p+q$, respectively. In particular the half-difference of the source ranks is $e(p,q) = n_0 + q_j - p_j + p - q$. The two splitting characters stay fixed along each Witt tower, since a hyperbolic plane changes each rank by two. Their quotient at every stage is $\nu_{ij,\boldsymbol{\mu}}$. Stage and splitting characters.

Let $D_{p,q}(\boldsymbol{\mu})$ be the quotient of the triangle's nonzero comparison constants at $\boldsymbol{\mu}$ and $\boldsymbol{\mu}_0$. This quotient is independent of the choice of common polynomial pair representing the comparison. In fact, if the two matched operators are $\lambda_{\boldsymbol{\mu}} J_{\boldsymbol{\mu}} p_*$ and $J_{\boldsymbol{\mu}} p'_*$ for fixed polynomial maps and injective $J_{\boldsymbol{\mu}}$, the proportionality of the two values on a fixed polynomial pair implies

$$p_*(P) = \frac{c_{\boldsymbol{\mu}}}{\lambda_{\boldsymbol{\mu}}} p'_*(P'). \tag{18}$$

For $p'_*(P') \neq 0$ this forces $c_{\boldsymbol{\mu}}/\lambda_{\boldsymbol{\mu}}$ to be independent of the tuple. This lets the successor obtained at one step be used as the input of the next. Comparison of stage constants.

On the actual row grid, the top equality from the matched, regularized theta integrals takes the form

$$\frac{c_{ij,\boldsymbol{\mu}}}{c_{ij,\boldsymbol{\mu}_0}} = a^{\mathrm{top}}_{\boldsymbol{\mu}} D_{p_j,q_j}(\boldsymbol{\mu}), \qquad a^{\mathrm{top}}_{\boldsymbol{\mu}} \in \iota_K(K^\times).$$

The two hyperbolic-plane comparisons, on the same stages and with the stage measures, give

$$\begin{aligned} D_{p,q}(\boldsymbol{\mu})\mathcal{R}_{\boldsymbol{\mu}}(e(p,q)+1) &= b_{p,q,\boldsymbol{\mu}} D_{p,q-1}(\boldsymbol{\mu}) && (q>0),\\ D_{p,0}(\boldsymbol{\mu}) &= b_{p,0,\boldsymbol{\mu}} \mathcal{R}_{\boldsymbol{\mu}}(e(p,0)) D_{p-1,0}(\boldsymbol{\mu}) && (p>0). \end{aligned}$$

The powers in the finite factors of the two comparisons are $-e(p,q)$ and $1-e(p,0)$, respectively. The L-functions initially omit places above the fixed base level. At its auxiliary prime the denominator character is ramified and the numerator is unramified. At the type primes the denominator has conductor exponent two and the numerator has exponent at most one. Thus the quotient is ramified above every omitted place, so both step identities use the finite primitive L-functions on their indicated arguments. Top comparison · Step on the smaller source · Step on the larger source · Euler factors at the base level.

The field assertion about $b_{p,q,\boldsymbol{\mu}}$ has an elementary local calculation once the two matched step identities hold. In the split chart at v, write ν_v for the quotient character on that factor and h_v for its conductor exponent. The local factor in the step, in the convention using ν_v^{-1} on units, is

$$\nu_v(s_*) q_v^{-h_v s_\varepsilon} \nu_v(\varpi_v)^{h_v} \sum_{u \in (\mathcal{O}_v/\varpi_v^{h_v})^\times} \nu_v(u)^{-1} \psi_v(u/\varpi_v^{h_v}),$$

where s_* is the fixed sign and s_ε is $-e(p,q)$ or $1-e(p,0)$ in the two respective steps. An unramified twist at a base place leaves the Gauss sum on units fixed, and the change of the uniformizer factor is a character value in the field. At the remaining places the relevant conductor exponent is at most the specified depth. The additive character is trivial on $\mathcal{O}_v$, so the phases have order dividing $q_v^{h_v}$. These roots and the local character values belong to the coefficient field of the grid. The two fractional-ideal factors in the step are finite products of local uniformizer values and their inverses. The change of a translating vector contributes a finite-order twist value in the field. Taking ratios on the matched nonzero tests proves $b_{p,q,\boldsymbol{\mu}} \in \iota_K(K^\times)$. The measure factor in each step is common to both tuples and cancels. This deduction uses only the field part of the local calculation. Gauss factors · Base places · Auxiliary places.

At stage $(0,0)$, transport the reached triangle to the standard definite triangle used for the rational class sums. First carry the Hermitian spaces by their isometries, then move the self-dual lattices and their frames. The common Haar and real changes multiply the comparison constants by a tuple-independent scalar. In the finite transport the two vacuum factors scale by $r_{\boldsymbol{\mu}}$ and $r'_{\boldsymbol{\mu}}$. The change in the comparison constant is $r'_{\boldsymbol{\mu}}/r_{\boldsymbol{\mu}}$, with ratio to its base value in $\iota_K(K^\times)$. The frame changes have the corresponding property from the finite character quotients. Writing $\lambda_{\boldsymbol{\mu}}$ for the factor on this standard triangle gives

$$D_{0,0}(\boldsymbol{\mu}) = a^{\mathrm{bot}}_{\boldsymbol{\mu}} \frac{\lambda_{\boldsymbol{\mu}}}{\lambda_{\boldsymbol{\mu}_0}}, \qquad a^{\mathrm{bot}}_{\boldsymbol{\mu}} \in \iota_K(K^\times).$$

Lattice and vacuum transport · Transport to the standard triangle.

Run the smaller-source steps first. The successive arguments on the left are $n_0+1,\dots,n_0+q_j$. After these steps run the larger-source steps, whose arguments on the right are $n_0+q_j,\dots,n_0+q_j-p_j+1$. Multiplying the identities cancels each intermediate comparison quotient and yields (17), with $a_{ij,\boldsymbol{\mu}} = a^{\mathrm{top}}_{\boldsymbol{\mu}} a^{\mathrm{bot}}_{\boldsymbol{\mu}} \prod b_{p,q,\boldsymbol{\mu}}$, the product over these steps. Since all L-values are nonzero there, the ranges cancel completely when $p_j = q_j$, and leave exactly $\mathcal{R}_{\boldsymbol{\mu}}(n_0+1)$ on the left when $p_j = q_j - 1$. In the product of square constants used for the downward coordinates, each remaining Hecke factor has exponent -1 and is indexed by its cell (i,j). Descent identity · Comparison on the same grid.

2.3.10 The real quotient for the square

At one real place use the oriented Hermitian–skew-Hermitian Schrödinger models of the two-leg kernels. The two Hermitian source signatures are $(a-p,p)$ and $(q,a'-q)$. In the convention that the skew-Hermitian target form is multiplied by i, its signature is $(R,S) = (a+a'-p-q,\,p+q)$. Write m_α, m_β for the exponents in the embedding of the triangle model, so the characters of $\mathbb{C}^\times$ have factors $(z/|z|)^{m_\alpha}$ and $(z/|z|)^{m_\beta}$. Put $\kappa_\infty = m_\alpha - m_\beta$. The grid gives $\kappa_\infty > a+a'$ and κ_∞ even. The latter follows because a, a' have the same parity and the restrictions of the characters to $\mathbb{R}^\times$ are the corresponding powers of the sign character.

The polynomial oscillator action in a signed orthonormal frame uses the operators

$$(B_\ell, A_\ell) = \begin{cases} (x_\ell, \partial_\ell), & \text{the two frame signs agree,} \\ (-\partial_\ell, x_\ell), & \text{the two frame signs differ.} \end{cases}$$

The matrix unit indexed by ℓ, ℓ' acts by $B_\ell A_{\ell'} + \frac{1}{2}\delta_{\ell,\ell'}$ before adding the splitting-character trace. This is the Fock action obtained by intertwining the prescribed Heisenberg operators with the Gaussian line. On compact diagonals it gives $(\frac{R-S}{2},\dots,\frac{R-S}{2};\frac{S-R}{2},\dots,\frac{S-R}{2})$ on the source vacuum, and the analogous target vacuum weight. These are the signs of the oscillator used in [93, (3.3)–(3.4)]. For the large source the splitting pair is $(\beta,\beta/\alpha)$, with the first entry used on the target. For the small source it is $(\alpha,\alpha/\beta)$. The source splitting twists therefore make invariance equivalent to the $\det^{\kappa_\infty/2}$ and $\det^{-\kappa_\infty/2}$ coinvariants of the respective untwisted oscillator actions. The target trace shifts are $m_\beta/2$ and $m_\alpha/2$.

The ordered pair of signatures

$$((a-p,p),(a'-q,q))$$

defines one θ-stable parabolic $\mathfrak{q}$ of $\mathfrak{u}(R,S)_{\mathbb{C}}$. We have $|R-S| < \kappa_\infty$, $R \geq a-p, a'-q$, and $S \geq p, q$. Paul–Trapa's Table 1 and Theorem 4.1(1),(2) give the two untwisted small theta lifts on this same ordered Levi. With the genuine character convention of their Section 2.8, they are

$$A_{\mathfrak{q}}(\overset{\kappa_\infty/2}{\det} \boxtimes 1), \qquad A_{\mathfrak{q}}(1 \boxtimes \overset{-\kappa_\infty/2}{\det}).$$

Tensoring with a character of the full group commutes with cohomological induction, by tensoring the defining complexes and restricting the character to the Levi. Thus the two representations agree after applying their respective target splitting twists, since the two resulting ordered pairs of exponents are $(m_\alpha/2, m_\beta/2)$. The parity of the ranks is the same on both legs, so the metaplectic genuineness conventions in this comparison agree. The theorem supplies nonzero irreducible quotients, including the compact cases and the cases $p = 0$ or $q = 0$. [93, §2.8, Table 1 and Theorem 4.1(1),(2)]

Choose surjective quotient homomorphisms from the two Fock polynomial modules to this representation, and tensor them over the real places to obtain q_A, q_B. Howe's polynomial version of duality gives a unique maximal proper submodule in each source coinvariant module. These maximal submodules are the kernels of q_A, q_B. For unitary groups the compact groups $\mathrm{U}(R) \times \mathrm{U}(S)$ are connected, and each vector is compact finite. A Lie-stable subspace is therefore stable under the compact group, by exponentiation on the finite-dimensional compact orbit of each vector. The resulting irreducible real module and both quotient maps can be chosen once for the prescribed frames and infinity types. The finite twist leaves those types fixed. [97, Theorem 2.1]

The polynomial quotients are the compact-finite parts of the continuous irreducible quotients of the ordinary real Schwartz models. In fact the signed frames give compatible compact subgroups for the oscillator, so the polynomial modules are the modules in the algebraic correspondence. Coincidence with the smooth correspondence on both unitary pairs gives nonzero continuous maps to the source character tensored with the smooth globalization of the computed target module. Their restrictions to the polynomial space are nonzero by density of the Hermite functions in the Schwartz topology. Polynomial multiplicity one then allows their normalization to the chosen q_A, q_B. The one-dimensional source twist and the target twist enter these continuous actions by the same characters as in the polynomial computation. The comparison uses polynomials with fixed finite tensors, and continuity on real Schwartz space is used to select an initial polynomial on which the large integral has nonzero value. [94, Proposition 5.1] [97, Theorems 1 and 2.1]

To apply these maps to the grid, identify the polynomial of the first arrow with the first target block in the orthogonal sum of the two exterior vertices. On the second arrow, before concatenating columns, transpose the matrix of Fock variables. In the large route the transpose has factor $\epsilon_i \epsilon'_j$ on the variable with indices i, j in the original arrow, where ϵ_i, ϵ'_j are its target and source frame signs. In the small route the original second-arrow source frame is negated, so the factor is the product of its signed frame signs. Indeed the Hermitian adjoint conjugates the coefficients and multiplies the transposed matrix unit by this sign product. In the large route the second conjugation in the signed model restores the original coefficients. In the small route pullback to the negated second-arrow space gives the signed transpose directly. These identities hold first for the Gaussian and then for each creation operator by differentiation of the Heisenberg translation. Orthogonal direct sums give the product of the two Gaussian polynomials in the two sets of columns. Let ι_A, ι_B be the resulting maps from the two route polynomial tensor spaces to the two Fock polynomial spaces. They are onto by the signed variable permutations and the polynomial tensor product in disjoint variables. The identities hold with the finite tensors fixed throughout, including finite sums of pure tensors. Large-route Fock substitution · Small-route Fock substitution

The comparison relation

$$q_A(\iota_A P) = q_B(\iota_B P')$$

has onto projections and is proper, since the common real module is nonzero. In any product realization of an irreducible adelic quotient, the polynomial maps at fixed finite input vanish on the corresponding local

kernels. Consequently the two maps on such an input have the form $q_A(\iota_A P)\otimes r_A(\varphi)$ and $q_B(\iota_B P')\otimes r_B(\varphi')$, in common product coordinates. This factorization follows by expressing the real-isotypic module as its irreducible real factor tensored with the span of the intertwining maps into the quotient. Indeed evaluation of this tensor product is onto and injective, since the real endomorphism ring is $\mathbb{C}$. Schur's lemma in this form applies to the countable-dimensional real polynomial module. Factoring the fixed infinity-type maps

2.3.11 Transport of the quotient across the two polarizations

For each leg, the rational Darboux transformation of the square intertwines the source and target actions on its algebraic Schwartz space. The transpose and sign on the small-source Darboux model give the opposite polarization. Its source relations are the augmentation differences for that transported action and the Lie derivatives for the actual small-source splitting character. The polynomial quotient maps on the two models have the same kernel after transport. In detail, take a finite tensor stage containing the exceptional places and all finite components that depart from the vacua. In the Hermitian model the sum of the local quotient kernels is the kernel of the tensor of the local surjections. Each local kernel is the unique maximal proper submodule containing the source relations. At the real place this follows by polynomial Howe duality, and at a finite place by the corresponding irreducible quotient and uniqueness for the smooth Weil representation. A tensor product of these finite-length coinvariant modules has only one irreducible quotient. To see uniqueness, use a composition series on the factors. An irreducible quotient of the product is the exterior tensor product of simple factors. Restricting to each leg gives a map from that local coinvariant module to a sum of copies of its simple factor, so it must kill the corresponding local kernel. The sum of those local kernels already has simple quotient. For split factors one may use the maximal-kernel assertion on the product directly. The transported kernels are compatible with passing to a larger finite tensor stage, since the extra vacuum maps are equal and nonzero. Comparing the two surjective algebraic quotient maps on each leg now identifies the kernels. The real Darboux operator takes the Gaussian and its creation operators to those in the corresponding transported Cartan frame. The finite intertwiners act on the corresponding finite Schwartz spaces; they introduce no new places away from the compatible lattice vacua. The smaller-leg comparison also applies to arguments with its prescribed full-rank finite support.

If both pairs of quotient maps use a common target, the two legwise identifications differ by a single target intertwiner. Its graph is stable for the real Lie algebra and the finite target group. It is therefore multiplication by a nonzero scalar in the common irreducible target coordinates. This supplies the scalar when transporting the Hermitian-triangle equality to the square, using the unchanged Schwartz tests in each Darboux model. Kernel on the first polarization · Kernel on the opposite polarization

2.3.12 Constant terms on the chosen hyperbolic planes

At an upper stage write the source ranks as $a > a'$, the target rank as $d = a + a'$, and $e = (a - a')/2$. Write T_A, T_B for the larger and smaller source integrals. In the mixed-model convention the smaller source has its form sign reversed, so the signatures are $(a - p, p)$, $(q, a' - q)$, and $(d - p - q, p + q)$. The rank hypotheses $e > 2(p + q) + 1$ and $a' > 2(p + q) + 1$ hold throughout the rectangle of stages. Indeed the two differences are bounded below by $n_0 - 2(p_j + q_j) - 1$ and $r - 2(p_j + q_j) - 1$, which are positive by $k \geq m + 3$ and $t \geq r + 5$. The infinity-type weight difference is $2a_\tau > d$, with a_τ as in the critical-value calculation below. Use the actual splitting characters α, β and set $\nu = \alpha/\beta$. The planes have isotropic bases (x, y) in the source being reduced and (x', y') in the target, with the prescribed pairings equal to 1. The orthogonal complements carry the lower lattices and Cartan involutions. The removed basis vectors are the two outer vectors of the corresponding block at the exceptional places. These conditions are the lattice and frame conditions used in the evaluation maps below. Adapted hyperbolic planes.

For the target parabolic of Ex', take the constant term by averaging over its rational unipotent quotient with total mass 1. In the mixed Schrödinger model the trivial Fourier character cuts out the isotropic

vectors in the source. The zero vector gives the target-only evaluation E_0 at the origin of the removed mixed coordinates. It yields the lower theta integral if the source is unchanged. When the source also loses a plane, both nonzero shifts and their reduced evaluation maps can be described on the smaller tensor space. Denote these by $u_A(l) = l\,x \otimes y'$, $u_B(l) = \bar{l}\,y \otimes x'$, and E_A, E_B, respectively. They are evaluation at zero after Fourier transform on the isotropic coordinates complementary to the corresponding shift. The mixed-model Fourier expansion follows by Poisson summation on the removed coordinates, with the self-dual measures for the fixed additive character. Its zero orbit and its nonzero orbit have the two Levi characters

$$\alpha(t)|t|_E^{a'/2}, \qquad \beta(t)|t|_E^{a/2}.$$

Since the real exponents differ, these two components are unique. Indeed, comparing at $t = 1$ and at an idele of norm other than 1 determines the values of both components on the lower Levi. Mixed evaluations · Two Levi characters.

For an isotropic line in a source of rank n, the determinant of conjugation by $t \in \mathbb{A}_E^\times$ on the unipotent radical has absolute value $|t|_E^{n-1}$. The E-linear coordinates contribute $|t|_E^{n-2}$, and the center, which is one-dimensional over F, contributes $|N_{E/F}t|_F = |t|_E$. The rational nonzero isotropic vectors form one orbit by Witt's theorem. Unfold its sum by the stabilizer of x, with quotient integration in the order unipotent, lower source, ideles, and compact group. This gives the factor $|t|_E^{1-n}\,d^\times t$. For the larger source one uses $t = l^{-1}$, and E_A has weight $|t|_E^{(d-2)/2}$ and splitting character $\nu^{-1}(t)$. The Mellin weight is thus $\nu(l)|l|_E^{a-1-(d-2)/2} = \nu(l)|l|_E^{e}$. For the smaller source the second polarization uses $t = l$, and E_B has weight $|t|_E^{d/2}$ and character $\nu(t)$. The Mellin weight is $\nu(l)|l|_E^{d/2+1-a'} = \nu(l)|l|_E^{e+1}$. The compact average is taken with the fixed product of normalized compact measures. For $J = A, B$, the resulting operator is

$$C_J\Phi = \int_K \int_{\mathbb{A}_E^\times} E_J\rho(u_J(l))\omega(k)\Phi\,\nu(l)|l|_E^z\,d^\times l\,dk, \qquad z = e\ (J = A), \quad z = e+1\ (J = B).$$

All source and target Haar changes in this unfolding depend on the geometry and the fixed additive character. The fixed rank bounds give $z \geq 2$, so the finite product here is in its absolute convergence range. The exchange of the global kernel sum and integral also uses the stage's absolute theta estimates, including the smaller-source support condition at the extra split place v_1. Line-parabolic quotient measure · Larger-source mixed weight · Smaller-source mixed weight · Absolute unfolding and compact average.

For the smaller source the switch to u_B must be made before unfolding. To see the correction, expand the theta sum in the two polarizations and average over their common isotropic coordinates. The intersection is spanned by $x \otimes x'$, $x \otimes W_1$, and $X_1 \otimes x'$, where X_1, W_1 are the two complements. Orthogonality for the mixed-model Fourier characters leaves the indices $\xi x \otimes y'$ on one side and $\xi y \otimes x'$ on the other, with $\xi \in E$. The nonzero rational vectors of a fixed isotropic line are reindexed on the second side by $\xi \mapsto \bar{\xi}^{-1}$. Subtracting the zero indices leaves a multiple of the unshifted E_A term, since $E_B\Phi = 0$ for the smaller-source tests. The unshifted correction lies in the $\alpha(t)|t|_E^{a'/2}$ component. It does not change the $\beta(t)|t|_E^{a/2}$ component used in this comparison. At v_1, the upper small input is in the split-coordinate vanishing subspace used in the smaller-source estimate. Its condition is preserved by the target action and by the projection associated to E_B, so the lower input in the cone integral still satisfies it. Fourier restriction to each rational line · Poisson correction and split components · Smaller-source support and quotient at the extra split place.

2.3.13 Local scalars on the matched quotients

Away from the ramified places, let $\mathfrak{b}, \mathfrak{c}$ be the line ideals for x, x' on the two adapted planes. Write $\mathfrak{d}_E^{-1}$ for the absolute codifferent and use a bar for conjugation of ideals. The two line ideals for y, y' are $\overline{\mathfrak{b}}^{-1}$ and $\mathfrak{d}_E^{-1}\overline{\mathfrak{c}}^{-1}$ in the Hermitian and skew-Hermitian lattice conventions. Thus the two shift conditions on l cut out the ideals

$$\mathfrak{a}_A = \mathfrak{b}\mathfrak{d}_E^{-1}\overline{\mathfrak{c}}^{-1}, \qquad \mathfrak{a}_B = \overline{\overline{\mathfrak{b}}^{-1}\mathfrak{c}}.$$

For the component at a place w of E these conditions on the standard vacuum are $\operatorname{ord}_w(l) \geq m_w$, where $m_w = \operatorname{ord}_w \mathfrak{a}_J$. With $\operatorname{vol}(\mathcal{O}_{E_w}^\times) = 1$ integration gives

$$\sum_{h \geq m_w} \nu_w(\varpi_w)^h q_w^{-zh} = \nu_w(\varpi_w)^{m_w} q_w^{-zm_w} L_w(z, \nu_w).$$

The ideal norm powers are common geometry factors across twists. The character evaluates to $\nu(\mathfrak{bc})$ for the first cone and $\nu(\mathfrak{bc})^{-1}$ for the second. This uses $\nu \circ \mathrm{conj} = \nu^{-1}$ and triviality on the codifferent factor. In this setting the relative discriminant is the unit ideal, so the absolute codifferent is extended from F, on whose ideles ν is trivial. At places omitted from the vacuum calculation the character is ramified over every place of E, so the Euler product in these factors is the finite primitive $L(z, \nu)$. Line ideals and their evaluations · Reassembly of the Euler product.

At an exceptional split place put $k = F_v$ and identify $E_v = k \times k$ using the selected place w, with character $(\vartheta, \vartheta^{-1})$ and conductor exponent $h > 0$. Use the measures $\operatorname{vol}(\mathcal{O}_k) = 1$ and $\operatorname{vol}(\mathcal{O}_k^\times) = 1$ for the additive and multiplicative groups. The local additive character is unramified. The two standard sections in the matrix model have their respective square blocks in $\mathrm{GL}(\mathcal{O}_k)$, weighted by ϑ^{-1} of the determinant. All entries in the other columns lie in $\varpi^h \mathcal{O}_k$. In the common quotient they are both the parabolic section supported by $Z \in \varpi^h \operatorname{Mat}_{a' \times a}(\mathcal{O}_k)$ on the opposite cell $\left(\begin{smallmatrix} 1 & 0 \\ Z & 1 \end{smallmatrix}\right)$, where they have value 1. Indeed the diagonal block minors there are identities and the remaining columns are Z or its transpose with sign. Off the open cell the two complementary projected block minors are singular, so both sections vanish. Parabolic covariance then gives equality at every group element. This verifies the common nonzero test in both the upper and lower quotients. Two sections on the opposite cell and its complement · Transport to the fixed frames.

The partial Fourier evaluations on these sections remove two pivot positions from the same square block. In the cone evaluation the determinant on the chart is minus the product of the two removed pivots times the determinant of the lower block. The translation sets one pivot to the first coordinate s of l. The Fourier variable pairs with the other coordinate t of l. The minus sign in the determinant and the Fourier sign on that coordinate cancel under the simultaneous change of the integration variable. The other integrated entries are one row of the untouched block or one column after transposition. Their total volume is $q_v^{-hn_{\mathrm{other}}}$, where $n_{\mathrm{other}} = a'$ for the larger source step and $n_{\mathrm{other}} = a$ for the smaller source step. Exactly the same volume multiplies the plain evaluation on the other section. Set $D_\vartheta(x) = 1_{\mathcal{O}_k^\times}(x)\vartheta(x)^{-1}$ and $\widehat{D}_\vartheta(t) = \int_k D_\vartheta(y)\psi_v(ty)\,dy$. After the common volume is canceled, the cone multiplier is

$$\int_{k^\times \times k^\times} D_\vartheta(s)\widehat{D}_\vartheta(t)\ \vartheta(s)\vartheta(t)^{-1}|st|^z\ d^\times s\ d^\times t.$$

The integral defining $\widehat{D}_\vartheta(t)$ is zero for $\operatorname{ord}_v t < -h$, by additive orthogonality on cosets modulo ϖ^h. It is zero for $\operatorname{ord}_v t > -h$, by multiplication by a unit in the last conductor subgroup on which ϑ is nontrivial, using all units when $h = 1$. For $u \in \mathcal{O}_k^\times$, partition into the residue classes of the units to get

$$\widehat{D}_\vartheta(\varpi^{-h}u) = q_v^{-h}\vartheta(u) \sum_{x \in (\mathcal{O}_k/\varpi^h)^\times} \vartheta(x)^{-1}\psi_v(\varpi^{-h}x).$$

Substitution in the multiplier gives $q_v^{h(z-1)}\vartheta(\varpi)^h$ times this Gauss sum before its factor q_v^{-h} and its u-character. Equivalently, the multiplier is exactly $\varepsilon_v(1-z, \vartheta, \psi_v)$ in the finite convention of the comparison. It is nonzero by Fourier inversion and the character's transformation law on the shell. Removed pivots and surviving row or column · Epsilon convention · Smaller-source volume · Larger-source volume.

To use these scalar calculations after taking a constant term, the plain and cone maps must descend to the common upper and lower target quotients. They intertwine the target line parabolic with the Levi character selected above. At an unramified place the scalar comparison on the spherical line then holds on its parabolic span. Indeed the parabolic times the maximal compact gives the full target group, the

compact translates of the vacuum are scalar, and its nonzero image is cyclic in the irreducible quotient. At a ramified split place the equivariant maps are proportional. The split matrix argument discards the two removed columns by the unipotent relations and a unit on which ϑ is nontrivial. The surviving full-rank stratum has the twisted coinvariants of the transitive space of full-rank rectangular matrices. Integrating over the invertible source block identifies this representation, up to a target determinant twist and inverse transpose, with normalized parabolic induction of the unitary determinant character and the trivial representation. Bernstein's irreducibility theorem and the countable-dimensional form of Schur's lemma give proportionality. The nonzero matched matrix sections computed above determine the scalar. At v_1 in the smaller-source step one also needs descent from the full Schwartz space to the quotient. The change of variables in the cone integral kills the augmentation relations from the source action, so the descent agrees with the one obtained using the small support subspace, which still surjects to that quotient. Cyclicity of the spherical comparison · Ramified Levi restriction · Descent of the small cone at the extra split place.

The real quotient maps are proportional on the polynomial vectors with the prescribed Fock data. In the highest-weight model the Levi weight and the compact diagonal weight on the deleted pair of coordinates determine the possible image of the upper highest-weight vector up to one scalar. The unipotent covariance equations and the mixed-root equations then propagate that scalar from the highest-weight vector to the whole polynomial module. For the maps at hand, nonvanishing of the plain map follows by surjectivity of the mixed evaluation and the lower quotient maps. Nonvanishing of the cone map uses a Schwartz bump at infinity in the removed variables, followed by polynomial density on the prescribed quotient. The real cone and evaluation maps, their quotient normalizations and their characters at infinity are common to all finite twists under comparison. Their ratio is therefore a common nonzero scalar. Highest-weight restriction · Nonzero quotient maps on the larger-source comparison · Nonzero quotient maps on the smaller-source comparison.

Upon tensoring, the vacuum scalar at each extra unramified place used to present a tensor cancels its inverse from factoring out the Euler product away from that finite set. Thus for a matched polynomial pair (Φ_A, Φ_B) the reduced pairs are

$$\big(C_A\Phi_A,\ k_A E_0\Phi_B\big), \qquad \big(k_B E_0\Phi_A,\ C_B\Phi_B\big)$$

in the larger and smaller source steps, respectively. Here k_J is a common nonzero geometric and real factor times the corresponding line-ideal character factor, the product of finite epsilon factors, and $L(z, v)$. This assertion uses the quotient maps of the prescribed test models on every tensor factor. Their product on an upper tensor equality gives the indicated lower tensor equality, and linearity gives the statement for the finite sums making up the matched pair. Assembly of the normalized larger-source maps · Assembly of the normalized smaller-source maps.

Using the matching constant. Let c, c_1 be the upper and lower matching constants on the specified quotients. For the smaller removal the $\beta(t)|t|_E^{a/2}$ coefficients of the larger and smaller upper integrals restrict to $C_0\,T_{A,1}(E_0\Phi_A)$ and $C_1\,T_{B,1}(C_B\Phi_B)$, respectively. The lower matching identity gives $k_B\,T_{A,1}(E_0\Phi_A) = c_1\,T_{B,1}(C_B\Phi_B)$. For the larger removal the matching coefficients are those of $\alpha(t)|t|_E^{a'/2}$. They restrict to $C_1\,T_{A,1}(C_A\Phi_A)$ and $C_0\,T_{B,1}(E_0\Phi_B)$, with $T_{A,1}(C_A\Phi_A) = c_1 k_A\,T_{B,1}(E_0\Phi_B)$. Here C_0, C_1 refer separately to the geometric constants for the two cases and are independent of the finite twist. Surjectivity of the target-only mixed evaluation on polynomial data and nonvanishing of the lower lifts give a test with nonzero chosen coefficient in each case. Canceling that coefficient gives

$$c\,k_B(C_1/C_0) = c_1 \quad \text{for the smaller removal,} \qquad c = (C_1/C_0)k_A c_1 \quad \text{for the larger removal.}$$

These equations compare actual matching constants from the chosen coefficient. The argument of Section 2.3.16 promotes the coefficient equality to equality on the matched polynomial tests. The nonzero

starting lift is the definite one in Section 2.3.22. Smaller-source nonzero matching coefficient · Larger-source nonzero matching coefficient · Induction using the end stages.

2.3.14 The normalized quotient at the inert places

Write L/k_v for the unramified quadratic extension at an inert place v. The triangle assumptions give odd residue characteristic, the conductor-zero additive character, and self-dual lattices. The exceptional set is split, so both splitting pairs are unramified here. The unramified characters α_v, β_v of $L^\times$ agree and have product one. Indeed a uniformizer of k_v is a uniformizer of L, and on it the two values are $(-1)^{a'}$ and $(-1)^a$. In particular the character on the source determinant in each Weil action is trivial.

The target is the quasi-split unitary group H of dimension $2m_v = a + a'$. Use the integral Siegel parabolic $P = MN$ and the stabilizer K of the target self-dual lattice. The decomposition $H = PK$ can be taken in an integral Witt basis for this lattice. Its Levi $M = \mathrm{GL}_{m_v}(L)$ acts on N by $B \mapsto tBt^*$ in Hermitian matrix coordinates (after multiplying by a fixed trace-zero unit if necessary). For diagonal t, each diagonal coordinate contributes $|N_{L/k_v} t_i|_{k_v}$. Each off-diagonal pair is one L-coordinate with multiplier $t_i\overline{t_j}$ and hence contributes $|N_{L/k_v} t_i|_{k_v}|N_{L/k_v} t_j|_{k_v}$. The elementary unipotents have determinant one in this action. Thus

$$\delta_P(t) = |\det t|_L^{m_v}.$$

Pass to the mixed Schrödinger model associated with the target Witt basis. Evaluation e_A at zero on the large source and e_B on the small source are invariant under their source unitary groups. The Heisenberg self-dual lattice lines on the two legs correspond to scalar multiples of the characteristic functions of the lattices in this model. In particular the evaluations on the given vacua u_A, u_B are nonzero. Normalize each evaluation by its value on the corresponding vacuum. The standard Schrödinger formulas give the two left P-characters of $f_D(h) = e_D(\omega_D(1,h)u_D)/e_D(u_D)$, $D = A, B$. They are

$$\chi_A(t) = \beta_v(\det t)|\det t|_L^{a/2}, \qquad \chi_B(t) = \alpha_v(\det t)|\det t|_L^{a'/2}.$$

In fact the Siegel unipotent acts at zero by the value one of its quadratic phase. The Levi acts by the indicated determinant character and absolute Jacobian factor. Both f_D are right K-invariant and $f_D(1) = 1$. Their product of left characters is precisely δ_P.

For Haar probability on K define

$$\ell_D(\phi) = \int_K \frac{e_D(\omega_D(1,k)\phi)}{e_D(u_D)}\,dk, \qquad \Omega_D(h) = \ell_D(\omega_D(1,h)u_D).$$

These are integrals over finite quotients on smooth vectors. The compact integral pairing $\int_K f_A(k)f_B(k)\,dk$ intertwines the right actions on the two induced representations, since the product of left characters is δ_P. Indeed this is the integral on $P\backslash H$ of the density of left weight δ_P. Its transformation law follows by changing variables in the Iwasawa formula. It follows by moving the right translation to the other argument that $\Omega_A(h) = \Omega_B(h^{-1})$. The inverse of a Cartan double coset KhK equals KhK, since the Weyl group acts by independent inversions of the uniformizer coordinates in the even-dimensional quasi-split unitary group. Thus $\Omega_A = \Omega_B = \Omega$, with $\Omega(1) = 1$. The two normalized spherical coefficients

For an arbitrary vector ϕ, put

$$F_{D,\phi}(h) = \frac{e_D(\omega_D(1,h)\phi)}{e_D(u_D)}.$$

Let $\chi = \alpha_v = \beta_v$. The images of the two maps $F_{D,\bullet}$ are precisely the evaluation representations $R(V_A, \chi)$ and $R(V_B, \chi)$ in the degenerate principal series, in the normalization of [47, §1]. To apply Proposition 1.4 there, the size of the Siegel Levi is m_v in our notation and the small source dimension is $a' < m_v$. The large source is isometric to the small source plus $m_v - a'$ hyperbolic planes. Indeed every self-dual Hermitian lattice for the unramified extension has an orthogonal integral basis with unit norms: find a vector of unit

norm in the nondegenerate reduction, split off its line, and repeat. Each norm unit can be made one using surjectivity of the norm on units. A hyperbolic plane also has a self-dual lattice, and the dimensions of V_A and of V_B plus those planes agree. The parity condition $\chi|_{k_v^\times} = \epsilon_{L/k_v}^{a'}$ holds on both sources. Yamana's Proposition 1.4(1),(2) identifies $R(V_B, \chi)$ with the irreducible small-source coinvariant module. Parts (4),(5) give a surjection $R(V_A, \chi) \twoheadrightarrow R(V_B, \chi)$. The full large-source coinvariant module has only one irreducible quotient by local Howe duality. Its maximal kernel is thus the kernel of the composed surjection. [47, Proposition 1.4(1),(2),(4),(5)] [98, Theorem 1.2]

For each leg send a smooth vector ϕ to $h \mapsto \ell_D(\omega_D(1,h)\phi)$. On the small leg this is a nonzero map from the irreducible $R(V_B, \chi)$, since $\Omega(1) = 1$. On the large leg, pair $F_{A,\phi}$ against $R(V_B, \chi)$ using the compact integral above. The pairing is nonzero on f_A, f_B, giving 1. Thus the resulting map to the smooth contragredient of $R(V_B, \chi)$ is onto and has the unique maximal kernel. Evaluation of this dual pairing on right translates of f_B separates the quotient, since the translates of the nonzero vector f_B span $R(V_B, \chi)$. In particular the large coefficient map has precisely that same kernel. Both coefficient maps now have irreducible image and send their prescribed vacua to the function Ω. Hence the images agree as the span of the right translates of Ω. This normalizes the two maps in the exact spherical model on all smooth inputs. The restricted tensor product uses this common vector at every inert place. The common coefficient representation

For the global square-integrable comparison, the resulting local kernel is killed by the theta map into L^2. With the other factors fixed, this local map factors through the admissible finite-length source coinvariants. Its image carries the positive invariant quotient form induced by L^2. Such a smooth finite-length module is semisimple. Indeed for any stable submodule M and any vector fixed by a compact open J, orthogonally project the vector to the finite-dimensional space M^J. By J-averaging the remainder is orthogonal to all of M. This gives an invariant direct complement. The unique irreducible quotient and local multiplicity one now identify the square-integrable image as the same quotient or zero. This also proves the required vanishing of the maximal local kernel at the restricted spherical place.

2.3.15 Growth bounds for the triangle tests

The source integrals and the constant terms below use an absolute kernel estimate uniform in the real vector for fixed finite data. Write p, q for the two source real Witt indices. The imposed rank gap is $e = (a - a')/2 > 2(p + q) + 1$. For two Schwartz tests on a Schrödinger space use the Heisenberg coefficient

$$B_{\Phi_1,\Phi_2}(z) = \int (\rho(z)\Phi_1)(u)\overline{\Phi_2(u)}\, du.$$

Identify the rational symplectic space with the E-matrix space of size n by d, where n is the source rank. Expanding the product of the two rational theta sums and applying Poisson summation in one variable expresses it as the sum of the coefficients over rational z, with the fixed additive covolume factor. Indeed, organize the double sum by the difference of its two indices and expand the sum in the second index by Fourier inversion. The central Heisenberg phase for rational indices is trivial. Weil covariance replaces the coefficient of the acted tests at z by that of the original tests at $g^{-1}z(h^{-1})^{\mathrm{t}}$. For fixed finite data the finite coefficient is supported in a compact set, since it is an integral of compactly supported locally constant functions against an additive character. Its real coefficient has Schwartz decay of any order, bounded by a product of continuous Schwartz seminorms of the real tests. This follows by exchanging derivatives and polynomial weights in the displayed integral, or by taking its partial Fourier transform in the integration variable. The rational coefficient majorant · The Heisenberg coefficients

For unrestricted tests use the original test in both inputs. For the restricted source at v_1, first write its local test as a finite sum supported in compact open patches D_i of the split full-rank matrix locus with $D_i + D_i$ still full rank. Such a patch about Y exists because $2Y$ has full rank and full rank is open. For one patch reflect the adelic test $u \mapsto \Phi_1(-u)$ at all places in the second input. It has the same theta sum, since reflection commutes with the Weil action and permutes the rational summation indices. In

the split model the Heisenberg translation in one component of z shifts the matrix variable, and the other component gives a phase. Nonzero support of the product in the coefficient implies that the translation belongs to $D_i + D_i$, since the second test is reflected. The translated component therefore has full row rank. This implies full row rank of the rational matrix in the sum, since the matrices multiplying it in Heisenberg covariance are invertible. Bound the original restricted test by summing the bounds of its finitely many patches. The restriction here is at the actual split place used in the cone calculation. Parity in the restricted square bound · The full-rank support patches

Use adelic Siegel sets on the source and the target through Borel's reduction for reductive groups over a number field (§§3.4–3.5, 5.1–5.2, and 5.6). In rational Witt bases write their nonnegative decreasing diagonal parameters as $s_1, \dots, s_l$ and $t_1, \dots, t_{l'}$, repeated with the opposite signs and zeros in the full bases. Bounded shifts in the chambers affect the estimates by constants. Moving compact positive unipotent factors to the right of the diagonal keeps them in a bounded set by contraction on the positive chamber. The parameters act by the same positive scales at all infinite embeddings of the number field, with the remaining factors in fixed compact sets. Let $S = \sum_i s_i$, $T = \sum_j t_j$, and $D = [E : \mathbb{Q}]$. Absorb the compact factors and the rational basis changes into the constants and the fixed fractional ideal containing the matrix entries in the coefficient sum. Each entry has dilation $e^{w_{ij}}$, with w_{ij} the sum of its signed source and target parameters. For any fixed fractional ideal $\mathfrak{a}$ in E, lattice counting in its Minkowski embedding gives

$$\sum_{x \in \mathfrak{a}} (1 + e^{w} \|x\|_\infty)^{-N} \ll_N \max(1, e^{-Dw}) \qquad (N > D).$$

This follows by counting points in balls using a bounded fundamental parallelepiped and summing over dyadic annuli. Apply a sufficiently large coefficient seminorm and split its decay among the entries. The negative weights satisfy $\sum_{i,j} \max(0, -w_{ij}) \le dS + nT$. Thus the kernel is bounded by $C \exp(D(dS+nT)/2)$.

In the supported sum, the row whose source dilation is e^{s_1} has some nonzero rational entry. The norm of a nonzero element of the fixed fractional ideal has a positive lower bound in the Minkowski embedding. After dilation this costs at least a constant times $e^{s_1 - t_1}$ in the real coefficient, interpreting $t_1 = 0$ if there is no target parameter. Reserve an additional Schwartz power to obtain an extra factor $e^{-M \max(0, s_1 - t_1)}$ in the bound for the kernel, for any prescribed M. This gives arbitrary source decay at the expense of a larger target exponent because $S \le l s_1$ and $t_1 \le T$. The anisotropic source case has an empty parameter list and a compact quotient. Rational diagonal weights · The source and target Siegel estimates

For the larger source the Siegel density in the source parameters is, up to bounded factors,

$$\exp\left(-D \sum_{i=1}^{l} (n - 2i + 1) s_i \right).$$

The decreasing parameters give $\sum_i (n-2i+1)s_i \ge (n-l)S$. The hypothesis $d + 2l < 2n$ therefore makes the unrestricted bound integrable in the source parameters. It holds for the larger source of the triangle and for its target-only lowered lift, since $l \le p$ and $e > 2(p+q)+1$. For the smaller source the supported bound gives an integrable source bound by its arbitrary decay. This is also valid after a target-only reduction onto supported data.

The resulting source integral of the absolute kernel grows at most as $e^{C_1 T}$ for a fixed C_1. This is a polynomial height bound. To see that the reduction controls the target height even before reduction, take a nonzero rational leading isotropic vector x in the target Witt frame for a reduced representative $\gamma h = a(t)k$. The adelic product norm of $h^{-1}\gamma^{-1}x = k^{-1}a(t)^{-1}x$ is $O(e^{-Dt_1})$ by compactness of k. The rational vector $\gamma^{-1}x$ has product norm at least one. Multiplying by the matrix norm of h proves that e^{Dt_1} is bounded by a constant times the adelic height of h, in a compatible matrix height convention. Since $T \le l' t_1$, this gives the required polynomial bound. The fixed finite tests contribute only constants. On compact target sets the same bounds give an integrable source majorant times a continuous real seminorm. They yield continuity in the real Schwartz test and justify differentiation in the target real action. Every real derivative is another real Schwartz test with the same finite data. In particular the height exponent

works uniformly for all derivatives. A compact unipotent average preserves these estimates. Domination in the unrestricted source range · Uniform seminorm domination and integration

2.3.16 The triangle identity from a Levi coefficient

Write V_A, V_B, W for the three spaces at one triangle, put $d = a + a'$ and $e = (a - a')/2$, and write $G = U(W)$. At each real place the signatures, in the chosen form conventions, are $(a-p, p)$, $(q, a'-q)$, and $(d-p-q, p+q)$. The two Weil splittings use $(\beta, \beta/\alpha)$ and $(\alpha, \alpha/\beta)$, in the order source character, target character. In particular the integrated source representations are trivial in these splittings. The rank and weight bounds used below are $e > 2(p+q)+1$, $a' > 2(p+q)+1$, and $m_\alpha - m_\beta > d$ at all real places, where m_χ is the exponent of $z/|z|$ in the chosen embedding. Let $\mathscr{D}$ be the space of differences $T_A(\Phi) - T_B(\Phi')$ of the two theta integrals on $[G] = G(F)\backslash G(\mathbb{A}_F)$, with polynomial real and Schwartz finite data. Impose the prescribed support at v_1 on the B tests. The two source tests may be chosen independently. In particular $T_A(\Phi) - cT_B(\Phi')$ belongs to $\mathscr{D}$ for every complex scalar c. Triangle hypotheses and splitting characters · Space of differences

For a rational isotropic line L in W, let P_L be its stabilizer and N_L its unipotent radical. The line constant term f_L is the average over $[N_L]$ with mass one. With a chosen dual isotropic line the Levi is $\mathrm{Res}_{E/F}\mathrm{GL}_1 \times U(W_1)$, where $W_1 = L^\perp/L$. The two line coefficients have the unnormalized characters $\alpha(t)|t|_E^{a'/2}$ and $\beta(t)|t|_E^{a/2}$. The modulus of P_L on this factor is $|t|_E^{d-1}$, so in the normalized constant term the two exponents are $(1-a)/2$ and $(1-a')/2$, respectively. The coefficients are uniquely determined by f_L, by evaluating at two different norms, and are left N_L-invariant. For the latter assertion, compare the two evaluations after replacing the argument by ng. The torus normalizes N_L, so the constant terms are unchanged in both evaluations.

Suppose that for every matched polynomial pair the selected coefficient of $T_A(\Phi) - cT_B(\Phi')$ vanishes on the embedded $U(W_1)$. Here the selected coefficient is the β coefficient when $q > 0$, and the α coefficient when $q = 0 < p$. A right translate by the target real maximal compact and by an arbitrary finite target element is still a matched polynomial pair. The real and finite quotient maps intertwine these actions, the real compact preserves the polynomial space, and the smaller test support at v_1 is preserved by the target action. Right translation commutes with the line constant term and with projection onto each of its two characters. Thus the coefficient vanishes on $U(W_1)K$, where K is the real maximal compact times the finite adelic group. Adelic Iwasawa decomposition $G(\mathbb{A}_F) = P_L(\mathbb{A}_F)K$ and the covariance of the coefficient now give its vanishing at every argument. This uses the coefficient identity for all matched pairs, including the translates of the given pair. Iwasawa decomposition and stability of matching · Propagation from the Levi

We use three consequences of the mixed-model line expansion for the other coefficient. When $q = 0 < p$, the β coefficient vanishes on all of $\mathscr{D}$. The definite smaller source has only the zero orbit, which gives its α coefficient. The larger source's β coefficient comes from the zero orbit and is a lift from V_A to W_1. Its target has real negative index $p-1$, and $(d-2)+2p < 2a$. The real vanishing calculation below applies to this convergent lift. When $p = 0 < q$, the α coefficient vanishes on all of $\mathscr{D}$. The definite larger source has only the β coefficient, and the smaller source's zero-orbit lift to W_1 has target negative index $q-1 < q$. Use the restricted support at v_1 for this lift. If its target-only evaluation does not already have that support, one uses a representative of the upper test with the same integral whose evaluation has the lower support. This is the local quotient support calculation at v_1. At this place the source character is nontrivial, and source augmentation relations within the prescribed support integrate to zero.

In the split matrix chart at v_1 a supported test is a Schwartz function on the open full row-rank locus. In coordinates for the removed plane the target-only evaluation is, up to its fixed nonzero measure factor, the map

$$E_0\phi(Y_1) = \int_{F_{v_1}^n} \phi(0, Y_1, z)\, dz.$$

The first coordinate is set to zero by the trivial Fourier character, and the dual coordinate is integrated by the partial Fourier transform. The image of full-rank tests is the Schwartz space on matrices Y_1 of rank at least $n-1$. Indeed for every such Y_1 one can choose a column z completing it to rank n. Small compact open boxes about these matrices give product charts lifting test functions, and a compact open partition handles every Schwartz function on this locus. Source actions intertwine the evaluation, since the Jacobian in z changes the source determinant power from $d/2$ to $(d-2)/2$. In the rank $n-1$ quotient the source stabilizer has arbitrary determinant. The nontrivial unitary source character on a uniformizer rules out its twisted coinvariant by the rank argument below. Hence a finite sum of source augmentation terms can be subtracted from $E_0\phi$ to obtain full-rank support. Lift every term of that subtraction to the full-rank upper test space using the surjectivity just proved. Subtract the corresponding source augmentations above, which preserve the upper integral since the tests and their translates lie in the supported subspace. The adjusted upper test has the required lower support after evaluation. Lowering the split support

When $p, q > 0$, choose a plane removal on the larger source and the target, keeping V_B fixed. For every $k \in K$ the restriction of the α coefficient at $h_1 k$ to $h_1 \in U(W_1)(\mathbb{A}_F)$ belongs to the lower difference space, with ranks $a-2, a'$ and indices $p-1, q$. The nonzero orbit on the larger source gives its cone evaluation. The cone integral converges at e, preserves polynomial real data through its compact average, and sends a restricted tensor to a lower restricted tensor times a scalar. The zero orbit on the smaller source gives its target-only evaluation with the same support representative just described. All these assertions concern separately chosen tests on the two sides. The rank and weight inequalities persist after this larger-source removal. The two absent coefficients and the lower difference · Changing the representative within the smaller support

For the assertion about the real cone, use the orthogonal symplectic decomposition into the lower tensor space and its complement. The adapted Cartan involutions preserve this decomposition. In compatible Fock coordinates, the Gaussian is a product and each polynomial vector is a finite sum of products of polynomials in the two sets of variables. The removed-coordinate shift acts on the complement. The mixed evaluation commutes with the lower Heisenberg action, so on each product it preserves the lower vector up to the scalar of the complementary test. Indeed a continuous operator commuting with the real Schrödinger phase multipliers has distribution kernel equal to a multiplication operator on the diagonal. Commuting also with translations forces that multiplier to be constant. The compact source orbit of the upper polynomial lies in the space of polynomials of bounded degree times the Gaussian, which is finite dimensional. All the reduced shifted vectors therefore lie in a fixed finite-dimensional space of lower polynomials. Absolute integration in the cone preserves this space. The finite local integrals are Schwartz by compact support in the evaluated variables, with the unchanged lower vacuum outside a finite set after its scalar has been factored out. The absolute convergence of the idele product at $e \geq 2$ permits summing the reduced tensors. This gives the lower polynomial tests used in the induction. Finite-dimensional lower polynomial space

The real weight obstruction in the zero orbit. Consider a convergent theta integral from signature (P, Q) to (R, S) with $\min(R, S) < \min(P, Q)$. The determinant splitting on the source has exponent m, with $|m| > R + S$. Fixing the finite test and the target point gives a continuous functional on the real Schwartz space, invariant under the real source. Restrict to the Fock polynomials at one real place and fix the vectors at the other places. For a rank-one trace-one matrix vw^{t} in one compact source block, the complexified infinitesimal action, after the constant term is removed, is the derivation

$$D = \sum_{i,j,c} v_i w_j \begin{cases} x_{ic}\partial_{jc}, & c \text{ has the same sign as the block,} \\ -x_{jc}\partial_{ic}, & c \text{ has the opposite sign.} \end{cases}$$

The shift removed from this operator is half the difference between the numbers of same-sign and opposite-sign target columns. Source invariance says that the functional intertwines D with the scalar $-m/2$ minus

that shift. This scalar has sign opposite to m. If it is positive, choose the block with fewer same-sign target columns. For an arbitrary linear polynomial ℓ, choose $w \neq 0$ annihilating the coefficient columns of ℓ of that sign, and choose v with $w^{\mathrm{t}}v = 1$. Such w exists because the source block has larger dimension. The action of D on the space generated by the orbit of ℓ has weights $0, -1$. The powers of ℓ have only nonpositive weights, so the functional kills them. For a negative required scalar, choose the block with fewer opposite-sign columns and exchange the roles of v, w. The weights of the corresponding powers are then nonnegative. Powers of linear forms span the homogeneous polynomials by polarization. Density of the polynomial vectors in the real Schwartz space gives the asserted vanishing. This calculation uses the actual source splitting on each lower zero-orbit lift, whose exponent is $m_\alpha - m_\beta$ or its negative. Rank-one action on powers of linear forms · Passage to the invariant Schwartz functional

Iteration of an α constant term. Suppose $q > 0$ and the full line constant term of a difference has only its α coefficient. Witt's extension theorem and left rational invariance move this property to any rational isotropic line. When choosing a plane adapted to a new lattice frame, conjugate the real Cartan involutions and right-translate the target argument accordingly. On the source, a right change of variables in the integral carries the Fock data to the conjugate Cartan model. The finite support condition stays at the same place. Thus the resulting function is a difference of the same kind and retains its one-character line constant term. With $p = 0$, the α coefficient is zero by the calculation above. If $p > 0$, take the larger-source removal. For each $k \in K$ write $f_{1,k}$ for the resulting lower difference on $U(W_1)$.

Choose a second rational isotropic line in W_1. The two successive averages equal the average over the radical of the parabolic of the flag formed by the first line and their two-plane U. That radical is also N_U followed by the upper triangular unipotent of the $\mathrm{GL}_2(\mathbb{A}_E)$ factor in the Levi of the two-plane. Both decompositions are semidirect products over F, so Fubini with probability quotient measures identifies the averages. Write $C(g)$ for this flag average. Fix all other data, and let t be the idele equal to 2 at the complex places and 1 at finite places. Let a_2 act on the two-plane as $\mathrm{diag}(t, t^{-1})$. The upper α law and the two lower laws imply

$$C(a_2^i g) = A(g) + \lambda^i B(g), \qquad \lambda = \alpha(t)\beta(t^{-1})|t|_E^{(a'-(a-2))/2}, \qquad |\lambda| < 1 \quad (i \in \mathbb{Z}).$$

This identity holds for every g. Indeed, write $g = nm(s)j(h_1)k$ in the first Iwasawa decomposition. The outer unipotent drops out of the first average, and the lower function is $\alpha(s)|s|_E^{a'/2}R(h_1)f_{1,k}$, with R denoting right translation. Apply the two lower laws to its line average. The α factors at t, t^{-1} cancel and the two different norm powers give the displayed λ. In the rank-one argument below the same result holds on every right translate, so $C(a_2 g) = C(g)$. The relation $(\lambda - 1)B(g) = 0$ gives vanishing of the lower β coefficient, at every argument. Thus each $f_{1,k}$ has only its α coefficient. Induction on p, applying line cuspidality and orthogonality below at each stage, makes $f_{1,k} = 0$. Iwasawa decomposition then gives a zero full line constant term of the upper difference. The induction applies to the entire difference space and needs no equality on lower matched pairs for these auxiliary removals. Two successive constant terms · Rank-one restriction of the flag average · Induction on the larger-source index

Here is the rank-one argument in a form that also applies to the rank-two Fourier calculation below. Let F_2 be a smooth, finite-level function on $\mathrm{SL}_2(E)\backslash\mathrm{SL}_2(\mathbb{A}_E)$ of uniform moderate growth, meaning the same height exponent works for all its real right derivatives, with constants depending on the derivative. Suppose its upper unipotent average C_2 obeys $C_2(a_2^i x) = A_2(x) + \lambda^i B_2(x)$ with $0 < |\lambda| < 1$. All right derivatives of the average have bounded growth on the sets $a_2^i K_2$ for $i \geq 0$ and compact K_2. To see this, differentiate the recurrence, or express its two coefficients in terms of the values at $1, a_2$; compactness bounds those values uniformly on K_2. Rank-one adelic reduction covers the quotient by $a_2^i K_2$, $i \geq 0$, for one compact K_2. This is the rank-one Siegel set, with the norm-one diagonal directions reduced modulo units and the remaining positive radial direction discretized by powers of 2.

For this compact K_2 and the fixed finite level choose a compact set of representatives of $E\backslash\mathbb{A}_E$ with finite parts in a sufficiently small compact open additive subgroup. Additive approximation gives these

representatives, with bounded archimedean parts obtained modulo a fractional ideal. The finite parts act trivially at the chosen level after conjugation by K_2. The remaining, archimedean unipotents conjugate across a_2^i by multiplication with 2^{-2i}. Thus comparison with the unipotent average by the mean value estimate bounds a derivative of F_2 at $a_2^i k$ by a constant plus 2^{-2i} times the bound for its first further derivatives on $a_2^i K_2'$, with K_2' compact. The estimate holds for every compact K_2 and all derivatives, since real differentiation preserves the finite level. If every derivative is $O(2^{Mi})$, this improves their common exponent to $\max(0, M-2)$. After finitely many applications F_2 is bounded. Its average is then bounded for negative powers also, forcing $B_2 = 0$. With $C_2 = 0$ the same argument gives boundedness directly. In the application, F_2 is the restriction of the N_U-average of a fixed right translate of f to the embedded $\mathrm{SL}_2(\mathbb{A}_E)$. The average is invariant under $\mathrm{SL}_2(E)$, which normalizes the rational unipotent. The theta estimates give uniform moderate growth because for fixed finite data the exponent is independent of the real Schwartz vector, and real derivatives act on that vector. The compact unipotent integral and the algebraic group embedding preserve this condition. Growth reduction for the rank-one unipotent average · The adelic rank-one reduction · Uniform moderate growth of the tested theta integrals

2.3.17 Line and parabolic constant terms

Let $f \in \mathscr{D}$ have zero constant term along one rational isotropic line of W. Witt's theorem and rational invariance give the same vanishing for every line. For a rational isotropic j-plane U, write f_U for the constant term along the parabolic stabilizing U. Choose a dual frame for U, and write $m_U(s)$ for the resulting embedding of $\mathrm{GL}_j(\mathbb{A}_E)$ in the Levi. Proceed by induction on j. For each target argument h the function $s \mapsto f_U(m_U(s)h)$ is left $\mathrm{GL}_j(E)$-invariant. Its constant terms along the proper maximal parabolics stabilizing prefixes of the chosen frame are zero. Indeed the average over N_U and then the indicated rectangular upper unipotent is the average over the radical of the flag parabolic. One can take the average for the smaller isotropic plane first in the same radical. All quotient measures here have mass one. Consequently the lower dimensional vanishing supplies all those constant terms.

An invariant function on $\mathrm{GL}_j(\mathbb{A}_E)$ with these zero averages is determined by its generic Whittaker integrals on all right translates. For this implication, expand successively along the upper unipotent column groups of dimensions $j-1, j-2, \dots, 1$. At each stage rational $\mathrm{GL}_r(E)$ acts transitively on nonzero characters of the vector quotient $(E\backslash\mathbf{A}_E)^r$. Fix the representative nonzero on its last entry. The zero Fourier term at stage r is an integral of the constant term for the rectangle $(r, j-r)$, with the already selected characters on the remaining upper triangular group. It vanishes by the previous averages. Thus if the last integral, with a fixed nontrivial character on successive simple roots, is zero on every translate, Fourier uniqueness on each compact vector quotient gives vanishing backwards through the expansions. This uses only uniqueness of Fourier coefficients of continuous functions on compact quotients. Whittaker testing of constant terms

We show that the required Whittaker integrals are zero on f_U. Choose a finite place v split in E and write $k = F_v$, ϖ for a uniformizer and Q_v for the residue cardinality. At a chosen factor of $E_v = k \times k$, choose the split Schrödinger chart with source rank n and matrix variable $Y \in \mathrm{Mat}_{n,d}(k)$. A local functional ℓ is obtained by varying the v-component of a fixed tensor and taking its source theta integral, the N_U-average and the Whittaker integral. The global integrals are linear on this test space, and the unipotent integrations are compact. For the smaller source take $v \neq v_1$, so this variation keeps the prescribed support. In the split chart the source and target actions are

$$\omega(g,1)\phi(Y) = \chi_W(\det g)|\det g|_v^{d/2}\phi(g^{\mathrm{t}}Y), \qquad \omega(1,h)\phi(Y) = \chi_X(\det h)|\det h|_v^{n/2}\phi(Yh),$$

where χ_X, χ_W are the two characters of this source in the chosen splitting. Invariance under the source and covariance of the Whittaker integral give the laws

$$\ell(\phi(AY)) = \xi(\det A)\ell(\phi), \qquad \xi = (\chi_W|\cdot|_v^{d/2})^{-1},$$
$$\ell(\phi(Y(1+e_i u^{\mathrm{t}}))) = \psi_v(u_{i+1})\ell(\phi), \qquad u_l = 0 \text{ for } l \leq i, \quad 1 \leq i < j.$$

Here the frame occupies the first j columns. A column shear has determinant one, preserves U, is trivial on $U^{\perp}/U$, and induces the indicated upper triangular unipotent on the frame. On the other split component it acts by inverse transpose, which fixes the dual-component part of U. This proves the displayed laws also when the shear changes columns outside the frame. Conjugating to an arbitrary rational frame gives the same laws in its coordinates. Split actions and determinant powers · The global functional in a split chart

Write y_i for the columns of Y. The first shear law forces ℓ to be supported on dependence of y_1, y_2. At a matrix with these columns independent, take w with $w(y_1) = 0$ and $w(y_2) = 1$, and choose t with $\psi_v(t) \neq 1$. The operation of left multiplication by $1 - t y_1 w$ combined with right multiplication by $1 + t e_1(wY)$ fixes this matrix and has character $\psi_v(t)$ on the functional. The matrices in the operation are fixed at the chosen point. Both multiplications are unipotent and commute. Their combined linear action preserves a lattice, as is seen by taking a basis adapted to the resulting nilpotent operator. Arbitrarily small compact open balls about the point are therefore preserved. The value of the functional on each such ball is zero. A finite disjoint compact open refinement on the support gives the claimed vanishing away from the dependence locus. Equivalently, the functional factors through restriction to this closed locus. For locally constant compactly supported functions, restriction to a closed subspace is onto by compact open extension, and its kernel consists of functions supported in the complement.

When $j \geq 3$, on the dependence locus with $y_1 \neq 0$ write $y_2 = s y_1$. Adding $t y_2 - s t y_1$ to the third column fixes the matrix and has character $\psi_v(t)$ by the second shear law. The compact open argument kills this stratum. On the remaining closed locus $y_1 = 0$, adding $t y_1$ to the second column fixes every point and again has nontrivial character. The functional is zero. Hence the rank at least three Whittaker integrals of each theta image vanish separately. The adjacent simple-root calculation

For $j = 2$ choose v outside all the finite supports of the tests, the frame denominators and the nonintegral places of the given argument h. Take it split with unramified splitting characters. There are infinitely many such places. The local vacuum is a scalar multiple of the integral matrix characteristic function. In the chosen chart the frame change and the v-component of h are integral with integral inverses, so right multiplication by them preserves that function. The element $b_i = \operatorname{diag}(\varpi^i, \varpi^{-i})$ at the chosen factor of E_v, with all other components equal to one, has determinant one in $\mathrm{GL}_2(\mathbb{A}_E)$. Its action on the vacuum moves the column bounds to $(-i, i)$, with no Weil determinant factor.

Here is the calculation for these bounds. Let $B(p_0, u_0)$ be the value of ℓ on the box with columns one and two in $\varpi^{p_0}\mathcal{O}_v^n$ and $\varpi^{u_0}\mathcal{O}_v^n$, respectively, and all other columns integral. Put $D(p_0, u_0) = B(p_0, u_0) - B(p_0 + 1, u_0)$ and $R = \xi(\varpi) Q_v^{n-d}$. Let c_v be the least integer such that ψ_v is trivial on $\varpi^{c_v}\mathcal{O}_v$. Then

$$\begin{aligned} B(p_0, u_0) &= 0 && (u_0 - p_0 < c_v), \\ D(p_0, u_0) &= R D(p_0 + 1, u_0) && (u_0 - p_0 - 1 \geq c_v), \\ D(p_0, u_0) &= Q_v D(p_0, u_0 + 1) && (u_0 - p_0 \geq c_v). \end{aligned}$$

For the first equality take a shear of column two by column one with parameter in $\varpi^{c_v - 1}\mathcal{O}_v$ and nontrivial character. It preserves the box. For the other equalities partition the first-column shell by the first row whose valuation is p_0. The corresponding test with all other entries of the first column at depth $p_0 + 1$ occurs with weight Q_v^{n-i} for pivot row i, by row shears distributing the later entries among the residue classes. These shears preserve all the other column bounds, and have determinant one. Scale the pivot row by ϖ. In the new box the pivot is at $p_0 + 1$, the other first-column entries still have depth $p_0 + 1$, and the other entries of the pivot row each gain a depth. Restricting the other first-column entries to depth $p_0 + 2$ gives a factor Q_v^{n-1} by the same row shear count. Loosening the extra depth in the pivot row in the other columns contributes $Q_v^{-(d-1)}$ in this expression, by the column shears. The shear affecting column two has a parameter of depth at least $u_0 - p_0 - 1$, so its character is one. Together with the determinant factor this proves the second equality. To prove the third, loosen the column-two entry of the pivot row from depth $u_0 + 1$ to u_0 using the column shear with parameter of depth $u_0 - p_0$. This contributes Q_v. On

the dependence locus the other column-two entries automatically have depth at least $u_0 + 1$ because the other first-column entries have depth at least $p_0 + 1$. Thus the resulting pivot test agrees with the one in $D(p_0, u_0)$ on the support of ℓ.

At $u_0 - p_0 = c_v$ one has $D(p_0, u_0) = B(p_0, u_0)$. The last two recurrences give $D(u_0 - c_v, u_0) = Q_v R D(u_0 + 1 - c_v, u_0 + 1)$. Sum the shells, using the second recurrence at each depth down from $u_0 - c_v$. The resulting value at $(-i, i)$ is a scalar multiple of

$$y_i(R) = (Q_v R)^{-i} \sum_{j=0}^{2i-c_v} R^j,$$

where the sum is zero for $2i < c_v$. For the two sources the parameters are $r_v = \nu_v(\varpi) Q_v^e$ and r_v^{-1}, where $\nu = \alpha/\beta$. The relation $y_i(r_v^{-1}) = r_v^{c_v} y_i(r_v)$ holds by reversing the sum, and $|r_v| = Q_v^e > Q_v$. Hence $|y_i(r_v)|$ tends to infinity. The same profile holds for finite sums of pure tensors and for their difference, since the same split place and frame can be used. In particular the Whittaker coefficient along $b_i h$ is $C_h y_i(r_v)$, for a scalar C_h independent of i. The two shell recurrences and their boundary · The reciprocal parameters at the same split place

On the other hand $s \mapsto f_U(m_U(s)h)$, restricted to $\mathrm{SL}_2(\mathbb{A}_E)$, has zero upper unipotent average by the line case and the same flag calculation. It is bounded on the $\mathrm{SL}_2(E)$ quotient by the rank-one argument above, which applies since it is smooth of uniform moderate growth and finite level. Its Whittaker integrals have absolute value bounded by the same bound, with mass-one unipotent measure. It follows that $C_h = 0$, and taking $b_0 = 1$ gives the required vanishing at h. This completes the induction on j. Every proper rational parabolic of $U(W)$ stabilizes a nontrivial isotropic flag. Its constant term is an iterated average including one of these plane averages. Thus f is cuspidal. Induction on the Witt index of the averaging plane

2.3.18 Orthogonality of cuspidal differences

Let $f \in \mathscr{D}$ be cuspidal. It is smooth, of finite level, and of uniform moderate growth. The last assertion follows by applying the theta kernel estimates to derivatives of the real Schwartz vector with its finite tensor unchanged. Cuspidality and uniform moderate growth give rapid decrease on Siegel sets. In the usual proof, a zero unipotent mean bounds the value by derivatives along the radical by integration on its compact quotient. Conjugating these derivatives across the positive radial part of a Siegel set gains an inverse power of a large simple-root coordinate. Choose the radical of the maximal parabolic for that coordinate. Each root of this radical contains that simple root. Applying the same estimate to the derivatives and using their uniform moderate growth gives any inverse power. The finite level lets one take the finite unipotent coordinates in a compact open set acting trivially. The bounded parts of the Siegel sets and the finitely many cusp representatives only change the constants. Thus the rapid decrease is a bound in the quotient height, the infimum of the adelic heights of rational translates. This also gives integrability after multiplying by an automorphic function with a fixed polynomial height bound. Growth and rapid decrease on the triangle

For $X = V_A$ or V_B, put $H_X = U(X)$ and define

$$A_X(f; \Psi)(g) = \int_{[G]} f(h) \overline{\theta_X(\Psi; g, h)}\, dh.$$

For fixed g the kernel has polynomial height growth in h, locally uniformly for g in compact sets. Rapid decrease gives absolute convergence, continuity and rational left invariance of A_X. The theta bounds for the two tested source integrals give a polynomial height bound on the source integral of the absolute kernel. Thus Fubini yields, with the given source Haar measure,

$$\langle T_X(\Psi), f \rangle_{[G]} = \int_{[H_X]} \overline{A_X(f; \Psi)(g)}\, dg$$

for the tested Ψ. We will show that for every Schwartz test the adjoint function is constant on $[H_X]$. Fubini for the adjoint pairings

Work at a finite place $v \neq v_1$, outside the ramification set, whose underlying rational prime splits completely in E. The value $v_v(\varpi)$ at either split factor is different from 1. In fact, for the corresponding degree-one prime $\mathfrak{w}$ of E choose b with $\mathfrak{w}^h = (b)$. The (2h)-th power of the unramified value, by the Hecke character relation on b, is, up to inversion, $\prod_\tau (\tau(b)/\overline{\tau(b)})^{m_\alpha(\tau)-m_\beta(\tau)}$. All embeddings are in the Galois field E. The conjugate prime ideals of $\mathfrak{w}$ are distinct, and all the displayed exponents are nonzero. Taking a valuation at a prime in this product rules out a root of unity. This argument applies at every such place, even though the coefficients of f need not be vacuum vectors there. Nontrivial uniformizer values at degree-one places

We use two facts about the split matrix model. First, on $\mathrm{Mat}_{N,d}(F_v)$ a functional with right covariance $\ell(\phi(\cdot h)) = \eta(\det h)\ell(\phi)$, where $\eta(\varpi) \notin \mathbb{R}_{>0}$, vanishes on each stratum of rank less than d. For a matrix of rank r, the orbit of right translation is the space of surjections onto its image, with stabilizer containing a GL_{d-r} on its kernel. The stabilizer therefore contains an element of determinant ϖ. On a transitive orbit, a relatively invariant distribution is possible only if the character on the stabilizer equals a ratio of positive Haar modular characters. One proves this by pulling the distribution back to the group by averaging test functions along stabilizer fibers, where the only relatively invariant functional is weighted Haar integration. On the rank stratum the image varies in a Grassmannian, and local choices of a frame give product charts with the orbit as fiber. Compactly supported locally constant test functions in each chart are finite sums of products. The stabilizer criterion therefore applies on the whole stratum. The filtration by closed rank bounds proves the vanishing, using restriction of Schwartz functions to the closed stratum. On square matrices the only surviving contribution is the open orbit GL_d. Its functional is a multiple of Haar integration weighted by $\eta(\det h)^{-1}$. In particular left multiplication on the test variable has the same covariance as right multiplication. The rank filtration for the local distribution

Second, the analogous source coinvariant on n-by-d matrices is represented by full row-rank tests. On the other strata the left stabilizer contains a determinant-ϖ element, and the same argument applies to $\xi = (\chi_W|\cdot|_v^{d/2})^{-1}$. On the full-rank stratum, integrating the test over the fibers of the row-space map with weight ξ^{-1} gives the coinvariant. Indeed each fiber is freely transitive under $\mathrm{GL}_n(F_v)$, and on compact open charts weighted Haar integration is the one-dimensional twisted coinvariant on that fiber. Local sections of the row-space map and a compact open partition lift every smooth section over the Grassmannian. With $Y_0 = (I_n\ 0)$, the resulting map in the target group model is

$$\phi \longmapsto \left[h \longmapsto \chi_X(\det h)|\det h|_v^{n/2} \int_{\mathrm{GL}_n(F_v)} \phi(AY_0h)\xi(\det A)^{-1}\, dA \right].$$

For a block lower triangular element with diagonal blocks (A, D) the multiplier is

$$(\chi_X/\chi_W)(\det A)\chi_X(\det D)|\det A|_v^{-(d-n)/2}|\det D|_v^{n/2}.$$

The two powers form the square root of the parabolic modulus. The two sources thus give the normalized inductions with unitary determinant characters α on the block of size a and β on the block of size a', in the two block orders. Each inducing representation is irreducible and unitary, so Bernstein's irreducibility theorem applies. The two irreducible inductions have the same class by commutativity of normalized parabolic induction in the Grothendieck ring of the general linear groups. Denote their common local module by P_v. [82, §8.2, Corollary (a)] [83, Theorem 1.9] Using the common split quotient in a cuspidal difference

For this place fix the other tensor factors in a finite expression for f as a difference. The local orbit maps of the two source integrals factor through P_v by source invariance. Each image is simple or zero, and their finite sum is a semisimple isotypic module. The submodule generated by the cuspidal vector f is therefore a finite sum of copies of P_v inside the space of rapidly decreasing functions. Consequently f

is a finite sum of values of local equivariant maps from the Schwartz space for either source, each with rapidly decreasing outputs. This step allows one to form the adjoint pairing with the opposite source locally for every varied test, with convergence for each of those pairings.

For A_X use the equivariant maps from the other source X', of rank $n' = d - n$. Fixing all places of the adjoint test except v, the local pairing is bilinear after conjugating the test for X. Stack the n-by-d and n'-by-d matrices to view it as a functional on square matrices. Simultaneous target right translation preserves the pairing. The Weil target scalars on the two tests are $\overline{\chi_X}|\cdot|_v^{n/2}$ and $\chi_{X'}|\cdot|_v^{n'/2}$. After removing these scalars, the functional has right covariance $\eta = (\chi_X/\chi_{X'})|\cdot|_v^{-d/2}$. Since $\chi_X/\chi_{X'} = \chi_W$ is unitary and nontrivial on the uniformizer, $\eta(\varpi) \notin \mathbb{R}_{>0}$. By the square matrix calculation the same left covariance holds on the first block, the other block fixed. This cancels exactly the conjugated source factor $\overline{\chi_W}|\cdot|_v^{d/2}$. Right translation of g by any element of $H_X(F_v)$ consequently leaves A_X unchanged. The square matrix calculation uses exactly d rows from the two blocks together. Cancellation of the adjoint Weil factors

The rational group and the groups at these degree-one places generate a dense subgroup of $H_X(\mathbb{A}_F)$. On $\mathrm{SU}(X)$, strong approximation off one such place applies because the group is simply connected, absolutely almost simple, and split and noncompact there. For the determinant, use a section of the norm-one torus S obtained by scaling a rational anisotropic line in X. The compact group $[S]$ is generated densely by the same local determinant groups. Otherwise there would be a nontrivial continuous character trivial on them. Its pullback by adelic Hilbert 90, $x \mapsto x/\bar{x}$, would be a nontrivial unitary Hecke character ω of E trivial at the degree-one places outside a finite set. The quotient $\zeta_E(s)/L(s,\omega)$ then has an absolutely convergent, nonzero Euler product near $s = 1$, since the remaining infinite set consists of primes of norm at least the square of their rational prime. This contradicts the simple pole of ζ_E and holomorphy of $L(s,\omega)$ there. This proves density on the determinant as well. Local right invariance implies local left invariance because conjugation by a fixed adelic element preserves each single-place subgroup. Continuity and rational left invariance now make $A_X(f;\Psi)$ constant. [84, Theorem 7.12] [32, §4] Strong approximation and the norm-one determinant · The degree-one determinant argument

Eliminating the adjoint pairings. Suppose first that X is isotropic. The rank-one constant-term formula for a theta lift of a cusp form shows that the constant term of $A_X(f;\Psi)$ along an isotropic line of X transforms under its idele torus by $\overline{\chi_W(t)}|t|_E^{d/2}$, in the conjugate Weil splitting. In the mixed model the central unipotent integral retains isotropic vectors of W. The nonzero orbit unfolds against the constant term of f along the stabilizer of an isotropic line, and its integral vanishes by cuspidality. The zero vector gives evaluation at zero in the removed coordinates, on which the dilation has the indicated weight. This is the computation of the line case in Wu's Proposition 3.1, with the two spaces interchanged and the conjugate Weil representation used for the test. Its proof uses cuspidality and rapid decrease of the input to unfold the nonzero orbit, so it applies directly to the cuspidal difference. A constant function equals its mass-one line average and obeys this law only if zero. Thus $A_X = 0$, and the adjoint identity gives $\langle T_X(\Psi), f\rangle = 0$ on the tested data. [85, Proposition 3.1 and (3.4)]

Suppose instead that X is definite, in a nondefinite triangle. Its adelic quotient is compact. Suppose a tested vector Ψ_2 has $\langle T_X(\Psi_2), f\rangle \neq 0$. Then the constant $A_X(f;\Psi_2)(1)$ is nonzero by Fubini. On the doubled group $U(W \oplus (-W))$ use the Siegel–Weil section obtained from $\Psi_1 \otimes \overline{\Psi_2}$, with arbitrary polynomial Ψ_1. Its inducing character is χ_X and its normalized parameter is $s_0 = (n-d)/2 < 0$. The tensor product of the original Weil model and its conjugate cancels the source splitting character and identifies with the restriction of the doubled model after multiplication by $\chi_X(\det h')^{-1}$ on the second target factor. Products of the theta sums therefore give the doubled kernel, with this determinant correction, at $\iota(h,h')$. The anisotropic Siegel–Weil formula in the unitary case $n < d$ gives a holomorphic value at s_0, equal to twice

the doubled theta integral with source quotient measure of mass one. If $v_X = \mathrm{vol}([H_X])$ for the source measure in T_X, this gives

$$J_{s_0}(h) = \frac{2}{v_X} A_X(f;\Psi_2)(1) T_X(\Psi_1)(h),$$
$$J_s(h) = \int_{[G]} E(s, \iota(h,h'))\, \chi_X(\det h')^{-1} f(h')\, dh'.$$

The identification uses the compatible doubled splitting. Compactness of $[H_X]$, polynomial kernel bounds in the other variable and rapid decrease justify both integrations in the identity. The Eisenstein series is the usual flat extension of the section at s_0. Its polynomial real test is finite under the doubled maximal compact in the corresponding Fock model, and its finite data have finite level. [19, §3.1, Theorem 7.1 and §11.3] [86, Theorems 2.1 and 2.2] The definite adjoint identity and its measures

The functions J_s have zero constant terms. To see this first in the half-plane of absolute convergence, unfold the Siegel sum against f. An isotropic maximal subspace of $W \oplus (-W)$ transverse to the two copies of W is the graph of an isometry and gives the open orbit. For any other subspace M, the intersection $U = M \cap (0 \oplus W)$ is a nonzero isotropic subspace. Indeed the two intersections with the coordinate copies have equal dimension, since isotropy makes each projection into the perpendicular of the intersection on its side. The second projection is $U^\perp$, by the same dimension count. The unipotent radical stabilizing U acts trivially on $U^\perp/U$ and therefore preserves M. Its determinant on M is one. Its average in the unfolding is thus the zero parabolic constant term of f. The open orbit gives

$$J_s(h) = \int_{G(\mathbb{A}_F)} F_s(\iota(1,z))\, \chi_X(\det z)^{-1} f(hz)\, dz,$$

with F_s the inducing section. Averaging in h gives zero along each proper parabolic, again by cuspidality of f. The section on the open orbit is bounded in a right half-plane by a large negative power of the matrix height of z. In the graph coordinates, the Siegel height is the inverse of the Plücker norm of the graph. Its local factors are comparable to the products of the larger of 1 and each Cartan singular value, using both components at split places. These products bound the matrix norms of z and z^{-1}. Cartan integration on each local group gives polynomial volume growth for these heights, and the finite factors converge once the power is sufficiently large. The closed-orbit interchanges follow by majorizing the defining series with the positive series of the real part of the parameter and pairing its polynomial growth with rapid decrease.

Meromorphic continuation of Eisenstein series with values in the space of smooth functions of uniform moderate growth continues each constant term of J_s. Pairing with the rapidly decreasing f after a restriction along $\iota(h,h')$ is a continuous linear functional on each fixed growth space. This remains true after the additional compact unipotent average in h. The constant terms are consequently zero also at the holomorphic point s_0. The argument works with the actual (possibly non-quasi-split) unitary group in the two copies of the doubling embedding. It uses the Siegel parabolic of the split doubled space. [87, Theorem 2.3(1)] The open and nonopen doubling orbits · Continuation of the zero parabolic means

The nonzero adjoint constant would thus force every $T_X(\Psi_1)$ in the polynomial image to be cuspidal. Remove a plane from the other (isotropic) source and the target, leaving X unchanged. Assume the strong induction for the triangle identity has supplied a nonzero lower polynomial lift for this definite source. This is a smaller value of $p+q$. The target-only evaluation is onto on polynomial data and gives the sole line coefficient of a definite-source integral, with a nonzero geometric factor. A preimage of the lower nonzero lift therefore has nonzero line constant term above, a contradiction. This use of nonvanishing works at the $p = 0 < q$ endpoint of the auxiliary induction on p as well. Its further plane removal still has smaller total index than the triangle under consideration. It follows that the pairings to the definite-source image vanish on the tested data. Contradiction using the lower definite-source lift · The lower-stage nonvanishing premise available in the induction

Conclusion of the comparison. The two pairings of a cuspidal difference with its defining lifts vanish, so its squared Petersson norm is zero. Smoothness and the full support of the quotient measure give pointwise vanishing. Thus equality of the selected line coefficient on the embedded Levi for all matched pairs implies equality of the two theta integrals with that scalar on every matched polynomial pair. The converse is obtained by taking the line average and projecting to the selected character. The proof handles the remaining constant-term contributions on the full space of differences before pairing in the cuspidal subspace. In particular the final difference has zero residual and continuous projections in $L^2([G])$. The nonzero lower coefficient specifies the common scalar with the normalization of the cone calculation. To pass to the matched Schwartz pairs, approximate their real tests by simultaneously matched polynomial pairs with the finite tensors fixed, retaining the smaller-source support. The real Schwartz continuity of both integrals for those finite tensors gives the same identity in the limit, pointwise on $[G]$. The positive norm conclusion · Passage to the matched Schwartz class

2.3.19 Regularization and the scalar of the square

The comparison on supported smaller-source tests extends to arbitrary finite tensors through a finite sum of source translates with nonzero total coefficient. For any finite collection of tensors and source translations, choose an additional completely split place v outside their exceptional places and the ramification set. In the chosen split Schrödinger chart the smaller source rank is $n = a'$. The finite vacuum is a scalar times $\mathbf{1}_{\mathrm{Mat}_{n,d}(\mathcal{O}_v)}$, and the translations at v preserve this vacuum. Let q_v be the residue cardinality and let U range over the subspaces of k_v^n, with L_U their inverse images in $\mathcal{O}_v^n$. For codimension c, choose g_U with $g_U^{\mathrm{t}} L_U = \mathcal{O}_v^n$ and $\det(g_U)$ of valuation $-c$. Under the source Weil action, the translated indicator is $\mathbf{1}_{\{Y\,:\,\mathrm{im}\,Y \subset L_U\}}$ times $\chi_W(\varpi_v)^{-c} q_v^{cd/2}$, where χ_W is the unramified source-action character.

Multiply that translate by $(-1)^c q_v^{c(c-1)/2} \chi_W(\varpi_v)^c q_v^{-cd/2}$ and sum over U. The lattice of subspaces has Möbius function $(-1)^c q_v^{c(c-1)/2}$ on an interval of codimension c. This follows by summing these values over the interval using the Gaussian binomial identity, whose recurrence counts subspaces according to their intersection with a fixed hyperplane. Summing over U containing the image of Y modulo ϖ_v therefore gives the indicator that this image has dimension n. The sum Z_v of source action operators has augmentation

$$d(Z_v) = \prod_{j=0}^{n-1} \left(1 - \chi_W(\varpi_v) q_v^{j-d/2}\right) \neq 0.$$

The inequality $n < d/2$ and unitarity of χ_W give nonvanishing. This produces the required smaller-source support and hence the Schwartz continuity and integrability of the regularized integral. In a translate of any one of the finite tensors, the v-component of the original translation can be dropped because it fixes the vacuum. The other components commute with Z_v. The finite operator on the actual tensor model · Simultaneous source translations.

In the common target quotient, applying Z_v to the smaller-source vector multiplies its image by $d(Z_v)$ because each summand is a source translate. For a matched pair in the same quotient, the normalized regularized integral is $d(Z_v)^{-1} I_B(Z_v \Phi_B)$ in the chosen measure. The factor from the supported comparison is independent of the auxiliary split place used for this calculation. Apply both finite operators at once to a test with nonzero common quotient and nonzero image under the larger-source integral. Their augmentations are nonzero. The resulting test is supported at both places, so evaluating the two equalities on this pair gives the same scalar. This uses the common quotient maps for the two supported realizations. If the original smaller-source integral already converges absolutely, changing variables by each source translate gives $I_B(Z_v \Phi_B) = d(Z_v) I_B(\Phi_B)$ directly. The integrals of finite source augmentation differences are zero after the same operation, by the preceding commutation and change of variables. This proves the regularized identity from the supported comparison whenever its two quotient maps and its common nonzero image are available. Augmentation and comparison · Agreement with the convergent integral.

The relative scalar from changing the reference sections of the triangle to the actual square can be checked before factoring the triangle L-values. For a fixed square let q_A, q_B be its maps to the same irreducible real polynomial module, with the two infinity types fixed across the tuples. For the two legs, let $r_{A,\mu}, r_{B,\mu}$ be the finite maps in the product quotient. The images of the two real and finite pairs are $q_A(P) \otimes r_{A,\mu}(\Phi_A)$ and $q_B(P') \otimes r_{B,\mu}(\Phi_B)$. Such a product description follows from the real quotient once each finite datum gives a real intertwiner out of the same irreducible module. Take the span of those intertwiners as the finite factor. Schur's lemma and simplicity identify its evaluation tensor with the span of the images. In the countable-dimensional real module, the scalar form of Schur's lemma follows since an endomorphism with no eigenvalue would give an uncountable linearly independent family of resolvents of a fixed nonzero vector. Indeed, clearing the denominators in a nonzero finite relation between resolvents gives a nonzero polynomial in the endomorphism annihilating that vector. The polynomial is a product of invertible shifts times a nonzero scalar, a contradiction. The real module and the two polynomial maps come from the same infinity types, so the factorization may be used with fixed real coordinates. Common real quotient · Product quotient.

For the reference triangle sections Φ_A^0, Φ_B^0 the finite maps are proportional with ratio γ_μ. Choose P, P' with the same nonzero real quotient, and with the reference smaller-source integral nonzero after identifying the triangle coordinates. Write its triangle polynomial coordinate as $p'_F(P')$, and write $p_F(P)$ for the larger-source coordinate. The actual quotient Haar measures differ from the triangle measures by factors s_A, s_B. If b_μ is the regularized integral scalar on the product quotient, injectivity of the triangle realization gives

$$s_A \lambda_\mu p_F(P) = b_\mu \gamma_\mu s_B \, p'_F(P').$$

Since the triangle polynomial maps and the chosen vectors are independent of μ, the ratio $m_{\mathrm{ref}} = b_\mu \gamma_\mu s_B / (s_A \lambda_\mu)$ is independent of μ.

On the actual finite flip tests the required normalized finite comparison reads

$$s_A r_{A,\mu}(\Phi_A) = \kappa_\mu \gamma_\mu s_B r_{B,\mu}(\Phi_B).$$

The two prescribed model identifications use one common Haar and real factor, the finite twist on the moving frame and the local characters of the complementary-frame diagonal units. In those identifications

$$\kappa_\mu = C \, \xi_\mu(y) \left(\prod_{v \in S_\mu} \nu_{\mu,v}(x_v) \right)^{-1},$$

where $C \neq 0$ is common, y is the prescribed moving idele, x_v is the fixed local frame unit and S_μ is the tested split support. The relative scalar $\kappa_\mu / \kappa_{\mu_0}$ is in $\iota_K(K^\times)$. Indeed $S_{\mu_0} \subseteq S_\mu$. On the old places the ratio of the two $\nu_v(x_v)$ is the finite twist value. At the added places the additional factors are values of the reference finite local characters on units and of the finite twist. The coefficient field includes these values on the chosen support, and $\xi_\mu(y)$ is a finite twist value in the same field. This argument checks the field consequence of the frame identity for the actual square. The finite frame and Haar factors · Local unit character ratios · The moving idele.

In the tied real relation the quotient maps are normalized by equality. The finite frame identity and the scalar of the reference triangle give

$$c^0_{ij,\mu} = b_\mu \kappa_\mu \gamma_\mu s_B / s_A = m_{\mathrm{ref}} \kappa_\mu \lambda_\mu.$$

The witness and the real comparison below give this same ratio for the original square constants. In particular a_μ^{top} is $\kappa_\mu / \kappa_{\mu_0}$. Scalar extraction on one nonzero triangle vector.

2.3.20 A nonzero evaluation on the prescribed square tests

Fix a tuple and a square. Use its two middle-vertex quotient measures, and place the Hermitian triangle on the same source lattices and the sum of the two exterior lattices. Write I_A for the large-source integral and T for the exceptional set of the placed triangle. The bijection from the selected split places above T to T uses the components prescribed by the grid. The designated auxiliary split place v_0 of the square belongs to T. The place used to regularize the small integral lies outside T.

First choose a polynomial Fock vector P and a finite pure tensor φ, with the prescribed vacua outside T, for which $I_A(P \otimes \varphi)$ is nonzero as a function. The nonzero Hermitian-triangle comparison gives a nonzero value of the large integral on the Schwartz space. Hermite polynomial density and its continuity with finite data fixed then give a polynomial and a finite pure tensor. To restrict the exceptional support to T, argue by contraposition. At any $v \notin T$ the common vacuum image is nonzero and cyclic in the irreducible local quotient. For fixed other factors the large integral depends only on the local quotient image. This follows by using the same matching small-source vector for two tests with the same quotient image. That small-source vector can be chosen with the required full-rank support at the good split regularizing place. Replacing a vacuum by a local target translate translates the resulting function. Consequently if the function vanishes with the vacuum it vanishes with every local vector. Induction on the finite set of exceptional factors outside T proves the claimed support reduction. Polynomial nonvanishing with exactly the permitted exceptional set

The evaluation may be taken in the block subgroup $\mathcal{H}$ of the target, the product of the two exterior unitary groups. In fact, $\det(\mathcal{H}(\mathbb{A}_F))$ is the full norm-one idele group. Use the action by the given norm-one idele on a nondegenerate line in a positive-rank exterior group and the identity on its complement. Every target element is therefore sh with $s \in \mathrm{SU}(\mathbb{A}_F)$ and $h \in \mathcal{H}(\mathbb{A}_F)$. At the specified split place v_0, strong approximation gives density of $\mathrm{SU}(F)\mathrm{SU}(F_{v_0})$ in $\mathrm{SU}(\mathbb{A}_F)$. If all evaluations on $\mathcal{H}$ were zero for all allowed φ with P fixed, the value at γkh would also be zero. The rational factor γ drops out by invariance of the theta distribution. Moving k to the right of h changes only the factor of φ at v_0. Continuity in the target then gives zero at every sh, contrary to the first step. There are thus φ and $h = \mathrm{diag}(g, g_1) \in \mathcal{H}(\mathbb{A}_F)$ with nonzero value. [84] Restriction to the exterior block subgroup

At each $v \in T$ replace its local test by one coming from the actual flip, and correct the evaluation by a block translate. For this, take a small-source preimage Ψ of the given large-source local quotient class in the complementary Schrödinger chart. The quotient in these charts is the space of the complementary zeta sections. We use the following consequence of the split frame calculation. Let w be the chosen split component above v. In the large-source chart the placed mixed-coordinate change has the form $Z \mapsto g_0 Z c_0$, followed by the prescribed column permutation and completion isomorphism. In the chosen completion, write C_1 for the frame of the first exterior vertex and C_2 for the negative of the frame of the second exterior vertex. Write b_W, b_A for the triangle target and large-source bases, γ_A for the large middle grid frame, and J_A for its Hermitian Gram matrix in the ambient coordinates. Then the matrices, with all entries transported by the chosen completion isomorphism, are

$$c_F = b_W^{-1}\,\mathrm{diag}(C_1, C_2), \qquad c_0 = c_F^{-\mathsf{T}}, \qquad g_0 = -\gamma_A^{\mathsf{T}} J_A^{\mathsf{T}} b_A.$$

The row change is integral invertible by self-duality: the pairing matrix of two integral bases of the same self-dual lattice is unimodular. The partial Fourier transforms on the complementary blocks cancel, so the image of the large flip test in this chart is the primitive-frame section of the first partial Fourier transform of the small test in the placed coordinates. This uses the inverse Fourier sign on the complementary block and the old middle level's inverse-transpose Haar measure. Normalize the mixed model by this Fourier inversion so that the chart identity has prefactor one. The integral row frame and the two mixed charts · The chosen column frame and its sign

Let J be the prescribed first column block of the exterior sum and choose a large integer M. Put

$$D_M = \mathrm{diag}(\varpi^M I_J, \varpi^{-M} I_{J^c}).$$

For each exterior block, use the unitary element with component $C_j(D_M|_j)^{-\mathsf{T}}C_j^{-1}$ on the chosen half and inverse adjoint on the other half. The two components give an element of $\mathcal{H}(F_v)$. In the large-source Schrödinger chart it acts by right substitution by $c_0 D_M c_0^{-1}$ with a nonzero scalar r. Indeed the operator of this substitution and the given Weil operator induce the same conjugation on the two Heisenberg generators. Their ratio is a nonzero scalar by Stone–von Neumann. Put $N = c_0 D_M$. Use for the first partial Fourier transform of the small test, in the transported coordinates, the function

$$\Xi(Y) = \frac{\nu(\det g_0)\nu(\det N)|\det g_0 \det N|_v^{(a+a')/2}}{r\,u\,s_\nu}\,\Psi(YN^{\mathsf{T}}).$$

Here u is the frame-permutation factor $(\nu(\det Q)|\det Q|_v^{(a+a')/2})^{-1}$ for the prescribed permutation Q, and $s_\nu > 0$ is the ratio of the transported inverse-transpose old-level Haar measure to the volume-one maximal-compact Haar measure on the small-source linear group. All these factors are nonzero. Indeed the frame-permutation factor follows from replacing the completion h by Qh and using its determinant weight. For the two moved frames the large zeta section acquires the inverse of the numerator above. The integrand on the small side is evaluated at $YN^{-\mathsf{T}}$; this cancels N^{T} in Ξ. The factors r, u, s_ν then restore the zeta section of Ψ in its original measure and frame. The equality of sections is an equality of local quotient classes.

For the depth check, choose M so that the support of $\Psi(Yc_0^{\mathsf{T}})$ lies in the box $\varpi^{-M}\mathcal{O}_v$ and it is invariant under translation by the box $\varpi^M\mathcal{O}_v$. Take $n = \max(2M, \mathrm{depth}_v)$. The pullback by D_M is supported in the lattice with J-columns in $\varpi^{-n}\mathcal{O}_v$ and J^c-columns integral. It is invariant by ϖ^n times this lattice: the changes in the arguments of $\Psi(Yc_0^{\mathsf{T}})$ have valuations at least M and $n - M$ respectively. The inverse finite Fourier transform on J gives an integral test Φ_w of depth n. The additive character and the local transported lattice use the same dual pairing as the split transfer. Applying that transfer to Φ_w, then taking the block translate, recovers the original local class in the required quotient. Local quotient reach by admissible flipped tests · Depth, frame movement and normalizing scalar

Replace a local component at a time. If ϕ_v^{new} is the unshifted flip datum, its image after action by $k_v \in \mathcal{H}(F_v)$ equals the old quotient image. Right covariance gives the same nonzero value with the new tensor at hk_v. The new evaluation remains block diagonal, and subsequent replacements leave earlier finite components in their new unshifted forms. This produces the local symbols Φ_w of each square and a nonzero evaluation on the pure tensor determined by them. The depths in this existence argument may exceed the minimum depth of the row tests.

To identify this tensor with the actual two-arrow tests, use the split idempotent which is one at the chosen components over T and zero at the other components. The local tensor decomposition gives a finite sum τ on the large route and a finite sum τ' on the small route, having symbols $\mathrm{Flip}_w \Phi_w$ and Φ_w on the prescribed halves. Off T use the specified lattice vacua. In the large Hermitian Schrödinger model, put the local mixed-polarization intertwiners in the exact frames of these halves. Each takes the local vacuum to its lattice indicator times its value at zero. This value is nonzero. The common Gaussian at infinity is multiplied by the product of these values so that the two models agree on the initial vacuum. Applying translations and multiplication operators of the two complementary local Lagrangians then preserves the equality. They act by the same Heisenberg phase in both models, and their translates of the local vacuum span the local Schwartz space. Thus the equality for the local symbols on the chosen halves extends to the actual two-arrow tensor datum on the full adelic space. This proves the equality for the Gaussian with exactly these finite factors. The route Fock substitutions are surjective and commute with the creation operators, so we can realize the specified P on τ with the finite factors unchanged. The real relation has onto projections, which supplies a matching small-route polynomial on τ'. Gaussian with the finite flip tensors · All Fock polynomials on those same finite tensors

The product of the two large-route kernels for this datum is the large Hermitian theta kernel times the nonzero splitting prefactor at $\mathrm{diag}(g, g_1)$. In the compatible frames the source conjugation pushes the

selected grid quotient Haar measure to the triangle measure. The integral over a different fundamental domain agrees by the rational invariance of the summed kernel. We obtain

$$\int_{[U_A]} \ker_A(\Theta_A; g, g_1, u)\, du \neq 0$$

on the matched route tests Θ_A, Θ_B just constructed. Take a finite-source translate regularizer z for the corresponding small-route test, with nonzero augmentation $d(z)$. Its normalized integral is the one in the square comparison. Hence

$$0 \neq d(z) \int_{[U_A]} \ker_A(\Theta_A; g, g_1, u)\, du = c_{ij,\boldsymbol{\mu}} \int_{[U_B]} (z \ker_B)(\Theta_B; g, g_1, u)\, du.$$

This proves nonvanishing both of the square constant and of the translated small-route evaluation, in the measures used in the matching identity. The selected square evaluation · Nonzero square constant

2.3.21 The constant for the real relation of the row tests

Keep the real relation $\mathscr{R}$ used in the grid row tests. The normalization argument also uses the relation $\mathscr{R}_0$ given by equality in the two fixed real quotient maps. Both relations have onto projections. They use the same real vacua in every arrow, so the condition of having paired finite flip symbols is the same for both. Write $c_t = c_{ij,\boldsymbol{\mu}}$ for the square constant on $\mathscr{R}$, with $t = \boldsymbol{\mu}$, and c_t^0 for the regularized comparison constant on $\mathscr{R}_0$.

To see why the local finite matching supplies one scalar for all paired tests, fix t and a product coordinate system with the real maps q_A, q_B of the relation being used. For one nonzero real quotient vector x, choose a large polynomial with image x. The pure-tensor matching gives a small polynomial with image x' such that

$$x \otimes r_A(\tau) = x' \otimes r_B(\tau')$$

for every finite pair with the prescribed flip symbols. The small polynomial can be chosen once here, since it is chosen in the onto real map before the matching finite tensors are inserted. Apply a complex-linear functional on the real quotient module taking x to 1. This gives $r_A(\tau) = \lambda_t r_B(\tau')$, where λ_t is its value at x'. It is independent of the paired finite test and its depth. Consequently the product quotient satisfies

$$p_A(\Theta_A) = \lambda_t p_B(\Theta_B)$$

on every test in the real relation. For the tied relation $\mathscr{R}_0$, the same argument uses the placed Hermitian triangle and its reading in the prescribed frames. Multiply the comparison scalar of the common product quotient by λ_t. The kernel realization has a common splitting prefactor, and it transports the two source measures to the given grid measures. This gives c_t, respectively c_t^0, in the regularized comparison. Uniform finite scalar for the original real relation · The tied comparison scalar in the actual measures

For fixed exterior arguments let $T_B(z)$ be the integrated translated small test when integrable. The normalized value $J = T_B(z_0)/d(z_0)$ is independent of the chosen good-place regularizer z_0, and $T_B(z) = d(z)J$ for every z with an integrable translated kernel. Indeed place an auxiliary regularizer w of nonzero augmentation at a good split place outside the support of z_0 and outside the finite exceptional set of the finite translation sum z. Strip the compact factors of z outside its finite exceptional set, preserving its augmentation and translated function. The resulting sum commutes with w. Change variables in the middle quotient and use absolute convergence for the two orders of each product of translates to get

$$d(w)\, T_B(z) = d(z)\, T_B(w), \qquad d(w)\, T_B(z_0) = d(z_0)\, T_B(w).$$

The latter identity uses disjoint support for w, z_0. Cancelling the two nonzero augmentations proves the assertion. Thus the original scalar identity may be applied to the regularizer used for the tied relation.

One normalized small-source value for every integrable translate · The original relation for all convergent comparisons

For each paired finite test, the large integrated kernel at fixed exterior arguments is a linear functional on the large-route real polynomial tensor space. The common nullspace, over all such tests and arguments at any one t, is

$$\mathcal{K} = \ker(q_A \iota_A).$$

Indeed the large kernel is the evaluation of the large-source lift on $q_A(\iota_A P) \otimes y$, times its nonzero splitting prefactor, where y is the fixed finite quotient image for the test. The quotient map of the theta integral is either zero or injective by irreducibility. The nonzero grid evaluation just constructed gives one finite test at each t with $y \neq 0$ and a nonzero integral map. If all paired grid evaluations vanish for a given P, the block-translation and local replacement arguments applied to this fixed P give vanishing of its global theta integral on that test. Injectivity and cancellation of y give $q_A(\iota_A P) = 0$. Conversely that quotient equality gives vanishing for every paired large test. Thus the nullspace is the same for all t, and it is proper. Common large-route nullspace

Choose $(P, Q) \in \mathscr{R}$ and $(P_0, Q) \in \mathscr{R}_0$. Use the very same finite tensors and quotient measures for the two large tests and their common small test. Take the regularizing finite translate supplied by the comparison on $\mathscr{R}_0$. For a large-route evaluation functional L, both identities now use the same normalized translated small-route value J, giving

$$L(P) = c_t J, \qquad L(P_0) = c_t^0 J.$$

It follows that $c_t^0 P - c_t P_0 \in \mathcal{K}$. If $c_t^0 = 0$, every P_0 would belong to $\mathcal{K}$ by the onto projection of $\mathscr{R}_0$. The nonzero evaluation thus gives $c_t^0 \neq 0$. Choose one pair as above with $P_0 \notin \mathcal{K}$, and keep it fixed for all t. The class of P in the quotient by $\mathcal{K}$ is then c_t / c_t^0 times the fixed nonzero class of P_0. Consequently

$$c_t = \ell_{ij} c_t^0, \qquad \ell_{ij} \in \mathbb{C}^\times$$

with ℓ_{ij} independent of the tuple. Transport between the two real relations

In the ratio $c_t / c_{t'}$ the factor ℓ_{ij} cancels. The reference triangle's coefficient formula is applied to the c^0 constants with the measure changes already included.

For a listed smooth cusp form, apply the regularized integral comparison inside the exterior integral in the prescribed quotient measures. The joined integrals on the chosen path are absolutely convergent after the finite-source translates, and the translated small lift is multiplied by the augmentation of the regularizer by change of variables. Cancel its nonzero augmentation. This gives

$$\mathcal{L}_{\mathrm{right,down}}(\Theta_A) = c_t \, \mathcal{L}_{\mathrm{down,right}}(\Theta_B).$$

For a square inside a longer path, decompose its finite test with the two exterior paths fixed. Each summand has two finite symbols related by the prescribed flip, and the corresponding Fock polynomials meet the original real relation. Linearity and the same convergence bound at that prefix give the identity for their sum. Replacing the small route by the large route therefore multiplies the expression by c_t^{-1}. Cancellation after integrating the source translates · Use of the inverse constant along the path

2.3.22 The definite triangle

At the bottom use the three standard Hermitian spaces of ranks $A_0, B_0, A_0 + B_0$, their self-dual lattices and their prescribed finite sections. The two sources and the common target are definite at every infinite place. The larger-source jet uses the negative-sign Fock convention. Put $h_0 = (A_0 + B_0)/2$ and $e_0 = (A_0 - B_0)/2$. Then $A_0 > B_0 > 0$ and $A_0 < A_0 + B_0 \leq 2A_0$. The target has Witt index zero, so there is no smaller global member of this Witt tower. The doubling parameter for the larger source is $s_0 = B_0/2$. At the unramified finite places the constant source vector gives an exponent string $(A_0 - 1)/2 - k$, $0 \leq k < A_0$. The local

standard factor in the doubling convention uses $\nu_{ij,\boldsymbol{\mu}}^{-1}$. Consequently at $s = s_0$ the arguments in that factor are

$$s_0 + \tfrac{1}{2} + \tfrac{A_0-1}{2} - k = h_0 - k, \qquad 0 \le k < A_0.$$

The exceptional set here consists of the base-level places and the places under the finite twist. At the new places $\nu_{ij,\boldsymbol{\mu}_0}$ is unramified and the local component of $\xi_{\boldsymbol{\mu}}$ is ramified. Ramification occurs on both split factors, since the character is trivial on the diagonal. The base-level factors are ramified as well, so the omissions remove only trivial primitive factors. Since $\nu_{ij,\boldsymbol{\mu}}^{-1} = \nu_{ij,\boldsymbol{\mu}} \circ c$ for CM conjugation c, reindexing the remaining Euler factors by $w \mapsto \bar{w}$ gives the primitive functions for $\nu_{ij,\boldsymbol{\mu}}$. Analytic continuation gives $\prod_{s=1-e_0}^{h_0} \Lambda_{\boldsymbol{\mu}}(s)$ at the doubling point. This is the finite factor in the anisotropic Rallis formula [19, S§7.1, 11.3–11.7]. Standard triangle · Exceptional places and exponent string · The inverse square character in the standard factor.

For the larger-source finite section one can check the nonzero exceptional doubling contribution using its support. At infinity take the product of the powers of the top A_0-minors of the Fock matrices. The powers are $a_\tau - h_0$, which are nonnegative integers by the grid gap. Right substitution acts on each such minor by the determinant of the source matrix. The vacuum exponent and the source splitting character contribute the inverse determinant power, so this section is fixed at infinity. At each exceptional split place write Z for the coordinate matrix in the polarization given by the selected place of E over v. The first section has integral Z with its top A_0-minor invertible over $\mathcal{O}_v$. If Z and $Z g_v^{-1}$ both meet this condition, their top minors differ by g_v^{-1}, which then belongs to $\mathrm{GL}_{A_0}(\mathcal{O}_v)$. The other split component is its adjoint inverse, by unitarity. Thus outside the lattice stabilizer the section and its translate have disjoint support in this polarization. Their local inner product is zero. Equivalently one can apply a finite Fourier projector on a common finite quotient of the two supports. On the lattice stabilizer the determinant character of the section cancels the splitting factor, so the section is fixed. At these exceptional places, the local coefficient in the doubling integral is therefore supported on the compact stabilizer and is the nonzero squared norm of the vector there. The Siegel section height is one on this stabilizer. Its integral is the Haar volume times the squared norm. The other places contribute the spherical factors above. In particular, with the prescribed doubling and section matching, the larger-source lift is nonzero by the norm formula and the nonvanishing of the Hecke product. Doubling calculation on the prescribed section · Norm on the definite triangle.

In the common oscillator quotient, take a fixed pair of maps p_*, p'_* on the two archimedean polynomial spaces, both surjective onto the quotient $\mathcal{P}$. They depend on the infinity types, which are fixed across the tuples. For each tuple the matched finite data lie on the same line in the common finite quotient. Thus after taking the definite theta integrals the two maps factor as $\lambda_{\boldsymbol{\mu}} J_{\boldsymbol{\mu}} p_*$ and $J_{\boldsymbol{\mu}} p'_*$. The kernel of $J_{\boldsymbol{\mu}}$ is a subrepresentation of the irreducible quotient. Nonvanishing of the larger-source lift gives $\lambda_{\boldsymbol{\mu}} \neq 0$ and injectivity of $J_{\boldsymbol{\mu}}$, and $p'_* \neq 0$ then gives nonvanishing of the other lift. This step uses the local matching on the prescribed finite data, the polynomial quotient and the global integral comparison on those quotients. Prescribed finite sections in the common quotient · Polynomial and finite factors · Nonvanishing of both integrals.

For each source leg, choose a compact open subgroup of its finite adelic group fixing the sections for every tuple. At level given by this subgroup the class set is finite by definiteness. Integration over the adelic quotient is a single Haar volume times a sum of compact real averages at class representatives, each divided by the size of the representative's finite stabilizer. The compact average of the Fock action is the projection to the invariant subspace after twisting by the determinant character. To see its algebraicity on a homogeneous polynomial space, restrict to the dense subgroup of compact unitary matrices with algebraic entries. The invariant space is the intersection of the kernels of $\rho(g) - 1$, and the complementary space is the span of their images. Finitely many of these matrices suffice to give the same intersection and span. Both spaces are defined over algebraic numbers, so the projection onto the first along the second has algebraic coefficients. Averaging over the central circles leaves a single multidegree at the infinite

places, fixed across the tuples. In the smaller-source leg the surviving determinant semi-invariants are spanned by products of maximal minors. Definite class formula · Real averaging in a fixed multidegree · Projection on the smaller source.

On the larger-source leg the finite Fourier projection expresses the kernel on the surviving multidegree as a finite algebraic linear combination of jets of finite Heisenberg translates of a standard vacuum. The CM jet comparison divides them by a common nonzero scalar, the vacuum comparison constant times the product of the local CM periods to the surviving degrees, with conjugation for the negative sign. On the smaller-source leg one uses the CM jet comparison for the surviving products of minors, again with a common period. The two Haar volumes are included in the two periods. At the target use a compact open subgroup fixing the finite vectors for all tuples. Its finite class set gives a finite list of evaluations on the surviving monomial spaces. Changing representative acts by polynomial substitution in the rational Hermitian matrices through the chosen embeddings, together with their determinant characters. These coefficients are in a finite number field. Thus adjoining the finitely many normalized class coefficients at the target representatives gives the field of the two class sums. This field is included at the same stage as the Hecke-ratio field in the construction of the tested coefficient field. The CM jet comparisons here must be taken with the stated vacuum and polynomial normalizations. Larger-source kernels · Fixed-degree CM period · Minor jets · Coefficient field of both sums.

Evaluation on the finite adeles distinguishes the resulting functions. In fact, $U(F)$ is dense in $U(F_\infty)$. The Cayley transform identifies the open subset where -1 is not an eigenvalue with an open subset of the space of skew-Hermitian matrices. The F-rational matrices are dense there by simultaneous approximation at the infinite places, and this open subset is dense in the real unitary group. For $g = (g_\infty, g_f)$, take $x_\ell \in U(F)$ with $x_{\ell,\infty} \to g_\infty$ and evaluate at $x_\ell(1, x_{\ell,f}^{-1} g_f)$. Left rational invariance and continuity show that a function vanishing on finite adeles vanishes at g. The same argument applies to linear combinations of the two operators.

The quotient of the bottom constants in the prescribed field. Fix the bottom triangle Y whose normalized class sums were adjoined to K for this cell. Write $I_{\boldsymbol{\mu}}^{Y}$ and $(I_{\boldsymbol{\mu}}^{Y})'$ for its two integrals, with the common quotient maps p_*, p'_* above. Use the product P of the powers of the top minors used to prove nonvanishing and take P' to be the product of the complementary minors, with the same powers. Both are polynomials with integer coefficients in the prescribed Fock variables. In the standard compact frames, the complementary minor identity on a unitary matrix relates the two column blocks by conjugation, the determinant, and the sign of the ordered complement. On the surviving determinant types this gives the intertwiner of the rectangular Schur spaces in the two quotient maps, with the determinant twist on the second block. Schur's lemma gives

$$p_* P = \rho\, p'_* P' \neq 0$$

for a scalar $\rho \in \mathbb{C}^\times$ that includes the sign. The spaces, frames and degrees at infinity are fixed throughout the tuples, so ρ is fixed. Since $J_{\boldsymbol{\mu}}$ is injective, $(I_{\boldsymbol{\mu}}^{Y})'(P')$ is nonzero. Choose a finite target point $b_{\boldsymbol{\mu}}$ at which it is nonzero. With the two fixed class periods Ω, Ω', the two values

$$\frac{I_{\boldsymbol{\mu}}^{Y}(P)(b_{\boldsymbol{\mu}})}{\Omega}, \qquad \frac{(I_{\boldsymbol{\mu}}^{Y})'(P')(b_{\boldsymbol{\mu}})}{\Omega'}$$

belong to $\iota_K(K^\times)$. Their quotient is $\rho\lambda_{\boldsymbol{\mu}}\Omega'/\Omega$. Dividing by the corresponding quotient at the base tuple shows that $\lambda_{\boldsymbol{\mu}}/\lambda_{\boldsymbol{\mu}_0} \in \iota_K(K^\times)$. Products of complementary minors · The common quotient for any prescribed bottom triangle.

Now let X be the bottom triangle used in the square and let $c_0(\boldsymbol{\mu})$ be its comparison constant. Both triangles have the same three ranks, with identity Gram matrices in the rational coordinates and the canonical definite frames at infinity. Identify their three spaces by these coordinates. At each selected split factor, move the finite frame β_Y to β_X by $\beta_X \beta_Y^{-1}$. This matrix and its inverse are integral. On the conjugate

factor take the adjoint inverse, giving a unitary move that preserves the self-dual lattice. At all other places the move is the identity. Write g, g', h for the resulting two source moves and the common target move at the given tuple. Let x, x', y be their determinant ideles on the selected split halves, extended by 1 elsewhere. In a split polarization the normalized Weil action on a moved section multiplies the pullback by $\chi_T(x)^{-1}\chi_S(y)$, for source and target characters χ_S, χ_T. If α, β are the two parity-twisted characters of this cell, the larger and smaller leg multipliers are therefore

$$\kappa = (\beta/\alpha)(x)^{-1}\beta(y), \qquad \kappa' = (\alpha/\beta)(x')^{-1}\alpha(y).$$

The finite tests have the same defining support in the new frames. Changing variables by $g^{-1}, (g')^{-1}$ in the two source integrals preserves quotient Haar measure, and both target arguments undergo the same right translation by h^{-1}. Consequently the transported comparison constant is $\lambda_{\boldsymbol{\mu}}\kappa/\kappa'$. In this quotient the common parity character cancels. For $\nu = \mu_{0,j}\xi_{\boldsymbol{\mu}}/\chi_{m-i}$ it gives

$$u_{\boldsymbol{\mu}}^{\mathrm{fr}} = \kappa/\kappa' = \frac{\nu(x)\nu(x')}{\nu(y)} \in \iota_K(K^\times).$$

Indeed the local values of these three characters are in the prescribed coefficient field. For a determinant in $E_w^\times$ its character value equals the value on a sufficiently close element of $E^\times$, since the local characters are smooth. Only the finitely many selected split halves contribute. Identifying the standard spaces · Integral moves on selected halves · The two split Weil multipliers · Common right translation · Membership of the frame factor in the prescribed field.

With the frames aligned, the prescribed vacua span the same vacuum line on each leg. Replacing either normalized vacuum multiplies its section by one fixed scalar, since the section operator for each tuple uses the same support, conductor and Weil action. The source measures on the two aligned presentations also differ by fixed factors. These changes only multiply the comparison by a nonzero constant independent of the tuple. Applying (18) on X to the transported pair and to the fixed pair defining c_0 yields

$$\frac{c_0(\boldsymbol{\mu})}{c_0(\boldsymbol{\mu}_0)} = \frac{u_{\boldsymbol{\mu}}^{\mathrm{fr}}}{u_{\boldsymbol{\mu}_0}^{\mathrm{fr}}}\frac{\lambda_{\boldsymbol{\mu}}}{\lambda_{\boldsymbol{\mu}_0}} \in \iota_K(K^\times). \tag{19}$$

Fixed vacuum factors · Fixed Haar factors · Transport between prescribed triangles · Bottom constant ratio.

2.3.23 Hecke values in the square comparison

Write E/F for the CM extension in the grid. At cell (i, j) the two splitting characters carry the same parity twist by η. Their quotient, with $\xi_{\boldsymbol{\mu}_0} = 1$, is

$$\nu_{ij,\boldsymbol{\mu}} = \frac{\mu_{0,j}\xi_{\boldsymbol{\mu}}}{\chi_{m-i}} = \nu_{ij,\boldsymbol{\mu}_0}\xi_{\boldsymbol{\mu}}.$$

It is unitary and trivial on $\mathbf{A}_F^\times$. In the CM type $\phi_\tau = \overline{\mathrm{emb}_\tau}$ the infinity character at τ is $(z/\bar{z})^{a_\tau}$, where $a_\tau = e_\tau(j+1)+\mathrm{cen}_\tau(m-i)$ is the sum specified by the signed weight grid. The grid gap gives $a_\tau > N_{\mathrm{amb}}+1$, where $N_{\mathrm{amb}} = 1 + r + m(t + r + 6)$ is the ambient rank bound. The finite-order character $\xi_{\boldsymbol{\mu}}$ has trivial infinity component, so these exponents also apply at the base tuple. Square character · Infinity type and finite twists · Weight bound.

Here $\Lambda_{\boldsymbol{\mu}}(s)$ is the finite primitive Hecke L-function. Its factor at w is $(1 - \nu_{ij,\boldsymbol{\mu},w}(\varpi_w)(\mathrm{N}w)^{-s})^{-1}$ if the character is trivial on the local units, and is 1 otherwise. This convention uses each tuple's own conductor. The archimedean factor is $\prod_\tau \Gamma_{\mathbb{C}}(s+a_\tau)$, with $\Gamma_{\mathbb{C}}(u) = 2(2\pi)^{-u}\Gamma(u)$. The dual factor at $1-s$ has arguments $1 - s + a_\tau$, so the critical integers are

$$1 - \min_\tau a_\tau \le s \le \min_\tau a_\tau.$$

In particular, $1 \le s \le N_{\mathrm{amb}} + 1$ is in this interval. For the two diagonal ranks write $R_{ij} = r + jt + i(r+6)$. Both t and $r+6$ are odd. Hence $n_0 = (t + r + 6)/2 \ge 1$, and $n_0 + q_j \le R_{ij} + n_0 = (R_{ij} + R_{i+1,j+1})/2 \le N_{\mathrm{amb}}$. In an even cell $q_j \ge 1$, so this bounds the remaining argument $n_0 + 1$ as well. The finite factors converge absolutely to a nonzero product at every argument $s \ge n_0 + 1$ used in (17). Primitive Euler factors · Rank identities.

The triangle norm calculation also uses nonvanishing at integers from $1 - e_0$ to the half-sum of the diagonal ranks, where $e_0 = n_0 + q_j - p_j$ is half the difference of the two bottom ranks. For a unitary Hecke character with nontrivial restriction to the ideles of norm one, Tate's theorem gives an entire completed L-function and relates its value at s to the inverse character at $1 - s$ by a nonvanishing epsilon factor. The completed L-function of the trivial character has a simple pole at $s = 1$. [32, Chapter XV, §4] These assertions apply to $\nu = \nu_{ij,\boldsymbol{\mu}}$, ν^2 and ν^{-1}, none of which can be a power of the norm because each has unequal paired infinity exponents. At the boundary, let S contain all ramified finite places. For real $\sigma > 1$, the Euler-product expansion and the inequality $3 + 4\cos\theta + \cos(2\theta) = 2(1 + \cos\theta)^2 \ge 0$ give

$$\left|\zeta_E^S(\sigma)^3 L^S(\sigma, \nu)^4 L^S(\sigma, \nu^2)\right| \ge 1.$$

The zeta factor has a simple pole at 1 and the last factor is holomorphic there. Thus $L^S(1, \nu) \ne 0$. Reinstating the unramified factors in S preserves nonvanishing. The same argument applies to ν^{-1}. If $1 - e_0 \le s \le 0$, then $1 - s \ge 1$ and both infinity factors in the functional equation are finite and nonzero, since $e_0 \le N_{\mathrm{amb}} < \min a_\tau$. The conductor power and root number are nonzero. This gives nonvanishing throughout the integer range for the triangle. Hecke functional equation.

For algebraicity, use Blasius's critical-value formula for an algebraic Hecke character χ of E with unequal infinity exponents at every complex place. At a critical integer s it gives

$$L(s, \chi) \sim_{\overline{\mathbb{Q}}} d_F^{1/2} (2\pi i)^{[F:\mathbb{Q}]s} p(\chi^{-1,c}, \Phi_\chi).$$

Here L is the finite primitive Euler product, c is complex conjugation of E, d_F is the absolute discriminant, and p is the CM period. The CM type Φ_χ takes the embedding with the smaller infinity exponent from each pair. The relation $\sim_{\overline{\mathbb{Q}}}$ means equality up to a nonzero algebraic factor in the chosen complex component. For a fixed CM type Φ, the CM periods multiply under product of characters with critical infinity type. This follows from the diagonal map of the Shimura torus to its square. [33, 35] [34, §1.2, Proposition 1.1 and Theorem 1.1]

Apply this with $\chi = \nu_{ij,\boldsymbol{\mu}}$ and $\chi_0 = \nu_{ij,\boldsymbol{\mu}_0}$. The conjugate inverses equal the original characters, since $\chi(x)\chi(x^c) = \chi(\mathrm{N}_{E/F} x) = 1$. They have the same CM type Φ, the conjugate of the type $(\phi_\tau)_\tau$ used to write the positive exponents. Take $d \ge 1$ such that $\xi_{\boldsymbol{\mu}}^d = 1$. For both characters all positive powers have distinct paired exponents $ha_\tau, -ha_\tau$ and the same Φ. Multiplicativity and $\chi^d = \chi_0^d$ give

$$\left(\frac{p(\chi, \Phi)}{p(\chi_0, \Phi)}\right)^d \in \overline{\mathbb{Q}}^\times.$$

Taking the actual quotient of the two complex periods gives an element of $\overline{\mathbb{Q}}^\times$, since it satisfies the corresponding polynomial of degree d. The common discriminant and powers of $2\pi i$ cancel in the two critical-value formulas. Thus $\mathcal{R}_{\boldsymbol{\mu}}(s) \in \overline{\mathbb{Q}}^\times$ for every $1 \le s \le N_{\mathrm{amb}} + 1$.

The coefficient field is chosen using the field generated by these ratios for the finitely many tuples, cells and integers. A twist key specifies the finite-order character and its prescribed primes. A grid key specifies the auxiliary prime, its selected place and the grid characters, all over the fixed host. For every realizable pair of keys take a basis of the compositum of the ratio field with the fields for the finite local coefficients and other grid data. The selection imposes the union of these finite bases before realizing the tested grid, using the common family thresholds. One uses a finite Galois field containing these algebraic numbers and its involution is the restriction of complex conjugation under the chosen inclusion ι_K. Transport between records with the same keys identifies $\mu_{0,j}$, χ_{m-i} and the matched finite-order twists,

and hence their primitive L-functions. The realized ratios therefore belong to $\iota_K(K^\times)$, in particular the factor $\mathcal{R}_{\boldsymbol{\mu}}(n_0+1)$ in each even cell. The finite field of ratios · Twist keys · Grid keys · Field constraint on the actual grid · The common thresholds · The complex embedding and involution.

2.4 Bessel periods and Poincaré series

The bottom inclusion $X_0 \subset Y_0$ has codimension $2r_S+1$, where $r_S = m/2-1$. The Bessel subgroup $\mathcal{S}_w$ for w is the semidirect product of the unipotent group N_w^{un} of a length-r_S isotropic flag with the copy of $U(X_0)$ stabilizing the flag and a marked vector of negative norm. Write ψ_w for the generic character of N_w^{un} determined by the flag and vector. The pairing with a bottom row form f_0 is

$$\int_{[U(X_0)]} \phi(h) \int_{[N_w^{\mathrm{un}}]} f_0(nh)\,\overline{\psi_w(n)}\,dn\,dh.$$

The product contracts the chosen coefficient vectors in the bilinear convention, using the prescribed quotient Haar measures. This is the Gan–Gross–Prasad Bessel period for a split odd-dimensional orthogonal complement. In the invariant-functional convention, $\overline{\psi_w}$ is the one-dimensional factor on the Bessel subgroup, trivial on the diagonal $U(X_0)$. Integrating first on the unipotent quotient gives the displayed Fourier coefficient. [22, §§12, 23] The ambient Y_0 is common to all words. The construction scales a reference marked vector of norm $-T_{\mathrm{ref}}$ by λ_w^{-1}, with $\lambda_w\overline{\lambda_w}T_w = T_{\mathrm{ref}}$, to obtain one of norm $-T_w$. The character, stabilizer and column level for w enter its period separately.

At a real place the compact column type ρ_w has parameter $\alpha_1 > \cdots > \alpha_k$. The row discrete series σ_∞ has decreasing entries $(\beta_1,\dots,\beta_{k-1},-b_0,\dots,-b_{m-1})$. The β_i carry positive signature labels and the $-b_i$ alternate, starting with a positive label. The column and row entries have half-integral differences and satisfy $\alpha_i > \beta_i > \alpha_{i+1}$ and $\alpha_k > -b_0$, as in the local interlacing conditions. Write $S = \mathcal{S}_w(\mathbb{R})$, embedded in the product of the row and column groups by (hn,h), and extend ψ_w trivially across the column group. Jiang–Sun–Zhu's multiplicity-one theorem for real unitary Bessel groups says that, for irreducible Casselman–Wallach representations τ,λ of the two groups and a generic Bessel character ξ, the continuous intertwining space $\mathrm{Hom}_S(\tau\widehat{\otimes}\lambda,\xi)$ has dimension at most one. Here λ is inflated across the unipotent radical. [30, Theorem A]

Smooth matrix coefficients of real discrete series with compact center belong to the Harish–Chandra Schwartz space. Both σ_∞ and the compact group representation $\rho_w^\vee$ are square-integrable with compact centers, so the same assertion holds for their product. [39, §7.2] For an odd-dimensional orthogonal complement that is the sum of hyperbolic planes and a nondegenerate line, Beuzart-Plessis shows that Harish–Chandra Schwartz functions on the product group restrict to $L^1(S)$. For an irreducible tempered product $\tau\boxtimes\lambda$, his continuous sesquilinear Bessel form on smooth vectors is nonzero exactly when $\mathrm{Hom}_S(\tau^\infty\widehat{\otimes}\lambda^\infty,\xi)\neq 0$. For Schwartz matrix coefficients, inversion in the unimodular subgroup gives the ordinary integral of $\langle v,(\tau\boxtimes\lambda)(s)u\rangle\overline{\xi(s)}$ in our pairing convention. [28, §6.2, Lemma 6.5.1(ii), §7.1 and Theorem 7.2.1] In the prescribed Haar measures its application to the actual compact column type gives, entry by entry, the absolutely convergent integral

$$\mathcal{B}_\infty(u,v) = \int_{U(X_0)(\mathbb{R})\times N_w^{\mathrm{un}}(\mathbb{R})} \big\langle v,\sigma_\infty(hn)u\big\rangle\,\overline{\psi_w(n)}\,\rho_w(h)^{-1}\,dh\,dn.$$

Xue's real Bessel criterion applies to generic packets for a unitary pair of ranks d and $d+2r_S+1$ whose complement has r_S hyperbolic planes and a marked line of sign $(-1)^d$, with the associated generic Bessel character. For discrete-series parameters $\varphi_{\mathrm{row}} = \bigoplus_a \chi_a$ and $\varphi_{\mathrm{col}} = \bigoplus_i \mu_i$, write e_a, e_i' for the component-group generators of these characters. The continuous Bessel Hom space has dimension one exactly when both packet characters satisfy

$$\eta_{\mathrm{row}}(e_a) = \epsilon(\chi_a\otimes\varphi_{\mathrm{col}},\psi_{\mathbb{C}}) \quad\text{and}\quad \eta_{\mathrm{col}}(e_i') = \epsilon(\varphi_{\mathrm{row}}\otimes\mu_i,\psi_{\mathbb{C}}) \qquad \text{for all } a,i.$$

The epsilon factors use $s = 1/2$ and the trace-zero character $\psi_{\mathbb{C}}(z) = \exp(2\pi i \operatorname{Tr}_{\mathbb{C}/\mathbb{R}}(iz)) = \exp(2\pi(\bar{z} - z))$, which also fixes the even-rank Whittaker normalization. In odd rank d' use the unique datum on $U((d'+1)/2, (d'-1)/2)$. [36, §§1.1–1.2, Theorem 1.1] The row and dual compact column lie in generic discrete-series packets. For $d = k$ even, negating both Hermitian forms gives the required positive marked line. For $r_S > 0$, scalars t_a on the flag basis and $\overline{t_a}^{-1}$ on its dual scale the consecutive character coordinates by t_a/t_{a+1} and the last by t_{r_S}. They fix the column group and marked line and realize the prescribed generic character by conjugation.

For an irreducible discrete series of $U(p, q)$, with distinct Harish–Chandra weights in $\mathbb{Z} + (p + q - 1)/2$, Atobe attaches the character $\chi_x(z) = (z/|z|)^{2x}$ to each weight x. In his normalization, a weight at decreasing position j has component sign $(-1)^{j-1}$ for a positive signature label and $(-1)^j$ for a negative one. Negating the form flips each label, and passing to the dual negates the weights, reverses their order and preserves the signature label attached to each negated weight. In even rank, changing to the other Whittaker datum reverses the component sign at every generator. [37, §2.2, Theorem 2.1(4),(6),(7) and Theorem A.4] For $a \in \mathbb{Z} + \frac{1}{2}$, the archimedean root number of $(z/|z|)^{-2a}$ at $s = 1/2$ for $\psi_{\mathbb{C}}$ above is $\operatorname{sgn}(a)$. Taking $a = -x$ gives $\epsilon(\chi_x, \psi_{\mathbb{C}}) = (-1)^{[x>0]}$. [38, Proposition 2.1]

To match the even-rank datum, one can test the positive-definite pair $U(d+1) \supset U(d)$. Write any compact dual column weights in decreasing order as $-\gamma_d > \cdots > -\gamma_1$. Choose the compact row weights $\delta_1 > \gamma_1 > \delta_2 > \cdots > \gamma_d > \delta_{d+1}$ in the other parity. This is the Harish–Chandra form of the compact branching rule, so the Hom space has dimension one. Xue's epsilon product on the generator of $-\gamma_i$ then has i negative factors, and Atobe's positive signature label at position $d + 1 - i$ gives the same sign $(-1)^{d-i} = (-1)^i$. This distinguishes the two even-rank normalizations. Both sources use the odd quasi-split form with one more positive than negative direction. [36, Introduction and Theorem 1.1] [37, Appendix A.3–A.4] Thus after the form reversal the packet signs at β_i, $-b_j$ and the dual column entry $-\alpha_i$ are $(-1)^i$, 1 and $(-1)^{i-1}$, respectively. Multiplying the root numbers for the summands counts positive row-minus-column differences, respectively $k - i$, 0 and $i - 1$. Since k is even, both equations of the criterion hold. By the nonvanishing equivalence above and smooth cyclicity, there are u, v among the translates of any unit vector v_0 in the lowest K-type with $\mathcal{B}_\infty(u, v) \neq 0$.

For the real orbital norm, fix the section on the row group used in the Gram estimates. It transforms on the left by ρ_w and ψ_w. It has positive finite norm on the Bessel quotient with the specified quotient measure and is reproduced by right convolution with $d(\sigma_\infty)\overline{\langle v_0, \sigma_\infty(g)v_0\rangle}$, where the formal degree uses the same ambient Haar measure. These properties identify the real test used in those estimates. Their proof in the release uses a direct nonvanishing calculation on translates of v_0. Direct nonvanishing argument using the chosen section. Its comparison with the chosen local factor then factors $\mathcal{B}_\infty(u, v)$ as $\xi \mapsto \langle \ell_w(v), \xi\rangle \ell_w(u)$ and identifies the local factor as a multiple of evaluation by the nonzero Bessel functional ℓ_w. This calibration enters the proof of the full Gram estimates in Section 3.

At the fixed bottom row level and type, let $\mathcal{H}$ be the finite-dimensional Hilbert space containing the Bessel Poincaré series for all words. From $\phi \in H_w^{\mathrm{full}}$, the construction forms a seed $\Phi_{w,\phi}$ on $G_0 = U(Y_0)$ with the Bessel character and the specified local factors. Let x_w be its chosen finite placement, K_2 the ambient finite level with character χ_2, and W_w the real section just described. Write $T_\phi(\xi; 1, h_f)$ for the column coefficient evaluated at h_f on the vector ξ in the real type. The seed has finite support modulo the Bessel group in $N_w^{\mathrm{un}}(\mathbb{A}_{F,f})\, U(X_0)(\mathbb{A}_{F,f})\, x_w K_2$, with the column group in its specified embedding. For $y_f = n h_f x_w k_2$ it has value

$$\Phi_{w,\phi}(y_\infty, y_f) = \psi_{w,f}(n)\chi_2(k_2)\, T_\phi(W_w(y_\infty); 1, h_f).$$

It is zero off that support. The finite level relation for the construction is as follows. If $x_w^{-1} n h_f x_w \in K_2$, then $x_w^{-1} h_f x_w \in K_2$ and the level character of the first product is $\psi_{w,f}(n)\chi_2(x_w^{-1} h_f x_w)$.

Here is the finite congruence calculation for the chosen flag. Let $g = g_0 g_1$ be the modulus, Λ_f the finite adelic lattice of Y_0, K_Λ its integral stabilizer, and e_a, f_a the rational isotropic vectors with $\langle f_a, e_b\rangle = \delta_{ab}$. The finite column group fixes these vectors and the marked z_w. The rational lattice splits along the

hyperbolic spaces and their orthogonal complement. Its coefficients on the e_a's are $I_a = \mathfrak{d}^{-(r_S-a)}$ for indices $0 \le a < r_S$, where $\mathfrak{d} = \mathfrak{d}_{F/\mathbb{Q}}\mathcal{O}_E$. The coefficients on the f_a's are I_a^{-1}. The orthogonal projection onto the complementary space preserves Λ_f. If $nh_f \in K_\Lambda$, apply this projection to $nh_f y$ for y in the complementary lattice. As h_f preserves the complement and $n-1$ maps it into the span of the e_a's, the projection gives $h_f y$. The hyperbolic coordinates are fixed by h_f. The inverse of nh_f has the form $n'h_f^{-1}$ because the column group normalizes the unipotent group. Applying the argument to the inverse gives $h_f \in K_\Lambda$ and then $n \in K_\Lambda$. Since $x_w \in K_\Lambda$ and $K_2 \subset K_\Lambda$, this applies to the overlap above. The rational flag and its lattice · Projection to the complementary lattice.

Write e_a^0, f_a^0, z^0 for the integral frames of K_2, and $\beta_i, \beta_i^\vee$ for the ordered residual vectors and their duals there. The equations of K_2 modulo $g\Lambda_f$ are triangular on the e_a^0's, have increment in their span on z^0, and permit on β_i increments in that span and the preceding β's. Use the given simultaneous frame approximations. At each prime dividing g, e_a^0, f_a^0 reduce to e_a, f_a and $x_w z^0$ reduces to z_w. The element x_w fixes e_a, f_a, so the first two congruences hold for $u = x_w^{-1} n x_w$ by the flag conditions defining N_w^{un}. Moreover $\langle e_a, x_w\beta_i\rangle$ is divisible by g by the marked-frame orthogonality. Subtracting the corresponding f_a-components gives a vector in the exact orthogonal to every e_a with the same reduction. Its increment under n lies in their span by definition of the unipotent group. This increment has integral coefficients, as the f_a's are integral dual vectors at these primes. Transporting back gives the third congruence for u, in fact with increment in the span of just the e_a^0's. At primes away from g only the integral stabilizer condition is needed. Hence $u \in K_2$. The triangular level · The transported flag, marked vector and residual congruences.

The numerator of the finite level character on u is the sum of the pairings $\langle f_a^0, u\ \mathrm{next}(e^0, z^0)_a\rangle$ and $\langle \beta_i^\vee, u\beta_{i+1}\rangle$ for existing successor indices. Here next uses the next flag vector, and the marked vector at the last position. Call the sum A. Write B for the corresponding rational flag sum $\sum_a \langle f_a, n\ \mathrm{next}(e, z_w)_a\rangle$, the numerator of the finite unipotent character. The residual pairings of A are zero modulo g by the congruences and the dual Gram matrix. Take a single integer C, prime to g, clearing the denominators in the given frame congruences. In particular $C(x_w f_a^0 - f_a)$ and $C(x_w\,\mathrm{next}(e^0, z^0)_a - \mathrm{next}(e, z_w)_a)$ lie in $g\Lambda_f$. Expanding each difference of pairings gives $C^2(A-B) \in g\widehat{\mathcal{O}}_E$. Also A is integral, and B belongs to $\mathfrak{d}^{-1}\widehat{\mathcal{O}}_E$. For the latter statement, $n-1$ preserves the finite lattice and its entry from $I_{a+1}e_{a+1}$ to $I_a e_a$ lies in $I_a I_{a+1}^{-1} = \mathfrak{d}^{-1}$. At the last index the marked vector is integral and the same conclusion uses $I_{r_S-1} = \mathfrak{d}^{-1}$. The additive character is trivial on this trace-dual ideal. It is therefore trivial on $A-B$ and on $C^2(A-B)/g$. A Bézout identity for C^2 and g gives triviality on $(A-B)/g$. Thus $\chi_2(u) = \psi_{w,f}(n)$. Now $x_w^{-1}h_f x_w = u^{-1}(x_w^{-1}nh_f x_w)$ is in K_2, and multiplicativity of χ_2 gives the stated level relation. The trace-dual ideal for the rational flag · and the character after transport.

For two decompositions $nh_f x_w k_2 = n'h'_f x_w k'_2$, apply this relation with $h = (h'_f)^{-1}h_f$ and $n'' = (h'_f)^{-1}(n')^{-1}nh'_f$, for which $x_w^{-1}n''hx_w = k'_2k_2^{-1}$. The column transforms under h with multiplier $\chi_2(x_w^{-1}hx_w)$, and the Bessel character on n'' is $\psi_{w,f}(n)/\psi_{w,f}(n')$. Thus the two expressions for the seed agree. The prescribed Poincaré seed · The finite level relation.

The archimedean Bessel covariance and the finite character formula are inverse on the rational unipotent by the product formula, and the rational action of the stabilizing column group agrees on the two coefficient realizations. The seed is therefore left invariant by $\mathcal{S}_w(F)$. Summing $\Phi_{w,\phi}(\gamma g)$ over $\gamma \in \mathcal{S}_w(F)\backslash G_0(F)$ gives the operator $\mathcal{P}_w\phi$ in the same-word and mixed Gram statements.

The orbital proof uses the weighted integrability of the chosen real section on the family with $r \ge 1801$ and successive real b-gaps at least $16r^{19}$. The real gap hypothesis. Set

$$H_w(x) = \inf_{s\in \mathcal{S}_w(F_\infty)} \prod_\tau \|s_\tau x_\tau\|_{\mathrm{op}}, \qquad N = 4r^2\sqrt{2q},$$

where q is the maximum real negative index of Y_0. The local argument in that proof gives

$$\int \|W_w(x)\| H_w(x)^N\, d\bar{x} < \infty$$

on the real Bessel quotient, and a bounded compactly supported right convolution reproducing W_w. The integrable matrix-coefficient projector above also reproduces W_w with the given Haar measure. The weighted estimate and reproduction use the actual frame and quotient measure.

Here is how the weighted bound gives absolute periodization. Write $G = G_0(\mathbb{A}_F)$, $\Gamma = G_0(F)$ and $\Delta = S_w(F)$. Use the compact convolution at infinity and the normalized K_2 average at finite places. They give a compact set $C \subset G$ and a constant a such that

$$|\Phi(x)| \le a \int_C |\Phi(xy)|\, dy, \qquad \Phi = \Phi_{w,\phi}.$$

For a compact set B define $m_B(x) = \#\{\gamma \in \Gamma : x^{-1}\gamma x \in B\}$. If the transporter from z to xC contains one element $\gamma_0 \in \Gamma$, every other such element satisfies $x^{-1}\gamma\gamma_0^{-1}x \in CC^{-1}$. Unfolding the nonnegative integral of $|\Phi|$ against this transporter gives

$$\sum_{\delta \in \Delta\backslash\Gamma} |\Phi(\delta x)| \le a\, m_{CC^{-1}}(x) \int_{\Delta\backslash G} |\Phi(z)|\, dz. \tag{20}$$

On the support of the seed, first take the quotient by the adelic stabilizing group. The finite coordinate is represented by $x_w K_2$. The rational quotient of the stabilizer is compact: the unipotent rational quotient is compact, and the column unitary group is definite. Thus representatives on the support have finite coordinate in one compact set and real coordinate $c_\infty x$, where c_∞ is in a fixed compact set and x can be chosen with $\prod_\tau \|x_\tau\|_{\mathrm{op}} \le 2H_w(x)$. The quotient measure is obtained by integrating first on these compact stabilizer fibers. The norm of the column coefficient is bounded on its finite class set. For every compact B, an element counted by m_B has finite entries in a fixed fractional ideal of E, and its entries at τ have absolute value at most $C_B\|x_\tau\|_{\mathrm{op}}^2$. Indeed it is a conjugate by $c_\infty x$ of an element of B, and the operator norms of x_τ and its inverse agree in the unitary frames. Counting integral lattice points in boxes in the r^2 complex matrices gives a constant times $\prod_\tau \|x_\tau\|_{\mathrm{op}}^{4r^2}$. Since $N \ge 4r^2$, this and the weighted bound make both $\int_{\Delta\backslash G} |\Phi|$ and $\int_{\Delta\backslash G} m_B|\Phi|$ finite.

Unfold one factor of the square of the absolute series on $\Gamma\backslash G$. Its integral is $\int_{\Delta\backslash G} |\Phi(x)| \sum_\delta |\Phi(\delta x)|\, dx$. Equation (20) makes this finite. Thus the periodization represents an L^2 function. It has the right character of K_2. The integrable archimedean projector can be interchanged with the periodization in L^2, using the L^2 absolute series and invariance of Haar. It fixes the result by the convolution of the seed, so the result lies in the selected Petersson space. In its pairing with a row form, absolute convergence and left rational invariance join the summation domains into the quotient by Δ. The seed's left equivariance there gives the Bessel integral with the indicated column coefficient. The absolute folding estimate · the realization in the selected Petersson space and the orbit unfolding.

In the direct application of the Gram bounds we take their operators on the prescribed columns. The change from the arithmetic column to the adelic column reads the same values on the transported finite cosets with the character of $x_w^{-1}h_f x_w$. These values determine the column by the class covariance. The invariant Hermitian coefficient pairings and class masses give the isometric ramified inclusion $J_w : \mathcal{M}_w \hookrightarrow H_w^{\mathrm{full}}$. The simultaneous comparison supplies the common row isometry. In the notation of that comparison, $s : E_{\mathrm{row}} \hookrightarrow \mathcal{H}$ is conjugate linear and isometric. The actual column levels and maps · and the corresponding bottom pairing. The conjugate-linear first-variable convention gives the bilinear form

$$\mathcal{E}_w(\phi, F) = \big\langle s(F), \mathcal{P}_w J_w \phi \big\rangle_{\mathcal{H}}.$$

Write π_w for the tensor product of the compact $U(k)$-types at the real places. Let $v_N(w)$ be the volume of the unipotent fundamental domain times the volume of the common ambient level, and $v_H(w)$ the volume of the real compact factor of the Bessel stabilizer. Compact Schur orthogonality with probability measure contributes $1/\dim \pi_w$. Indeed averaging a rank-one operator over that factor commutes with

π_w and has the trace of the original operator. Restoring the measure multiplies by $v_H(w)$. The quotient and level factors give $v_N(w)$. The full operators and reference measures are fixed by

$$c_0(w) = \frac{v_N(w)v_H(w)}{\dim \pi_w} > 0, \qquad L_w = c_0(w)^{-1/2}\mathcal{P}_w, \qquad K_{wv} = L_w^* L_v. \tag{21}$$

Put $d_w = \dim H_w^{\mathrm{full}}$ and $r_w = \dim \mathcal{M}_w$. The single-word estimate for K_{ww} and the value of r_w in Section 3 will give a nonzero adelic reference pairing.

2.4.1 The auxiliary grid and the detection problem

Write $Y_i = V_{i,0}$ for the first row of the auxiliary grid and C_i for its core spaces. The ranks are $r + i(r + 6)$ and $r + 1 + i(r + 5)$, respectively. The initial core is $C_0 = Y_0 \perp \langle 1 \rangle$. The inclusion $\iota : U(Y_0) \hookrightarrow U(C_0)$ fixes a norm-one vector z_0 in its orthogonal complement. The last core is identified with the definite column vertex. The last row and core form the Bessel pair used in (16). At a real place the initial row has negative index $m/2$. These groups and the test family are those of the auxiliary grid.

Fix the level, the Weil data and the quotient measures of the grid. Let t be a nonzero element of the bottom cuspidal multiplicity space at principal level and trivial finite level character. Fix a nonzero vector w of its prescribed smooth real representation. The weak separation assertion asks for a prescribed finite coset tuple Φ, a permitted column c at the last core, real Fock polynomials ρ and a real column covector ξ for which the last Bessel period of c_ξ and the first-row lift of $t(w)$ is nonzero. Here c_ξ is the scalar form obtained by evaluating the coefficient of c against ξ. The finite tuple may depend on t, w. Statements of core detection, core span and ascent.

2.4.2 The bottom corank-one period

AI-proposed candidate (detection at the prescribed principal level). For each size exponent ε and depth constant $C_F \geq 6510$, there is a common threshold $x_0(\varepsilon, C_F)$. For every permitted size datum of exponent ε and scale $x \geq x_0(\varepsilon, C_F)$, take any compatible row level and character family with an admissible base tuple, any auxiliary grid at constants $(1, C_F, b)$ for that family, and its admissible Weil data and row quotient Haar measures. Let M_Y and M_C be models of the prescribed bottom row and core discrete-series types. Let $t \neq 0$ be any cuspidal multiplicity map for M_Y at the full row principal subgroup on Λ_Y of modulus $(g_0 g_1)$, with trivial finite character. For every nonzero vector w in the tensor product of the smooth real modules of M_Y, there exist a cuspidal multiplicity map y for M_C of trivial finite character at $K_C = K(\Lambda_Y \perp \mathcal{O}_E z_0, \mathfrak{m})$, and a K_∞-finite vector v in its smooth tensor model, such that $a = y(v)$ satisfies

$$\int_{[U(Y_0)]} a(\iota h) t(w)(h)\, dh \neq 0. \tag{22}$$

Here S_0 consists of the chosen place v_0 over the least split auxiliary rational prime q_0 of the grid and the places over $g_0 g_1 p_{\mathrm{de}} p'_{\mathrm{de}}$. With one specified split prime $\mathfrak{w}_v$ over each $v \in S_0$, the core modulus is $\mathfrak{m} = \prod_{v \in S_0} \mathfrak{w}_v^{d_v}$, where $d_v = 2$ away from v_0 and

$$d_{v_0} = 1 + \left\lfloor \frac{C_F [F : \mathbb{Q}]^2 (r+1)^6 \log\bigl(b(r+1) g_0 g_1 p_{\mathrm{de}} p'_{\mathrm{de}} |\Delta_E|\bigr)}{\log q_0} \right\rfloor .$$

The modulus, row measure and real types are fixed before t, w. Quantifiers for detection at the grid level · Core modulus · Depth at the auxiliary place.

The real branching pair uses the dual of the row discrete series because the period is bilinear. Xue gives its local dimension criterion. [36, Theorem 1.1] The global corank-one period theorem concerns a fixed Hermitian Arthur parameter and uses central value nonvanishing. Its existence result with one rank fixed takes the larger-rank parameter as data. The displayed detection fixes each lower form at the stated

level and seeks an upper witness in a prescribed real type and level. [23, Theorems 1.1.5.1 and 1.1.5.3] The argument below proves this first step for the admitted grid. The next step is to span the prescribed core test space by descending chains, then ascend with the permitted row tests.

For $U(p,q)$ and the subgroup $U(p-1,q)$ fixing a positive line, He identifies the discrete part of the restriction of each discrete series. It is the direct sum of the discrete series of the subgroup whose Harish–Chandra parameters satisfy the signed interlacing rule, each with multiplicity one. [24, Theorem 0.1] To verify the signed interlacing, merge the core and row-dual parameters in decreasing order. The prescribed word is

$$a_0, \dots, a_m, \quad v_2, v_2 - \tfrac{1}{2}, \dots, v_k, v_k - \tfrac{1}{2}, \quad -b_1, \dots, -b_m.$$

The a_i are core entries with signs alternating from positive. All entries in the middle block are positive, with v_i from the row dual and $v_i - \frac{1}{2}$ from the core. The final block consists of row-dual entries with signs alternating from positive. The strict ordering uses $a_m > v_2$, $v_i - v_{i+1} \geq 1$ and $v_k + b_1 \geq 1$. Since m is even, a_m has positive sign. Thus two successive entries have the same sign precisely when they belong to different groups, including at the two block boundaries. This is the signed interlacing condition in He's theorem for deletion of a positive line. Merged real parameter word · Matching the word to the actual models.

Put $d = [F : \mathbb{Q}]$ and $n = r + 1$. In the adapted real frames, write $|z|_v$ for the positive Euclidean norm of a vector of C_0 at a real place v, normalized by $|z_0|_v = 1$. The needed coefficient k_∞ is the tensor of matrix coefficients of translated lowest-compact-type vectors in the specified ambient discrete series. It must reproduce the value of $f = t(w_*)$ for a fixed nonzero real row vector w_*. In the real Haar factors of dh, the required identities and bounds are

$$\int_{U(Y_0)_\infty} k_\infty(\imath h) f(xh)\, dh_\infty = c_\infty f(x), \qquad c_\infty \neq 0, \tag{23}$$

$$\int_{U(Y_0)_\infty} |k_\infty(g\imath h)|\, dh_\infty \leq e^{Q} |c_\infty| \prod_{v|\infty} |g_v z_0|_v^{-(2n+5)}. \tag{24}$$

The first variation along the row Lie algebra has an L^1-bound with constant $e^Q|c_\infty|$. Moving that variation to the coefficient in (23) gives the row Lipschitz constant $e^Q\|f\|_\infty$. The required uniform range is $Q = Cdn^4 \log b$ with an absolute C. The vector before translation must also have sufficient matrix-coefficient decay to form a cuspidal Poincaré series. The lower bound on the Hilbert projection and the transversal and first-variation bounds use the same translated vector. The reproducing identity follows from branching and Schur orthogonality.

At one real place let $H = U(Y_0)(\mathbb{R})$, let π be the ambient representation restricted to H, and let σ be the row dual. He gives a unique Hilbert copy $j\sigma$ of σ inside π, with j an isometry. Take the specified translated ambient vector v, with $j^*v \neq 0$ and $k_v(h) = \langle v, \pi(h)v\rangle$ in $L^1(H)$. Inner products are linear in the second argument. Let u_0 be a unit lowest-compact-type vector in σ. The real weight gaps of the row imply that its diagonal coefficient is integrable. Write $\mathcal{D}$ for the linear span of its H-translates, which is dense in σ. All coefficients with both vectors in $\mathcal{D}$ are integrable by left and right translation.

For $a, c \in \mathcal{D}$ let $e_{a,c}(h) = d_\sigma \overline{\langle a, \sigma(h)c\rangle}$, with d_σ the formal degree for the chosen Haar measure. Schur orthogonality gives $e_{a,c} * e_{a',c'} = \langle c, a'\rangle e_{a,c'}$ and $e_{a,c}^* = e_{c,a}$. On any unitary representation τ of H, integrate τ against these L^1 functions. If $\tau(e_{u_0,u_0})$ has nonzero image, choose a unit vector z in its image. The matrix-unit identities make $a \mapsto \tau(e_{a,u_0})z$ an isometry on $\mathcal{D}$. Left translation of e_{a,u_0} by $g \in H$ gives $e_{\sigma(g)a,u_0}$, so the isometry extends to a copy of σ in τ. Applied to the orthogonal complement of $j\sigma$, uniqueness shows that $\pi(e_{u_0,u_0})$ vanishes there. Factoring each $e_{a,c}$ through this projection gives the same vanishing for all matrix units. On $j\sigma$ they act by the usual rank-one operators.

Taking the coefficient against v gives

$$\int_H \overline{k_v(h)}\, \langle a, \sigma(h)c\rangle\, dh = d_\sigma^{-1} \langle a, j^*v\rangle \langle j^*v, c\rangle.$$

The integral extends continuously to all a, c by $k_v \in L^1(H)$. With $c = j^*v/\|j^*v\|$, this is the reproducing eigenvector with eigenvalue $\|j^*v\|^2/d_\sigma$. The vector is smooth since j^* is a bounded H-intertwiner and v is smooth. Transporting back to the row model conjugates the two coefficients in the bilinear convention. Multiplying the identities over the real places gives (23) in the prescribed measures.

The smooth tensor in the row multiplicity space is algebraic. Its map t extends to the Hilbert tensor, as follows. Freeze all but one real factor. Point evaluation is continuous on the remaining smooth Fréchet space by the defining Sobolev condition on multiplicity maps. Every such cuspidal value is a bounded function by rapid decay for finite-level smooth cusp forms of moderate growth, finite under the center. See Miller–Schmid [88, Theorems A and 2.23–2.24]. The uniform boundedness principle applied to all point evaluations therefore gives continuity of this one-variable map into bounded functions and hence into $L^2([U(Y_0)])$, since the quotient has finite volume. The pullback of the L^2 inner product is invariant on the compact-finite module. Such a form is a nonnegative scalar multiple of the model form: on each finite-dimensional compact isotypic space compare the forms by the Riesz operator, then invariance makes the combined operator an endomorphism of the irreducible Harish–Chandra module. Density of compact-finite vectors in the smooth topology extends the identity to the smooth factor. On pure tensors, applying the identity in each factor in turn gives

$$\|t(u_1 \otimes \cdots \otimes u_d)\|_2^2 = c_t \prod_j \|u_j\|^2.$$

Separate polarization gives the same scalar times the Hilbert inner product on all pairs of algebraic tensors. The scalar is positive for $t \neq 0$, since continuous functions zero in L^2 vanish. Thus t extends as a scalar isometry and intertwines the real group on the Hilbert tensor.

First apply the integrated Schur operator on the Hilbert space of the row representation and use the automorphic intertwiner. The pointwise integral against the smooth bounded row form is continuous by the L^1-bound on the kernel. By Fubini against compactly supported continuous functions on the quotient it agrees with the Hilbert-space integral. The resulting equality of continuous functions holds at each x. Quantitative real input · Integrable coefficients at the actual parameter gaps · Quantitative real projection · The prescribed real types and frames · Real reproducing normalization.

2.4.3 Quantitative real branching and the period kernel

Work at one real place. Write $G = U(p, q)$, $H = U(p-1, q)$, and $N = p + q = n$ for the ambient dimension. Use the adapted frame with signature matrix $J = \mathrm{diag}(\epsilon_1, \dots, \epsilon_N)$. The inclusion fixes the positive coordinate vector $z = e_{i_0}$. Let λ and μ be the ambient and row-dual Harish–Chandra parameters, respectively, in decreasing order with their signs. We use the following size properties of the grid:

$$\lambda_i - \lambda_{i+1} \geq 2, \qquad |\lambda_i|, |\mu_j| \leq B = b^2, \qquad \log(8NB) \leq 6 \log b.$$

Within either parameter the gap between entries of opposite signs is at least $1000N^3$. Indeed the a_i have adjacent gap at least $1000N^3$, as do the b_i, and a_m is separated from the middle block by at least $1000N^3$. In the row dual the positive middle block is separated from the first negative entry in the final block by the gap $b_2 - b_1$, in addition to $v_k + b_1 \geq 2$. The middle block in each parameter has successive spacing at least 2. The parameters are integral after subtracting the respective half-sum of positive roots. The signature includes at least one positive and one negative coordinate in each group. These are the real grid estimates at the chosen scale. Gaps and size of the actual real grid

Let $P = j^*$, and take unit vectors v_0 and u_0 in the lowest compact types of π and σ. Use Haar measures dg and dh in the two groups and denote the corresponding formal degrees by d_π and d_σ. Write $\|\cdot\|_F$ and $\|\cdot\|_{\mathrm{op}}$ for the Frobenius and Euclidean operator norms of matrices. All implicit constants below are absolute. Set $T = 8NB$ and $L = T^2$.

Derived vectors in the lowest compact type. Put $E_{ab} = d\pi(e_{ab})$ on the smooth space. Infinitesimal unitarity gives $E_{ab}^{*} = \epsilon_a \epsilon_b E_{ba}$ there. Let $C_G = \sum_{a,b} E_{ab} E_{ba}$ and let C_K be the same sum over equal-sign pairs. For any smooth vector y,

$$\sum_{a,b} \|E_{ab}\, y\|^2 = \langle y, (2C_K - C_G) y \rangle.$$

On the lowest type the highest weight is $\lambda + \rho_{\mathrm{nc}} - \rho_K$, with each coordinate bounded in absolute value by $B + N/2$. The space generated from this type by at most a derived factors is a quotient of a direct sum of its tensor products with $\mathfrak{g}_{\mathbb{C}}^{\otimes j}$, $j \le a$, as a compact module. The weights of the initial type lie in the convex hull of its compact Weyl orbit. Each adjoint weight has coordinates of absolute value at most 1. Hence each highest weight ν of the generated space satisfies $|\nu_i| \le B + N/2 + a$. The Casimirs have eigenvalues $\|\nu + \rho_K\|^2 - \|\rho_K\|^2$ and $\|\lambda\|^2 - \|\rho_G\|^2$, respectively. For $a \le N + 2$, the square root of the upper bound for $2C_K - C_G$ is at most L. Cauchy–Schwarz now gives

$$\|d\pi(X_1) \cdots d\pi(X_j) v_0\| \le L^j \prod_{i=1}^{j} \|X_i\|_F \qquad (j \le N + 2). \tag{25}$$

Derived vector estimate

A moment in the transverse coordinate. Let $F(g) = \langle j u_0, \pi(g) v_0 \rangle$. The required moment is

$$\int_G \|g^{-1} z\|^2 |F(g)|^2 \, dg \le T^{8N} \int_G |F(g)|^2 \, dg = T^{8N} / d_\pi. \tag{26}$$

Use $\Pi = \pi \otimes (\mathbb{C}^N)^*$ and put $x_a(y) = y \otimes e_a^*$. On coordinate tuples x its action is $(\Pi(g)x)_a = \sum_l (g^{-1})_{la} \pi(g) x_l$. Its invariant Hermitian form is $[x, y] = \sum_a \epsilon_a \langle x_a, y_a \rangle$.

The Harish–Chandra module of Π is the direct sum of the discrete-series modules π_i with parameters $\lambda - e_i$, in the same signed chamber. To see this, use the upper triangular θ-stable Borel in the ordered frame. The inducing character of π on its compact Cartan has weight $\lambda - \rho_G$. Apply the finite-dimensional tensor identity for cohomological induction to the Borel filtration of $(\mathbb{C}^N)^*$. Its graded characters have weights $-e_i$. Every shifted parameter is regular in the original chamber, since the ambient gaps are at least 2. Good-range concentration thus makes the degree-$\dim(\mathfrak{n} \cap \mathfrak{k})$ long exact sequences short exact. Each graded factor is the corresponding discrete series. Their infinitesimal characters are distinct, since each $\lambda - e_i$ remains decreasing and the tuples are different. The primary decomposition for the center splits the filtration. This uses the tensor identity and filtered complex in [78, Proposition 3.6 and (14)] and the good-range Borel construction in [77, Theorems 5.35, 5.99 and 8.2, and Chapter XI]. The smooth spaces give the same direct sum by Casselman–Wallach globalization [79, Theorems 1.1 and 10.6].

Write p_i for the projections and R_i for the maps onto the smooth unitary models of π_i. The summands are orthogonal for $[\cdot, \cdot]$. Indeed the total quadratic Casimir is symmetric for this form and its eigenvalues on the summands are real and distinct. On each summand the form is a nonzero real multiple of the unitary inner product, by irreducibility and Schur's lemma on compact-finite vectors. Nondegeneracy follows from the direct sum and orthogonality. Rescale R_i so that

$$[x, p_i y] = s_i \langle R_i x, R_i y \rangle, \qquad s_i \in \{1, -1\}. \tag{27}$$

The matrix operator $\mathcal{E} = (E_{ab})$ on tuples satisfies $C_\Pi = C_G I + N I - 2\mathcal{E}$, by expanding the action on the tensor with the dual standard representation. Thus its eigenvalue on p_i is $\lambda_i + (N-1)/2$. Let $Q_G(u)$ be the Capelli column determinant with column l shifted by $u - l + 1$. Its infinitesimal-character eigenvalue is $\prod_j (u + \lambda_j - (N-1)/2)$. Let $A(u)$ be the matrix whose (a, l) entry is the signed Capelli minor omitting row l and column a, with consecutive column shifts starting from u. The Capelli comatrix identity gives

$$A(u)(\mathcal{E} + (u - N + 1) I) = Q_G(u) I.$$

This is the specialization of the antisymmetrizer identity in [80, Proposition 2.7(4), Remarks 2.8(1) and 2.11] under $T(u) \mapsto I + \mathcal{E}/u$. Group the product by the row of its last factor, after fixing the column sequence

consisting of all columns except a followed by the final column. This gives the stated sign and Capelli shifts. The eigenvalue of $Q_G(u)$ also follows by evaluation on a highest weight vector: the diagonal contribution in column l is $u - l + 1 + \lambda_l - \rho_{G,l}$. The largest displaced column in each other summand gives a raising operator acting on the highest weight vector after scalar factors. Evaluating the comatrix identity on each eigenspace and taking $u_i = -\lambda_i + (N-1)/2$ gives

$$p_i = \frac{A(u_i)}{\prod_{k \neq i}(\lambda_k - \lambda_i)}.$$

The denominators have absolute value at least 1. By (25), each signed minor applied to v_0 has norm at most $(N-1)!(B+N+L)^{N-1} \le T^{4N}$. In particular, $|s_i|\|R_i x_l(v_0)\|^2 \le T^{4N}$ by (27) and the (l,l) minor.

The (i_0, i_0) minor is the Capelli determinant of the complement of z. On smooth vectors the bounded intertwiner P commutes with its derived action. Comparing the two Harish–Chandra shifts of rank N and $N-1$ therefore gives

$$P\big((p_i x_{i_0}(y))_{i_0}\big) = \gamma_i P y, qquad \gamma_i = \frac{\prod_j(\mu_j - \lambda_i + 1/2)}{\prod_{k \neq i}(\lambda_k - \lambda_i)}, \qquad |\gamma_i| \le T^N. \tag{28}$$

Put $\xi = ju_0$. The $L^1(H)$ matrix unit $e = d_\sigma \overline{\langle u_0, \sigma(\cdot)u_0\rangle}$ from the preceding argument acts on $\pi|_H$ as the orthogonal projection onto $\mathbb{C}\xi$. It acts on $\pi_i|_H$ as an orthogonal projection Q_i, by the self-adjoint idempotent identities. Choose a smooth vector η with $\langle \eta, \xi\rangle = 1$ and put $a = x_{i_0}(\eta)$. Since H fixes this coordinate, integrating the equality in (27) on $\Pi(h)y$ against $e(h)$ gives, for smooth tuples y,

$$\langle \xi, (p_i y)_{i_0}\rangle = s_i \langle b_i, R_i y\rangle, \qquad b_i = Q_i R_i a.$$

With $y = a$, (28) yields $s_i\|b_i\|^2 = \gamma_i$. This averaging takes place in the unitary Hilbert spaces of the summands. For each l we obtain

$$(g^{-1})_{li_0} F(g) = \sum_i s_i \langle b_i, \pi_i(g) R_i x_l(v_0)\rangle.$$

Schur orthogonality bounds the square integral of each summand by T^{5N}/d_{π_i}. Harish–Chandra's formal-degree formula on the same group and for the same Haar measure is proportional to $\prod_{a<b}|\lambda_a - \lambda_b|$, with a common factor for all these parameters. For the compact central factor, Haar integration of the squared coefficient contributes the same factor in every representation, since the central character is unitary. Decreasing one coordinate by 1 changes only $N-1$ factors and each ratio is at most 2, because the original differences have absolute value at least 2. Thus $d_\pi/d_{\pi_i} \le 2^{N-1} \le T^N$. [62, (3.10), (3.12) and Theorem (8.1)] Cauchy–Schwarz over i, followed by summation over l, introduces at most N^3 in the bound T^{6N}/d_π. This proves (26). Weighted Schur estimate

A translate normalized by its reproducing scalar. For a matrix unitary group D of matrix size r_D, let $B_D(\delta) = \{g \in D : \|g - I\|_F \le \delta\}$ and $D(R) = \{g \in D : \|g\|_F \le R\}$. Any Haar measure on it satisfies

$$\operatorname{vol} D(R) \le (1 + 2R^2/\delta)^{2r_D^2} \operatorname{vol} B_D(\delta). \tag{29}$$

Indeed $g^{-1} = Jg^*J$ gives $\|g^{-1}\|_F = \|g\|_F$. A maximal δ/R-separated subset of $D(R)$ gives a cover by its left translates of $B_D(\delta)$ and the Euclidean packing bound in real dimension $2r_D^2$ gives the number of points. With $\delta = 1/4$ and dyadic shells, (29) gives

$$\int_D \|g\|_F^{-u}\, dg \le 2\, 33^{2r_D^2} \operatorname{vol} B_D(1/4), \qquad u \ge 4r_D^2 + 1. \tag{30}$$

Let $Y = 2T^{8N}$. By (26), the region $S = \{g : \|g^{-1}z\|^2 \le Y\}$ carries square integral at least $1/(2d_\pi)$. This region is left H-invariant. A vector u of Hermitian norm one is the image of z under a matrix b_u with $\|b_u\|_F^2 \le N + 2\|u\|^2$. To construct it, take $u = u_+ + u_-$ in the positive and negative subspaces, with lengths

a_+, a_- satisfying $a_+^2 - a_-^2 = 1$. First map z to the direction of u_+ by an element of the compact group, then act in the span of the two directions by the block with diagonal a_+ and off-diagonal a_-. The compact map alone suffices when $u_- = 0$.

Apply this with $u = g^{-1}z$ for $g \in S$. Since $gb_u \in H$, every left coset in S has a representative of Frobenius norm at most $R_0 = (N + 2Y)^{1/2}$. Put $R_G = 2R_0$. The left translates of each representative by $B_H(1/4)$ belong to $G(R_G)$. Quotient integration with dg and dh bounds the quotient volume of $H\backslash S$ by $\operatorname{vol} G(R_G)/\operatorname{vol} B_H(1/4)$. It also identifies the square integral on each fiber with $\|P\pi(g)v_0\|^2/d_\sigma$, by Schur on H. Consequently some representative g_* has $\|g_*\|_F \le R_0$ and

$$c = \|P\pi(g_*)v_0\|^2/d_\sigma \ge \frac{\operatorname{vol} B_H(1/4)}{4d_\pi \operatorname{vol} G(R_G)}. \tag{31}$$

Here the factor 4 allows choosing a value at least half the average, and the fiber integral is independent of the representative.

For $\delta = (4L)^{-1}$, the logarithm of $g \in B_G(\delta)$ belongs to the real Lie algebra and has Frobenius norm at most 2δ. This follows from the convergent logarithm series and $g^{-1} = Jg^*J$. Integrating the derived orbit along the exponential and using (25) shows that $|\langle v_0, \pi(g)v_0\rangle| \ge 1/2$ on this ball. Schur on G gives $d_\pi \operatorname{vol} B_G(\delta) \le 4$. As $R_G \le T^{O(N)}$, (29) and (31) yield

$$\|g_*\|_F \le T^{O(N)}, \qquad c \ge \operatorname{vol} B_H(1/4)\exp(-C_0N^3\log T) \tag{32}$$

for an absolute C_0.

Decay of the lowest-type coefficient. For a group of matrix size M in this argument, let β be its Harish–Chandra parameter with opposite-sign gaps at least $1000M^3$. If τ is the representation and V is its lowest compact type, we claim

$$\|P_V\,\tau(g)|_V\|_{\mathrm{op}} \le \exp(100M^3)\|g\|_F^{-40M^2}. \tag{33}$$

Use the polar decomposition $g = k\exp X$, with X Hermitian and supported in the off-sign blocks. The compact factor is isometric and preserves V. The unitary group is connected reductive with finite center on its derived group. Sun's bottom-layer formula applies to the compression of $\exp X$ for the upper triangular θ-stable Borel and the Cartan character of weight $\beta - \rho_G$. The bottom type is the lowest compact type, with highest weight $\beta + \rho_{\mathrm{nc}} - \rho_K$, and occurs with multiplicity one. The weight is dominant for the induced compact ordering because the entries of β are decreasing with successive gaps at least 1. [81, Theorem 6.3 and (25)–(26)]

In the formula take $x = \exp(X/2)$. For l in the compact group, the map $H_{\mathfrak{q}}$ extracts the diagonal from the upper triangular factor of lx modulo the compact real form of the complex group. Thus $\widetilde{H}_{\mathfrak{q}}(lx)$ is the positive diagonal D in the factorization $\exp(lXl^{-1}) = uDu^*$ with u upper unitriangular. The scalar $(\delta_{\mathfrak{n}}\chi)(D)$ in Sun's formula has weight $\beta + \rho_G$, since $\delta_{\mathfrak{n}}$ contributes $2\rho_G$. If $D_i(\exp Z)$ denotes the principal minor on the coordinates $i+1, \dots, M$, this scalar is

$$\prod_{i=1}^{M-1} D_i(\exp Z)^{-(\beta_i - \beta_{i+1}+1)}, \qquad Z = lXl^{-1}.$$

Indeed the trailing determinant is the product of the corresponding diagonal entries of D, and $\det D = 1$. The other factor in the integral is the conjugate of the projection onto a highest weight line. With compact probability measure, its average is $I_V/\dim V$. It follows by positivity that the norm of the compression is bounded by the supremum of the displayed scalar, with the factor $\dim V$ in Sun's formula canceled by this average.

To bound a minor, view it as the expectation of e^U in the spectral probability measure of the exterior-power derived operator of Z on the unit wedge for the trailing set. If $t = \|X\|_{\mathrm{op}}$, then $|U| \le Mt$. Its mean is zero because the diagonal of Z is zero. Its second moment is $\sum_{a\le i<b} |Z_{ab}|^2$, as seen by replacing one

vector in the wedge. For $\theta = (\max(1, Mt))^{-1}$, use $e^v \ge 1+v+v^2/4$ for $|v| \le 1$, followed by $\log(1+h) \ge h/2$ for $0 \le h \le 1$. Convexity of the logarithm of the moment-generating function then gives

$$\log D_i(\exp Z) \ge \tfrac{1}{8}\theta \sum_{a \le i < b} |Z_{ab}|^2.$$

The sum of these variances with weights $\beta_i - \beta_{i+1} + 1$ has coefficient $\beta_a - \beta_b + b - a$ at $|Z_{ab}|^2$. Every nonzero entry has opposite signs, so the sum is at least $1000M^3\|Z\|_F^2/2$. For $t \ge 1$ the compression norm is therefore at most $\exp(-125M^2 t/2)$. For $t < 1$ it is at most 1 by unitarity. The inequality $\|g\|_F \le \sqrt{M}e^t$ proves (33). This argument also supplies the row integrability used above by (30), with $M = N-1$. Lowest-type power decay

Transversal integrals and first variation. Let $v = \pi(g_*)v_0$, $k(g) = \langle v, \pi(g)v\rangle$, $G_0 = \|g_*\|_F$, and $D_0 = 40N^2$. Conjugating (33) by g_* gives

$$|k(g)| \le M_0 \|g\|_F^{-D_0}, \qquad M_0 = \exp(100N^3) G_0^{2D_0}.$$

For $a \in G$, put $\ell = \|az\| \ge 1$. By the same block construction write $a = b_{az} h_1$ with $h_1 \in H$. Setting $x = b_{az} h$, we have $\|x\|_F \ge \ell$ and $\|h\|_F \le \sqrt{N+2}\,\ell\|x\|_F$. For $u = 4(N-1)^2 + 1$, this gives

$$\int_H |k(ah)|\, dh \le M_0 (N+2)^{u/2} \ell^{-D_0+2u} \int_H \|h\|_F^{-u}\, dh.$$

In particular $D_0 - 2u \ge 2N + 5$.

For X in the Lie algebra of H with $\|X\|_F \le 1$, let $A_h(t) = k(\exp(-tX)h)$. The bound for the second derived word in (25) and $\|\operatorname{Ad}(g_*^{-1})X\|_F \le G_0^2$ give $\sup_t |A_h''(t)| \le L^2 G_0^4$. For $|t| \le 1$ we also have $|A_h(t)| \le M_0 e^{D_0} \|h\|_F^{-D_0}$, since the operator norm of $\exp(tX)$ is at most e. Taylor's formula at $t = \|h\|_F^{-D_0/2}$ bounds

$$|A_h'(0)| \le (2M_0 e^{D_0} + L^2 G_0^4/2)\|h\|_F^{-D_0/2}.$$

Both integrals are bounded using (30). Dividing by the same c of (32), we get the stated transversal and first-variation bounds with a single $\exp(C_1 N^3 \log T)$. For the product over real places they give one $Q \le Cdn^3 \log b$, with $d = [F : \mathbb{Q}]$. The weaker exponent $Cdn^4 \log b$ in the depth argument also suffices. Using the reproducing identity at $x\exp(tX)$, translation puts the variation on the kernel as $k(\exp(-tX)h) - k(h)$. Its L^1-norm is at most $|t|$ times the first-variation bound, by integration along that curve and left invariance of Haar measure. This gives the stated Lipschitz bound for f and supplies the transversal estimate for the Poincaré seed. Transversal integration · Integrated first variation

2.4.4 The Poincaré test at the prescribed depth

For the level and the row quotient, put $A = r^2 d\, g_0 g_1\, |\Delta_E|^{1/[E:\mathbb{Q}]}$. The prescribed $\mathfrak{m}$ must allow a real number $0 < s \le 1/2$ with

$$e^Q A(2s)^{1/m} \le \tfrac{1}{2}, \qquad 2e^{Q+(6n+9)d} \le R, \qquad R = s^d \sqrt{\mathrm{N}\mathfrak{m}} > 1. \tag{34}$$

Take $s = \frac{1}{2}(2e^Q A)^{-m}$, so the first bound is an equality. With $L = \log b$, the size hypotheses give $d, m \le n$, $L \ge 1$, $n \ge C + 1$ at sufficiently large scale, and $\log A \le n^3 + 21L$. For the last inequality, $r^2 d \le n^3$, $\log(g_0 g_1) \le 20L$ and $\log |\Delta_E|^{1/[E:\mathbb{Q}]} \le L$. The prescribed depth at the chosen split auxiliary place, with $\mathrm{N}\mathfrak{w}_{v_0} \ge q_0$, gives $\log \mathrm{N}\mathfrak{m} > C_F d^2 n^6 L$. Hence

$$\log 2 + Q + (6n+9)d + d\log(1/s) \le (2C+42)d^2 n^5 L < \tfrac{1}{2}\log \mathrm{N}\mathfrak{m}.$$

This gives the remaining bounds in (34). The short-lattice estimate of Section 2.4.5 says that if some nonzero y of the transported rational row lattice $\mathfrak{a}\Lambda_Y(x)$, for a nonzero integral ideal $\mathfrak{a}$, satisfies

$$\prod_{v|\infty} |x_v^{-1} y|_v < s^d \sqrt{\mathrm{N}\mathfrak{a}},$$

then $|f(x)| \le e^{Q}\|f\|_\infty A(2s)^{1/m}$. The estimate uses reduction to a rational unipotent radical. Cuspidality makes the average along that radical zero and the small real representatives at the row level bound the variation away from $f(x)$. Small-lattice estimate · Depth of the actual grid.

Here is the nonvanishing argument from these inputs. Choose x with $|f(x)| > \|f\|_\infty/2$. Such an x exists since the nonzero intertwiner t and its irreducible real source give $f \neq 0$. The small-lattice estimate implies the reverse weak inequality for every $\mathfrak{a}, y$ at this x. By orthogonal addition of the norm-one line it follows that

$$\prod_{v|\infty} |\iota(x_v)^{-1} z|_v \ge R \quad (0 \neq z \in \mathfrak{m}\Lambda_C(\iota x)).$$

Indeed, if the Y_0 component is nonzero, its product gives the bound. A nonzero vector on the added line has coefficient in $\mathfrak{m}$, so the product is at least $\sqrt{\mathrm{N}\mathfrak{m}}$.

Use the seed $k = k_\infty 1_{K_C}$ in the Poincaré series

$$P_x(g) = \sum_{\gamma \in U(C_0)(F)} k\big(\iota(x)^{-1}\gamma g\big).$$

Here is the justification on the ambient smooth representation. At a real place the translate in the first coefficient vector changes the constant in (33). The bound with exponent $40n^2$ then holds when the second vector is either in the lowest type or its chosen translate. For a fixed second K_∞-finite vector an exponent $20n^2$ suffices. Indeed the lowest type generates the irreducible Harish–Chandra module under the derived action. For a derivative of order $j > 0$ of a translated lowest-type coefficient, use the values on a finite stencil in real exponential coordinates, at scale $h \le 1$, exact on polynomials of degree at most $2j-1$. Such a stencil is obtained by interpolation on a $2j$-point grid in each variable. The absolute value of the derivative is bounded by $C_j(h^{-j}\|g\|_F^{-40n^2} + h^j)$. The first term uses the power decay, uniform for each right shift in the stencil. The remainder uses derivatives of total order $2j$, bounded independently of g by unitarity and smoothness of the lowest-type vector. Taking $h = \|g\|_F^{-20n^2/j}$ proves the assertion for every fixed derived word and hence for each fixed compact-finite vector.

For any g, two distinct rational terms surviving the finite indicator have their quotient, in the order ab^{-1} for the real arguments a, b of their coefficients, in the discrete real image of

$$U(C_0)(F) \cap \iota(x_f) K_C \iota(x_f)^{-1},$$

conjugated by $\iota(x_\infty)^{-1}$. Choose $\delta > 0$ so that every nonidentity element of this subgroup at infinity has at least one component of Frobenius distance δ or more from the identity. In the box $\|a_v\|_F, \|b_v\|_F \le R_v$, the quantity $\|a_v b_v^{-1} - 1\|_F$ is at most $R_v\|a_v - b_v\|_F$. Thus the products of real balls of radii $\delta/(2R_v)$ centered at the distinct arguments have disjoint interiors. Volume comparison bounds the number of terms by $\prod_{v|\infty}(1+2R_v^2/\delta)^{2n^2}$. Dyadic summation with exponent $20n^2$ proves uniform boundedness and convergence of the series for every fixed compact-finite second vector and for the chosen diagonal seed. Convergence is locally normal: on compact right translates the matrix norms are comparable with constants independent of the rational term. The same bound for every fixed derived word allows differentiation term by term on the compact-finite module.

For a compact-finite second vector, unfold the constant term of this series along a proper rational parabolic. The quotient of its unipotent radical by the rational points is compact, and the absolute bound above permits unfolding. Each resulting integral over the adelic unipotent radical vanishes by the real discrete-series constant-term theorem at a real place. The theorem applies to the lowest-type vector and the compact-finite second vector. Translation of the first vector moves the translate into the argument of the matrix coefficient and leaves that theorem applicable. The level is K_C and the right compact and derived actions on the Poincaré series agree with those on the second vector.

These bounded, cuspidal functions lie in the cuspidal Hilbert space, since the adelic quotient has finite volume. The estimates for all fixed derived words identify differentiation with right differentiation in

L^2 as well, by dominated convergence. The resulting Harish–Chandra-module map pulls the invariant Hermitian form back to a scalar multiple of the given form, by irreducibility. It therefore extends to a bounded map from the ambient Hilbert representation. This map intertwines the real group. Indeed the compact-finite vectors are analytic for each real one-parameter subgroup. Their images have the same bounds on all powers of each infinitesimal generator, up to the single operator norm of the map. The intertwining follows first by power series and then along the entire one-parameter subgroup by analytic continuation. These subgroups generate the connected real group.

At the prescribed finite level, local Sobolev embedding on each real quotient chart gives the smooth representative of the image of every smooth vector. It also bounds evaluation at any fixed adelic point by a finite sum of Hilbert norms of derived vectors, with the order depending only on the real dimension. This proves the evaluation condition in the actual multiplicity space. The same bound at xg_∞ follows by evaluating $\pi(g_\infty)u$ at x. Commuting the fixed Sobolev operators with $\pi(g_\infty)$ via the adjoint action gives one moderate-growth exponent for all right derivatives, by applying the bound to each derived u. The center acts by the given infinitesimal character. Compact averaging approximates each smooth vector in every Sobolev seminorm, hence its image locally uniformly by the compact-finite images. The unipotent quotients in the constant terms are compact, so cuspidality persists pointwise. This gives a smooth cuspidal multiplicity map y on the algebraic tensor of the prescribed local models, of trivial character at K_C. Harish-Chandra, Theorem 18.1. Cuspidal Poincaré series at the prescribed level.

In the pairing with f, unfold using right cosets by $\iota U(Y_0)(F)$. For each coset represented by γ the finite part of the integral over $U(Y_0)(\mathbb{A}_F)$ has support either empty or a left translate of $K' = \iota^{-1}(K_C)$ in $U(Y_0)(\mathbb{A}_{F,f})$. The orthogonal lattice sum gives the row congruence modulo $\mathfrak{m}$. For a unitary automorphism of a self-dual lattice, the inverse is the adjoint. Taking the adjoint of the inverse congruence therefore gives the congruence modulo $\overline{\mathfrak{m}}$ as well. Each prime of E over g_0g_1 divides $\mathfrak{m}$ or its conjugate, and (g_0g_1) is squarefree in $\mathcal{O}_E$. Thus $\mathfrak{m} \cap \overline{\mathfrak{m}} \subset (g_0g_1)$ and K' lies in the bottom principal subgroup. For the stabilizer coset, left translation by x, K'-invariance and (23) give the term $\mathrm{vol}(K')c_\infty f(x)$. Real and finite factors of the stabilizer term · The bottom lattices and levels · Cuspidal Poincaré series and its multiplicity map.

For every other supported coset, put $z = \gamma z_0$. The support condition implies $z - z_0 \in \mathfrak{m}\Lambda_C(\iota x)$ by applying the congruence to z_0. The stabilizer property makes $z - z_0$ nonzero and distinct for distinct cosets. Since z has Hermitian norm one, $|\iota(x_v)^{-1}z|_v \geq 1$. The triangle inequality gives

$$|\iota(x_v)^{-1}z|_v \;\geq\; (1 + |\iota(x_v)^{-1}(z - z_0)|_v)/3.$$

By (24), the absolute tail is bounded by $\|f\|_\infty e^Q |c_\infty| \,\mathrm{vol}(K')$ times the sum of $\prod_{v|\infty}((1+|\iota(x_v)^{-1}y|_v)/3)^{-(2n+5)}$ over nonzero y of $\mathfrak{m}\Lambda_C(\iota x)$.

To bound this sum, group the vectors into boxes indexed by $j_v \geq 0$ with $2^{j_v} \leq 1 + |\iota(x_v)^{-1}y|_v < 2^{j_v+1}$. Set $D_{\mathrm{j}} = \prod_v 2^{j_v+1}$. In the norm on $\prod_v \mathbb{C}^n$ obtained by taking the maximum of the component norms scaled by 2^{-j_v-1}, the box lies in the unit ball and two distinct points are at distance at least $D_{\mathrm{j}}^{-1/d}$. This follows from the lower product bound for their difference. Packing disjoint balls of half that radius bounds the count by $(1 + 2D_{\mathrm{j}}^{1/d})^{2nd} \leq 3^{2nd} D_{\mathrm{j}}^{2n}$. The product of the quantities $1 + |\iota(x_v)^{-1}y|_v$ is at least R, so one power of that product supplies R^{-1}. Summing the remaining weights over the boxes gives

$$\frac{3^{(4n+5)d}\, 2^{2nd}}{R} \sum_{\mathrm{j} \in \mathbb{Z}^d_{\geq 0}} 2^{-4\sum_v j_v} = \frac{(3^{4n+5} 2^{2n}\, 16/15)^d}{R} \leq \frac{e^{(6n+9)d}}{R}.$$

This bound also justifies absolute unfolding. The last two inequalities of (34) make the tail at most $\|f\|_\infty |c_\infty| \,\mathrm{vol}(K')/2$, strictly smaller than the absolute value of the stabilizer term. The resulting period is nonzero.

Let v_* be the tensor of the translated lowest-type vectors in the diagonal seed. Use compact approximate identities of finite Peter–Weyl support for the product of the adapted maximal compact groups.

Their averages of v_* converge in the smooth topology. Averaging the corresponding right translates of P_x commutes with the locally normal sum and gives the series of these compact-finite second vectors. By continuity of evaluation and bounded continuity of P_x, the resulting pointwise identity is $P_x = y(v_*)$.

For an arbitrary nonzero w in the smooth row tensor, the period against P_x defines a continuous functional on the row Hilbert tensor. Indeed t is bounded into L^2, and $P_x \circ \iota$ is bounded on the finite-volume row quotient. The tensor of the local irreducible unitary representations is irreducible for the product real group, so the span of the real translates of w is Hilbert-dense. The functional is nonzero on w_*, hence nonzero on some real translate of w. Right invariance of the quotient measure moves the inverse translate to P_x. The translated ambient form still belongs to y at level K_C and is bounded and continuous. Average it on the right by finite Peter–Weyl approximate identities in the maximal compact subgroup of the prescribed ambient model. One may take products of nonnegative compact kernels with integral one. The resulting ambient vectors are compact-finite tensors, and the pairings converge by bounded convergence against the integrable function $t(w)$. Some compact-finite vector therefore gives (22). Passage to a prescribed initial real vector.

2.4.5 Reduction at a point with a short lattice

Write $d = [F : \mathbb{Q}]$, $g = g_0 g_1$ and $\mathrm{rd}(E) = |\Delta_E|^{1/(2d)}$. At an adelic point x, put $L = Y_0(E) \cap x_f \widehat{\Lambda_Y}$ in the finite adeles of the row space. At the real places give L the Euclidean norms from x_v^{-1} and the adapted frames. In these norms the original Hermitian form is represented by a self-adjoint involution J_v. The finite lattice L is self-dual for the original E/F-Hermitian form.

For an $\mathcal{O}_E$-lattice D of rank $k > 0$ with induced Hermitian norms, use the size

$$\mathrm{sz}(D) = \left(2^{dk} |\Delta_E|^{-k/2} \operatorname{covol}(D)\right)^{1/(2dk)}.$$

The covolume uses one complex embedding over each real place of F, and quotients use the orthogonal quotient norms. Write $\ell(D)$ for the minimum size of a nonzero sublattice and $h(D)$ for the maximum size of a nonzero torsion-free quotient. Put $\ell = \ell(L)$. For any nonzero integral ideal $\mathfrak{b}$ and $0 \neq y \in \mathfrak{b}L$, write $L \cap Ey = \mathfrak{a}y$. Saturation and projectivity over $\mathcal{O}_E$ give $\mathfrak{b}L \cap Ey = \mathfrak{b}\mathfrak{a}y$. Thus $1 \in \mathfrak{b}\mathfrak{a}$ and $\mathrm{N}\mathfrak{a} \leq 1/\mathrm{N}\mathfrak{b}$. The size of this line gives

$$\prod_{v|\infty} \|y\|_v = \mathrm{sz}(\mathfrak{a}y)^d (\mathrm{N}\mathfrak{a})^{-1/2} \;\geq\; \ell^d \sqrt{\mathrm{N}\mathfrak{b}}.$$

In particular, failure of the product lower bound with $s^d \sqrt{\mathrm{N}\mathfrak{b}}$ forces $\ell < s$. Size and quotient norms The line calculation.

Covering integral homomorphisms. Let D_1, D_2 be Hermitian $\mathcal{O}_E$-lattices of ranks $a, b > 0$, with $\ell(D_1) \geq l > 0$ and $h(D_2) \leq u$. Give the real space of complex-linear placewise homomorphisms from D_1 to D_2 its sum of Hilbert–Schmidt squares. The lattice $\mathcal{L} = \mathrm{Hom}_{\mathcal{O}_E}(D_1, D_2)$ has covering radius at most

$$\frac{ab\sqrt{d}}{2}\,\mathrm{rd}(E)\frac{u}{l}.$$

Indeed, the trace pairing of $\mathcal{L}$ with the reverse homomorphisms is $\mathrm{Tr}_{E/\mathbb{Q}} \operatorname{tr}(BT)$. Its trace dual is $\mathfrak{D}_E^{-1} \mathrm{Hom}_{\mathcal{O}_E}(D_2, D_1)$, as one sees in integral bases at the finite places. It follows that the Euclidean dual elements are $(2B_v^*)_v$ for B in this trace dual, because the field trace is twice the sum of the real parts at the chosen complex embeddings.

For such a nonzero B, of rank j, its image in $\mathfrak{D}_E^{-1} D_1$ has size at least $l/\mathrm{rd}(E)$. Indeed the different times the image lies in D_1, and multiplication by the different changes its size by $\mathrm{rd}(E)$. The size of $D_2/\ker B$ is at most u. Taking determinants of $B_v : D_{2,v}/\ker B_v \to \operatorname{im} B_v$ and comparing the two covolumes gives

$$\Big(\prod_v |\det B_v|\Big)^{1/(dj)} \;\geq\; \frac{l}{u\,\mathrm{rd}(E)}.$$

The dj nonzero singular values then give $\left(\sum_v \|B_v\|_{\mathrm{HS}}^2\right)^{1/2} \geq \sqrt{d}\, l/(u\,\mathrm{rd}(E))$ by the arithmetic–geometric mean inequality. Banaszczyk's transference bound $\mu(\mathcal{L})\lambda_1(\mathcal{L}^\vee) \leq \dim_{\mathbb{R}}(\mathcal{L}_{\mathbb{R}})/2 = dab$ proves the stated bound. For either rank zero the homomorphism space has radius zero. See Banaszczyk [53, Theorem (2.2)]. Integral homomorphism lattice · Application to the induced quotient norms.

The short isotropic step. Suppose $\ell < 1/2$. The Arakelov Harder–Narasimhan filtration of L has saturated steps and strictly increasing successive sizes. The normalized degree divided by rank equals $-\log \mathrm{sz}$. At a step R_i the largest quotient size is its last successive size, and the least size of L/R_i is the next successive size. These assertions follow by comparing the degree polygon with the chords ending or starting at the vertex for R_i. They also give that the first successive size is ℓ. See Bost–Chen [52, §§1.2–1.3].

Every step of last size smaller than $1/2$ spans an isotropic space for the original Hermitian form. Indeed this form induces an integral map $R_i \to \overline{R_i}^*$ of operator norm at most 1 at each real place, since J_v is isometric. The integral metric dual has least size $1/h(R_i)$ by degree additivity and duality. If the map were nonzero, its determinant on the quotient by its kernel would give a geometric mean singular value at least $1/h(R_i)^2 > 1$. This contradicts the operator norm bound. Consequently the number p of successive sizes below $1/2$ is at most the signature bound q, and $p \geq 1$. The bottom row has negative index $m/2$. Use $q = m$ as an upper bound for that index.

Write $s_1 = \ell < s_2 < \cdots < s_p < s_{p+1} = 1/2$ for these sizes with the cutoff appended, and put $\rho = (2\ell)^{1/q}$. The product of the p ratios s_i/s_{i+1} is 2ℓ. Some ratio is therefore at most ρ. For this step $R = R_i$ and $W = ER$, use $h = s_i$ and $l = s_{i+1}$. Then

$$h(R) \leq h, \qquad \ell(L/R) \geq l, \qquad h/l \leq \rho, \qquad h^2 \leq \rho/4.$$

The last bound follows from $h \leq \rho/2$ and $0 < \rho < 1$. The saturated gap used in the release.

A fundamental set for the unipotent radical. Put $P = L \cap W^\perp$ and $k = \dim_E W$. The integral form identifies L/P with $\overline{R}^*$. Indeed, self-duality gives all integral covectors on L, and $R \subset L$ is a saturated summand. At a real place the quotient norms agree because $V_v = W_v \oplus V_{0,v} \oplus J_v W_v$ is Euclidean orthogonal, with $V_{0,v}$ stable under J_v. In these coordinates the radical of the parabolic stabilizing W has Lie algebra of matrices of the form

$$\begin{pmatrix} 0 & A & Z \\ 0 & 0 & -J_{0,v}A^* \\ 0 & 0 & 0 \end{pmatrix}, \qquad Z^* = -Z,$$

where J_v identifies $J_v W_v$ with W_v. The first layer A maps $V_{0,v}$ into W_v. Its displayed lift has Hilbert–Schmidt norm $\sqrt{2}\|A\|_{\mathrm{HS}}$. The center consists of the Z blocks. All triple matrix products vanish.

The covering bound on $\mathrm{Hom}_{\mathcal{O}_E}(P/R, R)$ gives an integral approximation to the first layer, since $P/R \subset L/R$ is isometric onto its image and $\ell(P/R) \geq l$. Any such integral first-layer map extends to a rational element preserving L in the Lie algebra. To see this, extend it to an endomorphism D of L, with image in R and killing a complement of P. The adjoint $D^\dagger$ for the original Hermitian form preserves L by self-duality, kills P, and has image in $W^\perp$ because D kills R. Then $D - D^\dagger$ is the required lift. The radius for the first-layer error modulo the center, in the total Hilbert–Schmidt norm, is at most

$$\sqrt{2}\frac{(r-2k)k\sqrt{d}}{2}\,\mathrm{rd}(E)\frac{h}{l} \;\leq\; \frac{r^2\sqrt{d}}{8\sqrt{2}}\,\mathrm{rd}(E)\rho.$$

For the center, use the homomorphisms from L/P to R, for which the source has least size $1/h(R) \geq 1/h$. Approximate half the desired center and take $D - D^\dagger$. The adjoint is isometric in the chosen norms. The resulting radius is at most

$$k^2\sqrt{d}\;\mathrm{rd}(E)h^2 \;\leq\; \frac{r^2\sqrt{d}}{16}\,\mathrm{rd}(E)\rho.$$

First-layer approximation · Center approximation.

Let Y be a real logarithm in the radical. Apply the first radius bound to $Y/(2g)$. Write $Y = 2gY_0 + A + 2gZ_1$ with Y_0 rational and integral on L, Z_1 central and $\|A\|$ at most $2g$ times the first radius. Apply the second radius bound to $(2gZ_1 + [A, 2gY_0]/2)/g$. This gives a rational center Z_0 integral on L and a real central error Z' with $\|Z'\|$ at most g times the second radius. The two-step Baker–Campbell–Hausdorff identity gives

$$\exp Y = \exp(2gY_0 + gZ_0)\exp A \exp Z', \qquad \|A\| + \|Z'\| \le \frac{r^2\sqrt{d}}{4}\, g\,\mathrm{rd}(E)\rho.$$

The first exponential is congruent to the identity modulo g on L. Indeed its finite sum is $1 + 2gY_0 + gZ_0 + 2g^2Y_0^2$, and the same argument applies to its inverse. This gives the principal level on L even at primes over 2. Integral rational exponentials · Two-step fundamental set.

For $n \in N_W(\mathbb{A}_F)$, first move its finite component into $x_f K_Y((g))x_f^{-1}$ by left multiplication by a rational element of N_W. The polynomial logarithm identifies its finite-adelic points with those of its Lie algebra, where the approximation follows from the density of F in its finite adeles. Use the preceding real factorization on the remaining real component. Its rational factor has finite components in the same conjugate principal subgroup. Expressed in the frames transported by x, the outcome is

$$nx = \gamma x\Big(\big(\exp A_v \exp Z'_v\big)_v, k_f\Big), \qquad k_f \in K_Y((g)),$$

and the sum of the two real norms over the places is at most $r^2 d\, g\,\mathrm{rd}(E)\rho/4$ by Cauchy–Schwarz. For any cuspidal row f invariant under this finite level and with real Lipschitz constant L_f in the stated norms, average $f(x) - f(nx)$ over $N_W(F)\backslash N_W(\mathbb{A}_F)$ with probability Haar. Its second term averages to zero by cuspidality. If $\ell < s \le 1/2$, the variation estimate gives

$$|f(x)| \le L_f r^2 d\, g\,\mathrm{rd}(E)(2s)^{1/q}.$$

Combined with the line calculation, this is the short-lattice implication at the prescribed bottom level used in the Poincaré argument. The adelic point factorization · Integration of the variation.

2.4.6 The span of descending core chains

An admissible descending core chain is a sequence of rapidly decreasing forms a_i on $U(C_i)$ with $a_i = \Theta^{\mathrm{core}}_{i+1\to i}(\Xi_i)a_{i+1}$. Each Ξ_i lies in the lattice-model depth space of its arrow and the last a_m is c_ξ for a permitted column. We show that the K_∞-finite core forms of (22) lie in the algebraic span of the a_0 of these chains.

A common finite datum on each core arrow. Write $G_i = U(C_i)$, $n = r + 1$, and $h = r + 5$, so that $\dim C_i = n + ih$. Let J_i be the given finite core level and let S be the set of split places at its depth. Denote by χ_i the inverse of the grid splitting character with index $m - i$. Both splitting characters on the arrow oriented for the upward pairing are χ_i. On a chosen split half the integral variable for the descending arrow lies in the completed module

$$X_v = \mathrm{Hom}_{\mathcal{O}_v}(\Lambda_{i+1,v}, \Lambda_{i,v}).$$

The other half is its dual for the ψ-pairing. The two lattices are dual because the vertex lattices are self-dual and the trace coefficient at the chosen split place is a unit in these coordinates. Multiplying both lattice frames by their specified polarization ideals does not change the vertex stabilizers. When exchanging the legs of the pair use $x \mapsto x^*$, the actual Hermitian adjoint, and change the sign on the opposite leg. Adjoint reverses the two projections of the split algebra. It interchanges source and target norms, and the two sign changes in the symplectic form cancel. In particular it transports ψ, the self-dual measure, and

evaluation at zero. These observations identify the space of downward tests with the Schwartz space in the upward pairing using these frames.

Put $I_v = \varpi_v^{d_v}\mathcal{O}_v$ at $v \in S$. Choose $x_0 : \Lambda_{i+1} \to \Lambda_i$ surjective modulo $\mathfrak{m}$ on the chosen halves. To do so, choose any projection in bases of the two free completed modules at those places and lift their residues through $\mathrm{Hom}(\Lambda_{i+1}, \Lambda_i)/\mathfrak{m}\,\mathrm{Hom}(\Lambda_{i+1}, \Lambda_i)$. The homomorphism module is projective and tensoring with the finite quotient commutes with Hom. The Chinese remainder theorem thus gives a single integral x_0. In the split Schrödinger coordinate use $f_v = \mathbf{1}_{x_{0,v}+I_vX_v}$. Use the unit lattice vacuum at all other finite places and in the adapted real frames use the Gaussian. The finite split choice is the projection onto the selected coset by the characters of $I_v^{-1}X_v^\vee/X_v^\vee$: with the average normalized by the order of this finite group,

$$\mathbf{1}_{X_v}(x)\,\mathrm{avg}_{\lambda \in I_v^{-1}X_v^\vee/X_v^\vee}\,\psi_v\big(\langle x - x_{0,v}, \lambda\rangle\big) = \mathbf{1}_{x_{0,v}+I_vX_v}(x).$$

Thus it has trivial translation character on I_vX_v and trivial multiplication character on $X_v^\vee$. On the larger dual lattice it has the eigencharacter $\lambda \mapsto \psi_v(\langle x_{0,v}, \lambda\rangle)$. Exchanging the legs gives the downward depth condition of Dep with its selected half. Real polynomial creation operators commute with this finite operation. This constructs the same finite datum for every real polynomial. The completed lattices in the two orientations · Finite Fourier projection and the Fock vacuum.

The normalized split action on these matrices is pullback by the two-sided action $x \mapsto k_i x k_{i+1}^{-1}$. The scalar factors in the Schrödinger representation are a positive square root of the change in volume and, in the downward orientation $(S, T) = (C_{i+1}^-, C_i)$, the character $\chi_T(\det k_{i+1})^{-1}\chi_S(\det k_i)$. The absolute determinant is one on the level. Both determinants are 1 mod I_v in integral frames. The grid splitting characters are trivial on $1 + I_v$. In particular, all depth-fixed vectors are fixed under the two levels, since every (I_vX_v)-periodic function supported on X_v is fixed by these pullbacks. At an unselected split place the same computation uses $I_v = \mathcal{O}_v$ and unramified characters, and the vacuum is the line of $\mathbf{1}_{X_v}$. At a nonsplit place the lattice stabilizer acts in the normalized lattice model by pullback, fixing the basic lattice vacuum. Here the extension is unramified and the residue characteristic is odd, since the host is split at the places above 2. The standard splitting with the unramified parity characters agrees with the lattice splitting on the two self-dual lattice stabilizers. Changing those parity characters by the unramified splitting quotients multiplies the action by their values on unit lifts of the determinants, which are one. Such a lift exists by Hilbert 90, since a lift can be multiplied by an element of $F_v^\times$ to make it a unit in an unramified extension. At a dyadic split place use translations and modulations on the individual halves, as in the displayed split formula. This explains why the full 2Λ translation condition in the adelic lattice model gives the usual vacuum condition off the split halves and does not enlarge the split invariant space. The same argument applies to the first-row levels with χ_S, χ_T the corresponding row splitting characters, whose possible quotient by the unramified η preserves all these trivialities. Actual character and principal action · Application to both families · Dyadic split condition.

Let $b(g)$ be the oscillator coefficient of this tensor against itself, with any one of its real polynomial factors. At a selected place v, allowing arbitrary $g \in G_i(F_v)$, its coefficient vanishes unless the two cosets $x_{0,v} + I_vX_v$ and $g(x_{0,v} + I_vX_v)$ meet. The determinant multiplier never changes this support. Suppose they meet. Then $a = gb_0$ with a, b_0 both in the chosen coset. Choose a right inverse y to $x_{0,v}$ modulo I_v. The two integral matrices ay, b_0y are 1 mod I_v, so they have integral inverses. Hence $g = (ay)(b_0y)^{-1}$ belongs to $J_{i,v}$. Conversely on $J_{i,v}$ the coefficient is $\|f_v\|^2 > 0$. The other split half of g is forced by unitarity and the dual lattice, so this is the exact specified one-sided congruence subgroup, viewed inside the unitary group. Off S the finite datum is spherical. This proves support on $G_i(F_\infty)J_i$ inside the semilocal domain where the components off S lie in the spherical compact. Support read via the dual half eigencharacter · Use on the oriented core arrow.

The real polynomial space and the placed zeta integral. In the determinant-cover coordinates at a real place, subtract the common splitting center $a_i = -c_{m-i}$. In the orientation of the upward adjoint,

write the centered parameter of the smaller representation as $(A, B; C, D)$, with $A, C > 0$ and $B, D < 0$. The partner has centered parameter

$$(A,\ ((h-1)/2,\ (h-3)/2, \dots, -(h-1)/2), D;\ C, B).$$

An overall reversal of positive and negative forms uses the contragredient convention for both members. Every old entry has absolute value at least $(h+3)/2$. In the grid this follows by subtracting the floor of the midpoint of the new block. The distinct centers are spaced by gap, while the old entries lie in the base lists and the other blocks of length h. The formula for the signed partner exchanges the signs of the old entries on one side of this midpoint and inserts the new block in the positive list. To check the choice of list, the opposite insertion would put an entry of the block on both signs when compared with the grid partner. Distinct positive and negative entries of that partner are separated by more than $2\,\mathrm{bigN}^3$. The determinant exponents are integers on the original even-rank groups. After the shift all real parameters for this pair are half integral, as required on the equal-parity determinant covers. The actual centered signed parameters · Splitting centers and parity.

The displayed upward pair is the one in the real theta correspondence [95, Theorem 2.7(b)]. To check first occurrence in its tower use the real first-occurrence criterion [39, Theorem 6.1]. The centered half-integral smaller parameter has no central string, so in that notation $k_\pi = 0$ and $(r_\pi, s_\pi) = (|A|+|D|, |B|+|C|)$. The tested signatures on that tower are $(r_\pi + h + l, s_\pi + l)$. At $l = 0$ the condition on $C_\pi^{\pm}(h/2)$ is satisfied since these sets contain only entries in absolute value less than $h/2$. The condition for $t = h/2 > 0$ excludes $l < 0$. The reverse direction uses the complementary removal criterion [95, Proposition 3.4(ii)]. The inserted string in the larger parameter is maximal, since the neighboring half-integral entries of the string are absent. Deleting it gives the smaller signature and the prescribed parameter at first occurrence. A global smaller member in either direction has a smaller real member at each real place. The local vanishing there followed by the constant-term formula for theta lifts gives cuspidality of the two first-occurrence lifts. This uses the tower property in its constant-term form, including the vanishing of the nonzero rational orbits of smaller dimension by cuspidality of the input.

The local pairing for the actual adapted Fock model has a nonzero real integral [39, Proposition 7.1 and Remark 7.2]

$$\int_{G_i(F_v)} \langle \omega_v(b,1) P_1 f_{\mathrm{vac},v}, P_2 f_{\mathrm{vac},v} \rangle \overline{\langle \pi_{i,v}(b) u_1, u_2 \rangle}\, db$$

on polynomial and compact-finite data with this splitting and this signature. Here the inner products are linear in their first variable and $\pi_{i,v}$ is the member used in the adjoint convention. The oriented model agrees by the adjoint exchange and, if using the opposite sign of the form, the twist of Ichino's Remark 7.3. The polynomial space is unchanged by the unitary change from an adapted orthonormal frame to the oscillator coordinates. Polarizing successively in the oscillator slot and in the $\pi_{i,v}$ slot shows that the value is nonzero for some $P_1 = P_2 = P_v$, $u_1 = u_2 = u_v$. The conclusion uses absolute convergence at the critical point. For the section at s near $h/2$, the real doubling height weight is $H_v(b)^{s-h/2}$, with $0 < H_v \le 1$. The real doubled compact acts on a finite-dimensional span of the polynomial tensor. Iwasawa decomposition applied to its section at $h/2$ therefore bounds its coefficient by $CH_v^{(n+(i+1)h)/2}$. On $\mathrm{Re}\, s \ge 0$ the integrand is bounded by a constant times the absolute discrete-series coefficient. It is integrable by the integrability criterion for discrete-series coefficients [96, Theorem III.1], since the noncompact root gaps exceed $2\,\mathrm{bigN}^3$ and $\frac{1}{4}\sum_\beta |(\alpha, \beta)| = n + ih - 1$ for a root α of type A_{n+ih-1}, with roots of length squared two. This bound also yields a holomorphic integral on $\mathrm{Re}\, s > 0$. The product over real places uses the tensor of the P_v and the u_v, and its Haar normalization differs from their product by a positive scalar. Real product and the holomorphic strip · Height exponent of the actual doubled section.

Semilocal norm on a whole Hecke label space. Let $\mathcal{A}_i$ be the finite-dimensional space of cuspidal occurrences of the specified smooth real type with right J_i-fixed values and trivial finite character. Choose the multiplicity inner product by Schur's lemma from the Petersson inner product on values and the

invariant product on the common real type. Thus for $a \in \mathcal{A}_i$ and $u = \bigotimes_{v|\infty} u_v$ from the real integral, the lower coefficient of $a(u)$ at bk with $b \in G_i(F_\infty)$, $k \in J_i$, is $\|a\|^2 \prod_v \langle \pi_{i,v}(b_v)u_v, u_v\rangle$. The doubled height at k is one, as $\mathrm{diag}(k, 1)$ preserves the doubled lattice. The semilocal integral is consequently

$$\mathrm{vol}(J_{i,S})\Big(\prod_{v\in S} \|f_v\|^2\Big)\|a\|^2\, Z_\infty(h/2).$$

Here the spherical Haar factors off S have volume one, and any conversion from product to chosen Haar measure gives a positive constant. All factors on this line are nonzero for every nonzero a. The same finite datum and real witness can be used throughout the multiplicity space.

For the unramified coefficient use spherical unitary base-change labels over the places w of E outside S. The label B of a bottom constituent is a multiset of size n at each w. The elementary unitary Satake bound places its entries in the annulus $q_w^{-(n-1)/2} \le |\alpha| \le q_w^{(n-1)/2}$. With the standard normalization this is the convex-hull bound by the exponents of the trivial spherical representation. It follows by bounding every double-coset operator on a spherical unit vector by its Haar volume and applying the Satake transform to rays of dominant coweights. Spherical theta transfer adds at arrow j the string $q_w^{(h-1)/2}, \dots, q_w^{-(h-1)/2}$ times the common entry $\chi_j(w)$. The original label is unchanged here because the two characters are the same on each oriented core pair. For split v this is normalized parabolic induction with a character on the excess general linear factor. For inert v the same formula is read on base-change parameters from spherical Hecke matching [46, Proposition A.5]. Every inert place of the grid has odd residue characteristic.

At arrow i put $s_* = h/2$ and omit the places over S from all partial products. Its unramified numerator at s_* is

$$L^S((h+1)/2, B \otimes \chi_i^{-1}) \prod_{0\le j<i} \prod_{\ell=0}^{h-1} L_E^S(h-\ell, \chi_j\chi_i^{-1}).$$

The first Euler product is absolutely convergent and nonzero since $(h+1)/2 - (n-1)/2 > 1$. Every quotient character is unitary and nontrivial, since its real exponents differ. Tate's continuation and the nonvanishing at one for a nontrivial unitary Hecke character give the factors at $h - \ell = 1$. The remaining Hecke products are nonzero by absolute convergence. Removing the finitely many Euler factors multiplies by nonzero numbers, since the characters are unitary. The denominator is $\prod_{j=1}^{n+ih} L_F^S(h+j, \eta_{E/F}^{h+j})$ and is likewise finite and nonzero.

The Witt index of C_{i+1} is at most its real negative index, which is at most $m/2 < h$. The doubling Siegel–Weil integral is in the convergent range since $\dim C_{i+1} - \mathrm{Witt}(C_{i+1}) > \dim C_i$. The unfolding at s_* gives the squared norm of the upward lift as the semilocal integral times this quotient of partial products, times a positive measure factor. One can use the convergent-range formula [39, §8.2], with the semilocal coefficient of the actual occurrence. To see why this works for a sum of occurrences with one label, average the right translations off S between spherical vectors. On the whole label space each averaged operator acts as its scalar spherical eigenvalue. For a finite set of these places the coefficient therefore factors as the product of normalized spherical coefficients times the semilocal coefficient just computed. Unfold first in the half-plane of absolute convergence and then continue the common Euler quotient using the displayed products. The compact support at S and the real weighted bounds give the continuation of the semilocal integral through s_*. This argument treats cross coefficients between occurrences without needing to choose a global simple tensor for $a(u)$. Thus every nonzero a in a label descended from the bottom has nonzero upward lift with the placed tensor. Semilocal reduction in the core measure · The common Euler quotient · Application to the upward adjoint.

Algebraic span of the downward chains. Fix the finite datum on each arrow. The one-dimensional real Howe intertwining line in the two prescribed real types factorizes the downward kernel operator as $D_i \otimes T_i(P)$. Here $D_i : \mathcal{A}_{i+1} \to \mathcal{A}_i$ is one linear map on multiplicities and $T_i(P)$ is contraction of the nonzero real oscillator pairing against P. This factorization can first be taken on Harish–Chandra modules. The

global lift is in the target real isotypic component by the real correspondence and the first-occurrence constant terms. The fixed finite tensor preserves the two levels by the preceding calculation. The real contraction lands in the compact-finite module for compact-finite P and upper vectors by compact equivariance. Its image span as both arguments vary is nonzero. That span is stable under the lower Lie algebra, since differentiating its action differentiates P by a quadratic creation and annihilation operator and leaves a polynomial. It is stable under the lower maximal compact for the same equivariance. Irreducibility of the lower Harish–Chandra module makes this algebraic span the entire lower compact-finite space. Equivalently, apply this argument at each real place and take their algebraic tensor product.

Fubini for the double theta integral against two cusp forms makes the upwards operator $D_i^* \otimes T_i(P)^*$ in the chosen Petersson measures. The absolute convergence for this swap follows by multiplying the polynomial height bound for the theta kernel by the rapidly decreasing cusp forms on both quotients. Positivity of the squared norm just proved implies that $D_i^* a \neq 0$ for each nonzero a of a label transferred from the bottom, and the spherical relations put it in its transferred label space. Distinct labels stay distinct, since one can cancel the added common multiset at each unramified place. Decompose $\mathcal{A}_0$ into its mutually orthogonal simultaneous Hecke eigenspaces. Iteration shows that the adjoint of $D_0 \cdots D_{m-1}$ is injective on their direct sum. Thus $D_0 \cdots D_{m-1}$ is onto $\mathcal{A}_0$.

For $a \in \mathcal{A}_0$ choose one preimage a_m under this composite. Expand a given lower compact-finite vector as a finite sum of products $T_0(P_0) \cdots T_{m-1}(P_{m-1})u_m$, using the real span at each stage in turn. Each term evaluated at a is an actual sequence of descending lifts of $a_m(u_m)$, with the Schwartz test obtained by the adjoint exchange at each arrow. The intermediate terms are cusp forms with the prescribed levels. At the definite core C_m, the real irreducible type is the dual of the fixed finite-dimensional column coefficient, and the prescribed column space is the corresponding full space of equivariant coefficient functions at J_m. For clarity, if the coefficient is E_m and $a_m : E_m^\vee \to \mathcal{A}_{\mathrm{cusp}}(G_m)^{J_m}$, then $c(g) = \sum_\alpha a_m(e_\alpha^\vee)(g)e_\alpha$ is the column with $c_\xi(g) = a_m(\xi)(g)$. The expansion therefore uses only permitted terminal columns and the actual smaller Haar quotients on the way down. This proves the core span for the bottom space tested by the corank-one period. Stability of the polynomial span · Petersson adjoint with the actual kernels · Simultaneous multiplicity preimage and real span · Real dual coefficient at the definite endpoint.

Integrate the resulting finite expansion of the corank-one witness in (22). The core forms are bounded and the lower row form is integrable. At least one descending chain therefore gives a nonzero initial pairing. Fix that chain and its last column. The following finite support and ascent calculations choose the new row tests with these data fixed.

2.4.7 Finite support for the fixed ascent tests

In a transfer write $V = Y_{i+1}$, $W = Y_i$ and Λ_V, Λ_W for their prescribed lattices. The marked tuple x in W has rank j at stage $i = 2j + 1$ and rank $j + 1$ at stage $i = 2j + 2$. The even tuple has one final nonisotropic column. Let R_x be the subgroup carrying x to xu with u upper unitriangular. In rank zero $R_x = U(W)$. The point condition on a local component h is $h^{-1}xu = x + a$, with a in the first block A_1 of the actual mixed model. Write T_a for the corresponding map $V \to W$ obtained by extending a along the input isotropic block. Its columns satisfy

$$a_j = T_a e_j^V,$$

where e_j^V is the paired isotropic frame vector in V. The shift condition is triviality of ψ_E on the sum of the adjacent entries of strictly upper triangular N whenever $h^{-1}xN \in A_1$. Here ψ_E is the trace character used in the standard upper triangular entry character. At odd stages there is also the isotropic-block character condition. The localization of these conditions is by projection to each finite component. For the odd case, right invariance of the test gives the conditions on the full right level coset. The point, shift and isotropic-block conditions · Localization.

Here is the mixed-coordinate derivation. For $F_h = \omega(1, h)\Phi$, a nonzero flattening integral has a nonzero first mixed coefficient at some xu. For $a_2 \in A_2$ the dual-block translation by ha_2 acts there by the phase of

the pairing of xu with ha_2. By Weil covariance and the A_2 eigencondition on Φ, this is also the phase of the pairing of x with a_2. The difference $h^{-1}xu - x$ thus pairs trivially with A_2. Since A_2 has finite support, duality gives that its finite part belongs to A_1.

The flattening is the second mixed coefficient at the fixed y, integrated over u with weight $\overline{\psi_{\mathrm{uni}}(u)}$. Translation in the first isotropic block by xN replaces u by $u + N$. Additive translation of the strictly upper entries preserves the unitriangular Haar measure. Indeed ordinary group translation is triangular in these entries with all diagonal coefficients 1, so product measure is Haar. The change $u \mapsto u + N$ gives the multiplier $\psi_E(\sum_j N_{j,j+1})$. If $h^{-1}xN$ lies in A_1, Weil covariance says this same translation fixes F_h. Nonzero flattening forces the multiplier to be 1.

If $h\,\mathrm{inc}_X(\xi) = \mathrm{inc}_X(\xi')$ for $\xi \in B_1$, the B_1 eigencondition on Φ gives the phase at ξ for translation on F_h by $\mathrm{inc}_X(\xi')$. In the second mixed coordinate this translation is multiplication by the phase of the pairing of y with ξ'. The nonzero flattened value forces the two phases to agree. This gives the third condition, including its trivial form when $y = 0$. Right invariance under the actual finite level gives the same statements for hK. The two evaluations and the weighted Haar integral · Mixed-coordinate phases.

The selected split places. For $v \in S_0$ label the half at $w(v)$ by $+$, the other by $-$, and put $I = (\varpi^{d_v})$ there. The staircase ideal of an auxiliary column with zero-based index j in a block of c columns is

$$\mathfrak{a}^{c-j}(\overline{\mathfrak{a}}^{\,j+1})^{-1}, \qquad \mathfrak{a} = \mathrm{AA}\,\overline{\mathfrak{m}}^{-1}, \quad \overline{\mathfrak{a}} = \overline{\mathrm{AA}}\,\mathfrak{m}^{-1}.$$

The two completions of AA here are unit ideals. Their integrality follows from the excluded primes, and the integral element $|\mathrm{N}\mathfrak{c}_K|$ of AA is a unit at v. The upper block has $c+1$ columns when the lower one has c. At $+$ the j-th raw auxiliary vector of either lattice has coefficient ideal $(\varpi^{d_v(j+1)})$. At $-$ the lower and upper ideals are respectively $(\varpi^{-d_v(c-j)})$ and $(\varpi^{-d_v(c+1-j)})$. These are paired with the opposite columns by the antidiagonal Hermitian matrix of the auxiliary blocks. In the even lower block its middle diagonal entry has sign -1. Staircase ideals · The two local depth ideals · Unit coefficients at the selected places.

In both marked frames the scalar on the j-th raw auxiliary column is $\eta_j = (-\xi_0)^{j+1}$, including the last even column. Thus the marked columns in the $+$ lattice basis are normalized by $\varpi^{d_v(j+1)}/\eta_j^+$. By the depth condition $T_a^+\Lambda_V^+ \subset I\Lambda_W^+$, the same normalization sends a_j^+ into $I\Lambda_W^+$. At the other half T_a^- maps Λ_V^- into Λ_W^- by the underlying homomorphism lattice and the unit coefficient ideal. Consequently

$$\frac{\varpi^{-d_v(c-j)}}{\eta_j^-}\,a_j^- = \varpi^{d_v}\,T_a^-\left(\varpi^{-d_v(c+1-j)}a_j^{V,-}\right) \in I\Lambda_W^-,$$

where $a_j^{V,-}$ is the raw upper auxiliary vector. Pairing Λ_W^- with Λ_W^+ and absorbing the antidiagonal sign gives the normalized covector congruences. The normalized base columns and covectors are primitive coordinate frames in the two halves. The Gram matrix of the transformed tuple is unchanged, since h is unitary and u is upper unitriangular. In the odd stage the Gram matrix is zero. In the even stage the only possible nonzero entry is the last diagonal entry and the last row of u is fixed. The depth homomorphism lattice · The odd marked columns · The even marked columns · Application to the actual row lattice.

Write A, B for the resulting base inclusion and surjection and A', B' for the transformed pair. They satisfy $A' \equiv A$, $B' \equiv B \pmod I$ and $B'A' = BA$. With $L_A A = 1$ and $B'J' = 1$ over the local integer ring, set

$$T_0 = 1 + (A' - A)L_A, \qquad T = T_0 + J'(B - B'T_0).$$

Then $TA = A'$, $B'T = B$ and $T \equiv 1 \pmod I$. The inverse transpose on the other half is its unitary extension and preserves Λ_W^-. Thus one principal-level k satisfies $kx = h^{-1}xu$. It follows that $h = (hk)k^{-1}$ is in $R_x K_v$. The two-sided tuple matching · The one-sided modulus in the intrinsic level.

The places away from the depth. Put $\mathcal{O} = \mathcal{O}_{E,v}$, interpreted as the product of the two local rings at a split place. The depth modulus is a unit here. The completed A_1 condition means that T_a takes Λ_V into $\mathrm{AA}\Lambda_W$. For its j-th input line put $J_j = \{c : ce_j^V \in \Lambda_V\}$. The same coefficient condition on x_j corresponds to $\mathrm{AA}\Lambda_W$, as the ratio of the upper and lower staircase ideals in the same column is $\mathfrak{a} = \mathrm{AA}$ at v. This is the primitive-column property, with the scaled target lattice. For adjacent columns,

$$J_j J_{j+1}^{-1} = (\xi_0)\,\mathrm{AA}\,\overline{\mathrm{AA}} = \mathrm{Ann}_{\psi_E}(\mathcal{O}).$$

The first equality follows from the staircase formula and $\eta_{j+1}/\eta_j = -\xi_0$. The second is the trace-character normalization at v. These equalities are of completed ideals. Depth condition off the selected places · The scaled target of the homomorphism lattice · The adjacent ideal ratio and trace annihilator · The same ratio for the other parity.

Use the local Witt Iwasawa decomposition $h = n\,d(t)\gamma\kappa$ with t diagonal, γ in the unitary group of the Witt core, and κ in the full lattice stabilizer K_v. The unipotent n acts on the marked tuple by an upper unitriangular matrix. In the even case this includes the last column since n fixes it modulo the isotropic space. For each hyperbolic column j, multiply the point condition on the left by κ and take the projection to the j-th isotropic line in W. The projection preserves the lattice and annihilates the earlier marked columns. Upper unitriangularity therefore gives $t_j^{-1}x_jJ_j \subset \mathrm{AA}\Lambda_W$. The primitive-column property gives $t_j^{-1}J_j \subset J_j$, hence $t_j^{-1} \in \mathcal{O}$.

For $s \in J_jJ_{j+1}^{-1}$ use the adjacent matrix $t_j sE_{j,j+1}$, multiplied on the left by the upper unitriangular matrix of n. Its image under $h^{-1}x$ is $\kappa^{-1}x\,sE_{j,j+1}$. This extended map preserves the scaled target lattice on the source by the definition of $J_jJ_{j+1}^{-1}$. The sum of adjacent entries of the test matrix is $t_j s$. The shift condition and the displayed annihilator identity therefore give $t_j \in \mathcal{O}$. This treats every t_j in the even case and all but the last in the odd case. Odd Witt Iwasawa reduction · Projections for the even point and shift conditions.

In the odd case with $q > 0$ isotropic columns, let $\mathfrak{b}$ be the last isotropic coefficient ideal in W and $\mathfrak{c}$ the coefficient ideal of the distinguished line in the core of V. The staircase formula gives

$$\mathfrak{b} = (-\xi_0)^{-q}\mathrm{AA}^{q+1}\overline{\mathrm{AA}}^{-q}, \qquad \mathfrak{c} = (-\xi_0)^{-(q+1)}\mathrm{AA}^{q+1}\overline{\mathrm{AA}}^{-(q+1)}.$$

Thus $\overline{\mathrm{AA}}\mathfrak{b}^{-1} = (\xi_0^{-1})\mathfrak{c}^{-1}$. The terminal covector of the isotropic-block condition, multiplied before taking its adjoint by s, has image coefficient ideal $(\xi_0^{-1})\bar{s}\,\bar{\mathfrak{c}}^{-1}$ on the last dual isotropic line. That line has lattice ideal $\bar{\mathfrak{b}}^{-1}$ by self-duality. This adjoint belongs to the completed homomorphism lattice exactly when $s \in \mathcal{O}$. By right level saturation remove κ in the isotropic-block condition. The last coordinate of the unitriangular action of n is fixed and γ acts trivially on the isotropic block. The condition therefore says that the terminal covector multiplied before adjunction by $t_{\mathrm{last}} - 1$ annihilates the integral complementary isotropic block. It annihilates the entire homomorphism lattice since it only pairs with the projection to that block. Self-annihilation and the terminal covector calculation give $t_{\mathrm{last}} - 1 \in \mathcal{O}$. All t_j are now units, so $d(t) \in K_v$. The factor $n\gamma$ belongs to R_x in the odd case. The two terminal ideals · The terminal adjoint map · Using the terminal character and the self-annihilator.

In the even case apply the same point projection to the last column, projecting now to the Witt core W_0. If L_z is the completed lattice on its distinguished line, this gives $\gamma^{-1}(\mathrm{AA}L_z) \subset \mathrm{AA}\Lambda_{W_0}$, and cancellation gives $L_z \subset \gamma\Lambda_{W_0}$. The line lattice is self-dual. Orthogonal projection to it preserves $\gamma\Lambda_{W_0}$, since pairing with L_z is integral. Its orthogonal complement in $\gamma\Lambda_{W_0}$ is self-dual. At these places the étale extension is split or unramified. In the split case a self-dual lattice is determined by either half and its dual. In the unramified case an orthogonal unit basis is obtained by splitting a vector with unit norm and using surjectivity of the unit norm map. Consequently any two such complement lattices are related by an isometry. Extending this isometry by the identity on the line gives $\gamma = \gamma_0\kappa_0$ with $\gamma_0 \in R_x$ and κ_0 preserving the core lattice, hence also the full lattice under the Witt embedding. This gives $h \in R_xK_v$ in the even case. Splitting off the fixed line · The self-dual complement orbit.

At all but finitely many v, the given h_v already belongs to K_v, so choose $r_v = 1$ there. The local factorizations thus assemble in the restricted product to $h = rh_\infty k$ with $r \in R_x(\mathbb{A}_F)$ and k in the prescribed adelic level. This proves the finite orbit-support condition used in the unfolding. Assembly at the selected places · Putting the support into the odd and even transfers.

2.4.8 The depth blocks in the two mixed coordinates

The ordinary homomorphism lattice, before applying the depth condition, is the lattice of maps $\Lambda_V \to \mathrm{AA}\Lambda_W$. It is self-dual for the symplectic character. In integral dual bases the trace character pairs the matrix entries, and the product $(\xi_0)\mathrm{AA}\overline{\mathrm{AA}}$ is its coefficient annihilator. The same conclusion holds for the larger lattice used to define B_1, B_2. Its target lattice is shifted by $\mathfrak{m}\overline{\mathfrak{m}}^{-1}$, whose product with its conjugate is the unit ideal. This lattice contains the depth lattice Γ, since it has the same depth on the selected half and has the larger inverse-depth lattice on the conjugate half.

The upper and lower Witt projections preserve the completed integral lattices. In the auxiliary blocks these are the projections to the indicated coordinate summands with their staircase coefficients. Consequently they also preserve the ordinary, shifted and depth homomorphism lattices. For a self-dual homomorphism lattice, the two projected pieces of each hyperbolic pair are dual, since the pair forms a symplectic orthogonal summand with complement given by the other mixed blocks. The shifted lattice makes B_1, B_2 mutually dual in the second block. In the first block A_2 is defined as the annihilator of A_1, and finite local biduality gives the converse. These are compact open lattices in their finite coordinate spaces. Applying the projections decomposes Γ into components in the five prescribed blocks. The first dual component belongs to A_2 because the ordinary lattice is self-dual. The middle components belong to the shifted lattice as above, and the last is in the core depth lattice by its definition as the pullback along the core embedding.

For the eigencharacters on components coming from Γ, extension of the marked x as a first isotropic block map takes the integral source into the target scaled by $\mathfrak{a}$, the ratio of the two same-column staircase ideals. This is the ordinary homomorphism lattice off the selected places. At the selected places the map has the ordinary lattice on the $+$ half and the inverse-depth lattice on the $-$ half. These halves constitute the annihilator of Γ, since its $+$ half has depth and its $-$ half has the ordinary lattice. In the odd case the terminal ideal identity before evaluating away from the depth is

$$\overline{\mathrm{AA}}\mathfrak{b}^{-1}\mathfrak{m}^{-1} = (\xi_0^{-1})\mathfrak{c}^{-1}.$$

Thus the adjoint of the marked y has the same ordinary and inverse-depth pattern by the terminal covector calculation. It too annihilates the relevant projection of Γ. This proves that the two prescribed characters are trivial on the appropriate components of the depth lattice. The even stage has $y = 0$. The shifted target and exact mixed pieces · Odd projections and phase conditions · Even projections and phase conditions · The first marked character · The second marked character.

The core block has a nonzero vacuum because its depth lattice is contained in the ordinary self-dual homomorphism lattice in the residual symplectic space. In a symplectic integral polarization of that lattice, the characteristic function of one half in the Schrödinger representation is fixed by translations and by the paired character operators. At split places one may take the two split halves as the polarization. This also gives the required half-lattice invariance at the dyadic split places, and the translations by twice the lattice satisfy the same invariance. Taking the restricted tensor product of these local vectors and any nonzero real Schwartz vector gives an adelic vacuum. The odd residual vacuum · The even residual vacuum.

A nonzero flattened vector. The character in the flattening is trivial on the finite shifts satisfying $x(u-1) \in A_1$. In fact write J_j^W, J_j^V for the completed normalized coefficient ideals of the lower and upper frames. Before specializing to the places off the depth, the ratio is

$$J_j^W (J_{j+1}^V)^{-1} = (\xi_0)\bar{\mathfrak{a}}.$$

The $j, j+1$-entry of $u-1$ belongs to this ideal multiplied by the coefficient ideal of the scaled target. On the selected half this multiplier is the depth ideal, and the product is $(\xi_0)\overline{\mathrm{AA}}$. Here AA is a unit. On the other finite halves the multiplier is AA and $\bar{\mathfrak{a}} = \overline{\mathrm{AA}}$. In both cases the product is the trace annihilator $(\xi_0)\mathrm{AA}\overline{\mathrm{AA}}$. The adjacent entries therefore all have trivial local characters. This proves the stated character property. The odd finite shift character · The even finite shift character.

Given the mutually dual compact-open blocks A_1, A_2 and B_2, B_1, and a nonzero core vacuum ϕ_{core} in the fixed Γ_0-space, take in the two mixed coordinates a product $b_1(z)b_2(z')\phi_{\mathrm{core}}(w)$. The finite factors of the bumps are the indicators of $x_f + A_1$ and $y_f + B_2$, and their real-place factors are smooth nonnegative bumps positive at x_∞, y_∞. Fourier inversion in the two hyperbolic summands gives its theta vector, with the positive Fourier normalization for the prescribed measures. Finite translations by A_1, B_2 fix the bumps, and multiplication by their dual characters takes its constant phase at x, y on their supports. Thus the vector lies in the required mixed class.

The support of $b_1(xu)$ is compact in the unitriangular group since x has a rational left inverse. On its finite part the weight $\psi_{\mathrm{uni}}(u_f)$ is trivial by the preceding ideal calculation. Choose the real support sufficiently small so that the real part of $\overline{\psi_{\mathrm{uni}}(u_\infty)}$ is positive wherever $b_1(xu) \neq 0$. Such a support exists by the left inverse and continuity at $u_\infty = 1$. Then the flattening is a nonzero scalar times $b_2(y)\phi_{\mathrm{core}}$, with nonzero scalar by positivity of the real part of the integral on a neighborhood of the identity. The mixed class is stable under real Heisenberg translations because the prescribed phases are all finite. It is also stable under translations by the embedded core annihilator, which commutes with the outer hyperbolic blocks and with the translations that define the fixed core vacuum. This supplies the nonzero flattened test and the stability used by the Fourier and Moyal arguments, given the compact-open block duality and the fixed core vacuum. Positive real bump and the two Fourier inversions · Stability and induced core vacuum.

The core and line vacuum at each transfer. Write $\mathfrak{j} = \mathfrak{m}\overline{\mathfrak{m}}^{-1}$ and $\mathfrak{l}_\rho = \mathfrak{a}^\rho\overline{\mathfrak{a}}^{-\rho}$. In the raw line frame ζ_0 the central auxiliary line at stage 2ρ, for $\rho > 0$, has completed lattice $\mathfrak{l}_\rho\mathcal{O}_E\zeta_0$. The line vacuum space uses the target $\mathfrak{j}\mathfrak{l}_\rho\mathcal{O}_E\zeta_0$. Both ideals are inverse to their conjugates, so these are self-dual line lattices. At rank zero use the projection of the bottom core depth lattice to the unit line. It is contained in the ordinary line homomorphism lattice. This gives a nonzero line vacuum in every case by the same integral polarization as for the residual core.

Projection of the residual row depth map to the two core lattices preserves the ordinary map condition and the depth condition. This follows because their embeddings are the original orthogonal summands in the row lattices, with their prescribed lattices. In an even step the remaining target projection is onto the line of index $\rho + 1$. The containment in the shifted homomorphism lattice gives its target $\mathfrak{j}\mathfrak{l}_{\rho+1}\mathcal{O}_E\zeta_0$ with the ordinary coefficient factor AA. This is exactly the line vacuum condition in the even transfer. The line depth lattice · The even line target · The even core and line factors.

In an odd step take the Hermitian adjoint of the line component. The link of lines uses $\zeta_{\rho+1} = (-\xi_0)\zeta_\rho$, and its inverse on the line map is the Hermitian adjoint $\xi_0^{-1}\mathrm{starAdj}$ in the normalization of the Weil action. On adjunction the ordinary coefficient AA changes to $\overline{\mathrm{AA}}$, and

$$\overline{\mathrm{AA}}\,\mathfrak{l}_{\rho+1} = \mathrm{AA}\,\mathfrak{j}\mathfrak{l}_\rho.$$

This is the target required for the line vector at $\rho > 0$. For $\rho = 0$, using this equality with the ordinary original target gives the required depth on the selected half. Applying adjunction with the simultaneously allowed shifted original target $\mathfrak{j}\Lambda_{C_1}$ cancels $\mathfrak{j}$ in the displayed equality and gives the ordinary line map condition. The two conditions together are the rank-zero line depth lattice.

The link reverses the symplectic form. Complex conjugation of the line vector therefore transports its Heisenberg invariance. It also exchanges the two split halves in adjunction, so the one-sided translation condition transfers along the link. For the tensor of the core and line vectors, the Weil translation operators factor under the orthogonal decomposition. The two projection conditions just proved make this tensor

a vacuum in the flattened space. In the bottom rank-zero seesaw, the depth lattice on the target $Y_0 \perp L_0$ is the direct sum of the two projected depth lattices. The entrywise Heisenberg annihilator conditions thus restrict to each summand, and the tensor is a vacuum for the bottom core depth lattice. The adjoint line ideal · Adjoint normalization · The rank-zero depth restriction · Heisenberg invariance under the linked adjoint · The bottom split depth lattice.

2.4.9 Ascent with a prescribed descending chain

Write $[H] = H(F)\backslash H(\mathbb{A}_F)$. Work in the conjugated lattice realization of the row. Complex conjugation and reversing the sign of the first Hermitian form each reverse the symplectic phase. Their combination retains the prescribed standard additive character ψ, and $\psi_E = \psi \circ \mathrm{Tr}_{E/F}$. Write C_i for the fixed rapidly decreasing core forms, obtained by the actual descending lifts $C_i = \Theta_{D_i}(C_{i+1})$ with the prescribed core tests D_i. Write $G_0 = u_w$, and let G_{i+1} be a lift of G_i with one of the permitted tests of the i-th row arrow. The desired implication is the following. If the rank-zero pairing of G_0 and C_0 is nonzero, then the tests on the row arrows can be chosen so that the final Bessel period with C_m is nonzero. This assertion holds for each compatible grid, each prescribed descending sequence and every fixed w giving the initial pairing. In addition to the core forms, only the initial row and its partial lifts are required to be rapidly decreasing. From the first lift onward the row is fixed by its prescribed finite group. Indeed the resulting theta integral is fixed in its target variable by any element that fixes its Weil test on that side.

Put $\alpha_i = (\chi_{m-i}^{\mathrm{grid}})^{-1}$, and let η be the odd line splitting character. The ordered pair on the row arrow, with the larger Hermitian space first in the oriented symplectic realization $\mathrm{Hom}(V, W)$, is $(\alpha_i\eta, \alpha_i)$ for odd i, and $(\alpha_i, \alpha_i\eta)$ for even i. On the core arrow it is (α_i, α_i). These are the inverse characters of the un-conjugated data because complex conjugation changes every unitary scalar in the Weil action to its inverse. Orthogonal direct sum multiplies the character attached to the summands and leaves the opposite character unchanged. Thus on splitting off the one-dimensional summand its character is η, with the common opposite character on the same oriented space. For the other orientation its realization is transported by the prescribed Hermitian adjoint and complex conjugation. In particular, the additive characters on the central and side coordinates in the following computation come from the same adapted action. The conjugated Weil datum and the actual integrals.

Here is a way to check the simple action laws used below. With B the oriented symplectic form, use

$$T_w F(v) = \psi(B(w, v)/2)F(v + w), \qquad T_w T_t = \psi(-B(w, t)/2)T_{w+t}.$$

At a rational translate the lattice relation gives $\psi(B(v, w)/2)$ with the opposite sign, so their product has full symplectic phase after moving the translate past a Weyl operator. The adjoint identification of the linked line has $B'(Lx, Ly) = -B(x, y)$. Consequently evaluation of a mixed coefficient translated by that line equals the conjugate line coefficient, with its full central phase conjugated. At split finite places the same normalization, in matrix coordinates, is $\chi_W(\det k_V)^{-1}\chi_V(\det k_W)$ times pullback of the Schwartz function. The two determinant factors are on the two sides of the same action. At the other finite places the vacuum-normalized splitting gives the identity on the self-dual vacuum. The finite-coset proof gives the stronger depth, support and annihilator properties for these particular adapted lattices.

The two unfoldings with their test functions. In either step let V be the upper row space and W the lower. Write $V = X \oplus V_0 \oplus Y$ with X, Y dual isotropic, and $\ell = \dim_E X$. In the flag unipotent write $n = n_C(c)n_B(b)m(u)$, with $c^* = -c$, $b : V_0 \to X$, and $u \in U_\ell$ upper unitriangular. In the sign convention for B, the translate of the first mixed row by $n_B(b)$ has components $(x, xb, -xbb^*/2)$. Denote the first mixed transformation of a lattice vector by $F^X(x; t)$. Precisely, it is the evaluation at zero of the product of the Weyl operators for $j_Y(A)$, $j_X(x)$ and the middle vector $j_0(t)$, integrated over the rational quotient in A. This definition fixes the half phases. If the lower flag is nonempty, its analogous transformation is denoted $(F^X)^Y(y; t)$; its dual argument is inserted using the adjoint for B. Give all additive quotients

covolume one. Give $U_\ell(\mathbb{A}_F)$ the Haar measure for which $[U_\ell]$ has volume one. These choices realize the unipotent quotient Haar of the periods. With any prescribed positive quotient Haar proportional to them, both sides acquire the corresponding positive scale.

Fourier expansion of the evaluation at zero along the compact dual block gives the sum of $F^X(x;0)$ over rational frames x. On $n_C(c)$ its character is

$$\psi\big(B(j_X(x), j_Y(xc))/2\big).$$

The coefficient selected by the upper Bessel weight therefore has Gram matrix zero. For the upper Fourier–Jacobi weight the product of the conjugate line function and the conjugate adjacent-root character has central weight equal to the conjugate of this phase at ze_ℓ^t, where z is the lower distinguished line vector. In that case the selected Gram is $\mathrm{diag}(0,\dots,0,\langle z,z\rangle)$. The trace pairing on the space of skew-Hermitian matrices is nondegenerate, so integrating over $[c]$ imposes precisely the indicated Gram condition.

Only independent frames contribute. To see this, take the first dependent column j and a relation $xd = 0$ with $d_j = 1$ and $d_k = 0$ for $k > j$. For $j < \ell$, the stabilizer contains $1 + dse_{j+1}^t$, with $s \in \mathbb{A}_E$. Its adjacent-root weight is $\psi_E(s)$, while it fixes x. It also fixes the Jacobi line row ze_ℓ^t, and in the Bessel case it leaves the last-row side weight unchanged. With last Gram entry nonzero necessarily $d_\ell = 0$. In the other case with $j = \ell$, the side translation $b = sdz^*/\langle z,z\rangle$, with z now in the upper Bessel middle space, satisfies $xb = 0$ and has side weight $\psi_E(s)$. In either case the orbit integral contains the integral of a nontrivial character on $E\backslash\mathbb{A}_E$, hence is zero. For an independent frame the specified Gram gives a single rational orbit, by Witt's extension theorem applied to the isometry of the two ordered subspaces. Choose x_* to be the reference lower isotropic frame in the first step, and that frame followed by z in the second. Write

$$R = \{g \in U(W) \,:\, gx_* = x_* u_g,\ u_g \in U_\ell\}.$$

Its reductive part is the full lower middle group for the isotropic frame; with the last line appended it is the stabilizer of the distinguished line vector in that middle group. Its unipotent part is the lower flag unipotent. These descriptions follow by acting on the flag columns and on the orthogonal quotient of the isotropic span. In the appended-line case, the last adjacent entry of u_g is the lower side coordinate.

Let a be the upper core variable and embed it as $\iota(a)$ in $U(V)$. Let $\phi = C_{i+1}$, let $G = G_i$ with the conjugations of the upper period included in its weight, and set $F_{a,g} = \omega(\iota(a), g)\Phi$. Unfolding the independent rational orbit and changing the rational translates of u gives the upper period as

$$\begin{aligned} &\int_{R(F)\backslash U(W)(\mathbb{A}_F)} f_\Phi(g)\,dg, \\ f_\Phi(g) = G(g)\int_{[H]}\phi(a)\int_{U_\ell(\mathbb{A}_E)}\int_{[b]} &\vartheta_a(b,u)F^X_{a,g}(x_*u; x_*b)\,db\,du\,da. \end{aligned} \tag{A}$$

Here H is the upper core group and the unipotent group in the u-integral has the stipulated rational covolume. The two weights and the associated contractions are as follows.

1. In the Bessel-to-Jacobi unfolding the weight is $\overline{\psi_E(\sum_j u_{j,j+1})}\,\overline{\psi_E(e_\ell^t bz)}$. Put $y_* = -ze_\ell^t$ in the adjoint block. For $\ell = 0$ both frame blocks and their characters are trivial and these arguments are zero. The equality

$$\overline{\psi_E(e_\ell^t bz)}F^X(x_*u; x_*b) = \big(T_{j_{W,X}(b)}T_{\mathrm{incAdj}(y_*)}F^X(x_*u;\cdot)\big)(0) \tag{B}$$

is the Weyl law followed by the lattice relation at the rational vector $\mathrm{incAdj}(y_*)$. The two half phases add to the side phase. Thus integration of (B) over $[b]$ is the second mixed coefficient at y_*. Define $\mathcal{F}_{BF}\Phi$ by integrating the resulting mixed coefficient over u against the conjugate adjacent character; retain the remaining core variable.

   For n in the lower flag group, substitution of $nx_* = x_*u_n$ in u changes the adjacent factor by $\overline{\psi_E(\sum_j (u_n)_{j,j+1})}$. The second adjoint argument changes from $v = \mathrm{incAdj}(y_*)$ to $n^{-1}v$. Its projection to the remaining line block is the line Weyl translate. Moving the remaining rational block

through the first evaluation gives the scalar $\psi(B(v, n^{-1}v)/2)$. These are respectively the flag factor, the line side action and its central scalar in the lower Jacobi datum. For h in the lower middle group and a in the upper core, the remaining action is $\omega_0(a, h)$. This follows by commuting the two embedded middle actions with the mixed transformations. Equivalently, the new line action is the Weil action on the lower Jacobi line brought across the Hermitian adjoint; conjugation on the line and the sign in y_* are essential in this identification. For a tested tensor $D_i \otimes L$ on the remaining core and line, the integral over $[R]$ in (A) is consequently the lower Fourier–Jacobi period with core $\Theta_{D_i}(\phi) = C_i$ and the transported line test L. The product formula follows by inserting the tensor and integrating $\phi(a)\omega_0(a, h)D_i(0)$ in a.

2. In the Jacobi-to-Bessel unfolding let $Y_a = \omega_{\mathrm{line}}(a, 1)\Phi_0$ be the fixed upper line test after the upper core action. The weight is
$$\vartheta_a(b, u) = \overline{\psi_E(\sum_j u_{j,j+1})}\,\overline{Y_a(ze_\ell^t b)}.$$
For the first $\ell - 1$ rows of b the integral is the second mixed coefficient at zero. The last row maps isomorphically over E to the line block. Thus the remaining integral is
$$\int_{[\mathscr{W}_{\mathrm{line}}]} \big(\omega_0(a, (h, 1))\Xi\big)(0, t)\overline{Y_a(t)}\,dt, \qquad \Xi = \mathcal{F}_{FB}\Phi, \tag{C}$$
where $\mathcal{F}_{FB}$ is the weighted double mixed coefficient at $(x_*, 0)$, and h fixes the appended line. The rational line isomorphism carries the mass-one measure on the first compact quotient to the mass-one measure on the second. Substituting u_g as above produces the lower Bessel character: its last adjacent coefficient is the side term $e^t b z$, and the preceding adjacent coefficients are those of the lower flag. The second evaluation at zero has no remaining central phase. This verifies the lower mean as the Bessel integral with the handed form $h \mapsto \int_{[H]} \phi(a)\,(\mathrm{C})\,da$. For $\Xi = D_i \otimes \Phi_0$, factor the Weil action on this tensor. The line inner product in (C) is then $\|\Phi_0\|^2$ by unitarity, and the handed form is $\|\Phi_0\|^2 C_i(h)$.

First mean and normalized character comparison · Second mean with the actual line test.

Absolute convergence and the topology of the tests. The lattice realization is obtained from the adelic Schwartz realization by the Weil–Brezin transform: in a polarization it is the sum of $\phi(x + \lambda)$, $\lambda \in X(F)$, against the ordinary character of λ in the dual coordinate, with the half phase of the Weyl law. Conversely Fourier coefficients in that compact dual coordinate recover ϕ. On fixed finite support and level, differentiation on the real compact quotient gives Schwartz seminorms of the inverse, and summing against rapid decay gives seminorms of the forward transform. In particular the mixed transforms here are restrictions and partial Fourier transforms on adelic Schwartz space. For fixed finite datum the u-integral defining a flattened test is continuous in these seminorms. The map $u \mapsto x_* u$ is a proper affine linear embedding on upper-entry coordinates with a left inverse from the full rank of x_*; its real integral is bounded by integrating a Schwartz majorant in those coordinates. The finite part ranges over a compact set for a fixed finite test.

For absolute convergence of (A), first sum the absolute values of the mixed evaluations on the original compact flag quotients before unfolding. A Schwartz seminorm bounds the real sum over rational x by summation of $(1 + \|x\|)^{-B}$ on a lattice for B larger than its real dimension; a second power bounds the lattice-point count after increasing denominators. In detail, choose a product finite symplectic lattice L_0 small enough that its Weyl operators fix the original Schwartz vector. After acting by g the vector is fixed under gL_0. Locally $\varpi_v^{d_v} L_{0,v} \subset g_v L_{0,v}$ with $\prod_v q_v^{d_v} \ll H(g)^A$, by the entries of the inverse action on this lattice. The dual split coordinates of this smaller lattice force support in a denominator lattice with index polynomial in that product, and its translation coordinates give constancy on cells of inverse

polynomial volume. The finite supremum of the transformed vector is therefore bounded by a power of $H(g)$ times its fixed L^2 norm. At infinity use a Cartan decomposition in the real symplectic group. The two compact factors act by a bounded family on each Schwartz seminorm, with finitely many input seminorms. The diagonal part acts by a determinant factor and linear coordinate changes, which bound its Schwartz seminorms by powers of matrix height. These arguments give a polynomial bound in the upper core and lower quotient heights, uniform for a or g translated in a fixed compact set. For the Jacobi case the absolute value of the upper line function is periodic in the line translation and obeys the same bound uniformly on its compact quotient. The estimates descend to quotient heights: rational changes of the middle bases permute the frame sum and preserve the mass-one measures on the additive quotients, so one may take reduced representatives for a and g.

The same estimate applies when restricting g to R and when integrating the means for a compact family of right translates. The heights of $n\imath(b)$ and b are comparable up to powers on these quotients. One can see the inequality in the forward direction by reducing b in its middle group and then choosing n in a compact fundamental set for the rational unipotent. In the reverse direction use the wedge e_X of the rational basis of the isotropic flag. On $e_X \wedge V_0$, the unipotent action is the identity, and the embedded middle acts by $e_X \wedge v \mapsto e_X \wedge bv$. For every nonzero $v \in V_0(F)$ and every rational γ in the ambient group, the vector $(\gamma n\imath(b))(e_X \wedge b^{-1}v) = \gamma(e_X \wedge v)$ is nonzero rational. Its product norm is bounded below by the product formula on any nonzero rational coordinate. Thus the product norm of $b^{-1}v$ is bounded below by a fixed inverse power of $H(\gamma n\imath(b))$. Take the infimum over rational γ. Mahler reduction for the middle lattice bounds its quotient height by a power of the reciprocal of this shortest-vector bound. The same argument with the rational middle vector in the direct summand applies when the distinguished line is fixed. This comparison is invariant under changing the chosen compact set of right translates. Choosing the rapid decay powers of ϕ and G larger than the polynomial powers thus bounds the absolute integrals by an integrable majorant on the core and lower quotients, which have finite volume; the unipotent quotients are compact. In (A) one can equivalently unfold the nonnegative sum first, by Tonelli, obtaining exactly the integral of the absolute independent-orbit integrand. These arguments prove the two unfoldings, their disintegrations over $[R]$, continuity of the $[R]$-mean in the real right translate, and continuity as a linear functional of the remaining flattened Schwartz vector. They also justify passage from finite sums of real Fock tests to Schwartz tests by the convergence of the Hermite-function expansion in the Schwartz topology.

From a global zero to the mean on the finite coset. Fix the preceding row, together with a nonzero upper line test in the Jacobi-to-Bessel step. Assume all the permitted new row tests, with the allowed upper line test in the Jacobi case, give zero upper period. The finite Heisenberg Weyl operators on the depth space permute its actual coset basis, up to unit scalars. Explicitly, in a split polarization write the depth lattice as $IX \oplus X^\vee$. Its annihilator has split coordinates $X \oplus I^{-1}X^\vee$. Translation by $x \in X$ sends $\mathbf{1}_{\kappa+IX}$ to $\mathbf{1}_{\kappa-x+IX}$; the dual coordinate acts on that coset by the character evaluated at κ. The Weyl half phase is a constant. The conjugated model conjugates these scalars. The rational coset labels give all finite residues by the Chinese remainder theorem. Equivalently the allowed lattice vectors are fixed by T_{2u} for every depth-lattice vector and by T_u on both split halves. This is the same condition at dyadic split places, where requiring only the first invariance would lose the coset bound. Off the depth primes the space is the self-dual vacuum line. Thus the coset tests span the full finite tested space. The Hermite expansion extends the assumed vanishing to all real Schwartz tests in that finite space. In particular it holds on the family of tests with the finite support and annihilator properties of the finite-coset proof, a family closed under real Weyl translation.

Modulate Φ by any Weyl operator on the first dual block supported at infinity. The mixed-coordinate Weyl law multiplies the integrand before the a, u, b integrations in (A) by

$$\psi_\infty\big(B(j_X((g^{-1}x_*u)_\infty), j_Y(c))\big). \tag{D}$$

All real frequencies occur by nondegeneracy. The absolutely integrable unfolded integrand thus defines a finite complex measure whose pushforward under $(g, a, u, b) \mapsto (g^{-1}x_* u)_\infty$ has every Fourier coefficient zero. For a compact set Q_∞ in the real lower group, the subset $Q_\infty^{-1} x_* U_\ell(F \otimes \mathbb{R})$ of the real frame space is closed. Indeed, from a converging sequence in it, take a convergent subsequence in Q_∞ and apply a left inverse of x_* to obtain a limit in the closed upper unitriangular group. The inverse image of this set under the displayed pushforward condition is exactly $g_\infty \in R_\infty Q_\infty$, independent of u. Fourier uniqueness gives zero integral over every such strip. This argument uses the joint frame $g^{-1}x_* u$ and the joint side and central weights of the single period.

Use now the finite support result for the flattened evaluation of the actual test: the integrand after the inner integrations in (A) is supported on $R(\mathbb{A}_F)U(W)(F\otimes\mathbb{R})J$ with J the lower right finite group. It is right J-invariant by the Weil action on the test and the preceding row invariance. Disintegrate the unfolded quotient along $R(F)\backslash R(\mathbb{A}_F)$. Its quotient Haar has the product description from the lower flag unipotent and its middle group; the determinant of the latter on the unipotent is one, as it is an F-rational character of the middle unitary group. Over the support there is a single finite orbit, whose positive measure is a constant multiple of $\mathrm{vol}(J)/\mathrm{vol}(J \cap R(\mathbb{A}_{F,f}))$. The vanishing on strips therefore implies that the real quotient integral of the $[R]$-mean on every compact set is zero. The mean is continuous, so its value at the identity is zero. By the computed covariance this is vanishing of the lower functional on the flattened vectors.

Using a prescribed lower test. The chosen compact support bump gives a nonzero flattened depth vector, and every translate by its actual annihilator is the flattening of another permissible Schwartz vector: extend the translate in the core block before the two mixed transforms, which commute with it. To see why these translates test the entire depth space, let f_0 be the nonzero flattened test and use the self-dual identity

$$\|f_0\|^2 f = \int_{\mathscr{W}(\mathbb{A}_F)} \langle f, T_v f_0 \rangle T_v f_0 \, dv. \tag{E}$$

In the Schrödinger realization (E) follows by Fourier inversion in the modulation coordinates followed by integration in the translation coordinates. Both sides have the same normalization of Haar, so there is no extra scalar. At fixed finite test data its coefficient has compact finite support. At infinity its coefficient is Schwartz, by repeated differentiation and integration by parts in the translation and modulation coordinates. Since the seminorms of $T_v f_0$ grow at most polynomially, (E) converges in Schwartz space. By applying the common lattice invariances to both vectors in the coefficient and commuting their Weyl operators, the coefficient vanishes off the annihilator, including the two split-half conditions. Applying the continuous lower mean to (E) now shows that it vanishes on every lower depth vector.

In the first unfolding insert the fixed tensor $D_i \otimes L$ using the given lower line test transported by the linked line map. In the second insert $D_i \otimes \Phi_0$ with the very same upper line vector in its contraction. The split depth block decomposition and the conjugate adjoint matching ensure these are vectors of the flattened depth space. By (B) and (C), respectively, vanishing would give the lower period equal to zero, or $\|\Phi_0\|^2$ times that period equal to zero. The upper line test in the second case can be any nonzero prescribed depth test. This proves the contrapositive at each step. Since finite cosets and real Hermites span densely in the topology used in (A), at least one individual permitted coset and polynomial succeed on the new arrow. Earlier choices and the automorphic input remain fixed.

At rank zero $C_0 = Y_0 \perp L_z$, and $Y_1 = C_1$. The chosen line space of the Jacobi datum is $-L_z$, so a product test in the core descent has row factor D and line factor $\overline{L}$, with L on the Jacobi line space. For this test, use the split Weil action in the two blocks to write the bottom pairing as

$$\begin{aligned}\int_{[U(Y_0)]} \int_{[U(Y_1)]} & C_1(a)\, \theta_D(a,b)\, \overline{\theta_L(a,1)} G_0(b)\, da\, db \\ &= \int_{[U(Y_1)]} C_1(a)\, \overline{\theta_L(a,1)} G_1(a)\, da, \qquad G_1(a) = \int_{[U(Y_0)]} \theta_D(a,b) G_0(b)\, db.\end{aligned}$$

This is the rank-zero Jacobi period, with the prescribed upper core measure and lower row measure. If the bottom core measure differs from the bottom row measure, their constant ratio is positive. The equality is justified by the same Schwartz and rapid decay bounds. Finite sums of these tensor tests are dense in the depth space on the split block, by the finite residue basis and real Hermite expansion. Consequently some permitted first row test and linked line test give a nonzero Fourier–Jacobi period. The alternating transfers then retain each lower test. Since $m = 2\rho+2$ is even, the last period is Bessel. Rank-zero splitting and descent input.

At the last step the auxiliary isotropic frame is $e_j = (-\xi_0)^j e_j^{\mathrm{grid}}$, $1 \le j \le \rho$, and the marked line vector is $(-\xi_0)^{\rho+1} z^{\mathrm{grid}}$. The stabilizer and flag subgroup are the original ones. For the simple-root coordinates on all consecutive columns, including the final line coordinate, this changes the matrix coefficient by $(-\xi_0)^{j+1}/(-\xi_0)^j = -\xi_0$. The grid Bessel character uses $\psi_E(\xi_0 \sum_j s_j)$ in these unscaled coordinates. The auxiliary one is therefore its conjugate. Applying conjugation to the terminal auxiliary integral recovers B^+ in the original model up to the positive Haar ratio of the top definite quotient. The flag quotients on both sides have volume one, because the rational frame changes identify those quotients. This proves that the terminal integral in the ascent detects the prescribed B^+. Endpoint character · Endpoint pairing and Haar.

2.4.10 The real Bessel line in the arithmetic tests

Here is the real identification on the row tower. At a given real place let P_i be the unshifted signed row list, put $s = r + 6$ and $\delta_i = -1/2$ for odd i and 0 for even i. The actual row has the dual of $P_i + \delta_i$. For the arrow of index i, write $c = c_{m-i}$. In the original action the two splitting exponents, in the order larger, smaller, are $c + \delta_i$, $c + \delta_{i+1}$. They are opposite-parity because s is odd. The larger space in this orientation has negated form. After matching the two centers, reversing the signature signs below the center, and adding the centered string

$$((s-1)/2, (s-3)/2, \dots, -(s-1)/2)$$

in the negative list, one obtains precisely the parameter of the larger negated row on the determinant cover, before taking the dual. In fact every old entry obeys $x + (s+1)/2 \le c$ in P_i. Thus every old entry reverses its sign label, which is exactly the signature flip from negating the form, and the new block agrees with P_{i+1} with its shift. The inequality is the grid entry bound obtained by putting its new center above all old blocks and the old lists. The rank and parity shifts make each real infinitesimal character regular and in the discrete-series lattice for the corresponding row group. This is the signed pair in Paul, Theorem 2.7(b) with the displayed inserted string. For first occurrence, in the notation of Ichino, Theorem 6.1 the smaller centered parameter has no zero and $k_\pi = -1$, so we are in the odd difference case with $t = (s-1)/2$ and $l = 0$ in its tested tower. The old entries are outside the range for $C_\pi^{\pm}(t)$. Deleting one hyperbolic plane decreases l below zero, so the lower real lift vanishes. The global constant-term formula then gives cuspidality for these row lifts, and real Howe duality gives the specified target type. The same argument works after taking the original dual convention on both sides. Original row lists, shifts and centers · The corresponding real type for the placed data.

For each real place take these fixed smooth irreducible modules and the indicated top compact dual. The real polynomial row parameters ρ and a smooth bottom vector w map to the same vector $v(\rho, w)$ of this row module for all multiplicity vectors and all finite tests. Indeed real local Howe intertwining at each arrow has one line for the specified source and target types. Fix one intertwiner on the Fock polynomials at each real place. With all finite factors held fixed, write the kernel integral on the real isotypic target in its multiplicity space. Each component map of the source and real oscillator test is an intertwiner and hence a scalar multiple of the fixed map. The scalar can depend on the finite factors and the input multiplicity, and is independent of the real polynomial and input real vector. Use the smooth globalizations to extend this equality in the source vector: continuity of theta integration on smooth cuspidal vectors follows from uniform rapid decay seminorms and the polynomial bounds above. Iterating

gives $\theta_{\rho,\Phi}(tw) = F_{t,\Phi}(v(\rho, w))$. The finite multiplicity map $F_{t,\Phi}$ carries the remaining scalar and may be zero. The real column module is the tensor product of the local dual highest-weight modules, and each column map with fixed Ψ is equivariant on this module. These factorizations use the same real modules at every varying t, Ψ, Φ.

The iterated Bessel pairing for given maps is a continuous bilinear functional on these real modules. In the terminal period the top definite unitary group is anisotropic, so the space of evaluation of the two automorphic functions is covered by a compact set from that quotient and the flag-unipotent quotient. The integrals converge absolutely and are bounded by smooth seminorms on this compact set. A simultaneous right translation of the two arguments by the embedded top real group leaves the pairing invariant; right translation by the flag changes it by the fixed Bessel character, since replacing n by nn_0 in the weighted integral produces $\psi_{\mathrm{Bes}}(n_0)$ from the conjugate weight. For each real place it is generic in the sense of [30, Theorem A and §2.2]. All adjacent-root characters are nontrivial, and the side character is nontrivial on the distinguished nonisotropic line, including its prescribed nonzero scalar from the terminal frame. The theorem applies to the irreducible smooth Fréchet globalizations in question. Applying it with the other local vectors held fixed, successively at every real place, gives a single line of such bilinear functionals. More explicitly, choose pure tensors at which a nonzero functional is nonzero. At each place the local functional obtained by fixing those other factors spans its one-dimensional space. Proportionality then holds on pure tensors by induction on the number of varying factors and extends by continuity. Pulling back along $v(\rho, w)$ gives the same function $a(\rho, \xi)$ up to a scalar depending on t, Ψ, Φ, for fixed w.

If the input space is nonzero, bottom separation and the proved ascent give one triple with nonzero terminal period at some real pair. Choose that pair (ρ_*, ξ_*). For every other nonzero input, its triple supplied by the same separation has nonzero scalar on the common function, and hence the same pair (ρ_*, ξ_*) works. This choice depends on the grid, volumes and w and precedes the quantifier over the nonzero inputs. For an empty separation problem any real pair suffices.

For the final arithmetic-to-period implication, fix a finite row test Φ. The arithmetic list for this Φ runs through all pairs of a terminal double coset and a basis vector of the prescribed column coefficient space. The terminal group is definite, so the rational lattice stabilizers are finite. Its congruence level is principal at places above the two distinct type primes. A finite-order matrix acting trivially modulo one of those places has order a power of that residue characteristic. Indeed, in the filtration by deeper principal subgroups each successive quotient has that characteristic, and the intersection of the filtration is the identity. The two places therefore force every stabilizer at this level to be trivial.

In the terminal column level the finite right character is trivial. A tuple of values at the class representatives extends uniquely to a column by the rational coefficient covariance and right-level invariance. The class stabilizer calculation makes this extension well defined for arbitrary coefficient values. The arithmetic test attached to a representative and a basis vector has precisely that value at the representative and zero at the other classes. These tests thus span all columns of the specified compact dual type on the same level. In particular the column supplied by separation is in their span for this fixed Φ. Contraction by ξ_* preserves this linear combination. The arithmetic test index and coefficient read · Triviality of the actual class stabilizers and The span for fixed finite row test.

At this same real pair, absolute convergence gives linearity of both integrals in the column. If the arithmetic functionals vanish, their integral comparison gives vanishing on each column and therefore on the whole span. Here the volume and normalization factors in the comparison are nonzero. If starting from a preselected arithmetic real pair, apply the stipulated real-replacement property of that choice to the vanishing family before using the newly chosen pair. Equivalently one may arrange the arithmetic choice to use (ρ_*, ξ_*) from the outset. The base-choice conditions concern the Weil data, measures, nonzero w, and square constants, so this replacement preserves them. Injectivity at the good pair then yields the real-replacement property for that choice, since a simultaneous zero there forces the input to be zero. The separating nonzero period then gives the contradiction for every nonzero realizable input. Use of the arithmetic column span · Real replacement from an injective choice.

2.5 From top theta periods to bottom Bessel periods

The next argument concerns the simultaneous comparison stated at the end of the section. Its full base lift and nonzero reference are used to derive that comparison. The principal-level tests above and the arithmetic realizations below provide separate inputs to the descended pairing. At the selected translate, the period comparison uses the geometric embeddings, real types, twisting characters, levels and compatible local tests in its hypotheses. Here $r_S \geq 1$, so $m \geq 4$. In the argument establishing the comparison there are two surjectivity assertions. The projected rows descend onto E_{row}, and the base chain on the larger top space reaches every row in the full permitted bottom multiplicity space for the fixed level, character and discrete series. The second assertion is used for the nonzero reference in that argument. The first is also needed for descent in the arithmetic application below. For the theta norm, let A have rank n and let the target B have rank $M > n$ and global Witt index q. In the norm formula write χ for the character inducing the normalized Siegel section. It is χ_V in Gan–Qiu–Takeda and is the inverse of the character denoted χ_B in the release's Weil datum. Its restriction to $\mathbb{A}_F^\times$ is $\eta_{E/F}^M$. The doubled source is $A^\square = A \oplus (-A)$. Its Siegel parabolic stabilizes the diagonal isotropic subspace of dimension n. For a Schwartz test Φ in the doubled Weil representation $\omega^\square$ that is a polynomial Fock vector at infinity, let Φ_s be the standard Siegel section with value $\Phi_{s_0}(g) = (\omega^\square(g)\Phi)(0)$ at $s_0 = (M-n)/2$, in the normalized induction by $\chi \circ \det |\det|^s$. If $M - q > n$, Yamana's convergent Siegel–Weil formula gives

$$E_{\mathrm{Sieg}}(g, s_0; \Phi_{s_0}) = I(g, \Phi), \qquad I(g, \Phi) = \int_{[U(B)]} \theta^\square(g, h; \Phi)\, dh_1.$$

Here dh_1 gives the quotient volume one. The theta integral converges absolutely and the Eisenstein series is holomorphic at s_0. For $M > 2n$ the defining Eisenstein series itself converges absolutely there. [47, Theorem 2 and the convergence statements preceding it] The column tower has definite targets. For the upward adjoints of the row maps, the rank increase is t, which exceeds the negative index at each real place of the target and hence its global Witt index. For an anisotropic target and $s_0 > 0$, Gan–Qiu–Takeda's Theorem 7.1 gives the same identity. The factor in Tamagawa conventions is $\tau(U(B))/[E : F] = 1$ times $I(g, \Phi)$. [19, §3.1 and Theorem 7.1]

Let f_1, f_2 be cusp forms in an irreducible cuspidal representation π of $U(A)(\mathbb{A}_F)$. In the bilinear convention the lift with test φ_i is $\theta_i(h) = \int_{[U(A)]} \theta(g, h; \varphi_i) f_i(g)\, dg$. The doubling seesaw uses the doubled test $\Phi = \delta(\varphi_1 \otimes \overline{\varphi_2})$, where δ intertwines the doubled Weil representation on the two copies of $U(A)$ with the tensor product of the original representation and its dual twisted by $\chi \circ \det$, and $\delta(\varphi_1 \otimes \overline{\varphi_2})(0) = \langle \varphi_2, \varphi_1 \rangle$. In Tamagawa measures, in the convergent range where the lifts have Petersson pairings and the seesaw integrals are absolutely convergent,

$$\langle \theta_2, \theta_1 \rangle = [E : F]\, Z(s_0), \qquad Z(s) = \int_{[U(A)]^2} E_{\mathrm{Sieg}}((g_1, g_2), s; \Phi_s)\, f_1(g_1)\overline{f_2(g_2)}\, \chi(\det g_2)^{-1}\, dg_1\, dg_2.$$

The factor $[E : F]$ converts the volume-one theta integral to Tamagawa measure. In $n < M \leq 2n$, $M > n + q$, Gan–Qiu–Takeda's Propositions 11.1(ii) and 11.2 give this identity when the lower lifts of π in the target Witt tower vanish, so the lift under consideration is cuspidal. The lower-lift condition is vacuous for an anisotropic target. [19, §§11.3–11.5, Propositions 11.1(ii) and 11.2] For the selected chains the release uses the absolutely convergent seesaw with the Siegel–Weil value above. Its doubling integrand uses the conjugates of the bilinear inputs as its two arguments, so with inducing character χ_B^{-1} the scalar on the second copy is $\chi_B(\det g_2)$. In the prescribed measures the target volume is the actual quotient volume. The pairing uses L^2 bounds on both theta kernels and rapid decay of the source cusp forms. square-integrable kernels and Fubini for the norm

With factorized vectors, pairings and measures the doubling integral unfolds, initially for Re s sufficiently large, to

$$Z(s) = \prod_v Z_v(s), \qquad Z_v(s) = \int_{U(A)(F_v)} \Phi_{v,s}(g, 1)\, \langle f_{2,v}, \pi_v(g) f_{1,v} \rangle\, dg.$$

The local factors have meromorphic continuation. For unramified data the spherical pair has pairing one, the spherical section has value one on the hyperspecial compact, and the compact has volume one. In those normalizations $Z_v(s) = L_v(s + \frac{1}{2}, \pi_v \times \chi_v)/d_v(s, \chi_v)$. [19, §§11.5–11.6] For these spherical data at split and inert places outside the exceptional primes of Eischen–Harris–Li–Skinner's calculation, the denominator is

$$d_v(s, \chi) = \prod_{i=0}^{n-1} L_v(2s + n - i, \chi|_{F_v^\times} \eta_{E/F,v}^i).$$

Here $\eta_{E/F,v}$ is the local quadratic character, trivial at split places, and the numerator is the standard local factor of the base change of π_v twisted by χ_v. [20, Hypotheses 1.4.4 and §4.2.1, (53)] Since $\chi_v|_{F_v^\times} = \eta_{E/F,v}^M$, the characters in the denominator take values of absolute value one on a uniformizer. Every argument $2s+n-i$ is positive at real $s > 0$, so the local denominators are finite and nonzero there. With S_0 containing the archimedean places, the places over the auxiliary prime and all other places where the data fall outside this unramified calculation, the meromorphic expression for $Z(s)$ is

$$\frac{L^{S_0}(s + \frac{1}{2}, \pi \times \chi)}{\prod_{v \notin S_0} d_v(s, \chi_v)} \prod_{v \in S_0} Z_v(s).$$

The norm uses the value of the meromorphic expression as a whole at the holomorphic point s_0, with the prescribed measure conversion.

For the real local criterion, take A skew-Hermitian and B Hermitian, with $M \geq n$, and use the Weil representation of the pair with additive character $x \mapsto e^{-2\pi i x}$ and unitary splitting characters of exponents congruent to the two ranks. Given an irreducible tempered τ_∞ of $U(A)(\mathbb{R})$, Ichino's local form

$$(\varphi_1, \varphi_2, u_1, u_2) \longmapsto \int_{U(A)(\mathbb{R})} \langle \varphi_2, \omega(g)\varphi_1 \rangle \, \overline{\langle u_2, \tau_\infty(g) u_1 \rangle} \, dg$$

is absolutely convergent on Schwartz tests and smooth vectors. It is nonzero as a form exactly when the small theta lift of τ_∞ to $U(B)(\mathbb{R})$ is nonzero. One can test nonvanishing on polynomial Fock vectors and K-finite vectors. [39, §4.1, Proposition 7.1 and Remark 7.2] For the bilinear global inputs f_i, the local form uses $\tau_\infty = \pi_\infty^\vee$ and u_i dual to $f_{i,\infty}$ under the invariant Hermitian identification with the conjugate representation. On the first copy of the source group the doubled section at s_0 has value $\langle \varphi_2, \omega(g)\varphi_1 \rangle$ in the corresponding local normalization. The release proves the Hecke norm formula using the split unramified formula, its inert calculation and the convergent seesaw. Assembly of the Hecke norm formula. The surjectivity along the prescribed chains uses the norm in the specified quotient measures, the unramified factors and nonzero theta maps on the prescribed types and levels. With that surjectivity and the selected real, twisting, level and geometric hypotheses, the release uses rapid decay of cusp forms, discrete series, Harish–Chandra modules and strong approximation for the special unitary groups to choose the compatible top and bottom models.

In these models a relation between top periods is carried to the bottom in two stages. First the seesaw comparison at the last two levels uses the codimension-one top pairing. Then each two-level step uses a Bessel or Fourier–Jacobi coefficient of the corresponding lower row. Composing the two Weil kernels and unfolding the quotient integrals gives the lower coefficient at the prescribed translate. The ramified column conditions identify the needed coefficients and rapid decay justifies the integrations. For prescribed real tests and translates in a two-level step, one matching real test works for all tuples. The kernel expansion uses finitely supported frame coefficients chosen independently of the tuple, with the tuple dependence in the normalization and the prescribed vectors. These are the quantifiers in the test identities. Induction transports any finite vanishing linear combination of the normalized top periods to the corresponding combination of bottom coefficients at every permitted frame and real vector. This is the comparison by the last two levels and successive two-level steps.

For each word w, combine the invertible column change from the theta comparisons with the isometry identifying the chosen bottom columns with the columns of the Poincaré series. The composite multiplies

norms by a constant $r_{\mathrm{col},w} > 0$. Dividing by $r_{\mathrm{col},w}$ gives a unitary U_w on $\mathcal{M}_w$ and divides both bottom pairings by the same constant. For a permitted full bottom row T, choose a preimage under the base chain and let $p_{w,T}(\phi_0)$ be its normalized top period on the base column ϕ_0. Let $c_{\mathrm{top}}(w,\boldsymbol{\mu})$ and $c_{\mathrm{bot}}(w,\boldsymbol{\mu})$ be the products of the upper and lower factors in that comparison. The upper product includes the two column factors and the real comparison factor from each two-level step, and the factor for the top pairing. The lower product contains one factor for each downward step. The scalar in the period relation is

$$d_{w,\boldsymbol{\mu}} = s_w \frac{c_{\mathrm{top}}(w,\boldsymbol{\mu})}{c_{\mathrm{bot}}(w,\boldsymbol{\mu})} \neq 0,$$

where $s_w \neq 0$ comes from comparing the two column lifts and is independent of the tuple. For a tested row F_m in the projected top space, the relation

$$\alpha\, p_{w,T}(\phi_0) + \alpha'\, d_{w,\boldsymbol{\mu}} b_{w,\boldsymbol{\mu}}(\phi, F_m) = 0$$

gives a vanishing combination of the two bottom pairings with coefficients α, α'. The first uses T and $U_w\phi_0$, and the second uses the downward image of F_m and $U_w\phi$. The release obtains this implication first for the chosen top coefficient coordinates, including all choices of the two column inputs. The rule for $d_{w,\boldsymbol{\mu}}$ follows by dividing by the nonzero lower factors.

Let t_0 be the fixed finite element matching the bottom row level and character with those of $\mathcal{H}$, and let J_{row} be the common bottom row level subgroup in these coordinates. The construction starts with a reference placed and framed Bessel datum. Comparing its translate with t_0 and absorbing a central adjustment in the embedded column group places it in zt_0J_{row}, for one finite central element $z \in U(Y_0)(\mathbb{A}_F)$. The pairwise coset comparison then puts every word in that same coset. This gives the common placement. Matching at t_0 also gives matching at zt_0, since conjugation by z fixes the level subgroup and its character. Right translation by z and its inverse preserve the Petersson norm and distinguished row conditions, giving a unitary V on E_{row}.

Evaluation on the chosen real unit vector at zt_0, followed by complex conjugation, gives a conjugate-linear bijection from the full permitted bottom multiplicity space onto $\mathcal{H}$. This uses the matched level and character, the real discrete-series identifications and the real gap hypothesis. Write ev_z for this map with the bottom pairing normalization. For a full row T and a ramified column ψ the pairing is

$$\langle \mathrm{ev}_z(T), \mathcal{P}_w J_w \psi \rangle.$$

For the downward image of a projected top row, the matching to the distinguished space represents this pairing by $\mathcal{E}_w(\psi, V\iota_{\boldsymbol{\mu}}F_m)$, up to a fixed nonzero conversion factor in the row normalization. We use the same conversion on both terms of the bottom relation. The release compares these pairings in the chosen row and column coordinates.

Use the full-space diagonal estimate (56) and the ramified dimension bound (57) for these same full operators and inclusions. On $\ker \mathcal{P}_w$, the operator $K_{ww} - I$ acts by $-I$. Extending an orthonormal basis of this kernel to one of H_w^{full} gives

$$\dim\ker \mathcal{P}_w \leq \|K_{ww} - I\|_{\mathrm{HS}}^2 \leq e_0 d_w < r_w.$$

Indeed $r_w \geq (1 - 2f/b)d_w$, $d_w > 0$ and $e_0 < 1 - 2f/b$ at sufficiently large x, since $f \leq x^5$, $b = \exp(x^{\varepsilon})$ and $\sqrt{e_0} \leq \exp(-x^{\varepsilon/8})$. Choose $\psi_* \in \mathcal{M}_w$ with $y = \mathcal{P}_w J_w \psi_* \neq 0$. Take the full bottom row T_* with $\mathrm{ev}_z(T_*) = y$. Its pairing with ψ_* is $b_{*,w} = \langle y, y\rangle \neq 0$. Both the choice of T_* and its pairing use the full bottom space. See the construction of the reference pair.

Use the input $\phi_0 = U_w^{-1}\psi_*$ and the base lift of this same T_* to form the reference top period $p_{*,w} = p_{w,T_*}(\phi_0)$. If $p_{*,w} = 0$, the relation with $(\alpha,\alpha') = (1,0)$ would give $b_{*,w} = 0$. Thus $p_{*,w} \neq 0$. In the two-term relation take $\alpha = d_{w,\boldsymbol{\mu}} b_{w,\boldsymbol{\mu}}(\phi, F_m)$ and $\alpha' = -p_{*,w}$. The resulting bottom equality is

$$d_{w,\boldsymbol{\mu}}\, b_{w,\boldsymbol{\mu}}(\phi, F_m)\, b_{*,w} = p_{*,w}\, \mathcal{E}_w(U_w\phi, V\iota_{\boldsymbol{\mu}}F_m). \tag{35}$$

Here the fixed, nonzero conversion from the bottom row normalization is included in $p_{*,w}$. The equality holds on the compatible coefficient coordinates, which span the projected complex top row spaces. It therefore extends to those spaces by bilinearity.

For the base tuple, (35) makes the top period vanish on $\ker \iota_{\boldsymbol{\mu}_0}$ in its projected domain. Evaluate the base period on any preimage of $F \in E_{\mathrm{row}}$ to define $B_w(\phi, F)$. Surjectivity gives a preimage for every F, and kernel vanishing makes its choice immaterial. With $a_w = p_{*,w}/(b_{*,w} d_{w,\boldsymbol{\mu}_0})$ we obtain

$$B_w(\phi, F) = a_w \left\langle s(VF), \mathcal{P}_w J_w U_w \phi \right\rangle_{\mathcal{H}}, \qquad a_w \neq 0. \tag{36}$$

The same V applies to every word, with a_w and U_w chosen for each w. At another tuple the identity on its projected rows reads

$$b_{w,\boldsymbol{\mu}}(\phi, F_m) = \frac{d_{w,\boldsymbol{\mu}_0}}{d_{w,\boldsymbol{\mu}}} B_w(\phi, \iota_{\boldsymbol{\mu}} F_m). \tag{37}$$

2.6 The coefficient-field map and finite-place tests

2.6.1 The definite coefficient models

Write K for the coefficient field, $\mathcal{O}_K$ for its integers and ι_K for the chosen embedding into $\mathbb{C}$. Use the compact algebraic coefficient types at the top and on the definite column tower. In the chosen matrix coordinates these are polynomial representations with determinant twists, defined over K using the specified embeddings of E. At a definite vertex with level J, choose representatives c for $G(F)\backslash G(\mathbb{A}_{F,f})/J$. The stabilizer $\Gamma_c = G(F) \cap cJc^{-1}$ is finite. Evaluation at the representatives identifies coefficient-valued forms with $\bigoplus_c W_c^{\Gamma_c}$, where W_c is the algebraic coefficient space and the stabilizer action includes the level character. In fact the value at any other point is determined by covariance, and the stabilizer condition makes its definition independent of the presentation in that double coset. This description commutes with extension of scalars because the invariants are kernels of finite systems of equations over K. The integral model consists of the vectors whose value at each c belongs to its prescribed value lattice. The equations impose a lattice in the invariant space: it is a submodule of the direct sum of the value lattices and spans after clearing denominators.

For the principal-level space let Γ now denote the finite group of level changes. For the chosen type characters χ of Γ and the character α of a subgroup P, the good projector is the product of

$$\sum_{\chi} \frac{\chi(1)}{|\Gamma|} \sum_{g \in \Gamma} \chi(g^{-1})\, g, \qquad \frac{1}{|P|} \sum_{g \in P} \alpha(g)^{-1} g.$$

The first operator, by character orthogonality, selects the specified isotypes. The second is the projector on the α-equivariant level subspace by averaging. They commute since the first is central. With the class inner product the level actions are unitary, so the projectors are orthogonal. The class weights are $1/|\Gamma_c|$ when the level and compact real group have volume one. For the larger level use $|P|^{-1}$ times the principal class inner product: the orbit of a larger-level class contributes $|P|$ times its mass in the principal class sum by orbit–stabilizer. The level and type matrices lie in K with the stated character values. Their action can be read at every realigned class representative. If a representative translates to $cg = \gamma dk$, covariance reads the result as the value at d acted on by the coefficient representation of γ and the level character of k. The stabilizer condition makes this independent of the decomposition. This is also the finite matrix action on the class values. Thus for the prescribed transported frames and coefficient lattices, the projector image $A_{\boldsymbol{\mu},K}$ complexifies to the actual projected top space $\mathcal{V}^0_{\boldsymbol{\mu}}$. In this assertion the determinant twists, level characters and coefficient lattices are those on the same definite groups as in (12). Row model identities

Pull back the value lattice to the image of the projector by intersection. Equivalently, for a basis matrix $E_{\boldsymbol{\mu}}$ of that image and ambient class value lattice $L_{\boldsymbol{\mu}}$, its lattice in the basis coordinates is

$$M_{\boldsymbol{\mu}} = \{y \,:\, E_{\boldsymbol{\mu}} y \in L_{\boldsymbol{\mu}}\}.$$

Indeed every integral vector in the image has unique coordinates in $E_{\boldsymbol{\mu}}$, and their integrality condition is exactly the displayed one. This pullback uses the original value lattice even at primes dividing projector denominators.

The same class model gives rational projections onto the fitted auxiliary coordinates at the base level. Let E_J be the inclusion of the level coordinates in the full class coordinates and let $E_J^\flat$ be its inverse on its image, extended by the orthogonal projection there. Let V include the good eigenspaces at that level, with inverse $V^\flat$ on its image, again extended by orthogonal projection. Let Z include the selected complementary subspace of the good space and let P_Z be the orthogonal coordinate map to this complement, so $P_Z Z = I$. The stage projection from the level model is $P_Z V^\flat E_J$. Its right inverse is $E_J^\flat VZ$, since $E_J E_J^\flat VZ = VZ$ and $V^\flat V = I$. All matrices are over K for the class Hermitian form. The rational surjection maps the level value lattice onto a full lattice in these stage coordinates. The arithmetic fitted subspace is the kernel complement identified for the auxiliary theta map in §2.6.3. Right inverse of the stage projection.

2.6.2 Descent on the selected row spaces

In this argument the ranks are $n_j = r + jt$, with t odd and $t \geq r + 5$. The final row group is $U(X')$, identified by the prescribed isometry with $U(Y_m)$. All preceding row groups are $U(Y_j)$. For a tuple $\boldsymbol{\mu}$, write $\mu_{\boldsymbol{\mu},i} = \mu_{0,i}\xi_{\boldsymbol{\mu}}$ for its splitting characters before the parity adjustments. The finite-order character $\xi_{\boldsymbol{\mu}}$ is trivial at infinity and on the base field. Write S for the designated split places of the residue group and Q for the additional split places of the tuple. At Q the character $\xi_{\boldsymbol{\mu}}$ is tamely ramified and the base splitting characters are unramified. In the orientation $\mathrm{Hom}(Y_j, Y_{j+1})$ let (χ_S, χ_T) denote the two characters of this arrow. At step $i = j + 1$ they are

$$(\mu_{\boldsymbol{\mu},i}\eta^{-1}, \mu_{\boldsymbol{\mu},i}) \quad (i \text{ odd}), \qquad (\mu_{\boldsymbol{\mu},i}, \mu_{\boldsymbol{\mu},i}\eta^{-1}) \quad (i \text{ even}).$$

Here η is unramified at all finite places and supplies the half-integral change at infinity. Use the selected half of the split algebra to read all the matrices. The specified local matrices below are in the transported lattice frames. In particular the same description uses the top frame on X' and the final datum on $\mathrm{Hom}(Y_{m-1}, X')$. Tuple twist and parity · Levels and ramified tuple characters · Actual successive and final finite tests.

Finite type and measure. Use the chosen subgroup P with unitary character α in the common residue group $\Gamma = \prod_{v\in S} \mathrm{GL}_r(k_v)$. On each stage the level at S is the inverse image of the upper block parabolic for $(r, t, \dots, t)$, with first block in P. The character λ_j is α on the first block and the respective $\chi_1, \dots, \chi_j$ on the determinants of the later blocks. Write $a = n_j$, $b = a + t$. At S the finite test is supported on integral matrices reducing to $\binom{u}{0}$, with u in this lower level. Its weight is

$$\lambda_j(u)\chi_{j+1}(\det u)^{-1}. \tag{F}$$

Here χ_{j+1} on the residue field is the inverse of the restriction of $\mu_{0,j+1}$ on the chosen half. The tuple twist is unramified at these places. The base test off $S \cup Q$ is the unit self-dual vacuum. At a nonsplit finite place the given characters are unramified and this is the hyperspecial vacuum.

In the split Schrödinger model the source pullback is $X \mapsto Xg$, with unit factor $\chi_{j+1}(\det g)$ on integral units at S. Thus (F) transforms on the source by λ_j. For k on the target, the pullback is $X \mapsto k^{-1}X$, with unit factor $\chi_{j+1}(\det k)^{-1}$. Reducing $k = \left(\begin{smallmatrix} u_0 & * \\ 0 & d_0 \end{smallmatrix}\right)$, the change in the test is $\lambda_j(u_0)^{-1}\chi_{j+1}(\det u_0)$. Multiplying the unit factor gives $\lambda_j(u_0)^{-1}\chi_{j+1}(\det d_0)^{-1} = \lambda_{j+1}(k)^{-1}$. To see exact coefficient support, suppose X, Xg are both in the support. Write A, A' for their top a-row square blocks. Both are invertible integral, and $g = A^{-1}A'$ reduces to the lower level. Conversely on that level the coefficient is $\lambda_j(g)\|\Phi_v\|^2$. The doubled height is one there. With the conjugate Petersson input the two λ_j factors cancel. The semilocal contribution at

S is exactly the level volume times the product of the norms of the actual local tests. The residue weight · Target character law on the placed test.

Use the principal first-block level $P = \{1\}$. For each $\gamma \in \Gamma$ lift it to the same diagonal first block on the two stages. Acting simultaneously on both legs conjugates the first block in (F) and leaves all other diagonal determinants fixed. The two splitting determinant factors cancel. The test is therefore fixed by this simultaneous action. The lift and its Petersson adjoint intertwine the two actions of Γ, by the change of variable in the quotient integral. Both consequently commute with the common good central projector. The P, α test is the sum of the principal tests translated on the lower side by these lifts, with coefficient $\alpha(\gamma)^{-1}$, $\gamma \in P$. Indeed the term indexed by γ has first block γ^{-1} and the source splitting multiplier is $\chi_{j+1}(\det \gamma)$, as in (F). On the α-isotype the summed integral is $|P|$ times the principal integral before Haar conversion. The smaller-level Haar is $|P|^{-1}$ times the principal one, because the smaller level has index $|P|$ over the principal level. This cancellation applies on every stage where this first-block level changes, including the top restriction.

If dv_j is the underlying row measure in the period definitions, the normalized stage measure is

$$dv_j^{\boldsymbol{\mu}} = m_Y(\boldsymbol{\mu}, j)\, dv_j, \qquad m_Y(\boldsymbol{\mu}, j) = \prod_{v \in Q} [\mathrm{GL}_{n_j}(\mathcal{O}_v) : J_{(r,jt),v}]. \tag{G}$$

At $j = 0$ the suffix is empty and this factor is one, as it is for all stages of the base tuple. Multiplying the product of the two form evaluations by this row measure gives the Hilbert metric of the corresponding stage. The downward map integrates using $dv_{j+1}^{\boldsymbol{\mu}}$. Thus the adjoint integrates the conjugate kernel using $dv_j^{\boldsymbol{\mu}}$, by Fubini with the product of these two measures. This gives precisely the powers of m_Y in each direction, including the measure after the final isometry to X'. Stage metrics and row maps · Small-level Haar factor.

Doubling at a tuple prime. Here the full-rank support argument at S must be replaced by an integral on the suffix of the matrix. Fix $v \in Q$, omit v from the notation, and put $d = jt$. In the selected split coordinates the finite vector is

$$f(X) = \mathbf{1}_{\mathscr{X}}(X)\, \xi_{\boldsymbol{\mu}}(\det \overline{X}_{00}), \qquad \mathscr{X} = \left\{ \begin{pmatrix} A & B \\ C & D \end{pmatrix} \text{ integral} : A \in \mathrm{GL}_r(\mathcal{O}), C \equiv 0 \pmod{\varpi} \right\}. \tag{H}$$

In (H) the first row and column blocks both have size r. The finite level is the parahoric of shape (r, d) with character $\xi_{\boldsymbol{\mu}}^{-1}$ on the suffix determinant. This follows by applying the split Weil formula to (H). For instance the source pullback contributes $\xi(\det \bar{g}_{00})$, and the scalar of the Weil action contributes $\xi(\det \bar{g})^{-1}$. The same calculation on the target gives the inverse of its level character. This verifies level covariance on both sides of the actual step, including its top version. The additional-prime factor · The suffix determinant in the level character.

Let π_j be the representation on the conjugate automorphic input in the doubling integral. At $j = 0$ it is spherical at v. For a constituent obtained by the first j steps, π_j is the normalized parabolic induction of a spherical GL_r representation and a ξ-twist of a spherical GL_d representation. Indeed split type-II Howe correspondence adds the determinant character χ_S on the new GL_t block and twists the old representation by χ_S/χ_T. The latter is unramified at v. The successive suffix factors before taking the common twist out are unitary determinant characters. Normalized parabolic induction of irreducible unitary representations of split GL is irreducible. The suffix spherical representation at step j consequently has the Satake strings from the earlier steps. Mínguez, Corollary 6.3, Tadić, §0 and the unitary induction property.

Include the scalar $\chi_T(\det g)^{-1}$ from the Weil action in the group coefficient, so now use $\pi_j' = \pi_j \otimes \chi_T^{-1}$. It is the normalized parabolic induction of $\tau \boxtimes \sigma$ for the parabolic $P_{r,d}$ where τ is a ξ^{-1}-twist of a spherical representation and σ is spherical. Denote by $K = \mathrm{GL}_{r+d}(\mathcal{O})$ and $J = J_{(r,d)}$. The J-line needed in the group coefficient now has character $\xi^{-1}(\det \bar{g}_{00})$. Its dimension is one: take a principal-series presentation surjecting onto π_j' with the r characters of inertia type ξ^{-1} first and the remaining d unramified. Taking

the J-character is exact by averaging. In the induction, the double cosets for the upper Borel and J are the subsets of r diagonal positions assigned to the first block of J. On the intersection with the diagonal units the character can match only if this subset is precisely the r positions of type ξ^{-1}, since ξ is nontrivial on the residue units. That coset contributes at most one dimension. In the parabolic induced model a nonzero vector with this character is supported on $P_{r,d}J$, with its value at one the product of the two spherical vectors after their character twists. The equivariance prescriptions agree on $P_{r,d} \cap J$, so this defines the actual line. Normalize this value to have norm one in the inducing representation.

With $H(g)$ the normalized doubling height, the local integral, up to a positive scalar, is

$$\int_{\mathrm{GL}_{r+d}} |\det g|^{b/2} H(g)^{s-t/2} \langle \pi'_j(g)v, v\rangle \int f(Xg)\overline{f(X)}\, dX\, dg. \tag{I}$$

When $d = 0$, the full-rank block in (H) forces the support of the coefficient to be K, so (I) is a nonzero volume times two squared norms. Suppose $d > 0$. In the compact induced realization, write the group coefficient as $\int_K \langle v(kg), v(k)\rangle\, dk$. The compact variable is supported on J. Put $g' = kg$ and $X' = Xk^{-1}$ in (I), so that $Xg = X'g'$. The covariances of $f(X'k)$ and $v(k)$ cancel in the two conjugate slots. This leaves a positive compact volume times the same integral with $\langle v(g), v(1)\rangle$ in the integrand. It now only integrates g in $P_{r,d}J$.

For $X \in \mathscr{X}$, eliminate its first-column blocks by the target unit matrix

$$\begin{pmatrix} A^{-1} & 0 \\ -CA^{-1} & I \end{pmatrix}.$$

Then eliminate its first-row right block by the integral upper unipotent matrix on its right. These changes can be made in the inner g-integral: the target change multiplies $f(Xg)$ and $f(X)$ by the same unit scalar. The source change is in the upper unipotent of $P_{r,d}$ and is absorbed by left translation in g, using the left covariance of the induced vector and the trivial character on this upper unipotent. Doubling height is bi-invariant under K, by the lattice normalization on the doubled source. The remaining rectangular block $D_X = D - CA^{-1}B$ ranges over all integral matrices of size $(d+t)\times d$; the change from D is an integral translation and preserves Haar. In $P_{r,d}J$ use $g = m_0 uk$, with $m_0 = \mathrm{diag}(A_0, D_0)$, u upper unipotent, and $k \in J$. Its Haar integral is in the order $dm_0\, du\, dk$ up to a positive constant, since writing the parabolic element as $m_0 u$ absorbs the left modular density. In the test after the eliminations, A_0 is an integral unit and u is integral, while $D_X D_0$ must be integral. The k-character cancels against the character of f, and the A_0-character cancels against that of τ. The normalized inducing factor is $\delta_{P_{r,d}}(m_0)^{1/2} = |\det D_0|^{-r/2}$. Multiplying it by $|\det g|^{b/2}$ gives $|\det D_0|^{(d+t)/2}$. The doubled height of $\mathrm{diag}(I_r, D_0)$ is the doubled height of D_0, by the orthogonal product of the doubled lattice and the identity block.

The integral that remains is exactly the spherical split doubling integral for the pair $\mathrm{GL}_d, \mathrm{GL}_{d+t}$ with group coefficient of σ and integral-matrix vacuum. The calculation can first be made in the right half-plane of absolute convergence. At $s = t/2$ the local numerator is the Euler factor of σ at $(t+1)/2$. Every suffix Satake entry has modulus at most $q^{(t-1)/2}$, so the numerator has no pole or zero there. The spherical denominator is a product of nonzero finite Euler factors at positive arguments. For absolute convergence use the Cartan index $\mathbf{k} = (k_1, \dots, k_d)$. The oscillator coefficient is $q^{-(d+t)\sum_i |k_i|/2}$. The volume of its double coset is $O(q^{2\langle \varrho, \mathbf{k}\rangle})$, where ϱ is the half-sum of the positive roots of GL_d. The spherical coefficient is bounded by a fixed polynomial in $\sum_i |k_i|$ times $q^{-\langle \varrho, \mathbf{k}\rangle + (t-1)\sum_i |k_i|/2}$. This follows from the spherical-function formula, including the limiting polynomial factors at coincident Satake entries. [99, Proposition 4.1 and Theorem 4.2] Since $\langle \varrho, \mathbf{k}\rangle \le (d-1)\sum_i |k_i|/2$, their product has the exponential bound $q^{-\sum_i |k_i|}$ at the critical point. The doubling height changes this bound by at most a fixed exponential in $|\mathrm{Re}\, s - t/2| \sum_i |k_i|$, so it still sums on a neighborhood of that point. All omitted factors in the reduction are the positive volumes of the integral blocks. Hence the particular tested factor at every tuple prime is nonzero.

Real factors, labels and passage to the cuts. For the actual Y_j parameters the center is $-e_j$, with the parity shifts of χ_T and χ_S on the smaller and larger models. In the convention for the local correspondence before passing to the dual row, the centered smaller lists on the determinant cover have the form $(A, B; C, D)$, with the first and third positive and the other two negative. Centering the larger lists gives

$$(A,\ ((t-1)/2, \dots, -(t-1)/2), D;\ C, B). \tag{J}$$

The shift from $-e_j$ to the two actual centers is the half-step dictated by the corresponding rank parities. The sign labeling flips the old entries below the center and inserts the new block in the positive list. By the prescribed row interlacing the distance of any old entry from the center is at least $(t+3)/2$. In computing the real gap, opposite-sign entries lie in disjoint strips separated by two successive b-endpoints. Subtracting the block length t from the gap of these endpoints leaves a gap larger than the cube of the larger row rank. These numeric conditions are unchanged by the tuple twist. By Paul's Theorem 2.7(b) the centered parameters in (J) are the prescribed pair. The smaller centered parameter is integral with no zero. It has first occurrence in this tower by Ichino's Theorem 6.1 with $k_\pi = -1, t_{\text{Ichino}} = (t-1)/2, l = 0$; the prior centered entries are outside the two near sets. In the reverse direction one deletes the maximal central string, as in Paul's Proposition 3.4(ii). The global constant-term calculation gives cusp forms in the designated real stages, since the lower real members vanish. This applies with the definite final signature on X'. Centered signed lists and the gap · Real exponents of the placed tests.

Use Ichino's Proposition 7.1 on this pair, polarize in the two oscillator slots and then the two source slots, and use real Hermite density as in the core norm argument. This gives a joint real polynomial and compact-finite source vector at which the real norm integral is nonzero. The actual finite vector already fixed in (F) and (H) can be used with every real polynomial. The same discrete-series gap and doubled height estimate give holomorphy of this real integral on $\operatorname{Re} s > 0$.

Off $S \cup Q$, start from the unramified bottom label B and form the usual theta transfer at each adjacent step: multiply old entries by χ_S/χ_T and add the fixed length-t string centered at χ_S. These transfers are injective on labels, by cancellation of the added multiset and reversal of the common twist. At step i, after twisting the label by χ_T^{-1} in the doubling numerator, the old strings are centered at $\mu_{\boldsymbol{\mu},h}/\mu_{\boldsymbol{\mu},i}$ for $h < i$. This follows from the alternation of the two parity shifts: the product of the intermediate ratios is one or $\eta^{\pm 1}$ and cancels the endpoint shift of the string. The initial rank-r label is twisted by $\mu_{\boldsymbol{\mu},i}^{-1}$, with unitary values. Thus the unramified numerator at the critical point is the product of the initial Euler function at $(t+1)/2$ and

$$\prod_{h<i} \prod_{b'=0}^{t-1} L_E^{S\cup Q}(t - b', \mu_{\boldsymbol{\mu},h}/\mu_{\boldsymbol{\mu},i}). \tag{K}$$

The initial Euler product converges absolutely and is nonzero by the elementary spherical bound $q^{r/2}$ and $t \geq r + 5$. The factors of (K) at one are nonzero by Hecke nonvanishing, since the quotient characters have distinct real exponents. The other factors are in absolute convergence. Omission of the finite sets preserves these properties. At a tuple prime the local transfer is also injective, either by the multiset description of the split local correspondence or by canceling the added suffix. The archimedean types are fixed throughout. Cancellation of the parity factors · Continuation and the range of arguments.

The target Witt index is at most its negative real index, hence less than t. The normalized convergent-range norm formula therefore applies to these pairs. With the chosen real witness, its semilocal factor is a nonzero common scalar times the norm square of a lower multiplicity vector of one transferred label. At S this is the character cancellation on the whole prescribed finite isotype. At Q it follows from the one-dimensional line and (I). Away from these places average between spherical vectors and use that every spherical operator acts by its scalar on a common label. This verifies the product for arbitrary linear combinations within the occurrence space, without requiring the automorphic vector itself to be a pure tensor. Together with (K) it proves nonvanishing for every nonzero lower multiplicity vector of a transferred label.

To pass from the tested functions to the adjoint on multiplicities, write the real theta contraction as one fixed Howe map on each placed pair. Its values on varying polynomial and compact-finite upper vector span the smaller compact-finite real module: their span is nonzero and stable under its compact group and Lie algebra, since differentiating acts by quadratic oscillator operators on the polynomials. Then Schur's lemma factors the test integral as this real contraction and one map D_j on multiplicities. With the measures (G), Fubini identifies the upward test as the Petersson adjoint of this product. A nonzero tested lift on any lower multiplicity vector therefore implies its image under D_j^* is nonzero. The normalized norm formula gives a positive scalar norm square on each transferred label for the fixed real witness. Dividing out the norm square of its real image and polarizing shows that $D_j D_j^*$ is a positive scalar on the corresponding lower label. The real witness can depend on the arrow and label.

It remains necessary here to track the finite cuts. An upward adjoint sends the full lower level character to the actual upper level character by the covariance just computed. At S it also intertwines the common residue-group idempotent. For spherical Hecke projections, its covariance is the Satake transfer just computed, so it sends an eigenspace of label B into the eigenspace of the transferred label. These two families of projectors commute because the spherical operators are at disjoint places from the residue-group operators. The factor calculation at Q applies on the actual line of the corresponding transported representation. In particular a cut formed by the good central projector, the level-character projector and the successive spectral labels obtained by transfer contains each of these upward images. The adjoint for the cut is the same map on these inputs. A preimage of a vector with lower label is its adjoint image divided by the positive scalar above. Adding the finitely many labels gives surjectivity on the cut. Iterating yields the common bottom image for the normalized downward maps. Since the bottom group and its level and cut are unchanged as the tuple varies, this image is the same at every tuple. For the full base image take $Q = \emptyset$ and retain every permitted bottom label. The same induction applies before the good projector.

For the analytic row descent write $D_{0,j} = D_0 \cdots D_{j-1}$. Its kernel complement in the stage cut is the image of its adjoint, and the inverse there is

$$D_{0,j}^*(D_{0,j}D_{0,j}^*)^{-1}.$$

The middle operator acts on each selected bottom label by the product of the positive stage scalars. This gives the analytic rank and explicit preimages on the row chain. The full base chain.

2.6.3 The fitted auxiliary theta matrix at the base level

For each $0 \le i < m$ and each selected tuple $\boldsymbol{\mu}$, consider the auxiliary pair V_i, W_i, with V_i definite of rank $r + t$ and W_i of rank r. The record stage is $j = i + 1$. Write $T_{i,\boldsymbol{\mu}}^{\mathrm{raw}}$ for theta integration from the full definite class space at the tuple level to the row multiplicity space, and $T_{i,\boldsymbol{\mu}}$ for its good projection. Both use the phase-one normalization. The base level in this comparison is (P_0, ψ_0). Let ι_K be the prescribed complex coefficient embedding. The arithmetic realization used for taking adjoints and solving the relative matrix is at the prescribed infinite place v_0. If its representative is conjugate to ι_K, the complex source and the complex target, with their pairings and both group actions, are simultaneously conjugated. The phase-one condition makes the arithmetic lift in this realization exactly $T_{i,\boldsymbol{\mu}}$ with this scalar convention. The distinguished-place realization and its phase · The two coefficient readings · Transport back to the complex model.

The class evaluations in the two constructions. The definite source level $J_{i,\boldsymbol{\mu}}$ is the principal flag level for the tuple, including its character on the other blocks. Its compact real factor has volume one, and the finite Haar measure has volume one on $J_{i,\boldsymbol{\mu}}$. At a refined finite level $J' \subset J_{i,\boldsymbol{\mu}}$ let h_c be the prescribed double-coset representatives, with rational stabilizers $\Gamma_c = U(V_i)(F) \cap h_c J' h_c^{-1}$, where the finite-level condition is on the finite component. Lift the value of a class form f at h_c to the coefficient polynomial

$x_c(f)$ using the given semilinear coefficient map $c_{i,\boldsymbol{\mu}}$ in the placed frame. On a pure lowest-type tensor u and at every row point g, the evaluation is

$$T^{\mathrm{raw}}_{i,\boldsymbol{\mu}}(f)(u)(g) = \sum_c \frac{[J_{i,\boldsymbol{\mu}} : J']^{-1}}{|\Gamma_c|}\, \vartheta_{i,\boldsymbol{\mu}}\big(h_c, g;\ p_{i,\boldsymbol{\mu}}(u, c_{i,\boldsymbol{\mu}}(x_c(f)))\big). \tag{38}$$

Here $p_{i,\boldsymbol{\mu}}$ is the harmonic polynomial map followed by its tensor with the specified finite oscillator function. The quotient measure on a double coset contributes the volume of J' divided by the stabilizer order. At product levels this volume is the product $\prod_v [J_{i,\boldsymbol{\mu},v} : J'_v]^{-1}$. In the integrand the two level characters on J' are inverse, and real equivariance pairs the coefficient with the same harmonic vector at every point of the compact real factor. This proves the class formula. In a coefficient extension used for the complex class sum the polynomial $x_c(f)$ is extended with ι_K; the rational mass factor is unchanged. The sum uses the actual $(i, \boldsymbol{\mu})$ finite test and the same g on the two sides. Definite Haar convention · Refined level, placed coefficients and stabilizers · The lowest-type integral on the actual classes.

The finite actions can be compared before any choice of fitted basis. On the definite source at a residue-group element s use its prescribed lift k_V and right translation normalized by $\chi_{V,J}(k_V)^{-1}$, where $\chi_{V,J}$ denotes the extended flag character. At the row lift k_W the principal flag character is trivial. At a type place the common lifted block has the same residue s on each side. In the rectangular finite test, $\bar{k}_V^{-1}\left(\begin{smallmatrix} I \\ 0 \end{smallmatrix}\right) \bar{k}_W = \left(\begin{smallmatrix} I \\ 0 \end{smallmatrix}\right)$. The remaining determinant character of the finite oscillator action is $\chi_{V,J}(k_V)^{-1}$ on the simultaneous lift. The component lifts may be taken trivial at $Q(\boldsymbol{\mu})$. Right translation in the definite integral now gives equivariance with the row right-translation action. The quotient Haar measure is right invariant. The same equality on the lowest type extends to the whole real multiplicity map. With this source action the input at level (P_0, ψ_0) has residue character ψ_0^{-1}, since the source flag character divided by the enlarged level character is $\iota_K(\psi_0)$. On the row the level character is $\iota_K(\psi_0^{-1})$. Thus the projectors on the two sides are the transports of

$$e_X = \sum_s e_X(s)s, \qquad p_0 = \frac{1}{|P_0|}\sum_{s\in P_0} \psi_0(s)s$$

under their respective normalized actions. Here e_X is the good central selector of the chosen finite types. These are unitary group actions in the class and Petersson metrics. The projectors are self-adjoint and commute, so their product is an orthogonal projection. Equivariance intertwines each projector and its adjoint. The selector e_X acts at the residue group and the $c(\sigma)$ used to label a global spherical eigenspace acts through the Hecke operators. Change of variable and the inverse flag character · The inverse character on the row level · Compatibility with the good summands · The character in the auxiliary class projector.

How the record is constructed. Let $B_{i,\boldsymbol{\mu}}$ send coordinates to the class form in the prescribed full basis. For a level coordinate vector x the form used in the record is $f_x = B_{i,\boldsymbol{\mu}} E_J x$, where E_J is the prescribed inclusion of the level coordinates. The source identification J_{M_t} is obtained by sending the standard vectors of the record to these class forms and extending complex linearly. Its source Gram is the positive definite class Gram pulled back by E_J. In principal measures the class pairing of these vectors is the principal definite pairing. At the character level the source record pairing is this pairing divided by $|P_0|$. The division rescales the Gram and does not change the measure $\nu_{V_i,\boldsymbol{\mu}}$ in (38). On the row the Petersson pairing rescales by the same $|P_0|^{-1}$, with the fixed product of the lowest compact-type dimensions included in the multiplicity pairing at both levels.

The record's target space M_j is built by a Hilbert identification $\mathcal{E}_i$ to the good character-level row space, with $\mathcal{E}_i^* \mathcal{E}_i = |P_0| I$ for the principal row metric. The orthonormal coordinates used to construct M_j determine $\mathcal{E}_i$ and the lattice and basis pulled back along it together. Its lattice is the inverse image of the prescribed good integral forms. The auxiliary $\Theta_{P_0,\boldsymbol{\mu},j}$ is obtained as follows. Evaluate (38) on $J_{M_t}(\cdot)$, extend from the lowest type to the cuspidal row multiplicity space, apply e_X there and read the result via

$\mathcal{E}_i^{-1}$. Indeed the actual theta integral supplies an extension with the required row character by the finite covariance above. A row multiplicity map is determined by its value on the lowest type since the tensor of these types generates the product Harish-Chandra module and the multiplicity maps are continuous on its smooth globalization. Consequently the linear extension used in constructing M_j agrees pointwise with the same theta integral, with the same J', class coefficients and source Haar measure. In particular, for $I_{i,\boldsymbol{\mu}} = j_{M_t} \circ e_{M_t}$, with e_{M_t} the fixed identification of the level coordinates with the record's K-space $M_{t,K}$,

$$\mathcal{E}_i \Theta_{P_0,\boldsymbol{\mu},j}(I_{i,\boldsymbol{\mu}} x) = T_{i,\boldsymbol{\mu}}(f_x). \tag{39}$$

In the final record the source and target identifications are the same maps onto their images, and the construction inserts this very $\Theta_{P_0,\boldsymbol{\mu},j}$ as its auxiliary theta map. The two Hilbert pairings and the unscaled integration Haar · Source standard vectors and their class pairing · Construction by the lowest-type extension and the good projection · Insertion of the lift spaces, lattice and maps · The original identifications viewed as maps onto their images · Assembly of the auxiliary spaces and theta map.

The coefficient field of the auxiliary lift. Fix an auxiliary index $i = j - 1$ and the prescribed pair of ranks $r + t, r$, with $t \geq r + 5$. Use the same definite class space, finite test and source measure as in (38). Write $A_{i,\boldsymbol{\mu},K}$ for the full rational class space of the auxiliary definite group at its prescribed tuple level. For x in this space, denote the unprojected lift evaluated at the chosen unit lowest-type vector by $U_{\boldsymbol{\mu}}(x)$. We show that the multiplicity vector of the projected lift lies in the prescribed K-space $F_{i,K}$ of auxiliary cusp forms for every such x and every selected tuple. This includes the inputs before forming the complement to the kernel.

The base PEL marking has integer modulus $M = g_0 g_1$ and auxiliary marking parameter $A = 1$. Write $Q_{\boldsymbol{\mu}}$ for the actual marked twist places and put

$$p_{\boldsymbol{\mu}} = \prod_{q \in \{\#k(P)\,:\,P \in Q_{\boldsymbol{\mu}}\}} q, \qquad M_{\boldsymbol{\mu}} = M p_{\boldsymbol{\mu}}.$$

The residue degrees are one. Take the half ideal $I_{\boldsymbol{\mu}}$ from the product of the distinct chosen primes below $Q_{\boldsymbol{\mu}}$ in the host imaginary quadratic field and extend it to E. The one-sidedness condition at each rational prime gives $I_{\boldsymbol{\mu}} \bar{I}_{\boldsymbol{\mu}} = (p_{\boldsymbol{\mu}})$. The marked ideal on the cover is the base marked ideal times $I_{\boldsymbol{\mu}}$, and the cover still has $A = 1$. The field selection adjoins the finite CM Weber values at these moduli before choosing the tuple. It suffices to include the moduli formed from the subsets of the union of ramified places of the finite allowed twist set. Together with the selected roots of unity and character values, these choices give the marked models and their component fields inside the prescribed coefficient field. On each component the determinant character of the cover is well defined on the dictionary points because it is trivial on determinants of rational group elements and on the principal congruence subgroup at $M_{\boldsymbol{\mu}}$. Its finite-order values give the algebraic component function $c_{\delta,\boldsymbol{\mu}}$ over K. Its reading at a marked frame above q is $\xi_{\boldsymbol{\mu}}(\det_{\mathrm{root}} q)$, with the rational translate from the dictionary understood. The multiplicative constant in this reading is one. The half ideal and the exact cover modulus · Finite field selection before the tuple · Component fields at the selected moduli · The algebraic determinant function with constant one.

Here is the class section on this cover. Read x at the placed class representatives in its given coefficient frame and apply the fixed algebraic harmonic map to the resulting values. For each class c, use the split-supported lattice Λ_c, its allowed odd denominator N_c prime to $M_{\boldsymbol{\mu}}$, and the weight $w_c = \widehat{\phi}_{\boldsymbol{\mu}}$ on the actual marked quotient. The function $\phi_{\boldsymbol{\mu}}$ is the product of the prescribed type test and the finite twist test $\omega^{\mathcal{Q}}$ with inverse determinant character. Its values and its finite Fourier coefficients belong to K. At a framed point on the cover, the class jet is the sum of harmonic derivatives of the pulled-back torsion translates as in (45), with the enlarged marking in the Fourier sum. Its denominator is $m_c N_c^{d_i} |\mathcal{Q}_{\boldsymbol{\mu}}|$, where $\mathcal{Q}_{\boldsymbol{\mu}}$ is the finite space of marked labels for all source columns. The Haar class weight is $v'_{\boldsymbol{\mu}}/|\Gamma_c|$, and the algebraic class scalar is the one attached to the same split-supported lattice. Multiply the weighted class sum by

$c_{\delta,\boldsymbol{\mu}}$. The change of a lift of the marking scales the character function and the finite Fourier trace by inverse factors. The jet transforms by the prescribed Hodge weight. Thus the product descends through the marked cover to a section $s_{\boldsymbol{\mu}}(x)$ over K. The harmonic contractions remove the change of polarization gauge, and the last positive-block weight kills the constant term on each isotropic line as in the primitive calculation below. Hence this is a cusp section of the fixed base row model. Finite test in the actual class frame · The normalized short jet · The class weights and determinant function · Descent of the product on framed charts · The character and cuspidal class section for each tuple.

The normalized reading of this section must use the component factors already installed in $F_{i,K}$. In that installation the empty-tuple cover has $p_{\boldsymbol{\mu}_0} = 1$, $I_{\boldsymbol{\mu}_0} = \mathcal{O}_E$ and $c_{\delta,\boldsymbol{\mu}_0} = 1$. The construction initially uses an auxiliary finite extension to put the reference datum and all selected marked data in one model. It multiplies every other cover function of reading $a'\xi_{\boldsymbol{\mu}}$ by a/a', where $a = 1$ is the reference constant. Denote its cusp section for x by $s^{\mathrm{g}}_{\boldsymbol{\mu}}(x)$. The installed data keep this normalization after passage to the fixed space of algebraic forms. They read, at every adelic quotient point q,

$$U_{\boldsymbol{\mu}}(x)(q) = L_i c_i^{\mathrm{fr}}(q)\ \mathrm{rd}_q\ s^{\mathrm{g}}_{\boldsymbol{\mu}}(x), \qquad L_i = (2\pi\xi_R)^{-d_i/2}.$$

Here rd_q uses the intrinsic Hodge frame from the common adelic dictionary and includes its prescribed determinant factor. The scalar $c_i^{\mathrm{fr}}(q)$ is the installed phase-one component factor. The level move and root outside that factor are both one. These choices are part of the construction of the geometric half and its fixed cusp basis. The identity reference cover · The unit reference constant · The common cover constant in the merged data · Installation of the component constants and the algebraic basis · The identity move and root for every tuple.

To compare with the section over K, extend the same source coefficient isometry to each geometric coefficient field. At each class representative its value is that of x in the fixed class embedding. Pull the residual level action back from the same reference marking. These choices leave the finite test, adelic point and lowest-type vector in $U_{\boldsymbol{\mu}}(x)$ fixed. Each decomposition uses the prescribed source measure and its class weights $v'_{\boldsymbol{\mu}}/|\Gamma_c|$ for its representatives. The class calculation then gives for both geometric readings s of this input

$$U_{\boldsymbol{\mu}}(x)(q) = \lambda_i(q)\,\mathrm{rd}_q\ s, \quad \lambda_i(q) = k_i(-\delta_i)^{d_i} L_i c_i^{\mathrm{vac}}(q) \neq 0.$$

The finite Fourier expansion gives the common sign $\delta_i \in \{1,-1\}$. Multiplication by N_c in the harmonic derivative cancels $N_c^{-d_i}$. The even representatives contribute z_c/m_c, and the class scalar supplies the exact finite lattice index and the inverse determinant factor of the moved source vacuum. This leaves the displayed harmonic normalization. The unit harmonic scalar k_i is the same for the two fields because both coefficient maps send the same extremal tensor to the same nonzero harmonic polynomial before the Fock comparison. The vacuum factor $c_i^{\mathrm{vac}}(q)$ uses the base component of the same dictionary point, with the marked half and Hermitian frame transported by the level map. The twist of the analytic vacuum by $\xi_{\boldsymbol{\mu}}$ matches the value of $c_{\delta,\boldsymbol{\mu}}$ on that frame. Both cover constants are one. Cancelling $\lambda_i(q)$ for each q proves

$$\mathrm{rd}_q\ s^{\mathrm{g}}_{\boldsymbol{\mu}}(x) = \mathrm{rd}_q\ s_{\boldsymbol{\mu}}(x).$$

This equality includes zero lifts. The analytic vacuum and harmonic factors are used in the comparison inside the prescribed reading of a cusp section. Retie using the same coefficient isometry and level action · The class index and character · The harmonic and positive normalization factors · Matching the harmonic map on the underlying class coefficient · The determinant at the same marked frame · The same nonzero multiplier at every quotient point.

The algebraic cusp sections of the base model commute with field extension. On each affine chart of the reduced framed scheme the closed points have finite residue fields over the ground field. Their maximal ideals intersect in the nilradical, so a regular function vanishing at all algebraic geometric points is zero. The same reasoning over an algebraic closure, followed by scalar extension, shows that finitely many such framed evaluations separate the complexification of the finite-dimensional cusp space. This

also proves uniqueness in the class section comparison. Evaluation coordinates at those points may use their residue fields. The coefficients of $s_{\mu}(x)$ are in K by descent and scalar extension on the fixed cusp space. Apply the installed reading to a base-cusp basis expansion of this section. Its value is precisely $L_i c_i^{\mathrm{fr}}(q)\,\mathrm{rd}_q\, s_{\mu}(x)$ at the designated lowest-type vector. Equality on every quotient representative extends to the whole automorphic function by left rational invariance. The lowest-type evaluation is injective on the multiplicity space by the positive multiplicity norm on the unit tensor. Thus the unprojected lift is the image of a vector in $F_{i,K}$. Separation by algebraic geometric points · Independence of evaluations · Reading the fixed cusp basis · Injectivity at the unit lowest-type vector.

If $e_X = \sum_{g\in G} e_g[g]$ is the good projector, its operation in the coefficient model is $\sum_{g\in G} e_g \chi_i(g)^{-1}\rho_i(g)$. Here ρ_i is the action on the fixed algebraic cusp basis and χ_i is its finite determinant multiplier. The inverse scalar cancels the scalar of the finite level action in the adelic reading. The character and all the projector coefficients are in K. Hence the same conclusion holds for the actual projected lift, including the scalar conjugation at the fixed place when the complex model uses the conjugate form. For either face b, the target vectors $w_{0,i,b}(e_k)$ read the prescribed K-basis of the good character subspace of $F_{i,K}$. The factor $\sqrt{|P_b|}$ in its common-space inclusion cancels the inverse factor in its Hilbert character-level basis. The coordinate functionals of the fixed basis of $F_{i,K}$ are now the separating tests. They give K-values on both the full source lifts and these fixed target vectors. Projecting the unprojected algebraic vector · Reading the fixed target frame and cancelling the index · The full coefficient coordinates and the two required families of values.

The fitted projection and the arithmetic columns. The coordinate functionals on the fixed rational cusp space give a K-matrix A of separating tests of $T_{i,\mu}$ on the whole definite class space. Let G be the K-Hermitian Gram of the full class space, up to an overall positive real factor. Then extension of scalars identifies $\ker_K A$ with the kernel of $T_{i,\mu}$ and identifies $(\ker_K A)^{\perp_G}$ with $\operatorname{im} T_{i,\mu}^{*}$. In fact, the first assertion follows from row reduction and separation, and the second follows by applying the same extension to the linear equations for the orthogonal complement. This argument uses the source Gram.

Intersect the adjoint image with the good and level-character images. This is the specified Hermitian complement, and we use the chosen record basis for its rational form. Since the orthogonal projectors commute with $T^{*}T$, they preserve its kernel and complement, and the intersection is the complexification of that same rational intersection. In the fixed class coordinates the basis has matrix VZ. Write $\Pi_Z = P_Z V^{\flat} E_J$. Orthogonal projection of $E_J x$ to the good level complement is $VZ\Pi_Z x$. The difference is killed by $T_{i,\mu}$: both vectors are in the character level, their difference pairs to zero with the good component of the adjoint image, and the good projection already fixes the image of T. Therefore

$$T_{i,\mu}(B_{i,\mu}E_J x) = T_{i,\mu}(B_{i,\mu}VZ\Pi_Z x).$$

This identifies the complement in the very class model used to construct M_t.

The W_i coordinate basis is the prescribed basis $b_{i,k}$ of the inverse-character good part of the rational cusp forms. Denote its complex realization by $J_{W,i}$. In the arithmetic definition of the relative matrix its columns are first realized with the factor $|P_0|^{-1/2}$ used for the smaller-level pairing. The embedding into the principal multiplicity space multiplies by $|P_0|^{1/2}$, so its k-th column is precisely $J_{W,i}b_{i,k}$. In the record $j_{W,i}$ is the inverse image of this very basis under $\mathcal{E}_i$. Thus their evaluations at every point agree, with no relative phase from level conversion. The definition of $U_{i,\mu}$ solves for the columns of $T_{i,\mu}B_{i,\mu}VZ$ in this basis using the prescribed embedding. Its rationality uses the absolute field normalization of the tests at the same cusp forms. In particular, the matrix is fixed by

$$J_{W,i}\left(\sum_k (U_{i,\mu})_{kh} b_{i,k}\right) = T_{i,\mu}(B_{i,\mu}VZe_h)$$

for every fitted basis index h. Applying this equality with coefficients $\Pi_Z e_a$ gives, for every standard level basis vector e_a,

$$\Theta_{P_0,\boldsymbol{\mu},j}(I_{i,\boldsymbol{\mu}} e_a) = j_{W,i}\big(U_{i,\boldsymbol{\mu}}(P_Z V^{\flat} E_J) e_a\big)\,. \tag{40}$$

Indeed both sides have the same image under the injective $\mathcal{E}_i$ by (39). Semilinearity with the same ι_K then gives the identity for every level coordinate vector. This argument keeps the prescribed integral lattices and all coefficient readings in place. Definition of the rational adjoint-image coordinates · The orthogonal complement killed by the same lift · Definition of the relative lift matrix · The prescribed good basis and its norm rescaling · The identity on the actual level coordinates.

The positive adjoint norm for this auxiliary pair in §2.6.5 gives surjectivity onto its good principal row space for the selected tuple. Averaging a preimage by $e_X p_0$ therefore gives a preimage at the good character level. The restriction of $T_{i,\boldsymbol{\mu}}$ to the fitted complement is injective and has this good image. It follows that $U_{i,\boldsymbol{\mu}}$ is square and invertible, and $E_J^{\flat} V Z U_{i,\boldsymbol{\mu}}^{-1}$ gives the explicit small-level preimages of the prescribed row coordinates. In particular this is the projection attached to M_j and $j_{W,i}$. On the base level, the separate comparison of finite types in (172) and (173) identifies the dimension with D_0, the rank of the bottom part of the row chain on that level. The coordinate right inverse · The type-block partition · The rank on the recorded blocks.

2.6.4 The common downward coordinates

Return to the descent from X' through the selected Y_j to Y_0. Section 2.6.2 gives the common image on the analytic cuts for these normalized maps. To identify the arithmetic kernel complement in the original class model, use the principal-level test identities, separation, and the rationality of $V_{\boldsymbol{\mu},a}$ and the square ratios. The base row space has positive dimension in the application. Surjectivity there and separation give a nonzero left side in (16) for a rational input, since the rational inputs span. Thus $\kappa \neq 0$. Multiply each ψ_a by $\iota_K(e_{\boldsymbol{\mu}_0}) c^{\square}_{\boldsymbol{\mu}_0}/\kappa$. The resulting fixed linear functionals still separate. Their values on every rational downward image are over K, since by (16) they are

$$\frac{\iota_K(e_{\boldsymbol{\mu}_0}) c^{\square}_{\boldsymbol{\mu}_0}}{\iota_K(e_{\boldsymbol{\mu}}) c^{\square}_{\boldsymbol{\mu}}}\, V_{\boldsymbol{\mu},a}(x).$$

This uses the square ratios before constructing a common field model of the row space.

For each tuple the kernel in the complexified top space is the kernel of the matrix of these tests. All its entries are over K, so this kernel is the complexification of its rational kernel $U_{\boldsymbol{\mu},K}$. The rational adjoint image is the orthogonal complement of $U_{\boldsymbol{\mu},K}$ for the class Gram matrix. Indeed the Gram is over K with the involution induced by conjugation at ι_K, and the equations for orthogonality extend to the complex orthogonal complement. Include the central good projector and impose either of the two prescribed level characters. In the principal model write the latter projector as

$$e_{b,\boldsymbol{\mu}} = \frac{1}{|P_b|} \sum_{g \in P_b} \psi_b(g)^{-1} \rho_{\boldsymbol{\mu}}(g), \qquad b \in \{0, s\},$$

where $\rho_{\boldsymbol{\mu}}$ is the determinant-corrected finite right action. The two actions and their covariance were computed in the level-change calculation. Their complex bottom averages intertwine the same downward maps. In particular their image on the common principal row space equals the image of each corresponding top cut. These averages on the bottom preserve its K-form: for a rational base preimage, multiplication by an element of the residue group can be performed on that preimage before descent. Unitarity gives the analogous action on the adjoints. The adjoint complement on the cut is thus the intersection of the principal adjoint complement with the good character space. Its Gram in the character-level normalization is a positive common multiple of the restricted principal Gram, so its orthogonal complement is unchanged. This applies both to the base level used for the auxiliary stage matrices and to the level of the periods indexed by words.

In the good coordinates of the level under consideration, let H_μ be the class Gram and Z_μ the prescribed K-basis matrix of the adjoint complement. The orthogonal projection to the complementary coordinates is

$$P_\mu = (Z_\mu^* H_\mu Z_\mu)^{-1} Z_\mu^* H_\mu.$$

The indicated Gram matrix is invertible by positivity at the chosen embedding. Thus $P_\mu Z_\mu = I$ and the complex downward map factors through these coordinates. Use the images of the base complementary basis as the basis of E_{row}, and write R_K for their K-span in these coordinates. The matrix of the fixed, normalized tests on this base basis has full column rank by separation, so select a square invertible minor over K. For the downward image of any other complementary basis vector, the coordinates in the base basis solve the linear system for this minor with right side in K. Hence their transition D_μ is over K, and is invertible by the equal-image and kernel-complement properties. Let j_{row} be the base complex realization and j_μ the realization of the projected top coordinates. We have proved

$$\iota_\mu(j_\mu y) = j_{\mathrm{row}}(D_\mu P_\mu y). \tag{41}$$

For each cut the base realization is the image of its prescribed complementary basis. Thus the transition in (41) is the unique relative matrix between these realized bases. This identifies the record on its own coordinates, including at the sharp level, even if the two levels were given different bases. Here and below j_{row} and j_μ on rational inputs are semilinear for ι_K. Tests on all rational inputs · Finite equivariance, the principal row and its two level cuts and the relative matrices in the prescribed bases.

For the metric at ι_K, use the bottom Gram matrix pulled back along j_{row}. If the basis is realized first on the principal level, divide its Gram by $|P|$ to obtain the metric at the character level. Indeed the Petersson inner product on the larger level uses volume one there, whereas its volume in the principal normalization is $|P|$. The two realizations can use conjugate representatives of the infinite place, which conjugates the pairing and preserves its norm. On a summand with $\iota_\mu \iota_\mu^* = \lambda I$, the downward map multiplies squared norms by λ on the kernel complement. Indeed a preimage there is $\lambda^{-1}\iota_\mu^* e$, whose squared norm is $\lambda^{-1}\|e\|^2$. The same argument gives the cross-pairing. If a fitted vector y_p has image in that summand, then $\iota_\mu^* \iota_\mu y_p = \lambda y_p$. Pairing this identity against any fitted vector z gives

$$h_{v_0}(D_\mu y_p, D_\mu z) = \lambda\, h^Q_{\mu,v_0}(y_p, z)$$

in the identified metrics at the chosen place v_0. Here $\lambda = \kappa_{v_0} \prod_j \lambda_{j,p,v_0}$ once the normalized step factors have been established.

At any other infinite place v, use the fixed base complement metric h_v^Q on the rational coordinates, with the prescribed positive scalars $\kappa_v, \lambda_{j,p,v}$. The underlying class Hecke operators are the commuting double-coset operators with their adjoints for the class mass. Integration of a right translate and change of variables give the adjoint by inverting the double coset and conjugating its character. The class-sum version is the same calculation on each source and target class. The adjoint matrix identities transport to v because the operator and level matrices have entries in the CM subfield, on which the field involution maps to complex conjugation at every embedding. The fitted class subspace is invariant under this family, including its adjoints. Thus its orthogonal projection commutes with the operators, and the induced operators on fitted quotient coordinates are commuting normal operators. Different joint spectral labels have disjoint joint eigenspaces, which are orthogonal for the quotient form h_v^Q. For a joint spectral label p let E_p be the projector to that summand. The prescribed row metric is

$$h_v(y, z) = \sum_p \kappa_v \left(\prod_j \lambda_{j,p,v} \right) h_v^Q(E_p y, E_p z).$$

Each term is positive semidefinite and the projections sum to the identity, so this gives a positive metric. Moreover for y_p in one summand and any z, it gives $h_v(y_p, z) = \kappa_v(\prod_j \lambda_{j,p,v}) h_v^Q(y_p, z)$. In particular

h_v dominates h_v^Q times the smallest of the positive factors, $h_v(y,y) \geq \kappa_v \min_p(\prod_j \lambda_{j,p,v}) h_v^Q(y,y)$. These identities give the dual norms and comparisons used in Section 4 once the selected theta norm scalars on the actual summands are determined. Principal-level and character-level metrics · Metrics at the other embeddings · Transport of adjoints and spectral orthogonality.

2.6.5 Petersson eigenvalues of the auxiliary row stages

Fix an auxiliary fitted row stage and a selected tuple $\boldsymbol{\mu}$. Its pair consists of the rank-r row space W and the definite space V of rank $n = r + t$, with $t \geq r + 5$. These ranks are fixed over the stages. Write $f = [F : \mathbb{Q}]$. First choose a coefficient embedding over the host imaginary quadratic field. The relative twist is finite order and trivial at infinity, so for this embedding its real test and real type are those of the base tuple. The positive norm calculation below applies to every selected tuple at the principal level. The quantitative bound needed for $\lambda_{j,p,v}$ concerns the base tuple. Fixed auxiliary row rank · Fixed auxiliary definite rank. Let T be the full downward lift from the source class space to the cuspidal row multiplicity space before the principal row projection. Use the Hermitian class-sum metric on the source. On the row use the Petersson multiplicity metric with finite principal-level volume one and real measures dg_ρ as specified below. For an auxiliary subgroup H_P of the row residue group, the character-level adjoint in these measures is $T^{*\mathrm{lev}} = |H_P|^{-1} T^*$ when the source retains its principal normalization. At principal level $H_P = \{1\}$. In (39) both pairings are divided by $|P_0|$, so their simultaneous rescaling leaves the adjoint of the restricted map unchanged. The factors $|P|$ in the passage from the later principal row metric to the character metric are handled by the prescribed simultaneous norm rescaling. The level adjoint · The multiplicity metric · The row metric conversion.

For a pure unit lowest-type vector $e = \otimes_\rho e_\rho$ and an orthonormal basis $(b_k)_k$ of the definite coefficient space C_V, let φ_k be the actual oscillator test associated to $e \otimes b_k$ by the joint-harmonic map. Its finite factor f_{fin} is the prescribed level test and its real factor is the polynomial applied to the prescribed Cartan vacuum. The global oscillator measure has rational-lattice covolume one. For y in the multiplicity space write $u_y = y(e)$ for lowest-type evaluation, and $\Theta^\uparrow(\varphi_k, u_y)$ for the scalar reverse lift to $U(V)$ with kernel $\theta(g,h;\varphi_k)$ paired against $\overline{u_y(g)}$. Put $d_K = \prod_\rho d_{K,\rho}$, where $d_{K,\rho}$ is the dimension of the lowest compact type at ρ, and put $d_V = \dim C_V$. Lowest-type evaluation satisfies $\langle y, y' \rangle = d_K \langle u_y, u_{y'} \rangle_W$ in compatible Hermitian conventions. Taking the adjoint inside the defining class sum therefore first multiplies it by d_K. Each source term is weighted by the inverse of the class stabilizer order. Expanding the coefficient pairing in $(b_k)_k$ gives

$$\langle T^* y, T^* y' \rangle = d_K^2 \sum_k \langle \Theta^\uparrow(\varphi_k, u_y), \Theta^\uparrow(\varphi_k, u_{y'}) \rangle_V .$$

Here the definite Haar measure gives the source level and the real compact factors total volume one, so integrating the coefficient pairing gives exactly the above class sum. Definite Haar measure · Evaluation and class-sum adjoint · The tested adjoint form.

Apply the convergent norm formula to these reverse lifts, with doubling source $W \oplus (-W)$ and target V. The forms u_y are cuspidal and K-finite, the actual φ_k are Schwartz with polynomial Fock part, V has Witt index zero, and $n > 2r$. With $s_0 = t/2 > r/2$, the ordinary Siegel Eisenstein series converges absolutely and equals the doubled theta integral in quotient-volume-one measures by the convergent Siegel–Weil formula. The ratio $\tau(U(V))/[E : F] = 1$ converts the formula of Gan–Qiu–Takeda to these measures. The tensor doubling map evaluates at zero by the oscillator scalar product in the covolume-one measure. The inducing character in the release convention is χ_V^{-1}, where χ_V is the V-splitting character denoted c_μ. The inverse character on the second copy in the doubling integral thus cancels the character in the tensor identity. Cuspidality kills the lower rational orbits. The open orbit unfolds to the integral of the oscillator matrix coefficient paired with the conjugate of the row matrix coefficient. The condition

$s_0 > r/2$ gives the absolute open-orbit unfolding. Rapid decay of the row cusp forms and compactness of $[U(V)]$ justify the kernel integrations in the seesaw. Thus with the actual measures the norm formula is

$$\langle \Theta^{\uparrow}(\varphi_k, u_y), \Theta^{\uparrow}(\varphi_k, u_{y'})\rangle_V = \operatorname{vol}_{v_V}[U(V)] \int_{U(W)(\mathbb{A}_F)} \langle \omega(g)\varphi_k, \varphi_k\rangle_{\mathrm{osc}} \cdot \overline{\langle R(g)u_y, u_{y'}\rangle_W}\, d\nu_W(g), \tag{42}$$

up to simultaneously transposing the arguments when using a conjugate-linear-first convention. Changing the Haar scale on $U(W)$ multiplies both sides by the square of that scale, and the factor on $[U(V)]$ gives the exact definite Haar convention. This is the $n > 2r$ absolutely convergent application of Siegel–Weil and the doubling open-orbit calculation. [19, §§3.1, 7.1, 11.3, 11.5, Theorem 7.1] Placed doubling integral · Actual definite and row measures.

The real coefficient in (42) at ρ is $H_\rho c_\rho(g_\rho)$, where

$$c_\rho(g) = \langle \pi_\rho(g)e_\rho, e_\rho\rangle, \qquad H_\rho = \operatorname{hook}(\mathbf{a}^{+}_{j,\sigma,\rho}) \operatorname{hook}(\mathbf{a}^{-}_{j,\sigma,\rho}).$$

The highest joint harmonic is a product of leading and trailing minor powers in disjoint sets of variables. The mixed Laplacians annihilate it because the used column sets are disjoint. The corresponding lowering operators annihilate the real highest vector in the Fock model. Its compact orbit maps equivariantly from the tensor of the lowest row type and the definite coefficient type. The lowest-weight module it generates has the row discrete-series parameter: the K_ρ-type fixes the holomorphic parameter, and the lowering relations give its cyclic lowest-weight model. Equivalently, its invariant scalar product is fixed by the lowest-type pairing using the adjoint raising and lowering operators and the Poincaré–Birkhoff–Witt ordering. Schur's lemma on this tensor of compact types fixes the constant from a unit pair mapping to the specified minor product. The Fock squared norm with $\|z^\beta\|^2 = \beta!$ is the Gaussian average of the squared modulus of the polynomial. For a partition κ of length at most p, the product of leading minors is the matrix coefficient on a unit highest vector of the Schur representation S_κ of GL_p. This follows by realizing the highest vector in the tensor of exterior powers with exponents $\kappa_i - \kappa_{i+1}$. Factor the square complex Gaussian matrix as $X = UR$ with U unitary and R upper triangular with positive diagonal. Conditional on R, U is Haar. The upper triangular factor acts on the highest line by $\prod_i R_{ii}^{\kappa_i}$, and compact Schur orthogonality contributes $1/\dim S_\kappa$. Successive orthogonal projections in Gram–Schmidt give independent squared diagonal entries with gamma distributions of shapes $p - i + 1$ and scale one. Hence the squared norm is

$$\frac{1}{\dim S_\kappa} \prod_{i=1}^{p} \frac{(\kappa_i + p - i)!}{(p-i)!} = \prod_{b\in\kappa} h(b)$$

by the Weyl dimension and hook product formulas. The trailing minors follow by coordinate relabeling. Disjoint variables give the product of the two hook products and each b_k has norm one. The other matrix coefficient in the integrand has the opposite conjugation. The resulting real factor is $H_\rho J_\rho$, where $J_\rho = \int_{U(W)(F_\rho)} |c_\rho(g)|^2\, dg_\rho$. Lowest-type Fock coefficient · Fock norm of the prescribed minors.

Let S be the finite places below the type and twist primes, and q_v the cardinality of the residue field of F_v. Each place of S is split, with exactly one chosen prime of E above it. On the chosen half the support of the test is integral, with reduction $\left(\begin{smallmatrix} A \\ 0 \end{smallmatrix}\right)$ in the n-by-r matrix coordinates, where A is in the specified subgroup $H_v \subseteq \mathrm{GL}_r(k_v)$ at a type place and in $\mathrm{GL}_r(k_v)$ at a twist place. The other half uses the conjugate-dual lattice Fourier model. Here k_v is the residue field and the multipliers on the support are unitary level characters. If the support and its translate meet, take a full-rank top block A there. Integrality of both A and Ag, with unit determinant reductions, gives $g \in \mathrm{GL}_r(\mathcal{O}_v)$. At type places A and Ag in H_v modulo the uniformizer give $\tilde{g} \in H_v$. The split dual frame gives the inverse conjugate at the other half, so the same conclusion holds in the unitary group. Thus the finite oscillator coefficient at v is supported on the actual row level J_v, there transforming by the same unitary character as the row functions. In the convolution the two characters cancel. Integration over J_v contributes the finite squared norm and the index of the

principal subgroup in J_v. The global index is $|H_P|$. At a tuple prime $v \mid Q(\boldsymbol{\mu})$ the determinant weight in the prescribed test is the inverse of the relative splitting character on row units. They cancel in the row action, so J_v is the full hyperspecial stabilizer and its character is trivial. The relative character has modulus one, so the squared norm of the test still counts every invertible top block with weight one. The finite theta test · Support intersection and the row level · Finite convolution on the actual Hecke part.

Outside S the self-dual vacuum is spherical and the stabilizer has volume one. The extension E/F is unramified at all finite places. For the tested joint character let $\beta_{w,1}, \dots, \beta_{w,r}$ denote the normalized standard Satake parameters at $w \mid v$. In the split chart they are the GL_r parameters at one prime and their inverses at the other. In the inert chart they are $\alpha_1, \dots, \alpha_{\lfloor r/2 \rfloor}$, their inverses, and a 1 if r is odd. The unramified character of the spherical section at w is χ_V^{-1}. With $q_w = \mathrm{N}w$ and ϖ_w a uniformizer, put

$$E_S = \prod_{w \nmid S} \prod_{a=1}^{r} (1 - x_{w,a})^{-1}, \qquad x_{w,a} = \beta_{w,a} \chi_{V,w}^{-1}(\varpi_w)\, q_w^{-(t+1)/2}.$$

Let $\epsilon_v = 1$ at a split place and $\epsilon_v = -1$ at an inert place. The spherical coset sum at v is the local factor of E_S multiplied by d_v^{-1}, where

$$d_v = \prod_{a=t+1}^{n} (1 - \epsilon_v^a q_v^{-a})^{-1}, \qquad d_{\mathrm{off}} = \prod_{v \notin S} d_v, \qquad d_S = \prod_{v \in S} d_v.$$

For the split computation the two determinant coset factors have inverse Satake lists and inverse character values, and the common density correction is $\prod_{a=t+1}^{n}(1 - q_v^{-a})$. For the inert computation let $z = \chi_V^{-1}(\varpi_v)$. The section on a graph of length ℓ in the doubled lattice has value $(z q_v^{-n})^{\ell}$. The spherical sum in this length variable $X = z q_v^{-n}$ has numerator $\prod_{i=0}^{r-1}(1 - (-q_v)^i X)$ and denominator the Euler polynomial on the list above at $q_v^{r-1} X$. Indeed the sum over the integral hyperbolic-frame Iwasawa cells gives the length series, with the i-th geometric factor carrying the residue norm $(-q_v)^i$. In this calculation the trace-zero rank-one term in the radical has one residue coordinate and the paired coordinates have two. This uses the unramified involution at the ordinary inert places of the family. Since $z = (-1)^n$ by $\chi_V|_{\mathbb{A}_F^\times} = \eta_{E/F}^n$, setting $a = n - i$ identifies the numerator with d_v^{-1} and $q_v^{r-1} X = z q_w^{-(t+1)/2}$. Equivalently the normalized section at $s_0 = t/2$ has denominator the factors at $2s_0 + r - i$ with inducing characters of sign $\epsilon_v^{n+i} = \epsilon_v^{n-i}$. [19, §11.6] Split spherical coset sum · Inert length series and the local factors · Comparison with the normalized base-change parameters.

In the finite summation one may replace the actual matrix coefficient by the normalized spherical coefficient times its value at the identity. Aggregate over each hyperspecial double coset using its Hecke eigenvalue on y; this works for y' in the full cuspidal multiplicity space. The principal coset support above and the good spherical sum therefore give the exact finite multiplier

$$|H_P|\, \|f_{\mathrm{fin}}\|_{\mathrm{osc}}^2\, E_S / d_{\mathrm{off}}.$$

The trivial unitary bound is $|\beta_{w,a}| \le q_w^{(r-1)/2}$. At a split place one can sum the normalized coefficient over the integral sublattice cosets of each determinant valuation in Hermite normal form. The generating series has roots the Satake parameters scaled by $q_w^{(r-1)/2}$, and its coefficients are bounded by the corresponding counts for the trivial representation. Its disk of convergence contains $|X| < q_w^{-(r-1)}$, which excludes poles there and gives the asserted bound. The inverse bound follows by applying the same argument to the dual. In the inert length formula, using the trivial spherical coefficient shows that the series of absolute values converges on $|X| < q_v^{-2(r-1)}$. Every unitary spherical coefficient is bounded by one, while the numerator of the rational expression has no zero in that disk. The absence of poles gives $|\alpha_i|$ and $|\alpha_i^{-1}|$ at most q_v^{r-1}, as required. Here the underlying unitary structure is the right-regular Petersson structure on the cusp space. The estimates also give absolute convergence of the placed coset sums and their Euler product at s_0. Split trivial Satake bound · Inert trivial Satake bound.

The definite volume is $[K_{\mathrm{lat}} : K_V]\,\mathfrak{m}(V)$, with $\mathfrak{m}(V)$ the self-dual definite class mass. At $v \in S$ the definite source level is the preimage of the parabolic residue matrices with diagonal blocks in H_v and $\mathrm{GL}_t(k_v)$, with a free off-diagonal block of size rt. On the chosen finite polarization, normalized additive lattice volume one gives the local squared norm $|H_v| q_v^{-nr}$. Thus

$$[K_{\mathrm{lat},v} : K_{V,v}]\,\|\varphi_{\mathrm{fin},v}\|^2 = \frac{|\mathrm{GL}_n(k_v)|}{|H_v|\,|\mathrm{GL}_t(k_v)|\,q_v^{rt}}\,\frac{|H_v|}{q_v^{nr}} = \prod_{a=t+1}^{n}(1-q_v^{-a}).$$

At good places both local factors equal one. This proves the exact cancellation $[K_{\mathrm{lat}} : K_V]\,\|f_{\mathrm{fin}}\|_{\mathrm{osc}}^2 = d_S^{-1}$. Combining the real and finite integrals with the adjoint identity and dividing by $|H_P|$ gives, on every nonzero joint Hecke summand,

$$TT^{*\mathrm{lev}} = \lambda I, \qquad \lambda = \mathfrak{m}(V) d_V\,\frac{|E_S|}{d_{\mathrm{off}}\,d_S}\prod_{\rho} H_\rho\, d_{K,\rho}\, J_\rho. \tag{43}$$

In fact, testing the adjoint identity on a nonzero diagonal vector shows E_S is positive real, since the other factors are positive and the absolutely convergent E_S is nonzero. One of the two factors d_K in that identity is absorbed in the multiplicity pairing, and summing the orthonormal basis gives d_V. Volume in the definite metric · Level index times squared norm · The scalar on the full Hecke part.

For each selected tuple, the good principal row space is the orthogonal sum of its joint Hecke eigenspaces away from S. The spherical operators there commute and taking the Petersson adjoint is induced by inverse double cosets and conjugation of the coefficients, so the sum exhausts this finite-dimensional space. On each nonzero summand (43) is positive. All finite places of $Q(\boldsymbol{\mu})$ contributed the compact factor computed above, and the remaining Euler product is in absolute convergence with nonzero factors by the unitary bound. The full adjoint is thus injective on the good row multiplicity space for every $\boldsymbol{\mu}$. Explicitly, for any y in a summand its preimage is the adjoint image divided by that summand's scalar. This gives the principal-level surjectivity used in §2.6.3. The estimates that follow use the base tuple. Auxiliary norm formula for every tuple · Positive adjoint on the whole good row space.

Both coefficient dimensions in (43) are at least one. The standard orthonormal lattice in the genus contributes $(|\mu(E)|^n n!)^{-1}$ to $\mathfrak{m}(V)$. For an integral unitary matrix, summing the squared moduli of a column over all embeddings gives $[E : \mathbf{Q}]$. Each nonzero entry contributes at least that much by the arithmetic–geometric mean inequality and its nonzero integral norm. Thus the stabilizer consists of monomial matrices with root-of-unity entries. The inequality $\varphi(u)^2 \geq u/2$, checked on each prime power, gives $|\mu(E)| \leq 8f^2$. This also proves $-\log \mathfrak{m}(V) \leq n\log n + 5nf$. For the finite factors, $|x_{w,a}| \leq q_w^{-7/2}$. Each $\log|1-x_{w,a}| \leq |x_{w,a}|$, and $\sum_w q_w^{-2} \leq [E : \mathbf{Q}]\sum_p p^{-2} < 2f$. Similarly $-\log(1-\epsilon_v^a q_v^{-a}) \leq 2q_v^{-a}$. As there are r factors at each v and $t+1 \geq 2$, the original bounds give $\log(|E_S|/(d_{\mathrm{off}}d_S)) \geq -10fr$. The orthonormal class mass · Uniform local factor bounds.

At ρ, dg_ρ is the Haar measure of identity density $2^{-r(r-1)/2}$ in the release's real matrix tangent coordinates. On the compact dual this has total mass $\mathrm{vol}_0\,U(r) = \prod_{a=0}^{r-1} 2\pi^{a+1}/a!$. The normalized compact-dual measure transfers the invariant identity density from the compact group of volume one. Consequently dg_ρ is $\mathrm{vol}_0\,U(r)$ times the normalized compact-dual Haar. Schur orthogonality for the genuine row discrete series gives $J_\rho = (d_\rho^D)^{-1}$, and Harish–Chandra's degree polynomial, in the measure just specified, is

$$d_\rho^D = \frac{1}{\mathrm{vol}_0\,U(r)}\prod_{a<b}\frac{|\Lambda_{\rho,a}-\Lambda_{\rho,b}|}{b-a}.$$

The central circle is compact and the central character unitary. The measure on the central circle scales the identity by the same factor on the group and on its compact dual, so the formula is for the full group $U(p,q)$. The actual $\Lambda_{\rho,a}$ are strictly decreasing and in $(r-1)/2 + \mathbf{Z}$. In particular the degree polynomial uses this shifted parameter. The labels over the host quadratic field enumerate the original decreasing row parameters. The spread is at most $d_\rho + n + 3$ with $d_\rho = |\mathbf{a}^+_{j,\sigma,\rho}| + |\mathbf{a}^-_{j,\sigma,\rho}|$. Indeed for $q > 0$ the first

positive and first negative Fock row lengths bound the extreme parameters by their defining affine offsets. In the zero-based stage notation, with stage offset e_j, the first Fock row lengths are the nonnegative parts of the floors of $\Lambda_{\rho,1} + e_j - (p-1) - (t+1)/2$ and $-(\Lambda_{\rho,r} + e_j + q + (t-1)/2)$, respectively. Adding the resulting upper bounds on the two affine expressions, with one added to each length for its floor, gives the parameter spread bound $d_\rho + n + 1$. For $q = 0$ the last parameter is $-b_j$ in the release parameter data, and $e_j - b_j \geq (t+3)/2$ gives the same bound from the positive length alone. Relaxing the spread by two yields the displayed size used in the degree estimate. [61, §41, Theorem 16] [62, (3.10), (3.12), Theorem (8.1)] Transfer of the real density · Actual parameter spread · Shift and degree of the selected representation.

The stage sizes give $d_\rho \leq d_{\max} = 3rb$, $L' = \log(b\ \mathrm{rd}_E g_0 g_1 n) \geq 5.5$, $\log n \leq L'$, and $\log(d_{\max} + 2) \leq 2L'$. The compact volume is at least r^{-r^2}, from the factorial product. Since $d_\rho + n + 3 \leq n(d_\rho + 2)$ for $n \geq 3$, one obtains

$$\log d_\rho^D \leq r^2 \log r + \frac{r(r-1)}{2} \log(n(d_{\max} + 2)) \leq \frac{5}{4} n^2 L'.$$

Here the denominators $b - a$ are at least one and $2r \leq n$. With $n \geq 8$, $f \geq 1$, the total loss from the mass, the finite factors and the real degrees is at most $n \log n + 5nf + 10fr + \frac{5}{4} f n^2 L' \leq 8 f n^2 L'$. It follows that

$$\log \lambda \ \geq \ \sum_\rho \log H_\rho - 8 f n^2 L'. \tag{44}$$

Quantitative stage bound.

The positive Fock hooks give exactly the base hook sum $\sum_\rho \log H_{j,\rho}^+$ at every host-field embedding: the place labels permute the tested weights and the row signature is constant along the real places. The negative hooks have positive integral values. Thus $\exp(\sum_\rho \log H_{j,\rho}^+ - 8 f n^2 L')$ is a common lower bound on TT^* at these embeddings and all nonzero Hecke summands. If P_{row} is the orthogonal principal projection, then $\|(P_{\mathrm{row}} T)^* y\|^2 = \|T^* y\|^2$ on $\operatorname{im} P_{\mathrm{row}}$. The same lower bound holds for the squared singular values of $P_{\mathrm{row}} T$ on its kernel complement.

2.6.6 The primitive class section of the base stage

The phase-one comparison uses the base tuple of the auxiliary pair V, W above. Fix its prescribed PEL row model (the moduli model with polarization, endomorphisms and level structure) over K, the automorphic bundle for its weight, and the framed theta cover with the two marked torsion levels. The algebraic class section is built on this cover, using the polarization line, its normalized theta structure, and the E-linear harmonic differentiation in the prescribed Hodge frames. At the primitive tuple the relative twist is trivial and its ramified set is empty. The level cover has degree one. Its integer multiplier is 1, its half-level ideal is $\mathcal{O}_E$, and its scalar and finite determinant correction sections are both 1. Thus both readings use the same marked half of the row lattice. Geometric data used in the phase-one construction · Identity cover and trivial relative determinant character · Scalar-one level cover · The selected size of the auxiliary definite rank.

Use the definite source class representatives at the refined level. Let Λ_c be the placed lattice at class c, Γ_c its finite stabilizer and v' the refined level volume, so that the mass of this term is $v'/|\Gamma_c|$. Write P_c for the joint harmonic polynomial of the coefficient vector at this class, of degree d_c. Let N_c be the prescribed odd denominator prime to the marked level. In the covering marked quotient Q_2 let $w_c(q)$ be the Fourier weight of the actual finite test and let $Y_{c,q}$ be its marked lift after multiplication by N_c in the refined lattice coordinates. This lift is read by inverting N_c on the marked torsion module. Multiplication by N_c is invertible there because the level and N_c are coprime. Let Z_c be the finite set of even representatives used in the projection of the standard vacuum to Λ_c. Write $z_c = |Z_c|$ and m_c for the index used to normalize the short value. These are the two indices in Section 2.6.7. The marked Fourier weights and the class rule · The class-scalar normalization.

In the fibre over a framed point, pull the theta section back along the isogeny $[N_c]$ and translate its n factors by the theta-group lifts prescribed by $(Y_{c,q}, u)$, $u \in Z_c$. Give the product of translates the

algebraic differential D_{P_c} at the zero section, using the prescribed invariant E-differentials. The maps of the abelian scheme, the finite theta-group action and the finite trace are algebraic. In characteristic zero, taking the indicated derivative is extraction of a jet coefficient in the formal group, multiplied by its factorial, so it is algebraic in the given frames. The harmonic projection removes the quadratic change of gauge. In fact, changing the E-linear Hodge splitting changes the polarization connection in the fibre by a linear differential form from the polarization pairing. In logarithmic coordinates, its change of gauge is $\exp(Q(z))$; Q is the sum of quadratic terms pairing a τ-direction with its conjugate direction, traced over the n source coordinates. In the prescribed frames these terms act on the polynomial by the contractions

$$\sum_{b=1}^{n} \partial_{\tau,i,b}\partial_{\bar{\tau},i',b}.$$

The formula for a jet in the changed gauge is obtained by expanding the exponential. For any polynomial P the extra terms are $D_{Q(\partial)^j P}G(0)/j!$, with $j \geq 1$. They are zero for the allowed joint harmonics. This proves that the differential reads the same holomorphic derivative in the prescribed fibre coordinates once the constant in the theta-line frame has been matched. The actual mixed contractions · The algebraic fibre derivative · Comparison in the normalized chart.

Let $F_{c,q,u}$ denote the resulting product of pulled-back translated theta sections before differentiating, including the specified theta-group phase. Write e_c for the source determinant coefficient. On the primitive level the class section is given in framed coordinates by

$$\sum_{c} \frac{v' e_c}{|\Gamma_c|\, m_c N_c^{d_c} |Q_2|} \sum_{q\in Q_2} w_c(q) \sum_{u\in Z_c} D_{P_c} F_{c,q,u}(0). \tag{45}$$

The sums and weights are those of the actual class input, so the expression is linear over K. The theta-isogeny formula identifies the pullback and the finite translations with the standard Heisenberg expansion in the compatible theta structures. It determines the relative term coefficients up to the common factor of the frame. [90, §1, Theorem 4] Multiplication by N_c in the fibre coordinate introduces $N_c^{d_c}$, which cancels its inverse in (45). The sum of the even representatives is z_c times the projection to the placed lattice vacuum by character orthogonality on the quotient defining Z_c. The remaining class index is therefore z_c/m_c, and the other two weights are exactly the inverse Fourier normalization and the class mass. This gives the same normalized derivative sum as in the analytic lift. The theta section of the row fibre is the same in all terms at the point in question. The short value and its lattice projection · Summation over actual source classes.

We describe the vacuum and the framing in the reading of (45). On an analytic component write $L = p_\Lambda(\Lambda)$ for the period lattice of the row and H for its positive Appell–Humbert form. The imaginary part of H is the trace of the given skew form and is unimodular on L. The integral half-root $w_0 + \bar{w}_0 = 1$ determines the semicharacter for both the marked theta-line generator and the standard oscillator vacuum. In a symplectic basis the quasi-periodicity determines all Fourier coefficients of the principal theta function from one coefficient. This gives the one-dimensional fibre vacuum. Identify the generator at a unit frame with its holomorphic theta function ϑ, normalized by

$$\int_{\mathbb{C}^g/L} |\vartheta(z)|^2 e^{-\pi H(z,z)}\, d\bar{z} = 1,$$

with probability Haar measure. The product of n such generators reads $C(x)$ times the classical oscillator vacuum $v_{\Lambda,x}$, with the Gaussian factor removed, at the unit frame x. The weighted probability norm of $\vartheta^{\boxtimes n}$ is also one. In the self-dual model the classical vacuum's weighted mean is independent of the unit frame. In symplectic coordinates, orthogonality in the real half leaves squared Gaussian coefficients on the other half. Their sum unfolds to a Gaussian integral whose covariance determinant cancels the Jacobian of the period coordinates. Consequently $|C(x)|$ is constant. The marked generator and its probability mean · Comparison of weighted and oscillator probability means.

The phase of $C(x)$ is constant as well. The marked half is represented by the lattice $L' = p_\Lambda(\mathfrak{a}\Lambda)$ for $\mathfrak{a}\bar{\mathfrak{a}} = (M)$. Since the trace dual of Λ is Λ, the trace dual of $\mathfrak{a}\Lambda$ for the form divided by M is $M\bar{\mathfrak{a}}^{-1}\Lambda = \mathfrak{a}\Lambda$. On $\mathbb{C}^g/L'$ the same theta function is a section of the M-th power of a principal theta line, up to a flat unitary character. Such a character has an M-th root, so it can be included in the principal line. Its M-th power is generated by sections for $M \geq 2$. Indeed, translates of a nonzero principal section by $a_1, \dots, a_M$ with $\sum a_j = 0$ multiply to a section of that power by the theorem of the square. One can choose the a_j so that each factor is nonzero at any prescribed point, by avoiding finitely many proper divisors in the parameter space of the first $M - 1$ translates. The space of sections of the power is the irreducible finite theta representation of dimension M^g. For a symplectic decomposition of M-torsion, one half acts with a full orbit of M^g distinct characters, permuted transitively by the other half. This proves irreducibility using the dimension from Riemann–Roch for an ample line on an abelian variety. The subspace of sections vanishing on all M-division points is invariant under the theta group and proper by generation at the origin, so it is zero. There is thus a marked torsion value of ϑ which is nonzero at every given fibre, with the torsion label allowed to vary. The product can use this same label in its n factors. [90, §1] The half-lattice form and division points.

Around a frame with a nonzero marked value, compare the two powers at that torsion label. Both evaluations are holomorphic in the domain coordinates, with the same frame weight. For the theta series this follows by writing the lattice sum in the holomorphic Appell–Humbert gauge, where it converges normally on compact subsets. The frame weight from the Gaussian vacuum is the determinant power $n/2$ in the prescribed convention. For the algebraic evaluation it is the corresponding Hodge determinant power. Taking the quotient cancels this determinant even when evaluating it in unit frames. The quotient is $C(x)$. It is holomorphic where the marked value is nonzero and has constant modulus, so it is locally constant. The domain of each component is connected, and these neighborhoods cover it. This gives a single nonzero constant C_C on each component C. Generator comparison in the unit frames.

For the adelic reading choose a finite representative placing the row lattice in C. The transported spherical vacuum reads $\kappa_C v_{\Lambda,x}$ times the prescribed determinant character $\mu(g)$. Indeed the finite lattice stabilizer acts on its one-dimensional vacuum by the splitting character. In the self-dual lattice realization the action before splitting is pullback, fixing the defining lattice vector. A change of finite representative for the same lattice therefore multiplies the reading by the same factor as $\mu(g)$. The real action moves the frame and has the determinant factor already in the analytic weight. The probability inner products agree under the self-dual lattice realization. One can integrate the same Weil vector on the adelic Heisenberg quotient, using the unit finite lattice and a fundamental parallelotope at infinity. The Gaussian-weighted real integral expressed in period coordinates has probability measure and is the product over the n source columns. In particular $|\kappa_C| = |C_C|$, so κ_C/C_C is a unit scalar common to all inputs and all points of C. Adelic vacuum, stabilizer and real frame · The quotient of the two component scalars.

The harmonic coefficient map and the normalized Fock map are maps of the same definite source type to the joint harmonic row type. They differ by a scalar u_j of modulus one because their distinguished extremal vectors have the same norm. Multiplicity one here can be read from the extremal polynomial, a product of flag minors in the positive and negative blocks. In the adapted normalized frame a homogeneous polynomial of degree d_j contributes the factor

$$u_j(-\delta_j)^{d_j}\left(\sqrt{2\pi\xi_R}\right)^{-d_j}, \qquad \delta_j \in \{1, -1\}.$$

The sign is that of the additive variable in the short Fourier series, and the positive factor comes from the normalized Gaussian derivatives. This factor, the character $\mu(g)$, and κ_C/C_C turn (45) into the analytic stage evaluation. Absorb the two unit scalars into its component frame constant. At the primitive tuple both the group element and the root of unity outside this component constant are 1, giving (47). Unit scalar for the harmonic coefficient map · Positive Lie factor and the short-series sign.

The class formula is a section over the coefficient field of the marked model. To see the gluing, change a marking in the framed model. The finite Fourier weights transform by the inverse character to the theta-

group action on the torsion sum, and the homogeneous derivatives transform by the specified weight of the Hodge frame. Thus the expressions on overlaps satisfy the transition rule of the automorphic bundle. Equivalently, on the fibre product of the frame charts the two readings are equal as sections of the pulled-back bundle. The charts are smooth in characteristic zero and equality may be tested on complex points. Faithfully flat descent gives the section S_x over K. Descent for one normalized class series.

The sections are cuspidal. For an isotropic row line, choose a symplectic basis adapted to it and use the Fourier series in the central line-unipotent coordinate. In a theta term its frequencies, one at each real embedding, are sums of squares of the leading row entries of the definite-source columns, including the fixed shifts. They vanish simultaneously only when all these entries are zero. In the positive row block the variables then lie in a hyperplane. Every polynomial in the specified highest-weight component is a sum of products of flag minors, and each product contains a maximal positive-block minor. The latter follows from positivity of its last weight. Indeed for the j-th stage the positive block has size $k + j$, and its last polynomial weight in the given convention is $\lfloor -b_{\tau,j} + e_{\tau,j} - (t+1)/2 \rfloor \geq 1$ by the prescribed gap at j. The maximal minor vanishes on the hyperplane. Therefore the zero frequency vanishes in every normalized class sum. The same Fourier calculation applies to every rational isotropic line. Any proper rational parabolic stabilizes an isotropic flag. Integrate its constant term first over the centre of the line unipotent to get zero. At an anisotropic row group there are no proper rational parabolics. Last positive weight at the actual stage · Adapted line expansion and the maximal minors.

On a toroidal compactification of the neat PEL model the cusp sections are global sections of the subcanonical extension of the automorphic bundle. The space is finite-dimensional over K, and extension of scalars commutes with its sections, by computing sections as the kernel of the restriction maps of a finite affine cover. Thus the class values give the K-linear coordinate map on the cusp basis. The lowest-type vector is cyclic, so equality of its evaluations gives equality of the multiplicity maps. Lan–Stroh, subcanonical cusp bundles on PEL models. The component constants at phase one · Comparison on the base class sections · The cusp-space realization.

The construction uses the PEL framing whose conjugate of the lattice module is the placed row lattice, with the prescribed polarization, marked module and coefficient isometry. These identifications are needed at every class representative of the selected stage, with the same phase-one choices before the conjugate comparison.

The base reading on a marked short section. In the common construction of the base and enlarged tuple data, retain the vacuum comparison $c_i(C)$ before multiplying by the harmonic phase. The base marking is the identity cover, so its determinant section is constant. Write $c_{0,i}$ for its value in the comparison frame. The common comparison multiplies the vacuum factor by $k_i(-\delta_i)^{d_i}/c_{0,i}$, using the harmonic map and additive sign of (45). The coefficient-field extension agrees with the original harmonic map on the base coefficients. The construction installs this product as the base component factor $c_{E,i}(C)$ of the row model. In particular, at the distinguished embedding and the chosen lowest-type vector, the reading of a section f in that model is

$$\mathcal{R}_i(f)(g) = L_i \, \frac{k_i(-\delta_i)^{d_i}}{c_{0,i}} \, c_i(C(g)) \, \mu_i(g) \, \mathrm{rd}_g(f). \tag{46}$$

Here $L_i = (2\pi\xi_R)^{-d_i/2}$, $C(g)$ is the component of the framed dictionary point, and rd_g uses its Hodge coordinates. The formula is the definition of the installed reading on the fixed principal cusp space. The component factor together with $\mu_i(g)$ is $c_i^{\mathrm{fr}}(g)$ above. The class calculation identifies these factors before choosing a finite class input. For either prescribed level cut, first apply the corresponding finite character projector in this same cusp space. Base restriction of the common datum · The phase multiplying the base vacuum comparison · Installation on the base components · Reading in the installed cusp space.

Use this same formula for the short section on the run's marked cover at Mp_o after descent to its base weight space. The level map pulls back the framed Hodge coordinates; the cover coefficient-field

isomorphism carries the coefficients of the finite harmonic polynomial in these same coordinates. Its complex embedding restricts to the distinguished embedding of the base field. Equivariance here follows by applying the coefficient map to the polynomial, the finite Fourier weights and the theta-group phases in the same finite trace. The adapted embedding fixes the small field containing the deletion character at p_o, so the reference character and its marking stay fixed while the extra ray-character values are read through the adapted embedding. The joint algebraic determinant coefficient on this lattice is the one in (48). The coefficient map and the marked fibre of the short rule · The adapted embedding fixing the small field.

2.6.7 Transport of the phase-one base stage

Fix one of the primitive base-tuple stages. Use the same definite source V, row space W, principal projection P_σ, and degree and coefficient fields as in the stage norm calculation. Write ι for the chosen coefficient embedding, k_{CM} for the host imaginary quadratic subfield, and K for the row coefficient field. The embeddings over the chosen inclusion of k_{CM} give the holomorphic framed models in the released comparison. They permute the embeddings of E in the chosen CM type, since that type consists of the extensions of this inclusion. The signatures and hence the positive hook sum are preserved under the permutation. The embedding $\sigma^+ = \overline{\sigma}$ over the chosen inclusion is the partner of any σ whose restriction to k_{CM} is the opposite one. Embeddings over the host field · The geometric embedding condition.

The fixed-basis comparison uses the algebraic cusp sections of the framed unitary moduli space of the row in this stage. Use the chosen weight and geometric cusp basis. For each rational class input x, set S_x to be its algebraic theta class section, a linear combination $\sum_\ell a_\ell s_\ell$ of the fixed cusp sections after base change to K. Both the action of the finite level group G on sections and this decomposition are over the field. The character of a level move in the adelic row comparison is denoted χ_W. It is the product of the prescribed residue-determinant characters, so has finite image in $K^\times$. The geometric row model reads a section s_ℓ at the fixed unit lowest-type vector as a function $J_\sigma(s_\ell)$, including the specified adelic and weight factors, in the cuspidal multiplicity model. The normalized base class identity of the phase-one realization is

$$P_\iota T_\iota j_\iota(x) = \iota(\zeta_0) P_\iota j_{W,\iota}(A_{z_0} S_x). \tag{47}$$

Here A is the geometric level action, z_0 a fixed central element of G, and $\zeta_0 \in K$ a root of unity. One may take both to be one in the phase-one base choice which absorbs the common vacuum phase into its base component factors. This comparison starts with the framed cusp sections, their linear decomposition over K, and the normalized phase-one class-section formula on the actual source classes. In that formula the theta generator, differential coordinates, cover marking and term volumes are part of the realization. The equality at the lowest-type vector determines the map because the vector is cyclic in the real representations and the multiplicity maps are equivariant. The algebraic class section and its base comparison · The normalized base choice · The phase-one base reading.

We compare the same section with the actual stage lift after an embedding σ over k_{CM}. Let M be the row marking modulus, M_2 the covering marking modulus, with $M \mid M_2$, and Q_2 the finite quotient of the marked half at level M_2, taken with one copy per source coordinate. Since the auxiliary base twist is trivial, its twisting support is empty. The finite base test is the characteristic function ϕ of the identity leading block and the zero complementary block. In particular its values are $0, 1$. Let $w_c(q)$, $q \in Q_2$, be its finite Fourier coefficients on the c-th source lattice class, with the inverse transform normalized by $|Q_2|^{-1}$. The pairings in the Fourier matrix have values in μ_{M_2}. An embedding of the algebraic class coefficient field has $\sigma(\zeta_{M_2}) = \zeta_{M_2}^{m_\sigma}$ with $(m_\sigma, M_2) = 1$. Extending σ to the covering coefficient field if necessary, the same formula holds for the chosen Fourier matrix. Consequently

$$\sigma(w_c(q)) = w_c(m_\sigma q)$$

for every c, q, with one multiplier m_σ for all classes. The actual comparison of the finite additive character with the dual marking has one sign $\epsilon \in \{\pm 1\}$. With the sign included in the first marking, replacing the labels q by $m_\sigma q$ is exactly the change to the marking ϵm_σ times the original marking. Indeed, the rational roots at label q are specified by a congruence modulo the lattice. Taking a representative of $m_\sigma q$ in the marked half is equivalent to multiplying a representative of q by m_σ, modulo M_2 on that half. The difference yields an integral lattice translation. The theta-group lifts use the same normalized signs, so the resultant sum is the reindexing of the finite translations. Multiplication by m_σ is a bijection of Q_2, with inverse given by a Bézout coefficient. Actual finite Fourier average and marking sign · The common Galois Fourier dilation · Dilation on the finite marking quotient.

The chosen half of the row level has ideal $\mathfrak{m}_b = \prod_{w \in P_g} w$, with $\mathfrak{m}_b \overline{\mathfrak{m}_b} = M\mathcal{O}_E$. The marked submodule before reduction is $\overline{\mathfrak{m}_b}$ times the row lattice. At each chosen type prime send the residue of ϵm_σ to the scalar matrix in $\mathrm{GL}_r(k_w)$. At the conjugate prime the unitary condition uses the inverse conjugate matrix. These scalars define a central element z_σ in the prescribed full product G of residue groups. Its action on the marked submodule agrees with multiplication by ϵm_σ modulo M. The extra marking at M_2 can be lifted after this move because $(m_\sigma, M_2) = 1$, using the compatible lift of the dilated marking just computed. In particular the element of G is independent of the source class and of the individual torsion label. The central element lifts adelically to a scalar on the selected split places, with inverse conjugate on the other half and the identity elsewhere. Denote its right translation on row forms by $R_\sigma(z_\sigma)$. The central level move · The compatible lifted marking · The central adelic lift.

The analytic and geometric evaluations now use the same finite class expression. Here are the scalar checks on it. Take the same source representatives h_c for an open subgroup $J' \subset J_V$, with class stabilizers Γ_c. Passing to the finer representatives gives $\mathrm{vol} = 1/[J_V : J']$; the class weights $\mathrm{vol}/|\Gamma_c|$ are rational. Let Λ_c be the conjugate of the corresponding finite lattice in the oscillator frames, I_c the index of its intersection with the standard lattice in Λ_c, and N_c the denominator in the class construction. Use m_c, z_c for the two integer cardinalities in the class term $\mathrm{mIdx}(N_c, r), \mathrm{zCard}(N_c, r)$. In the finite lattice projection $z_c = m_c I_c^r$. This is the index formula on the inclusion $N_c\mathcal{O}_E^n \subset \Lambda_c \cap \mathcal{O}_E^n \subset \Lambda_c$, applied in the row variables. The m_c normalizes the kernel average and z_c counts the translated terms. The differentiation of a test of total degree d on the rescaled fibre introduces N_c^d, which cancels the compensating factor N_c^{-d} in the class definition. In particular the holomorphic evaluation of S_x on the moved framed point has the same polynomial derivatives of the same $|Q_2|^{-1}$-weighted translates as the analytic class sum. This is the theta pullback for compatible theta structures, followed by the finite trace and the harmonic derivative in the marked coordinates. One can compute it by first expanding the finite theta-group projection as a sum of translations of the generator. The sum over the even representatives is z_c times the normalized lattice projection. Linearity of the derivative then leaves the above cardinalities and a common factor $\kappa'(g)$ from the n-fold theta generator on the row fibre. The generator is nonzero, and its conversion to the oscillator vacuum is an equality of theta lines at the same point. Thus $\kappa'(g)$ is common to all classes, polynomials and torsion terms at g. [90, §1, Theorem 4] Finer source class volume · The exact finite derivatives and their coefficients · The index ratio in the short theta value · Trace and harmonic differential in the theta comparison · Evaluation of the class section.

The remaining class coefficient is the determinant character of the source lattice move times a positive square root. To transport it, let $\mathfrak{d}_c/\overline{\mathfrak{d}_c}$ be the determinant ideal of the move, with $\mathfrak{d}_c$ and its conjugate coprime. This is obtained by taking the positive valuations of the determinant; the opposite valuations are their negatives by unitarity. Write y_c for the specified half determinant idèle, and write μ_σ, η_σ for the placed splitting characters. At the supporting primes outside the type primes one transports $\eta(\varpi_w)(\mathrm{N}w)^{-1/2}$, using the fixed algebraic readings of this product. Indeed the product over a principal ideal for the η-reading is the inverse of the product of the embeddings of its generator in the compatible CM type, by the half exponent of η. The conjugate embeddings transport this principal-ideal identity. Restoring the positive norm factors yields the unramified placed character η_σ. Its restriction to F is preserved, since the norm-normalized restriction at a split pair or an inert prime is rational. Thus the product

that transports at y_c is $\eta(y_c)^{-1}\sqrt{\mathrm{N}\mathfrak{d}_c}$, with the character and positive square root taken together. For μ, the prescribed readings on units and uniformizers are algebraic. Conjugating them and permuting the integral CM infinity exponents gives the placed character μ_σ. The character is trivial on principal idèles because their infinity factor is an algebraic product of paired embeddings with the prescribed exponents. The conjugated finite readings on units are roots of unity, and the values on uniformizers have absolute value one by raising their ideals to principal ideals, with the unit factors at the type primes included.

In the prescribed class coefficient $B(\mathfrak{d}_c)$ these readings give

$$\sigma(B(\mathfrak{d}_c)) = (\mu_\sigma \eta_\sigma^{-1})(y_c)\sqrt{\mathrm{N}\mathfrak{d}_c}. \tag{48}$$

The coefficient is defined by multiplying the corresponding readings at prime ideals. This applies to every nonzero determinant ideal prime to the bad modulus of the base datum. It therefore applies to the determinant ideal of an adapted short-section lattice supported off that modulus. The own-prime correction on its marked cover contributes the separate value of the deletion character on that ideal. The joint reading on ideals prime to the bad modulus · The same base reading for the short rule. The scalar in the primitive class section is $e_c = q_c/B(\mathfrak{d}_c)$, with $q_c \in \mathbf{Q}_{>0}$ and $q_c^2 I_c^r = \mathrm{N}\mathfrak{d}_c$. Therefore

$$\sigma(e_c)\,\frac{z_c}{m_c} = (\mu_\sigma \eta_\sigma^{-1})(y_c)^{-1} I_c^{r/2}.$$

The square root $I_c^{r/2}$ is exactly the conversion between the projection of a unit vacuum and the unit vacuum of the moved tensor lattice. The formula fixes the scalar of each source class, including when the index has odd valuation. The class-scalar normalization · Determinant ideal and the joint square-root reading · The transported class coefficient.

There are three remaining unit factors. First, in the Fock coefficient map the fixed rational highest vector goes to the prescribed minor product. In the comparison at σ both coefficient highest vectors are unit for the standard Hermitian form, and the same minor product is transported by relabelling its variables. The Cartan-frame change on the definite side is unitary. Thus the two maps have a scalar of absolute value one on the highest line. Since the coefficient type is irreducible and the polynomial maps intertwine its action, this scalar α_σ applies to the whole coefficient space. Second, the cover factor on the base tuple is a constant γ_σ, the embedding of the prescribed root of unity. Indeed its values on the marked framed points are the root of unity times the auxiliary twist character on the determinant, which equals one on the base tuple. These points uniformize the complex charts, so the cover section is that constant after base change. Third, on a component c of the row charts the two normalized vacua differ by a factor $c_{\mathrm{fr},\sigma}(c)$ times the specified determinant automorphy factor. In both descriptions the Gaussian squared norm on the lattice torus has mean one. The self-dual lattice gives a common one-dimensional theta space, hence $|c_{\mathrm{fr},\sigma}(c)| = 1$ at one frame of the component. For another frame in the same component, move first by a real unitary operator and then by the finite self-dual lattice stabilizer. Their oscillator multipliers are the determinant characters appearing in the adelic automorphy. The comparison with the transported framing therefore uses the same $c_{\mathrm{fr},\sigma}(c)$. In the geometric reading the fixed phase is $c_{E,\sigma}(c)$, also of norm one. Their quotient gives a function $u_\sigma(g)$ of absolute value one which depends only on the chart component. It is unchanged under the left rational group, the right real group, and the right principal level. For any unipotent radical N and any g, additive approximation along a unipotent filtration gives density of $N(F)$ in $N(\mathbb{A}_{F,f})$. Thus for $n \in N(\mathbb{A}_F)$ one can choose $\gamma \in N(F)$ with $g^{-1}\gamma^{-1}ng$ in the right real group times the principal finite level. It follows that $u_\sigma(ng) = u_\sigma(g)$, so multiplication by u_σ preserves the vanishing of the cuspidal constant terms. This multiplication is unitary, since Petersson integration is unchanged. The multiplier also passes through the unipotent averages. For a rational unipotent radical N, given $n \in N(\mathbb{A}_F)$ and g, choose $\gamma \in N(F)$ so that $g_f^{-1}\gamma_f^{-1}n_f g_f$ belongs to the principal level. Such a γ exists by additive approximation along a split unipotent filtration. The left rational, right real and right principal invariance of the multiplier then gives $u_\sigma(ng) = u_\sigma(g)$. Fixed highest vector and unit Fock

factor · The root of unity on the base cover · Vacuum movement by the determinant multiplier · A unit vacuum phase on each component.

In detail, let $D_{\sigma,c,x}$ denote the harmonic derivative with the transported polynomial of the class value x_c, and let v_c be the projection of the normalized real vacuum to the class tensor lattice. Write $\mathcal{H}(Y)$ for the normalized Heisenberg translation by the marked root Y. At the framed point of g the common expression is

$$C_\sigma(x,g)=\sum_c \sigma\left(\frac{\operatorname{vol} e_c}{|\Gamma_c|}\right)\frac{z_c}{m_c|Q_2|}\,D_{\sigma,c,x}\left[\sum_{q\in Q_2} w_c(q)\mathcal{H}(Y_c(q))v_c\right].$$

On the adelic side, evaluate $\alpha_\sigma T_\sigma j_\sigma(x)$ at the unit vector. This gives

$$c_{\mathrm{fr},\sigma}m_W(g)\ell\,\kappa'(g)C_\sigma(x,g),$$

with ℓ the prescribed real derivative conversion and m_W the determinant automorphy factor at σ. On the geometric side the value of $A_{z_\sigma}S_x$ at the frame is $\gamma_\sigma\kappa'(g)C_\sigma(x,g)$. The adelic reading $j_{W,\sigma}$ multiplies the last value by $c_{E,\sigma}m_W(g)\ell$. Using $u_\sigma=c_{\mathrm{fr},\sigma}/c_{E,\sigma}$ therefore gives

$$\alpha_\sigma\gamma_\sigma T_\sigma j_\sigma(x)=U_\sigma j_{W,\sigma}(A_{z_\sigma}S_x). \tag{49}$$

Here the equality is promoted from evaluated functions by the injectivity of lowest-type evaluation. The use of the same coefficient vectors x_c on both sides is through the fixed algebraic coefficient basis and its σ-linear realization, so the arithmetic input on the right remains $S_x=\sum a_\ell s_\ell$. The common section factors are $\sigma(a_\ell)$ after this realization. The level action is read with the same U_σ. In fact a level move $a\in G$ changes a section frame by the analytic theta-line scalar at that move raised to n. The vacuum comparison yields exactly the ratio of component constants before and after that frame change. The theta-line scalar cancels in this ratio and the remaining multiplier is the placed determinant value $\sigma(\chi_W(a))$. Equivalently, writing $B_\sigma=U_\sigma j_{W,\sigma}$, one has

$$B_\sigma A_a=\sigma(\chi_W(a))R_\sigma(a)B_\sigma.$$

Both identities use the same marked section evaluations. The chosen full conjugate-character lift may use a different orthonormal realization of its coefficient space. The two realizations are related by the isometry obtained from their common Gram matrix. The normalized oscillator vacua are on the same line and the joint-harmonic couplings agree up to a single unit scalar by their normalized extremal vectors. Thus this replacement affects (49) only through a fixed unit scalar. Level action and cancellation of the theta-line scalar · Equality of the two framed class sums · Fixed-basis identity at the chosen lowest-type vector · Unitary independence of the conjugate model.

Let $e=\sum_{a\in G}e_a a$ be the K-rational group idempotent defining the principal row projection on the fitted summands, so $P_\sigma=\sum_a\sigma(e_a)R_\sigma(a)$. On the geometric space put $P_K=\sum_a e_a\chi_W(a)^{-1}A_a$. The level law gives $B_\sigma P_K=P_\sigma B_\sigma$. Define the stage map over K by $\theta_K(x)=\zeta_0 P_K A_{z_0}S_x$. By (47) this realizes the actual projected map at ι, and it is unique because the row realizations preserve independence. Write $z=z_0z_\sigma^{-1}\in Z(G)$. Then

$$\begin{aligned}U_\sigma j_{W,\sigma}\theta_K(x)&=\sigma(\zeta_0)P_\sigma B_\sigma A_{z_0}S_x\\&=\sigma(\zeta_0\chi_W(z))\alpha_\sigma\gamma_\sigma\,R_\sigma(z)P_\sigma T_\sigma j_\sigma(x).\end{aligned}$$

The last coefficient has absolute value one. Right translation by z is an isometry for the fixed Petersson measure and commutes with the projection because z is central. This proves the required equality of target Gram forms on all rational inputs. The conjugate coefficient source realizations have their class Gram forms by the common coefficient pairing, with the rational class weights computed above. Scalar extension of θ_K, followed by the indicated invertible unitary transformations, shows that the kernel of $P_\sigma T_\sigma$ is the complex span of the realized rational kernel. Thus the fitted quotient is computed using the

same kernel complement. Rational projector and central-translated comparison · All embeddings of the phase-one stage · The complexified kernel.

At an embedding opposite on k_{CM}, complex conjugation of the real unitary representation at σ^+ replaces the highest type by the opposite type. It gives the antiunitary mate of the chosen unit lowest-type vector. The principal finite character is trivial, so taking complex conjugates of its automorphic values preserves its level and cusp conditions. For the fixed cusp basis the readings at the two embeddings are conjugate under this identification. This is the conjugate reading used in forming the rational row model. The multiplicity Gram is Petersson evaluation at the unit vector times the product of compact type dimensions. Those dimensions agree under complex conjugation, so the Gram forms are conjugate on basis vectors and hence on the entire rational row. The rational class-sum source forms are conjugate as well. Finally the comparison of principal and character volumes divides both the source quotient squared norm and the row squared norm by the same $|P|$. Thus the squared singular value on the fitted summand is measured by the stage norm estimate at σ^+, including at the embedding opposite on k_{CM}. Opposite embedding and principal level · Conjugate Gram on the fixed basis · Both principal-to-character rescalings.

The least stage factor. The isometries in this transport identify the rational kernel complement with that of the actual projected lift and carry the class-sum and Petersson forms to the prescribed metrics. Apply (44) in the model over the chosen CM type and use the conjugate metric at its partner. At the character-level comparison the principal factor divides both squared norms. On a nonzero fitted row summand insert a vector of nonzero quotient norm into the stage Gram equality, and divide by that norm. One obtains

$$\log \lambda_{j,p,\sigma} \;\geq\; L_{6,j}, \qquad L_{6,j} = \sum_{\rho} \log H^{+}_{j,\rho} - C_{\mathrm{sz}} f n^2 L', \qquad C_{\mathrm{sz}} \geq 8.$$

Empty summands contribute no constraint to the least eigenvalue of the actual quotient. The inequality is uniform in the fitted label p and the coefficient embedding σ. Consequently taking the minimum of the sums of stage logs over fitted labels and averaging over all coefficient embeddings preserves the lower bound $\sum_j L_{6,j}$ for the period-map height. Base projected lift · Transport of the conjugate row norm · Unitary comparison of the two Gram forms · Fitted eigenvalues away from the base place · Base hook term and the common loss.

2.6.8 The normalized theta germs

Let k_{CM} be the host imaginary quadratic field, with discriminant $-d_{\mathrm{CM}}$, and use its chosen complex embedding. The host curve ideal in the chosen family is the unit ideal, so the period lattice of the host elliptic curve is $\Lambda = \delta\mathcal{O}_{k_{\mathrm{CM}}}$, with $\delta^{12} = \Delta(\mathcal{O}_{k_{\mathrm{CM}}})$ in Weierstrass normalization, where $\Delta = g_2^3 - 27g_3^2$. Scaling gives $\Delta(\Lambda) = 1$. The CM j-invariant is algebraic, so $g_2(\Lambda)^3 = j/1728$ and $27g_3(\Lambda)^2 = g_2(\Lambda)^3 - 1$ are algebraic. The invariant differential dx/y on $y^2 = 4x^3 - g_2(\Lambda)x - g_3(\Lambda)$ pulls back to dz. This realizes the actual lattice, differential and complex coordinate on an algebraic CM curve.

Put $A = \mathrm{area}(\Lambda)$. The reduced odd theta function of the divisor $[0]$ is normalized by $\theta'(0) = 1$. Its multiplier on $\gamma \in \Lambda$ is the sign 1 for $\gamma \in 2\Lambda$ and -1 otherwise, times $\exp(\pi(z\bar{\gamma} + |\gamma|^2/2)/A)$. The entire function with this multiplier, divisor and derivative is unique, since the quotient of two such functions is elliptic and entire. This is the odd theta function used in the vacuum. Bannai–Kobayashi prove the following special case. For a CM elliptic curve over a number field with the period coordinate of its algebraic invariant differential, the reduced theta for $[0]$ normalized by derivative one has algebraic Taylor coefficients. Its reduced translate

$$U_a\theta(z) = e^{-\pi(z\bar{a}+|a|^2/2)/A}\theta(z+a)$$

has algebraic Taylor coefficients at zero for every $a \in \mathbb{Q}\Lambda$. Indeed their translation theorem works for a symmetric divisor, with algebraic germ at zero. Divide by a rational function defining the translated

divisor locally at zero. The normalized coefficients lie in a torsion field, and multiplying back by the rational germ retains algebraicity. The reduced and Mumford translations differ by a root of unity on torsion. [31, Example 1.9, Corollary 2.2 and Theorem 2.9] Thus the even theta $\theta_{\mathrm{ell}} = U_{\omega_0/2}\theta/(U_{\omega_0/2}\theta)(0)$ and all its reduced torsion translates have algebraic Taylor coefficients. The denominator is nonzero since the chosen half period is nonzero modulo Λ.

Here is the passage to the actual vacuum of the host E. Use the CM type over the embedding of k_{CM}. The host fractional ideal $\mathfrak{A}_0 = \mathfrak{b}^{-1}$ satisfies $\mathfrak{A}_0\overline{\mathfrak{A}_0} = \mathfrak{D}^{-1}_{E/k_{\mathrm{CM}}}$. The lattice Λ_M is the type embedding of $\delta\mathfrak{A}_0$. For the Hermitian form H obtained by dividing the standard form in the complex coordinates by A, its alternating form $E_H = \operatorname{Im} H$ on that lattice is

$$E_H(\delta x, \delta y) = \frac{2}{\sqrt{d_{\mathrm{CM}}}} \operatorname{Im} \operatorname{Tr}_{E/k_{\mathrm{CM}}}(x\bar{y}),$$

with the trace in the chosen embedding. An element pairing integrally with the lattice also pairs integrally against its ω-multiples, for $\omega = (1 + i\sqrt{d_{\mathrm{CM}}})/2$. Hence the trace pairs in $\mathcal{O}_{k_{\mathrm{CM}}}$. The trace dual calculation gives $\mathfrak{D}^{-1}_{E/k_{\mathrm{CM}}}\overline{\mathfrak{A}_0}^{-1} = \mathfrak{A}_0$. Testing real vectors against a rational basis of the lattice shows first that the dual vectors lie in the rational span, since the alternating Gram matrix is nondegenerate and rational there. This proves self-duality for E_H. The same calculation applies to the Hom lattice between two of the self-dual Hermitian class lattices: the Hom lattice is Hermitian self-dual, as can be checked in local bases using the two unimodular Gram matrices. The relative trace dual then introduces the same $\mathfrak{A}_0$.

The fourfold isometry g_C of the vacuum is induced by an isometry over k_{CM} from the fourfold trace form to the standard form. For existence, diagonalize the positive Hermitian trace form over k_{CM}. Each diagonal value is a positive rational number a, and the four-squares identity gives a rational matrix with transpose product aI_4. These matrices realize the isometry on the four copies of each diagonal summand. For the chosen isometry, its complex matrix has algebraic entries by expressing it in the type embeddings of a k_{CM}-basis. It takes rational period vectors to rational period vectors. Thus $L_1 = \Lambda^{4[F:\mathbb{Q}]}$ and $L_2 = g_C\Lambda_M^4$ are commensurable and self-dual for the same alternating form. The inverse relative different factor · The chosen fourfold isometry.

The common semicharacter of the two lattices is $\beta(x) = \exp(\pi i \sqrt{d_{\mathrm{CM}}}\, H(x,x)/2)$. Its exponent divided by πi is $E_H(\omega x, x)$, an integer on either lattice. Polarizing this expression shows that its change at $x + y$ is $E_H(x, y)$ modulo two, since $2E_H(\omega x, y) = E_H(x, y) + \sqrt{d_{\mathrm{CM}}} \operatorname{Re} H(x, y)$. On one standard elliptic factor, multiplication by ω has distinct eigenvalues zero and one on $\Lambda/2\Lambda$. This follows from its trace one and its even norm, using $d_{\mathrm{CM}} \equiv 7 \pmod 8$. The chosen ω_0 mod 2Λ is the sum of the two nonzero eigenvectors, since the images of $\omega_0/2$ under ω and $1 - \omega$ are nonzero modulo Λ. The even theta sign at γ is the odd sign times $(-1)^{E_H(\omega_0,\gamma)}$. This sign and $\beta(\gamma)$ both equal 1 on zero and on the eigenvectors, and -1 on their sum. This identifies the sign in the actual product F_{ell} of the θ_{ell} over the standard elliptic factors.

For two such self-dual lattices L_1, L_2, put $L_0 = L_1 \cap L_2$. The finite quotients $L_1/L_0, L_2/L_0$ have a perfect pairing through $e^{2\pi i E_H}$, because the annihilator in each lattice is the intersection. In particular they have equal order. The transport T_{12} is the sum of $\beta(x)U_x$ over L_2/L_0, divided by $\sqrt{[L_2 : L_0]}$. Here U_x is the reduced translation for H. The translation law $U_xU_y = e^{-\pi i E_H(x,y)}U_{x+y}$ makes this sum independent of representatives and invariant under the operators $\beta(x)U_x$ for $x \in L_2$. To compose with T_{21}, commute the L_1/L_0 operators past the L_2/L_0 operators on the source theta space. Summing the commutator characters over L_1/L_0 leaves only the identity coset of L_2/L_0. The two square roots of the index cancel the order of the surviving sum, so the transports are inverse. This also proves nonvanishing of the transported vacuum. The standard theta space is one-dimensional by the one-variable elliptic theta condition applied in each variable. After transport, pull back by g_C. The result belongs to the theta line for Λ_M^4. A nonzero slice shows that the theta space for Λ_M is nonzero. Two independent vectors there would give independent fourfold products. Thus the transported product is a nonzero scalar times a product of one vacuum at four independent arguments.

The current normalization specifies

$$\theta_M(z_1)\theta_M(z_2)\theta_M(z_3)\theta_M(z_4) = \kappa_M(T_{12}F_{\mathrm{ell}})(g_C(z_1, z_2, z_3, z_4)).$$

It can be imposed by taking a fourth root of the scalar in the one-dimensional space. Here κ_M is the prescribed finite product of powers of roots of $\varphi(\pi)/\varphi(\bar{\pi})$ for generators of powers of split prime ideals of k_{CM}, with φ its complex embedding. In particular κ_M is algebraic and nonzero. The right-hand side is a finite sum of products of reduced torsion translates of the elliptic theta after algebraic linear substitution, with algebraic coefficients. Composing simultaneously with the four reduced translations by a rational period vector a preserves this property. Indeed $g_C(a, a, a, a)$ is rational in the standard periods, and the translation law adds only roots of unity. Taking the coefficient of a multi-index α in each of the four independent argument blocks gives the fourth power of the Taylor coefficient of $U_a\theta_M$. Thus each such coefficient is algebraic, also when the value of θ_M at zero vanishes. Fourth-tensor normalization · Algebraic factor in the vacuum.

For a source and target class lattice N, N', take the product of θ_M in the matrix entries, then move its source lattice and target lattice by the same normalized Heisenberg averages. The two relative indices are finite since all lattices are in the same rational span. The average at each step and its torsion translates preserve algebraicity, by the translation law and the preceding result. The determinant factor $\chi_F(N, N')$ in the moved vacuum is $\chi(\det N)^{\dim N'/2}\chi(\det N')^{-\dim N/2}$, where the determinant ideals are measured from the standard lattices. In the normalized germ its reciprocal removes this factor. Thus this argument uses the stated normalization, including when changing the chosen ideal character in the paired construction. Moved vacuum and determinant factor · The normalized germ under the change of character.

2.6.9 Finite coset jets at the chosen points

In one step with ranks a, b, let $\mathcal{L}$ be the semilocal Hom lattice of the source and target class lattices at the tested places. The finite symbol Ξ is supported on $c_A^{-1}\mathcal{L}$ and constant modulo $c_B\mathcal{L}$. The arithmetic kernel sums its values over $c_A^{-1}\mathcal{L}/c_B\mathcal{L}$. To see the normalization of a term, write s for the lift of its coset in the period-coordinate lattice. For the translation construction of the indicator germ, let A_1, A_2 be the two finite sets of representatives in the Fourier sum and its dual translation sum, and let $\mathbf{1}_M$ be the subgroup indicator in that Fourier sum. If α_H is the quadratic phase $\exp(\pi i\sqrt{d_{\mathrm{CM}}}\,H(x,x)/2)$ and G the moved vacuum before the determinant factor, the normalized germ for this coset is

$$G_s = \sum_{u\in A_2} \alpha_H(s)\alpha_H(u)e^{\pi i E_H(s,u)}\left(\frac{1}{|A_1|}\sum_{v\in A_1}\mathbf{1}_M(v)e^{2\pi i E_H(v,u)}\right)U_{u-s}G.$$

Indeed Fourier transformation of the indicator gives the inner average times the square root of its volume, and the germ normalization contributes the reciprocal square root. The multiplier for $U_{-s}U_u$ is $e^{\pi i E_H(s,u)}$. The determinant factor $\chi_F(N, N')$ cancels in G_s. This proves the formula by finite character expansion with the normalized measures. Its summands are algebraic germs since the phases are roots of unity on rational period vectors and the translates of G have algebraic Taylor coefficients. Finite Fourier convention · Indicator germ expansion.

Reading the normalized real coefficients. Use the orthonormal definite frames on the source and target. Write z_τ for the real-place complex matrix coordinate in these frames and $\xi_R = |\tau(\xi_0)| > 0$ for the host form scale, which is constant on the CM type. With the prescribed additive sign the Gaussian is $g(z) = \exp(-\pi\xi_R\sum_\tau \|z_\tau\|^2)$. In a single entry of z_τ, translation by $t\alpha$, $t \in \mathbb{R}$, acts in the Weyl model with derivative of its phase equal to $2\pi i\xi_R\,\mathrm{Im}(\bar{\alpha}z)$. The derivative of the Gaussian contributes $-2\pi\xi_R\,\mathrm{Re}(\bar{\alpha}z)$. The sum is $-2\pi\xi_R\alpha\bar{z}$. Consequently the combination of the directions $1, i$, with coefficients $1, -i$, acts on $g(z)H(z)$, after division by $2\sqrt{2\pi\xi_R}$, as

$$g(z)(2\pi\xi_R)^{-1/2}\big(\partial_z - 2\pi\xi_R\bar{z}\big)H(z).$$

The holomorphic derivative commutes with multiplication by each $\bar{z}$. In a word of creation operators evaluated at zero every term involving one of those multiplication operators vanishes. This leaves the iterated holomorphic derivative with one factor of $(2\pi\xi_R)^{-1/2}$ for each operator. The value of the Gaussian itself at zero is 1.

With the zero finite-argument vacuum in the standard lattice model, $H(z) = c_{\mathrm{vac}} G(\delta z)$. Indeed rescaling by δ identifies the Gaussian-removed quasiperiodicity with that of the reference theta line of Section 2.6.8. Both nonzero vacua belong to that line. Finite Heisenberg translation and the normalized finite averages defining G_s preserve the common scalar c_{vac}. Thus the monomial of each multi-degree $(d_\tau)_\tau$ reads the derivative of G_s times

$$c_{\mathrm{vac}} \prod_\tau \left(\frac{\delta}{\sqrt{2\pi\xi_R}} \right)^{d_\tau}.$$

In the conjugate Fock convention this formula is conjugated. Division by the vacuum scalar and by these period powers, followed by multiplication by $\prod_\tau c_{D,\tau}^{d_\tau}$, is exactly the derivative normalization below. These factors must be used on the same real polynomial spaces and with the same finite vacuum identification as in the lifted coefficients. Directional Weyl derivative · Normalized creation at zero · CM coordinates and period power · Positive and conjugate Fock conventions.

For a multi-index α on the variables $Z_{\tau,i,j}$, the derivative used in the step is

$$J_s(\alpha) = \left(\prod_\tau c_{D,\tau}^{\sum_{i,j} \alpha_{\tau,i,j}} \right) \partial^\alpha G_s(0)$$

with the prescribed algebraic differential scalars $c_{D,\tau}$. The factorial converting a Taylor coefficient to a derivative is rational. To pass from algebraicity to the specified K, fix the modulus and the maximal degree for each word and step before choosing the final class data. Definite class finiteness gives finitely many self-dual lattice representatives in each rank. For each pair, every positive divisor of the modulus gives a finite Hom quotient and finitely many derivatives below that degree. Include these normalized values, the differential scalars and the embeddings of E in a finite extension of $\mathbb{Q}$. Transport to any representative in the genus is an algebraic linear substitution on the germs. Under this transport the determinant ideals change by the principal ideals of the rational frame changes, whose character values are algebraic products of their embeddings. The chain and product rules therefore express the transported derivatives in the same finite extension. In the family selection the modulus for a tuple and word is $c_A c_B$, where c_A is the product of the residue cardinalities on the selected row support to the power $m+1$ and $c_B = T_w c_A$. The local support set belongs to the finite list of subsets of the primes over the type and word primes and the conductors of the finite list of twists. The prescribed coefficient field contains these finite generators before the grid is realized. This gives membership of $J_s(\alpha)$ in $\iota_K(K)$ for the actual point computations. Normalized derivatives and degrees · Finite field through lattice transport · Actual moduli and field.

Here is the resulting coefficient calculation. Multiply the source coefficient polynomial $P(X)$ by the determinant powers $\prod_\tau \det(X_\tau)^{n_\tau}$, with n_τ the source determinant shift for the step. Substitute

$$X_{\tau,i,j} = \sum_{h=1}^{b} Z_{\tau,h,i} Y_{\tau,h,j}.$$

In the expanded polynomial replace Z^α by $\sum_s \Xi(s) J_s(\alpha)$. Multiply by the ratio of the inverse determinant character values of the source and target frames. The resulting polynomial in Y is evaluated with the target determinant shift divided out. The substitution comes from evaluating the source frame against the columns of the target frame via the Hom variable. Each monomial in the derivative variables has the prescribed degree at every embedding because the source polynomial has its prescribed homogeneous weight and the determinant shift adds an_τ. Consequently only the indicated jets enter. Their field of definition and the framed character readings give an operator over K. Polynomial substitution in the jet · Source and target determinant shifts · Normalization of the finite characters.

The determinant character readings may be checked on units and a uniformizer at each place of the support. At a type place they are the prescribed type characters on residue units and the fixed inverse reference and twist value on the uniformizer. At a twist place the reference is unramified and the twist is tame, so the residue-unit value is that of the inverse twist. The two rules agree with the value on every determinant by writing a local element as a uniformizer power times a unit. Off the support both characters are unramified, and the determinant contributes the inverse unramified value raised to the valuation of the local frame determinant. The factors at conjugate split places use inverse values because the characters are trivial on the ground-field ideles. For the spherical factors, in a self-dual lattice model with the indicated additive character, the Fourier transform of the lattice indicator is that same indicator. This follows by character orthogonality on the lattice and self-dual volume one. It also applies at the unramified inert places when both tests are in that self-dual model. Tame local readings · Determinant character and frame valuations.

2.6.10 The finite symbols at the coupled class points

Write $A = B = m + 1$ for the two depth exponents of the chosen tuple's local box and $M = m$ for its marking depth. At each selected split place P, set $q = q_P$, let $\mathcal{O}$ be the local ring and use q as uniformizer in the selected component. These primes have absolute residue degree one. The matrix of a Hom variable x in the two chosen frames is $U = (f')^{-1}xf$. At the primes of the word T_w use the indicator $\phi_T(x)$ that x maps the source local lattice into the target local lattice and that the target's marked residue covector pulls back to the source marking. These conditions are unchanged under a perturbation by T_w times the integral Hom lattice at these primes. Split places and their residue cardinalities · The framed matrix and the marked lattice indicator.

On an ordinary arrow $i < m - 2$, the symbol is $\phi_T(x)\prod_P \Xi_{i,P}(U)$. In $\Xi_{i,P}$, let a, b be the two definite ranks, and c', c the numbers of added source and target coordinates in the mixed model. Use zero-based indices. The enlarged matrix $X(U, \mathbf{t})$ has U in its first b rows and first a columns. A position (ℓ, s) is free if

$$s \geq a, \qquad (s - a = c) \text{ or } \quad (s > a \text{ and } (\ell < b \text{ or } \ell - b < s - a)).$$

The case $\ell - b < s - a$ is read for $\ell \geq b$. At a position in an added column which is fixed, put one when $\ell = b + s - a$ and zero elsewhere. The remaining positions below U are zero. Write $\mathcal{F}$ for the set of free positions, and $\mathfrak{u}(X) = \sum_{j=0}^{c-1} X_{b+j,a+j+1}$ for the entry sum in the unipotent character. With the free entries $\mathbf{t}$ ranging over $(q^{-A}\mathcal{O}/q^{B}\mathcal{O})^{\mathcal{F}}$, the normalized local value is

$$\Xi_{i,P}(U) = \mathrm{vol}(C_{\mathrm{be}} \cap N_1 \cap \mathrm{Stab}(\Lambda_Y))^{-1} q^{-B|\mathcal{F}|} \sum_{\mathbf{t}} (\omega\Phi_i)\big(X(U,\mathbf{t})\big)\,\psi_P\big(-\mathfrak{u}(X(U,\mathbf{t}))\big).$$

Here $\omega\Phi_i$ is the finite matrix symbol transported by the two prescribed extended-frame automorphisms using their splitting characters and determinant powers. The local volume is for the indicated target unipotent subgroup in the same measure as the integral defining the coefficient. If the supported mixed-model kernel in these coordinates is this integrand and is constant on its free cosets modulo q^B, partitioning the integral into these cosets gives the displayed factor $q^{-B|\mathcal{F}|}$. This calculation isolates the support and frame checks needed for the ordinary integrals. Extended frames and the free entries · Normalization of the ordinary local symbol.

On the penultimate arrow the extended source has one added column and the target has none. A coupled point includes a marked coset Y on the target of this arrow. In its frame at P the coset is $y' + q^{M_P}\mathcal{O}^b$, where $M_P = 0$ at a type prime and $M_P = M$ at a twist prime. At a type prime $y' \in q^{-1}\mathcal{O}^b$, and at a twist prime $y' \in q^{-M}\mathcal{O}^b$. The prescribed symbol is $\phi_T(x)$ times the product of $(\omega\Phi_{m-2})(U \mid y')$, with y' the added column. At a type prime the integral extended source basis scales that column by q. Thus a change of y' by an integral vector changes the matrix in that basis by q times an integral matrix. The source frame automorphism preserves the rescaled lattice, so the change after applying it is still divisible by q in

the rescaled entries. At a twist prime the source frame matrix has denominator exponent at most $m-2$, so a change by q^M times an integral vector changes the framed extended matrix by a matrix divisible by q. The symbol conditions read the integrality and the residues of these entries, and the frame multipliers are unchanged. This gives independence of the representative of Y. The same argument for changes of U by q^B times an integral matrix gives the modulus in the Hom variable. The additional column · The framed automorphism bounds · The marked cosets and penultimate symbol · The box and marking exponents.

On the last arrow both extended column counts are zero. Now $a = k + (m-1)(t+1) = r + (m-1)t$ and $b = k + m(t+1)$. Write U_{up} for the first $b-1$ rows of the framed matrix and ℓ_U for its last row. At the top point over a row point, the frame on the added line uses the scale π_{T_w}. Thus ℓ_U in the unscaled last-line coordinate includes the factor $\pi_{T_w}^{-1}$ when read in the coupled row frame. The prefix at this same point is read using the row frame. Let $\Phi_{\mathrm{top}}(U_{\mathrm{up}})$ be the corresponding ordinary or sharp type symbol, or the tame twist symbol at a twist prime. The local factor after the last partial Fourier transform is

$$\Phi_{\mathrm{top}}(U_{\mathrm{up}})\, q^{-M_P a}\, \mathbf{1}_{\ell_U \in q^{-M_P}(\mathcal{O}^a)^\vee}\, \psi_P(-\ell_U(y')),$$

where $(\mathcal{O}^a)^\vee$ is the integral linear dual and y' in the prescribed marked coset of the source. To verify the factor, integrate $\psi_P(-\ell_U(z))$ over $y' + q^{M_P}\mathcal{O}^a$. Translate by y'. With $\mathrm{vol}(\mathcal{O}) = 1$ and additive character of conductor $\mathcal{O}$, character orthogonality gives zero unless the dual-lattice condition holds, and otherwise gives the indicated volume times the phase. It also shows that the phase is independent of the representative y'.

The prefix on this arrow has integral support. In the type symbols all columns are among the first $r + (m-1)t = a$, so the defining entry conditions impose integrality on each prefix column. With no added columns its entry weights are unscaled. The tame twist symbol also imposes integrality on each entry. The prefix is constant under adding q^B times an integral matrix since its other conditions and characters read only residues. For the last row, $M \le A, B$, the type bound $qy' \in \mathcal{O}^a$ and $B \ge 1$ show that adding q^B times an integral covector to ℓ_U changes neither the dual-lattice condition nor the phase. On its support the phase argument has denominator exponent at most 1 at a type place and $2M$ at a twist place. These are within the prescribed character depths. Thus the terminal symbol has the required box and gives values in K. The volume factor is a unit at places of K over T_w since the primes in this row support lie away from T_w. The local formula agrees with the symbol at a top column point lifted from a row point. To check this, compare its top frame f_{top} with the block frame whose prefix is the row frame and whose added vector is π_{T_w} times the prescribed last basis vector. At a twist prime that block frame is the chosen top frame. At a type prime both frames are integral bases of the same lifted lattice and give the same prescribed residue frame of its decoration. Their transition therefore lies in $\mathrm{GL}_b(\mathcal{O})$ and is the identity modulo q. The integrality of the entire matrix, which is necessary for nonzero symbol at a type prime, is invariant under this transition. On an integral matrix its prefix residues are unchanged. Its last covector changes by an element of $q(\mathcal{O}^a)^\vee$, whose pairing with $y' \in q^{-1}\mathcal{O}^a$ is integral. Thus the terminal Fourier phase is unchanged too. The decorated top lattice at the primes over T_w is the same lattice in the two descriptions, so the factor $\phi_T(x)$ agrees. This proves the equality of the two terminal symbols at the lifted points, including their marks. The tame prefix · The type prefixes · The marked-coset transform · Same lattice and residue frame at the type places · The same symbol at a lifted row point · Support and the depth bound at the last step.

The character used to move the arithmetic vacuum can be changed on the class kernels if its restriction to norm-one fractional ideals agrees with the original character. In fact, for two self-dual lattices in the same Hermitian space, the relative determinant ideal $\mathfrak{d}$ satisfies $\mathfrak{d}\bar{\mathfrak{d}} = \mathcal{O}_E$. In local bases, duality conjugates and inverts the determinant ratio at the two conjugate places, which proves the ideal identity. Both relative determinant ideals of the source and target self-dual lattices therefore have the same character readings before and after the change. The vacuum scalar and the determinant normalization agree, so the germs G_s, the derivatives and the three types of finite kernels agree in the same frames. Norm-one determinant ideals · Compatibility of the three kernels with the character change.

2.6.11 The normalized step entry

Fix two class points p, p' in one of the selected definite steps, of ranks a, b. Let $\Xi_{p,p'}$ be the finite symbol just described in their frames. This means the term-level average in an ordinary step, the marked-column restriction in the penultimate step, and the marked-covector Fourier transform on the last step. At the lifted row points the last symbol is the coupled row symbol by the frame calculation above. All three use ϕ_T on the decorated T_w-lattice and the indicated tests at the row-support places.

Consider the normalized mixed-model theta kernel $\mathcal{T}_{p,p'}(z)$ before real differentiation. The finite support gives its decomposition into indicators of the classes of $c_A^{-1}\mathcal{L}/c_B\mathcal{L}$. The self-dual vacuum takes care of the spherical places and the Chinese remainder theorem gives the representative sum at the remaining places. In a local matrix space the frame change $U \mapsto f_{p'} U f_p^{-1}$ has Jacobian

$$|\det f_{p'}|^{a}\,|\det f_p|^{-b}.$$

Thus its half-density factor is the inverse square root of the local lattice-volume ratio. The character expansion of each indicator carries the square root of the same volume, so the two factors cancel. The remaining average has the denominator $|A_1|$ and the two Heisenberg translations with quadratic phase displayed in the formula for G_s. This gives the normalized finite expansion, with the vacuum determinant included,

$$\mathcal{T}_{p,p'}(z) = \chi_F(N, N') \sum_s \iota_K\big(\Xi_{p,p'}(s)\big)\, G_s(z).$$

Here G_s uses the same normalization for all cosets, and finite differentiation commutes with this sum. The normalization can also be checked by applying the two finite character projectors to a single coset indicator. The first projector uses the annihilator of $c_B\mathcal{L}$, so its average selects precisely the given coset. The second transports to the paired half of the period lattice. Multiplication of its two translations yields $e^{\pi i E_H(s,u)}$. Linearity then gives the expansion for $\Xi_{p,p'}$.

Let $M_j(p)$ denote the inverse value of the prescribed reference-and-twist character $\mu_j\xi$ on the determinant idèle of the framed class p. The unramified valuation and the ramified frame values in the preceding calculation give $M_j(p)$ in $K^\times$. At the conjugate half of a split place the local determinant valuation is opposite; triviality on ground-field idèles gives the corresponding inverse character. At an inert place the determinant has valuation zero. For the oscillator character on the relative determinant ideals use χ_ψ. The two complex class multipliers are

$$\rho_j(p) = \iota_K(M_j(p))\,\chi_\psi(\det N)^{-b/2},$$
$$\rho'_j(p') = \iota_K(M_j(p'))^{-1}\,\chi_\psi(\det N')^{a/2}.$$

Their product is $\iota_K(M_j(p)/M_j(p'))\,\chi_F(N, N')^{-1}$. This cancels the character of the moved vacuum before taking the normalized derivatives. Frame volume, half-density, and unit-volume kernel · Indicator decomposition on the test box · The two actual class multipliers · Cancellation with the moved-vacuum character.

Write $\mathscr{J}_{p,p'}$ for the K-linear coefficient contraction taking

$$Z^\alpha \longmapsto \sum_s \Xi_{p,p'}(s)\,\iota_K^{-1}\big(J_s(\alpha)\big)$$

on the degrees selected by the input weight. If P represents the input, write $P^\sharp(X) = P(X)\prod_\tau \det(X_\tau)^{n_\tau}$. In the fixed complex frames the normalized analytic entry evaluated at Y is

$$\iota_K\left[\frac{M_j(p)}{M_j(p')}\,\mathscr{J}_{p,p'}\big(P^\sharp(Z^{\mathsf{T}}Y)\big)\right](Y)\prod_\tau \det(Y_\tau)^{-s_{\mathrm{div},\tau}}. \tag{50}$$

Here the embedding is applied to the coefficients of the polynomial before evaluation at Y, and $s_{\mathrm{div},\tau}$ is the prescribed output determinant exponent. Indeed the real operator replaces each derivative variable by the

Gaussian-normalized creation operator. At zero its evaluation is $J_s(\alpha)$, by the calculation of the creation operator and the prescribed differential scale above. The source sharp factor increases the determinant exponent from the input coefficient convention to n_τ. The exponent of the determinant in the target coefficient model is $-s_{\mathrm{div},\tau}$, so removing the output shift gives exactly that displayed evaluation. The exponents and degrees are integral on these ranks. In particular the output polynomial has the required covariance over K: substitution in the equivariance identity for the complex map gives polynomial identities with coefficients in K, and the embedding is injective. Input sharp and output determinant exponents · Normalized derivatives in the coefficient contraction · Application to each actual stage representative.

Formula (50) therefore gives the entry on every selected pair with the specified tests and frame placements. Applying it at the penultimate coupled points and restricting it at the lifted top points proves the corresponding entry identities there. The local orbital realization of the opposite path in (15) is a separate calculation on its indicated term levels. For the definite top pairing, only the normalized test symbols and placements of the column chain enter the entry calculation.

2.6.12 Descent of the period

Write N_w for the prescribed column lattice before division. Its scalar extension $N_{w,K}$ has complexification $\mathcal{M}_w$. To read the top pairing, evaluate the normalized step kernel on each pair of the chosen definite class representatives, with the framed coefficient vectors. Denote its prescribed arithmetic value by $K_{j,K}(c, c')$ in these coordinates. Equation (50) identifies it under ι_K with the step in (12), including the source and target determinant factors. At the final two steps the entries are read at the coupled points, and at the final step only at the images of the top row points.

The corresponding class measures are those of the term levels. The step entries at the class representatives · The coupled and top entries.

Given these entries, start with the prescribed values $C_0(c) = \phi(c)$ after transport to the tuple's stage zero. Successively set

$$C_{j+1}(c') = \sum_{c} |\operatorname{Stab}_j(c)|^{-1} K_{j,K}(c, c') C_j(c).$$

If one refines a class into orbits at an intermediate level, the sum of the reciprocal stabilizer orders over a finite orbit decomposition is the orbit size divided by the acting group order. Hence a refinement by an index h multiplies the weighted sum of an invariant value by h. This accounts for the level volumes when using the coupled intermediate points. At a row class a, let $\widetilde{a}$ denote the top column point over a, with its prescribed frame transport and coefficient projection included in res_a. Pairing against the class value $x(a)$ gives

$$m_Y(\boldsymbol{\mu}, m) \sum_{a} |\operatorname{Stab}_Y(a)|^{-1} \big(\mathrm{res}_a\, C_m(\widetilde{a}), x(a)\big)_{\mathrm{alg}}.$$

The algebraic pairing is invariant under changes of representatives. Indeed the column and row carry inverse finite characters and the branching projection pairs dual coefficient actions. The scalar $m_Y(\boldsymbol{\mu}, m)$ is a product of indices. The class unfolding at each step and the entry identity imply by induction that the displayed sum maps to the actual top pairing under ι_K. All its operations are over K. The same entry identity enables the coefficient valuation for the actual period in Section 4.

Class sums at the coupled points · Invariant pairing and the top class sum. Top period entry at the chosen points.

Using the entry identity and (41), take the base K-valued period on any rational preimage of $F \in R_K$ under $P_{\boldsymbol{\mu}_0}$. The projected base map is onto, and the simultaneous theta–Bessel identity gives zero on its kernel. The embedded value determines the value in K, so the same kernel vanishes over K. This defines uniquely

$$B_{K,w} : N_{w,K} \longrightarrow R_K^\vee, \qquad B_w(\phi_{\mathbb{C}}, F_{\mathbb{C}}) = \iota_K\big(B_{K,w}(\phi)(F)\big).$$

The equality extends by bilinearity to the full complexifications.

For each finite place v of K, choose the tuple $\boldsymbol{\mu}(v)$ prescribed by the family. The choice is the same at conjugate places. At an untwisted place the base tuple has trivial $\xi_{H,\boldsymbol{\mu}_0}$. At a twisted place of residue characteristic p, the character $\xi_{H,\boldsymbol{\mu}(v)}$ is the product of the specified twist unramified above p and the auxiliary character supported at p_{de}, using p'_{de} if $p = p_{\mathrm{de}}$. The rational primes $p_{\mathrm{de}}, p'_{\mathrm{de}}$ are distinct. In both cases the character is unramified above p. Every prime Q in the tuple's $Q(\boldsymbol{\mu}(v))$ is a ramification place of this character. A prime above p cannot also lie above the chosen auxiliary prime, since these two rational primes generate the unit ideal. At the base the trivial character gives $Q(\boldsymbol{\mu}_0) = \emptyset$. In particular

$$Q(\boldsymbol{\mu}(v)) \cap \{Q \,:\, Q \mid p\} = \emptyset. \tag{51}$$

The selected finite-place tuple and the two distinct auxiliary primes. In the compatible rational top models, evaluate (37) on a pair of coordinates over K where the period is nonzero. The quotient of the top evaluation by the evaluation on the descended row gives $\zeta_{w,\boldsymbol{\mu}} \in K^\times$ with image $d_{w,\boldsymbol{\mu}_0}/d_{w,\boldsymbol{\mu}}$ at ι_K. If all evaluations on these coordinates vanish, both sides vanish on the complex span and we may choose scalar 1. In all cases the base scalar is 1 and

$$b_{w,\boldsymbol{\mu}}(\phi, F_m) = \iota_K(\zeta_{w,\boldsymbol{\mu}})\, B_w(\phi, \iota_{\boldsymbol{\mu}} F_m).$$

Use the support conclusion of the simultaneous theta–Bessel comparison at this particular tuple. Its relevant power is a nonzero rational number supported on its twisting primes. Applying v_v to the power identity gives $n\, v_v(\zeta_{w,\boldsymbol{\mu}(v)}) = 0$ for some $n > 0$. The possible scalar 1 also has valuation zero. Both conjugate places have the same residue characteristic, so the factor is a unit at both. Consequently the valuation of the tested top pairing equals the valuation of $B_{K,w}$ on the descended row.

Use the sharp level of the word periods. All integral test rows below lie in the good part $\mathcal{V}^0_{\boldsymbol{\mu}}$ on which the simultaneous identity is stated. For each tuple let $A_{\boldsymbol{\mu},K}$ be the rational top class space at that character level, $\mathscr{L}_{\boldsymbol{\mu}}$ its prescribed value lattice and $E_{\boldsymbol{\mu}}$ the inclusion matrix of its good coordinates. Thus $E_{\boldsymbol{\mu}}$ reads the same coefficients as $j_{\boldsymbol{\mu}}$ in (41). The lattice of integral top test rows in these coordinates is

$$M_{\boldsymbol{\mu}} = \{x \,:\, E_{\boldsymbol{\mu}} x \in \mathscr{L}_{\boldsymbol{\mu}}\}.$$

In the full principal class coordinates the identical condition uses the inclusion of the character level followed by $E_{\boldsymbol{\mu}}$ in the displayed read. A vector passes the condition precisely when its class values, in the prescribed frames, are integral. Multiplying a rational basis by a common denominator shows that the cut lattice is full.

Use the basis of the adjoint image and the projection $P_{\boldsymbol{\mu}}$ in (41). The lattice of projected coordinates is the image of $M_{\boldsymbol{\mu}}$ under $P_{\boldsymbol{\mu}}$. This claim uses the integral vectors before projection. Requiring their projected representatives to have integral class values would give a different condition at primes dividing projection denominators. In the prescribed base row coordinates $D_{\boldsymbol{\mu}}$ sends this image to

$$L^{\mathrm{row}}_{\boldsymbol{\mu}} = D_{\boldsymbol{\mu}} P_{\boldsymbol{\mu}} M_{\boldsymbol{\mu}} \subset R_K.$$

The kernel of the projection is killed by the actual descent and by the word period. Surjectivity on the rational class coordinates makes $L^{\mathrm{row}}_{\boldsymbol{\mu}}$ full. For each v the specified row gluing uses the image lattice for $\boldsymbol{\mu} = \boldsymbol{\mu}(v)$. Localization commutes with this matrix image, since each localized vector has the form x/d for a scalar d outside v. In particular the prescribed local generators are the actual descended integral rows:

$$R_v = (L^{\mathrm{row}}_{\boldsymbol{\mu}(v)})_v = \mathcal{O}_{K,v}\, D_{\boldsymbol{\mu}(v)} P_{\boldsymbol{\mu}(v)} M_{\boldsymbol{\mu}(v)}. \tag{52}$$

The sharp value lattice and the actual downward map · The cut in class coordinates · The projected-coordinate image and the prescribed localizations.

These localizations glue to one lattice R. To prove this, fix the base image lattice L_0. The tuple list is finite, so there is an integer $h > 0$ such that all the $L^{\mathrm{row}}_{\boldsymbol{\mu}}$ lie between hL_0 and $h^{-1}L_0$. Outside the

primes dividing h, they have the same localization as L_0. The finite $\mathcal{O}_K$-module $h^{-1}L_0/hL_0$ is the direct sum of its localizations at the primes dividing h, by the Chinese remainder theorem. In each summand select the image of the required local lattice. Their sum pulls back to the full lattice R with exactly these localizations. Membership in a full lattice is tested at all the localizations: if x does not belong, the proper ideal of $a \in \mathcal{O}_K$ with ax in the lattice lies in a maximal ideal where localization still excludes x. This also proves uniqueness.

Let $C = R^\vee = \mathrm{Hom}_{\mathcal{O}_K}(R, \mathcal{O}_K)$, viewed in $R_K^\vee$. Its localization is the dual of R_ν, which can be checked in a basis over the discrete valuation ring. A valuation bound on the top periods extends to evaluations on R_ν by the unit identity above and the ultrametric inequality, since the actual descended integral rows generate R_ν. When the local estimate of Section 4 puts the localized image in $\mathfrak{p}_\nu^{m_{w\nu}} C_\nu$, multiplication by the inverse prime power cancels it. Intersecting over all finite places, with $I_w = \prod_\nu \mathfrak{p}_\nu^{-m_{w\nu}}$ for the prescribed rounded gains, gives

$$B_{K,w}(N_w') \subseteq C, \qquad N_w' = I_w N_w \quad (w \in W). \tag{53}$$

The row lattice in (52) uses the sharp cut. At a twisted place the arithmetic covering in (175), for $1 \le j \le m$, uses the full auxiliary class space $M^L_{\boldsymbol{\mu}(\nu),j}$ of the base level above. Its source is the one-step class model of (40). The localized Hecke algebra in Section 6 uses the normalization of its actual split operators. Its maximal ideals index the chosen local records. Applying their primary projections to a spanning set of the twisted class lattice gives the projected coordinate generators of each local module. Their spans over the local fraction field are the primary summands, whose projections add to the identity. This proves the covering and the joint detection on the prescribed auxiliary space. The finite-place row lattice, the avoidance in (51), and this covering can therefore all be used with the same function $\nu \mapsto \boldsymbol{\mu}(\nu)$. Local record indices and projected lattices · All residual ideals and Covering on the normalized class space.

2.7 The matrices in common row coordinates

Choose an orthonormal basis (F_i) of E_{row}. The vectors $s(VF_i)$ form one orthonormal basis of the distinguished space $H_{\mathrm{row}} = s(VE_{\mathrm{row}})$. Complete it to an orthonormal basis of $\mathcal{H}$ and let Π be the orthogonal projection to these distinguished coordinates. In each orthonormal basis $(\phi_{w,j})$ of $\mathcal{M}_w$, the matrix J_w of $J_w U_w$ into an orthonormal basis of H_w^{full} has orthonormal columns. Write P_w for the corresponding matrix of $\mathcal{P}_w$. By (36) the i,j entry of $\Pi P_w J_w$ is $a_w^{-1} B_w(\phi_{w,j}, F_i)$. If Z is the coordinate matrix of (F_i) in the complexified rational row basis and X_w is that of $(\phi_{w,j})$ in the complexified rational column basis, expansion gives

$$\Pi P_w J_w = Z^{\mathsf{T}}\, \iota_K([B_{K,w}])\, (a_w^{-1} X_w). \tag{54}$$

The rational matrix uses the dual of the rational row basis on its target. The transpose is ordinary because the pairing is linear in the row vector. The invertible row change Z is common to the words, and X_w is invertible for each word. Choosing the conjugate representative of the infinite place conjugates the whole identity.

The common row change is needed when taking sums of images. For example, take $E_{\mathrm{row}} = \mathcal{H} = \mathbb{C}^2$ with s coordinatewise conjugation and two one-dimensional columns given by $z \mapsto (z, 0)$ and $z \mapsto (0, z)$. With common row coordinates the resulting functionals span the dual. Using the identity on the first row input and the coordinate swap on the second makes both functionals multiples of $F \mapsto F_1$. With (54), all blocks use the same rows. Section 3 estimates the full column pairings in $\mathcal{H}$ and then restricts to these common row coordinates. Its Gram normalization is the positive volume factor $c_0(w)$ of (21). The nonzero scalar a_w comes from the period comparison.

AI-proposed candidate (simultaneous matching of theta and Bessel periods). Fix a selected translate with the towers, levels and tests above and with bottom rank difference $2r_S + 1$, where $r_S \ge 1$. There

is one finite central placement zt_0 and one conjugate-linear row isometry $s \circ V$ for all words. For each w there are a unitary U_w and $a_w \in \mathbb{C}^\times$, independent of the tuple and the vectors, such that

$$b_{w,\boldsymbol{\mu}}(\phi, F_m) = \frac{d_{w,\boldsymbol{\mu}_0}}{d_{w,\boldsymbol{\mu}}}\, a_w \big\langle s(V\iota_{\boldsymbol{\mu}} F_m), \mathcal{P}_w J_w U_w \phi \big\rangle_{\mathcal{H}}$$

for every allowed tuple $\boldsymbol{\mu}$, ramified column $\phi \in \mathcal{M}_w$, and projected top row $F_m \in \mathcal{V}^0_{\boldsymbol{\mu}}$. A positive power of the displayed ratio is rational with prime support contained in the twisting primes of $\boldsymbol{\mu}$. In particular the identity matches the prescribed coefficient-field tests and their actual downward images at the base tuple and every finite-place tuple $\boldsymbol{\mu}(v)$.

The stated identity gives kernel vanishing on the projected base rows. Together with the arithmetic model and normalized comparisons in the preceding subsection it supplies $B_{K,w}$, the common block identity (54), and the local units used for (53).

For a definite unitary group and its codimension-one unitary subgroup, Michael Harris constructs a rational bilinear period on the coherent-cohomology models with compatible coefficient projection. The pairing takes values in the compositum of the two coefficient fields and is Galois-compatible with Tamagawa measure. [27, §2.5, Lemma 2.5.4] This gives the rationality principle for the top pairing once the lifts, coefficient models and measures are identified. The finite placement, identification with the prescribed Bessel vectors, and support of the twisting ratios require the comparisons across the two towers.

For comparison, fix $V_0 \subset V_1$ over E/F with complement given by hyperbolic planes and the norm line of E/F. Write P for the absolutely convergent integral of a cusp form over the Bessel subgroup quotient against its generic flag character, trivial on $U(V_0)$. For a cuspidal automorphic representation σ of $U(V_0) \times U(V_1)$ tempered at every place with discrete Hermitian weak base change Π, Beuzart-Plessis and Chaudouard give, for each nonzero factorizable $f = \bigotimes'_v f_v \in \sigma$,

$$\frac{|P(f)|^2}{\|f\|^2_{\mathrm{Pet}}} = |S_\Pi|^{-1} \mathcal{L}(\tfrac{1}{2}, \sigma) \prod_v \frac{P^\natural_v(f_v, f_v)}{\|f_v\|^2_v}.$$

Here S_Π is the component group of Π, the Petersson norm uses Tamagawa measure, and the invariant local norms are used in the matrix coefficients as well. The local subgroup measures multiply to Tamagawa measure. The form $P^\natural_v$ is the local Bessel functional on matrix coefficients divided by $\mathcal{L}_v(\frac{1}{2})$. Writing $d_i = \dim V_i$, the local factor is $\mathcal{L}_v(s) = L_v(s, \Pi_v) \prod_{j=1}^{d_1} L_v(s + j - \frac{1}{2}, \eta^j_{E/F,v}) / L_v(s + \frac{1}{2}, \Pi_v, \mathrm{As}')$, with the Rankin–Selberg numerator and Asai signs $(-1)^{d_0}, (-1)^{d_1}$ in As'. The global factor is continued from the product of the local factors. [26, §§1.3.1–1.3.2, §1.3.5 and Theorem 1.3.6.1]

The seesaw identity supplies the interchange of compatible theta kernels. The squared-period formula gives absolute squares on factorizable vectors of a fixed σ. The displayed matching identity gives a bilinear comparison in common coordinates across constituents and tuples, with the relative normalizations on the prescribed tests.

3 Relative traces and the span of the periods

Fix the family selected in Section 6.1, a scale x beyond its thresholds, and the record at one of the tied translates. Write W for its finite set of words. The two full Gram estimates and the multiplicity estimate are the inputs to this section's conditional argument. We use them to control the span of the periods after restricting each column space and projecting onto the common row space. The orbital calculations give context for the Gram estimates and the same-word estimate also explains the construction of the nonzero bottom reference pairing in Section 2.

3.1 Unfolding the full Gram operators

For each $w \in W$, use the Poincaré-series operator $\mathcal{P}_w : H^{\mathrm{full}}_w \to \mathcal{H}$ in (21). The ambient level and archimedean row type fix the same finite-dimensional Petersson space $\mathcal{H}$ for all words. Let $d_w = \dim H^{\mathrm{full}}_w$,

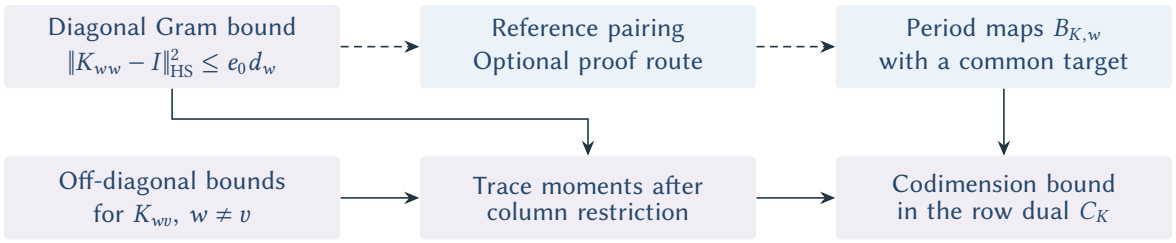

Figure 2. The dashed route gives the reference pairing for an optional proof of the granted simultaneous period comparison. There the diagonal and ramified rank bounds give $\dim\ker \mathcal{P}_w \le e_0 d_w < r_w$, and full row evaluation pairs with a nonzero image of a ramified column. To descend the granted comparison over K, use the compatible base pairing and the surjective projected row map over K. With the word and dimension estimates, both Gram bounds control the trace moments for isometric column inclusions. Using $J_w = J_w U_w$ matches those moments to the period maps. The common projection and the distinguished-row dimension bound give the corank over K.

$D = \dim \mathcal{H}$, and $R_{\mathrm{full}} = \sum_w d_w$. The coefficient representation π_w is the tensor product of the residual compact $U(k)$-types at the real places. Write $v_N(w)$ for the volume of the unipotent fundamental domain multiplied by the ambient level volume, and $v_H(w)$ for the volume of the real Bessel stabilizer. Then $c_0(w) = v_N(w)v_H(w)/\dim \pi_w$, with the measures of (21). Set $L_w = c_0(w)^{-1/2}\mathcal{P}_w$ and $K_{wv} = L_w^* L_v : H_v^{\mathrm{full}} \to H_w^{\mathrm{full}}$. The inner product is linear in its second argument.

The orbital formulas in the next paragraphs describe the operators in the two Gram inputs. Their quantitative estimates are used in the forms stated at the same-word estimate and the mixed estimate.

Put $G_0 = U(Y_0)$, where Y_0 has rank $r = k + m - 1$ over the CM field E, and write $\mathcal{S}_w$ for the Bessel subgroup attached to the flag and its distinguished vector for w. For $w, v \in W$, $\phi \in H_w^{\mathrm{full}}$ and $\psi \in H_v^{\mathrm{full}}$, the corresponding Poincaré series periodize the seeds $\Phi_{w,\phi}$ and $\Phi_{v,\psi}$, which are invariant under their respective left rational Bessel subgroup actions, over $\mathcal{S}_w(F)\backslash G_0(F)$ and $\mathcal{S}_v(F)\backslash G_0(F)$, respectively. Let Ω_{wv} be representatives of the rational double cosets $\mathcal{S}_w(F)\backslash G_0(F)/\mathcal{S}_v(F)$, and set $\mathcal{S}_\gamma(F) = \mathcal{S}_w(F) \cap \gamma \mathcal{S}_v(F)\gamma^{-1}$. At the chosen level Alpöge proves absolute convergence using the level conditions, neatness and real Whittaker decay. The sum of the orbital integrals with the product of seeds replaced by its absolute value is finite. The double-coset unfolding therefore gives the normalized identity

$$\langle \phi, K_{wv}\psi \rangle = \frac{1}{\sqrt{c_0(w)c_0(v)}} \sum_{\gamma \in \Omega_{wv}} \int_{\mathcal{S}_\gamma(F)\backslash G_0(\mathbf{A}_F)} \overline{\Phi_{w,\phi}(g)}\Phi_{v,\psi}(\gamma^{-1}g)\, dg.$$

All quotients carry the measures induced by the adelic Haar measure in the Petersson product. After unfolding the w periodization, the v sum is indexed by $\mathcal{S}_v(F)\backslash G_0(F)$. The cosets in the double coset of γ^{-1} are represented by $\gamma^{-1}a$ for $a \in \mathcal{S}_\gamma(F)\backslash \mathcal{S}_w(F)$. Integrating their translates joins the corresponding fundamental domains. Getz–Hahn use the same stabilizer reindexing for the geometric side of their relative trace formula with compactly supported tests. [40, §4C, (4C.8) and its unfolding, p. 114] Alpöge identifies the unfolded coefficients using compatible flag and lattice representatives. Separating the lattices and the finite level conditions gives $K_{wv} = \sum_t A_t$. In each coefficient block allowed by the level, A_t is the archimedean orbital operator multiplied by the finite character, the level volume and $(c_0(w)c_0(v))^{-1/2}$. The archimedean operator integrates $W_w(g) \otimes W_v(\gamma^{-1}g)^*$, where W_w, W_v are the coefficient vectors of the Poincaré seeds.

Each decorated flag comprises an ordered isotropic frame of length $r_S = m/2 - 1$ and a vector orthogonal to that frame of norm $-T_w$ or $-T_v$. Here m is even and $m \ge 4$ on the chosen large-rank family. Set $a = r_S + 1$, and write Z for the a-by-a matrix of cross pairings of the augmented frames. With e the last coordinate vector, their joint Gram matrix and its determinant are

$$M = \begin{pmatrix} -T_w ee^* & Z \\ Z^* & -T_v ee^* \end{pmatrix}, \qquad \det M = (-1)^a G, \qquad G = |\det Z|^2 - T_w T_v |\det Z_{a-1}|^2,$$

where Z_{a-1} is the leading $(a-1)$-by-$(a-1)$ submatrix and $|z|^2 = z\bar{z}$ is the CM norm. The determinant follows by swapping the block columns, taking the Schur complement of Z and using the cofactor

$e^{*}\operatorname{adj}(Z)e$. Polynomiality covers singular Z. When $G \neq 0$ the augmented frames span a nondegenerate space of dimension $2a = m$, since m is even. Its orthogonal complement has dimension $k - 1$.

Separate the contribution of the identity double coset, the other $G = 0$ terms, and the generic and nongeneric terms with $G \neq 0$. Write g_0, g_1 for the two designated auxiliary rational primes, which split completely in E. At places above these primes the finite level uses specific Whittaker characters. A nondegenerate pair is generic if its lattice in the joint span is generated locally by the two augmented frames, with the adjusted second frame at the places above the auxiliary primes. At the identity, the level calculation for $w = v$ matches columns and gives $v_N(w)$. The average of the unit coefficient vector's rank-one projector over the real stabilizer commutes with π_w, and its trace is $v_H(w)$. Schur's lemma gives $(v_H(w)/\dim \pi_w)I$, so the identity contributes I to K_{ww}. For distinct words the calculation at the identity gives zero by averaging between inequivalent coefficient types.

The unipotent average on the stabilizer quotient vanishes when the characters disagree on the intersection after transport by γ. For a compatible $G = 0$ orbit the geometric stabilizer acts by the full residual $U(k)$ on both coefficient spaces at each real place, up to fixed unitary conjugation. Schur's lemma forces the resulting intertwiner to vanish for $w \neq v$ because the chosen highest weights differ at every real factor. This is the release's vanishing of the degenerate terms for distinct words. The mixed block consists of the generic and nongeneric sums with $G \neq 0$ in the same weights and measures.

3.2 The single-word estimate

Put $f = [F : \mathbf{Q}]$ and $b = \exp(x^{\epsilon})$. The reference dimension is

$$D_{\star} = \frac{\mathfrak{d} V_G}{v^{\dagger}} > 0,$$

where $\mathfrak{d}$ is the product of the target formal degrees, V_G is the ambient adelic volume and $v^{\dagger}$ is the deep ambient level volume. These are the measures of the Bessel realization. For the common factor $\Gamma = \exp(C_{\Gamma} m^5 f \log b)$, with $C_{\Gamma} > 0$ fixed on the family, the column-size output of the selection in Section 6.1 is

$$C_0^{-1} \leq \frac{d_w \Gamma^2}{D_{\star}} \leq C_0 \qquad (w \in W), \qquad C_0 \geq 1. \tag{55}$$

The selected constant is $C_0 = 3$. The balanced dimensions here count the full level-zero columns before cutting by the word-prime characters. Since $D_{\star} > 0$, the lower bound makes every d_w positive. The selected rank inequalities give $m^4 \leq f \leq k \leq x^5$ and $k \geq 8 \log \Gamma$ at all sufficiently large scales.

The level index explains the balance. Let $d_w^{(0)}$ be the compact column dimension before marking a nonzero covector at each place above T_w, with the two auxiliary levels and the word's real type retained. The class cover has degree $(T_w^k - 1)^f$. Each local factor is the index of the stabilizer of a nonzero covector in $\mathrm{GL}_k(\mathbf{F}_{T_w})$. The definite class stabilizers are trivial by the two auxiliary reductions, as below, so this index multiplies the number of classes and hence the column dimension. The level index and compact rank. Write

$$Y_w = \left(\frac{D_{\star}}{d_w^{(0)} \Gamma^2} \right)^{1/(fk)}.$$

The word-prime interval of the chosen record is $Y_w \leq T_w \leq Y_w(1 + 1/(fk))$. The word-prime interval. Thus the exact ratio $d_w \Gamma^2 / D_{\star}$ is $(T_w/Y_w)^{fk}(1 - T_w^{-k})^f$. Its upper bound is $\exp(1)$ and its lower bound is $1 - f T_w^{-k} \geq 1/2$ for large x. The count uses the same $D_{\star}$ as the row projector below.

With the contribution I of the identity removed, there are two sums in $K_{ww} - I$. Alpöge bounds the squared Hilbert–Schmidt norm of the remaining degenerate sum by $d_w \exp(-kf \log k)$ and that of the entire nondegenerate sum by $\Gamma d_w^2 / D_{\star}$. The first estimate uses absolute orbital amplitudes. The second combines the scalar-pair and error estimates for equal words, in their prescribed levels and measures. For the scalar pairs, the local Bessel integral from Section 2 enters in both word slots of the real orbital

integral used with Siegel–Weil. Alpöge combines this scalar estimate with the other pair estimates. The exponent C_Γ absorbs their constants, and the rank satisfies $k \geq 8C_\Gamma m^5 f \log b$. The upper ratio in (55) gives $\Gamma d_w^2/D_\star \leq (C_0/\Gamma)d_w$. Combining the degenerate and nondegenerate estimates by the squared triangle inequality gives

$$\|K_{ww} - I\|_{\mathrm{HS}}^2 \leq e_0 d_w, \qquad e_0 = 2\exp(-kf\log k) + 2C_0/\Gamma. \tag{56}$$

Here $k \geq 8\log\Gamma$ and $\log\Gamma \geq C_\Gamma x^\varepsilon$, so $\sqrt{e_0} \leq \exp(-x^{\varepsilon/8})$ for large x, uniformly over the selected translates. **AI-proposed candidate (the same-word Gram defect).** Fix the target, precision and tied family with its levels, weight gaps and rank conditions, including (55) and $k \geq 8\log\Gamma$. At every selected translate beyond one common scale threshold, the full Gram operator of each word has mean squared defect

$$\frac{\|L_w^* L_w - I\|_{\mathrm{HS}}^2}{d_w} \leq 2e^{-kf\log k} + \frac{2C_0}{\Gamma}, \qquad \Gamma = \exp(C_\Gamma m^5 f \log b).$$

The normalization gives I on the identity double coset. The estimate uses absolute bounds for the remaining degenerate terms and the equal-word nondegenerate bound $\Gamma d_w^2/D_\star$. This nondegenerate calculation shares orbital estimates with the mixed-word proof below. The same-word statement controls the mean of $(\sigma_j^2 - 1)^2$ over the d_w singular values of L_w, including zeros. Mixed bounds alone allow mutually orthogonal images with each L_w twice an isometric inclusion, for which the mean squared defect is 9.

Wei obtains a growing matrix of qI_2-Fourier coefficients of the pI_2 Siegel Poincaré series, with main term $8I$ and an entrywise error decreasing faster than any power of the weight. [43, Lemma 2.3 and §3.1] The estimate above treats all full Bessel columns of one word in the unitary family, with the indicated dependence on its rank, field and level.

The period source $H_w^{\mathrm{act}} = \mathcal{M}_w$ is the ramified column space included by $\jmath_w$ in Section 2. At level zero the columns are sections over the definite class set, with one residual compact representation fibre per class. Since T_w splits completely in E, there are f places of F over it and the T_w-diamond group is $\mathsf{A}_w = \prod_{z|T_w} \mathbf{F}_{T_w}^\times$. It acts by right translation on the classes and carries the fibres isomorphically along this action. A diamond representative is an integral local dilation at the T_w-places and the identity at the auxiliary places. If it fixes a class, the double-coset equation identifies its product with a level element with a rational point preserving the lattice of that class. This rational stabilizer is finite by definiteness. At each auxiliary prime the reduction is upper unitriangular, and the preimage of that group is pro-g_i because the principal congruence kernel is pro-g_i. Thus the stabilizer element has order a power of each of two distinct primes. It is the identity, so the diamond representative is itself in the level and acts by the trivial multiplier at every T_w-place. Freeness at the level-zero classes.

On each free orbit fix an identification of the fibre at one class and transport it by the group action. This identifies the full column representation over the orbit with the regular representation of A_w tensored with that fibre. In each $\mathbf{C}[\mathbf{F}_{T_w}^\times]$ the kernel of the local averaging projector has dimension $T_w - 2$. The actual ramified space is the intersection of these local kernels. Its restriction to the orbit is the tensor product of the individual kernels with the fibre, since the f local projections act in distinct factors. Adding the ranks over the orbits therefore gives the first equality below. With $r_w = \dim \mathcal{M}_w$ and $T_w \geq b \geq 3$, Bernoulli's inequality gives the rest:

$$r_w = d_w\left(1 - \frac{1}{T_w - 1}\right)^f, \qquad 0 \leq d_w - r_w \leq \frac{2f}{b} d_w. \tag{57}$$

This is the ratio for the same level-zero columns counted in (55). Ramified rank formula.

The diagonal input gives another use of this calculation. On $\ker \mathcal{P}_w$, the operator $K_{ww} - I$ is $-I$. Consequently $\dim\ker \mathcal{P}_w \leq e_0 d_w$. The bound $f \leq k \leq x^5$ implies $e_0 < 1 - 2f/b$ for sufficiently large x, so some vector in $\mathcal{M}_w$ has nonzero image under $\mathcal{P}_w$. Pairing with that image in the full $\mathcal{H}$ supplies the nonzero bottom pairing.

In the construction of the bottom reference pairing, the earlier single-word estimate comes from the same local Bessel integrals and Siegel–Weil. In combination with the ramified rank bound, it puts the

squared diagonal defect strictly below the dimension of the ramified column space of each such bundle. The formal construction uses this inequality to obtain a nonzero image of a ramified column. For the conditional argument we use the period maps supplied by the stated comparison.

3.3 Squaring the orbital sum at finite level

Fix $w \neq v$ and let S be either of the two orbit families with $G \neq 0$. Group the terms by their two lattice representatives and the remaining index $i(t)$ in the finite Levi group. Call this triple the row $r(t)$. At a place of F above $g_0 g_1$ with residue size q, the group is $G_q = \mathrm{GL}_k(\mathbf{F}_q)$ with upper unitriangular subgroup U_q. Write $I = \prod_q G_q$, with all such products indexed by these places. The weighted vector v_t is the normalized real orbital operator in the Hilbert–Schmidt summand for its lattice pair, multiplied by the finite unipotent character and the square root of the product of the two lattice weights. Put $K_S = \sum_{t \in S} A_t$ and $V_i = \sum_{t \in S \colon i(t)=i} v_t$. In the chosen column representatives the unipotent factor of the finite character is constant for each lattice pair. The remaining factor is the character of $\prod_q U_q$, extended by zero, evaluated at the relative Levi index for the two column representatives. The representatives give free decompositions of I, and the Gram sums for each family are invariant under the corresponding left and right translations in I. These identities give the squared norm

$$\|K_S\|_{\mathrm{HS}}^2 = \mathrm{Re} \sum_{i,j \in I} \mathcal{G}_{ij} \langle V_i, V_j \rangle, \qquad \mathcal{G}_{ij} = \overline{\prod_q \gamma_q((j^{-1}i)_q)}.$$

Here $\gamma_q = \mathrm{ch}(\mathrm{Ind}_{U_q}^{G_q} \psi_q)$ is the Gelfand–Graev character, where ψ_q is the nondegenerate character obtained from the chosen nontrivial residue-field character on the sum of the simple superdiagonal entries. Let $\Gamma_1 = \prod_q \gamma_q(1) = \prod_q [G_q : U_q]$ and $P_g = \prod_q (1 + 3/q)$. The finite convolution gives

$$0 \leq \mathcal{G} \leq P_g \Gamma_1 \, \mathrm{Id}_{\mathbf{C}^I}, \qquad \|K_S\|_{\mathrm{HS}}^2 \leq P_g \Gamma_1 \sum_i \|V_i\|^2 = P_g \Gamma_1 \sum_{\substack{t,t' \in S \\ r(t)=r(t')}} \mathrm{Re} \langle v_t, v_{t'} \rangle.$$

For the inequality, use an orthonormal basis of the span of the finitely many V_i and apply the scalar matrix bound to each coordinate vector. For the last equality, within each i different lattice pairs have orthogonal summands. Absolute convergence then permits the expansion over ordered pairs in each row.

The character $\overline{\psi_q}$ is conjugate to ψ_q by a diagonal matrix whose successive diagonal entries have ratio -1. Consequently γ_q is real. For convolution with counting measure, character orthogonality and Frobenius reciprocity give the eigenvalue of $\overline{\gamma_q}/\gamma_q(1)$ on the irreducible ρ as

$$\frac{|U_q|}{\dim \rho} \dim \mathrm{Hom}_{U_q}(\psi_q, \rho).$$

For $\mathrm{GL}_k(\mathbf{F}_q)$ over a finite field and the nondegenerate ψ_q just defined, Gel'fand proves that the complex representation $\mathrm{Ind}_{U_q}^{G_q} \psi_q$ is multiplicity-free. For the mirabolic P_q, defined by $e_k^{\mathsf{T}} g = e_k^{\mathsf{T}}$ with e_k the last basis vector, $\mathrm{Ind}_{U_q}^{P_q} \psi_q$ is irreducible of dimension $[P_q : U_q] = \prod_{j=1}^{k-1} (q^j - 1)$. Thus for each irreducible complex ρ of G_q, the displayed ψ_q-multiplicity is 0 or 1. Frobenius reciprocity for $U_q \subset P_q$ gives exactly one copy of the mirabolic induction in $\rho|_{P_q}$ when the multiplicity is 1, and therefore gives the same lower bound for $\dim \rho$. [45, Theorem 0.5 and Proposition 1.2, pp. 17–18] With $|U_q| = q^{k(k-1)/2}$, the local eigenvalue is at most $\prod_{j=1}^{k-1} (1 - q^{-j})^{-1} \leq 1 + 3/q$ for $q \geq 3$. The product before inversion is at least $1 - 1/(q-1)$. Taking tensor products over the places yields the bound for $\mathcal{G}$.

For a pair of terms in the same row, consider the fixed flag $\mathfrak{f}$ and two moving flags $\mathfrak{f}'$, $\mathfrak{f}''$, set W_1 and W_2 equal to the orthogonal complements of the two joint augmented frames. Let β be the isometry comparing the moving flags and let X and X'' be the fixed and second moving isotropic spans. The comparison space is the subspace of $x \in W_1$ for which there exists $y \in W_2$ with $x - y \in X$ and $\beta x - y \in X''$. The corank s_0 is k minus its dimension, so $s_0 \geq 1$ because $\dim W_1 = k - 1$. To define a scalar pair, write $U' = X^+(\mathfrak{f}) + X^+(\mathfrak{f}')$

for the joint span of the first pair of augmented frames, and U'' for the analogous span for the second. In U', take lifts y and y' perpendicular to $X^+(\mathfrak{f})$ and $X^+(\mathfrak{f}')$, respectively. If $y - y' = \lambda z + \lambda' z'$ modulo the sum of the two isotropic spans, call (λ, λ') a pair of distinguished-vector coefficients. Compare this relation with the corresponding one for lifts $\tilde{y}, \tilde{y}'$ in U'' such that $y - \tilde{y} \in X$ and $\tilde{y}' - \beta y' \in X''$. A pair is scalar if it has corank one and the two coefficient relations agree for every such choice of lifts.

3.4 The leading contribution to the mixed Gram blocks

Keep $w \neq v$ fixed and set $D_{wv} = d_w d_v / D_\star$. The word primes T_w split completely in E and are congruent to 1 modulo $(g_0 g_1)^2$. We check their comparison with a fixed constant, as used in the mixed estimate. Write $j_0 = k - 2$. At a real place, the compact word parameters have two common endpoints and the j_0 interior entries $v_i - 1/2 - \Delta b_i$, where $b_i \in \{0, 1\}$ and $v_i - v_{i+1} = 4\Delta$. The upper endpoint A satisfies $A - v_1 \geq 4k\Delta$. The lower endpoint is $v_{k-1} - 1/2$ away from the designated top places, and $-b_{\mathrm{top}}$ with $b_{\mathrm{top}} \geq B_1/2$ at the top places. The selected weight data give $\Delta \geq 1$, $v_i \geq 0$ and $B_1 \geq 4kx\Delta$. Endpoints and the shifted grid · Endpoint inequalities · Weight size.

For a single word the interior log Vandermonde relative to the unshifted grid is

$$\sum_{h=1}^{j_0-1} \sum_{i=1}^{j_0-h} \log\left(1 + \frac{b_{i+h} - b_i}{4h}\right).$$

The linear term for lag h is bounded in absolute value by $\min(h, j_0 - h)/(4h)$, by cancellation of the overlapping sums of bits. For $|t| \leq 1/4$, integration of $1/(1+t)$ gives $|\log(1+t) - t| \leq (2/3)t^2$. The sum of these remainders is at most $(j_0/24)\sum_{h\geq 1} h^{-2} \leq j_0/12$ by integral comparison. The difference for two words on the interior is therefore at most $2j_0/3$ in absolute value. At the upper endpoint the log change of each gap is bounded by $\Delta/(4k\Delta)$. At the lower endpoint outside the top places the gap for interior entry i is at least $3(k-1-i)\Delta$, so the total log change there is at most $\sum_{j=1}^{j_0} 1/(3j)$. At a top place the interior entries are at least $-\Delta - 1/2$, so each gap to the lower endpoint is at least $k\Delta$ for large x, at total cost at most 1. The difference of the two endpoints is fixed. For large k the resulting logarithmic difference of the full compact Vandermondes is at most $3k/4$, at every real place.

The compact Weyl dimension formula gives $d_w^{(0)}$ as a common finite-level class factor times the product over real places of the compact Vandermondes, divided by the same Weyl denominators. This is d_w^* in the compact dimension formula of Section 4. Hence $\left|\log(d_w^{(0)}/d_v^{(0)})\right| \leq 3kf/4$. Thus $|\log(Y_w/Y_v)| \leq 3/4$. By the chosen prime interval, every $\log(T_w/Y_w)$ lies in $[0, \log(1 + 1/(fk))]$. Combining the bounds gives

$$|\log(T_w/T_v)| \leq \tfrac{3}{4} + \log(1 + 1/(fk)) \leq 1, \qquad \left|\log \frac{\mathrm{Vand}_z(w)}{\mathrm{Vand}_z(v)}\right| \leq \tfrac{3}{4}k \quad (z \mid \infty). \tag{58}$$

Here $\mathrm{Vand}_z(w)$ is the compact Vandermonde at z. This supplies $A_0 = \exp(1)$ for both types of ratio. The orbital estimates use these comparisons and (55) together with the prescribed gaps between archimedean weights and the local lattice models. These hypotheses allow Alpöge to apply the lattice-counting and orbital-decay estimates to each pair of words.

The leading term comes from the generic scalar pairs of corank one. At an ordinary finite place v, write $E_v = E \otimes_F F_v$, assumed split or unramified quadratic over F_v. Let ϖ be a uniformizer of F_v, $q = |\mathcal{O}_{F_v}/\varpi|$, and J_v the Gram matrix in an integral basis of a self-dual Hermitian lattice of rank r. For a nonsingular integral Hermitian matrix S of size ℓ, let c be the corank of its reduction. At congruence depth n, count matrices $X \in M_{r,\ell}(\mathcal{O}_{E_v}/\varpi^n)$ with $X^* J_v X \equiv S \pmod{\varpi^n}$ and with injective reduction modulo ϖ. In the split case injectivity is required in each of the two factors. For $r > \ell$ and $\ell + c \leq r$, the limit of this count divided by $q^{n(2r\ell - \ell^2)}$ is

$$\prod_{j=r-\ell-c+1}^{r} (1 - \xi^j q^{-j}), \qquad \xi = \begin{cases} 1 & \text{at a split place,} \\ -1 & \text{at an inert place.} \end{cases}$$

For the unramified quadratic field E_v, the self-dual lattice has Gram matrix I_r in an integral orthonormal basis. This is Hironaka's primitive density for that ambient Gram matrix. The c positive elementary-divisor exponents of S give denominator $\prod_{j=1}^{r-\ell-c}(1-(-q)^{-j})$ in her quotient of products. [48, Equation (3.1), p. 572, and Corollary 3.3(i), p. 576] At a split place, let X_1, X_2 denote the two reduced matrices after identifying the ambient form with the standard one. In $X_1^{\mathsf{T}} X_2 = \bar{S}$, with $\bar{S}$ the reduction in the first component, choosing X_1 gives factors $j = r-\ell+1, \dots, r$. The free block of X_2 must be injective on $\ker \bar{S}$, giving the c remaining factors. Surjectivity of $\dot{X}_2 \mapsto X_1^{\mathsf{T}} \dot{X}_2$ gives the same normalized count at every congruence level. In the release, self-duality and the local Gram matrices identify the orbital measures with this normalization.

In the scalar calculation $\ell = m = 2(r_S + 1)$ and $k \geq \ell + f + 5$. The first exponent in the density is $k - c \geq k - \ell$. At ordinary places the reference density is the product of the two single-frame factors. At the places above the auxiliary primes it uses one common single-frame factor. The ratio of the pair density to the reference factor is at most $\exp(4q^{-(k-\ell)})$ at each place. Let $s_1 = k - \ell \geq 2$. There are at most f places of F above each rational prime, so $\sum_{\mathfrak{p}} (\mathrm{N}\mathfrak{p})^{-s_1} \leq f \sum_{n \geq 2} n^{-s_1} \leq 3f\, 2^{-s_1}$. The upper bound E_{den} for the normalized density products, enlarged to at least 1, satisfies $1 \leq E_{\mathrm{den}} \leq \exp(12 f\, 2^{-(k-\ell)})$.

Use the convergent unitary Siegel–Weil formula with Y_0 of Hermitian rank $r > 2\ell$ over the quadratic CM extension E/F. Let G_ℓ be the unitary group of the split skew-Hermitian space of rank 2ℓ, with Siegel parabolic P_ℓ. Fix a nontrivial additive character ϑ of $\mathbf{A}_F/F$ and a Hecke character χ of E restricting to $\epsilon_{E/F}^{r}$ on $\mathbf{A}_F^{\times}$, where $\epsilon_{E/F}$ is the character of the quadratic extension. Write ω for the corresponding Weil representation on the Schwartz space of $(Y_0 \otimes_F \mathbf{A}_F)^{\ell}$. If J is the Gram matrix of Y_0, the form $2J$ in Liu's convention gives moment matrix $X^* J X$. In particular the Siegel unipotent element $n(b)$ acts by the multiplier $\vartheta(\mathrm{tr}(b X^* J X))$ for Hermitian b.

For any Schwartz Φ, set $s_0 = (r-\ell)/2$ and use the Siegel–Weil section $f_{\Phi,s}(g) = (\omega(g)\Phi)(0)\lambda(g)^{s-s_0}$, where $\lambda(g) = |\det a|_{\mathbf{A}_E}$ for the Siegel Levi coordinate a in an Iwasawa decomposition of g. The Eisenstein series and the theta integral below both converge absolutely at s_0, and for every $g \in G_\ell(\mathbf{A}_F)$

$$E_\Phi(s_0, g) := \sum_{\gamma \in P_\ell(F)\backslash G_\ell(F)} f_{\Phi,s_0}(\gamma g) = \int_{[G_0]} \sum_{X \in Y_0^{\ell}} (\omega(g)\Phi)(h^{-1} X)\, dh,$$

where dh has total mass 1 on $[G_0] = G_0(F)\backslash G_0(\mathbf{A}_F)$. [46, Theorem 2.1(1) and the Weil-representation formulas, pp. 858–859]

For any nonsingular Hermitian ℓ-by-ℓ matrix S over E, take the S-Fourier coefficient $E_{\Phi,S}(s_0, 1)$, integrating over the rational quotient of the Siegel unipotent with multiplier $\vartheta(-\mathrm{tr}(bS))$ and quotient volume 1. Taking this coefficient of the absolutely convergent formula above gives equality with the corresponding coefficient of the theta integral. Yamana gives this nonsingular coefficient identity for polynomial Fock sections, with factor 1 in this range since the Witt index of Y_0 is at most $r/2$ and $r > 2\ell$. [47, Theorem 2 and Lemma 3.1, pp. 237, 246] The theta coefficient selects the Gram fibre. On the Eisenstein side the nonsingular coefficient unfolds through the open cell. At s_0, the local section at $w_\ell n(b)$, where w_ℓ is the long Siegel Weyl element, is the Fourier transform of the pushforward of $\Phi_v(X)\, dX$ under $X \mapsto X^* J X$, times the local Weil constant in the action of w_ℓ. Applying $n(b)$ and evaluating that action at zero gives the integral with phase $\vartheta_v(\mathrm{tr}(b X^* J X))$. Fourier inversion for factorizable $\Phi = \bigotimes_v \Phi_v$ therefore gives

$$E_{\Phi,S}(s_0, 1) = \int_{[G_0]} \sum_{\substack{X \in Y_0^{\ell} \\ X^* J X = S}} \Phi(h^{-1} X)\, dh = \prod_v \alpha_v(S; \Phi_v),$$

where $\alpha_v(S; \Phi_v)$ is the value at S of the density of the pushforward of $\Phi_v(X)\, dX$ under $X \mapsto X^* J X$, using the self-dual measures for ϑ_v on the tuple space with the trace of the Hermitian form and on the Hermitian matrices with the trace pairing. The Weil indices have global product 1. The reductions of S can have different coranks at different finite places. Applying the formula to the scalar pairs uses the release's comparison of the prescribed test functions, stabilizer volumes and orbital measures.

In this normalization the signed sum over the scalar pairs is bounded by a sum of squared archimedean amplitudes. Here the real orbital calculation uses the chosen local Bessel integrals of both words. Positivity and Poisson summation bound that sum by $(1+\delta_{\mathrm{sc}})E_{\mathrm{den}}$ times the archimedean energy integral, with $0 \le \delta_{\mathrm{sc}} \le b^{-7}$. The inequality applies to the Γ_1-weighted ordered pairs in the finite convolution. After disintegration the formal degrees bound the archimedean integral by $1/\mathfrak{d}$. Write V_{wv} for the adelic volume of the joint stabilizer in the mass calculation, V_w, V_v for the two adelic Bessel stabilizer volumes, and s_w, s_v for their deep Bessel level volumes. Put $c_w = c_0(w)/v_{\mathrm{lev}}$, removing the common ambient level volume v_{lev} already supplied by the unfolded level integral. The product of the local measures, including the finite factor Γ_1 at the places above the auxiliary primes, gives $v^{\dagger}/(s_w s_v)$. The stabilizer volumes and numbers of columns satisfy

$$V_{wv} V_G = V_w V_v, \qquad c_w^{-1} = \frac{d_w s_w}{V_w}, \qquad c_v^{-1} = \frac{d_v s_v}{V_v}.$$

Consequently the scale multiplying the allowance is

$$\frac{V_{wv} v^{\dagger}}{s_w s_v c_w c_v \mathfrak{d}} = \frac{d_w d_v v^{\dagger}}{\mathfrak{d} V_G} = D_{wv}.$$

The signed sum over the generic scalar pairs is therefore at most AD_{wv}, where $A = (1+\delta_{\mathrm{sc}})E_{\mathrm{den}}$.

By counting lattices at fixed corank and using decay of the real orbital amplitudes, Alpöge bounds the absolute sums over pairs of corank at least two, and separately over generic nonscalar pairs of corank one, by $\epsilon_L D_{wv}$ in each case. The first sum includes both nondegenerate families. Here $\epsilon_L = \exp(-kf \log k/3)$, using $6\log(3A_{\mathrm{L}}) \le kf \log k$ for the fixed multiplicative constant $A_{\mathrm{L}} \ge 1$ in the lattice and amplitude bounds.

For nongeneric pairs of corank one, a failure of local generation decreases either the mass of the local configurations or the character sum. Summing the weights over the fibres of the local data bounds the signed sum by $\epsilon_N D_{wv}$, where

$$\epsilon_N = 2^{-(k-2r_S-5)} \exp(C_N m^5 f \log b)$$

for a fixed $C_N \ge 0$.

The generic row sum contains the scalar main term, the higher-corank error and the generic nonscalar corank-one error. Its associated operator has squared norm at most $P_g(A + 2\epsilon_L)D_{wv}$. The nongeneric part contains the higher-corank and signed corank-one errors, so its operator has squared norm at most $P_g(\epsilon_L + \epsilon_N)D_{wv}$. Hence

$$\|K_{wv}\|_{\mathrm{HS}}^2 \le D_{wv}\left[\sqrt{P_g(A+2\epsilon_L)} + \sqrt{P_g(\epsilon_L+\epsilon_N)}\right]^2.$$

The selected family has $m^4 \le f \le k \le x^5$, $T_w \ge b^{m/2}$, $b \ge 16$, $g_i \ge (kb)^8$ for $i = 0, 1$, and $k \ge \max(32, 4r_S + 12, 70\log b)$. Each g_i has f places of F above it, all of residue size g_i, so

$$\log P_g \le 3f(g_0^{-1} + g_1^{-1}) \le 6b^{-7}.$$

The size conditions also give $8C_N m^5 f \log b \le k$. Since $k - 2r_S - 2 \ge k/2 + 4$ and $k - 2r_S - 5 \ge k/2 + 1$, the exponents in the density and error estimates give $\log E_{\mathrm{den}}, \epsilon_L, \epsilon_N \le e^{-k/5} \le b^{-14}$. In particular $E_{\mathrm{den}} \le \exp(b^{-14})$. Substituting these bounds in the two square roots gives, for $b \ge 16$,

$$\|K_{wv}\|_{\mathrm{HS}}^2 \le (1+b^{-6})\,\frac{d_w d_v}{D_\star} \qquad (w \ne v). \tag{59}$$

3.5 Trace moments after restricting the columns

AI-proposed candidate (uniform multiplicities for unitary groups at finite level). Let $G = U(V)$ for a rank-r Hermitian space over a CM extension E/F unramified at every finite place, with a self-dual

lattice Λ. Put $f = [F : \mathbb{Q}]$ and $\mathrm{rd}(E) = |\operatorname{disc} E|^{1/(2f)}$. The real signatures are all (p, q), with $p + q = r$ and $q^2 \le r$. Let K_f be the finite adelic stabilizer of Λ and $K_f(g)$ its principal congruence subgroup modulo g. With the given finite Haar measure, take a locally constant self-adjoint convolution idempotent e of positive type, supported on K_f and bi-invariant under $K_f(g)$. For a rational prime $g_0 > 4r$ splitting completely in E, require $(x-1)^r \widehat{\Lambda} \subseteq g_0 \widehat{\Lambda}$ for x in its support, where $\widehat{\Lambda}$ is the finite adelic lattice. These are the finite projector conditions.

For fixed $A > 0$ and r beyond a threshold depending on A, assume $b \ge 2$, $1 \le g \le b^A$, $\log\log \mathrm{rd}(E) \le r$, $\log f \le \sqrt{r}$, and $\log\log b \le rf \log r/16$. At each real place v, let σ_v be a discrete series with strictly decreasing Harish–Chandra parameters $\lambda_{v,i} \in \mathbb{Z} + (r-1)/2$ and their signature labels. Require $|\lambda_{v,i}| \le b^2$ and, for every pair with opposite signature labels,

$$|\lambda_{v,i} - \lambda_{v,j}| \ge 8 \cdot 10^{14} (A+4)\, r^{19} f^5 \log(b\ \mathrm{rd}(E)).$$

Write $\sigma_\infty = \widehat{\bigotimes}_{v|\infty} \sigma_v$ and $R(e)$ for finite convolution in the right regular representation on $L^2([G])$. The count $D(e, \sigma_\infty)$ is the dimension of $R(e) \operatorname{Hom}^{\mathrm{bdd}}_{G_\infty}(\sigma_\infty, L^2([G]))$, which is finite. The estimate is

$$D_\star(e, \sigma_\infty) = \mathrm{vol}([G])\ \mathrm{Re}\, e(1) \prod_{v|\infty} d(\sigma_v),$$

$$|D(e, \sigma_\infty) - D_\star(e, \sigma_\infty)| \le \exp(-rf \log r/16)\, D_\star(e, \sigma_\infty).$$

The formal degrees use the real Haar measures in the quotient volume. The error is uniform in the fields, ranks, signatures and projectors satisfying these hypotheses. The release uses the counting statement with its lattice-counting input from the volume and kernel estimates. In this calculation the identity gives the main term. The other terms use large-gap matrix coefficient decay and bounds for the near unipotent terms and the deep part of the quotient.

Shin–Templier's quantitative Plancherel results and counting asymptotic apply to level and weight families of a fixed group over a fixed base field. [42, Theorems 9.16 and 9.19 and Corollary 9.25] The bound here is uniform as the unitary rank and field vary, and its large-gap condition permits small gaps within a signature block. It concerns the full projected multiplicity at the given real representation.

We apply the multiplicity input to the full row projector of the same record as $\mathcal{P}_w$. The finite quotient at level $g_0 g_1$ of the selected self-dual lattice stabilizer is $\prod_{i=0}^{1} \prod_{z|g_i} \mathrm{GL}_r(\mathbf{F}_{g_i})$ by the split reductions and the Chinese remainder theorem. The full row projector e is the tensor product of the finite character projectors on the prescribed unitriangular subgroups at these places, inflated with the principal congruence kernel of measure $\mathrm{vol}\, K_f(g_0 g_1)$. Equivalently e is the average with conjugate character over the deep ambient subgroup of volume $v^\dagger$. Its identity value is $1/v^\dagger$. With this reciprocal-volume normalization it is self-adjoint and idempotent, and it is of positive type as the matrix coefficient of the vector e in the regular representation. Its support reduces to a unipotent group at each auxiliary prime, so $(y-1)^r \widehat{\Lambda} \subseteq g_i \widehat{\Lambda}$ for every point y of the support. The condition also holds on the conjugate half of each split unitary reduction, which is the inverse transpose of the chosen half.

At the real places, evaluation of the bounded intertwiner at the chosen tensor of unit lowest-type vectors identifies the projected multiplicity space with the Petersson row model $\mathcal{H}$ in (21). This evaluation is injective by cyclicity of the vector in the irreducible tensor. For the converse, let π_∞ be the actual ordered tensor of real discrete series and u its chosen unit lowest-type tensor. In the right-translation convention use

$$q_\infty(g) = d(\pi_\infty) \overline{\langle u, \pi_\infty(g) u \rangle}.$$

The matrix coefficient is integrable at the chosen real gaps. Schur orthogonality, with this formal degree for the row Haar, gives $q_\infty^* = q_\infty$, $q_\infty * q_\infty = q_\infty$, and

$$q_\infty * \delta_g * q_\infty = \langle u, \pi_\infty(g) u \rangle q_\infty.$$

Thus if z is fixed by both the real line projector and $R(e)$, the prescription $\pi_\infty(g)u \mapsto R(g)z$ extends uniquely to a bounded intertwiner. Indeed the Gram of the image of any finite orbit combination is its original Gram multiplied by $\|z\|^2$. Its image is fixed by $R(e)$ because that finite projector commutes with the real action. This is inverse to evaluation at u, with the Petersson norm of z for the same Haar. The dimension D is therefore the multiplicity of exactly the full row projector to which the count applies. Row realization.

Here are the size checks for this application and for the restricted projectors below. The real parameters, in order, are a decreasing nonnegative block $v_1, \dots, v_{k-1}$, all with the positive label, and $-B_1, \dots, -B_m$ with alternating labels starting with the positive label. Thus the signature is $(k - 1 + m/2, m/2)$. The selected bound $m^2 \le k$ gives $(m/2)^2 \le r$, where $r = k + m - 1$. Write $L = \log b$ and Δ for the weight displacement. The selected inequalities for the weights in this real type give

$$0 \le v_j \le 4k\Delta, \qquad C_{\mathrm{gap}} mk\Delta \le B_1, \qquad B_m \le b^2, \qquad B_{j+1} - B_j \ge C_{\mathrm{gap}} r^{19} f^5 L.$$

Weight and level inequalities. The family comparison allows C_{gap} at least $\max(2, 16 \cdot 10^{14}(A + 4))$, with $A = 2C_g$ when $g_i \le b^{C_g}$. Then $4k\Delta \le B_1$, so all parameters have absolute value at most b^2. A gap between oppositely labelled negative parameters is at least one adjacent B-gap. An opposite-signature pair involving a v_j uses a $B_{2j'} \ge B_2$, so the same lower bound holds. The bound $\log \mathrm{rd}(E) \le L/C_{\mathrm{gap}}$ implies the required gap with $\log(b\ \mathrm{rd}(E))$. The selected levels are completely split with $(kb)^8 \le g_i \le b^{C_g}$, hence $4r < g_i$ and $g_0 g_1 \le b^A$. With $x^2 \le f \le k \le x^5$ and $x^4 \le k \le r \le 2k$, all three logarithmic size conditions in the count hold for sufficiently large x. The CM extension is unramified at finite places and the selected Hermitian lattice is self-dual, as in the ambient row construction. Count parameters for the common record.

The main term is therefore $D_\star = \mathfrak{d} V_G / v^\dagger$ in the Gram normalization. The multiplicity input gives $|D - D_\star| \le \exp(-rf \log r/16) D_\star$, in particular $D > 0$.

At each translate the estimates give

$$\begin{aligned} \|K_{wv}\|_{\mathrm{HS}}^2 &\le (1 + \eta_0)\frac{d_w d_v}{D_\star} && (w, v \in W,\ w \ne v), \\ \|K_{ww} - I\|_{\mathrm{HS}}^2 &\le e_0 d_w && (w \in W), \\ \sum_{w \in W} \|K_{ww} - I\|_{\mathrm{HS}}^2 &\le e_0 R_{\mathrm{full}}, && D \le (1 + \eta_1) D_\star. \end{aligned} \tag{60}$$

Take $\eta_0 = b^{-6}$, $\eta_1 = \exp(-rf \log r/16)$, and e_0 from (56). The three quantities η_0, η_1, $\sqrt{e_0}$ are each bounded by $\exp(-x^{\varepsilon/8})$ at every indicated translate above a common threshold.

AI-proposed candidate (a uniform mixed Gram bound). For the fixed target, precision and tied family, the mixed estimate (59) holds beyond one scale threshold at every selected translate, simultaneously for all distinct words w, v in its word set. It bounds $\|L_w^* L_v\|_{\mathrm{HS}}^2$ by $(1 + b^{-6}) d_w d_v / D_\star$ on the full columns, with the common Petersson target and reference dimension at that translate. For the selected seeds under the prime, weight-gap, lattice and rank conditions above, the degenerate terms vanish. The upper bounds use the signed sums for generic scalar pairs and nongeneric corank-one pairs, and absolute sums for the remaining errors.

Getz–Hahn's compact-test formula supplies the orbit reindexing in a simple relative trace formula. A quantitative precedent is Kowalski–Saha–Tsimerman's estimate for Fourier coefficients of genus-two Siegel Poincaré series over $\mathbb{Q}$. For fixed discriminant d their normalized error at indices L, M and scalar weight κ is $O_{d,\epsilon}(L^{1+\epsilon} M^{3/2+\epsilon} \kappa^{-2/3})$. [41, Proposition 3.3, arXiv v4] The unitary bound here uses the different local test and volume identifications above, including the reference dimension from V_G, $v^\dagger$ and $\mathfrak{d}$. Its uniformity covers the fields, ranks and word primes of the tied family. An aggregate trace moment alone allows two identical isometric blocks of dimension d inside a tight frame in dimension $D_\star = 2d$. Their mixed squared norm is d, greater than $(1 + b^{-6})d/2$.

The trace argument sums the mixed bounds and uses the independent single-word defect (56) to control the diagonal blocks. With the row and word counts these estimates give the moments for Waldron's inequality. The ramified compression gives (61), and (63) controls the rank after further column restrictions.

Write $Q_{\mathrm{full}} = \sum_w L_w L_w^*$. Cyclicity of trace gives

$$\operatorname{tr}(Q_{\mathrm{full}}^2) = \sum_{w,v} \operatorname{tr}(L_v^* L_w L_w^* L_v) = \sum_{w,v} \|K_{wv}\|_{\mathrm{HS}}^2.$$

Put $\delta_0 = \sqrt{e_0}$. The diagonal terms have $\|K_{ww}\|_{\mathrm{HS}} \le (1+\delta_0)\sqrt{d_w}$ by the triangle inequality. Summing the off-diagonal bound in (60) over ordered pairs $w \ne v$ and using $\sum_{w\ne v} d_w d_v \le R_{\mathrm{full}}^2$ yields

$$\operatorname{tr}(Q_{\mathrm{full}}^2) \le (1+\eta_0)\frac{R_{\mathrm{full}}^2}{D_\star} + (1+\delta_0)^2 R_{\mathrm{full}}.$$

Section 6.1 selects the word set with $|W| \ge C_0 \Gamma^2 x^{C_H}$. These are the same words and the same $C_0 = 3$ as in the full column balance. Word-count condition. Together with the lower bound in (55) this gives $R_{\mathrm{full}} \ge D_\star/\sigma$, where $\sigma = x^{-C_H}$ and $C_H > 0$ is fixed on the chosen family. This makes the diagonal term small compared with $R_{\mathrm{full}}^2/D_\star$.

For any isometric inclusions $J_w : H_w^{\mathrm{act}} \hookrightarrow H_w^{\mathrm{full}}$, define the restricted column maps A_w and their frame operator $Q_{\mathcal{H}}$ by

$$A_w = L_w J_w, \qquad Q_{\mathcal{H}} = \sum_w A_w A_w^* = \sum_w c_0(w)^{-1} (\mathcal{P}_w J_w)(\mathcal{P}_w J_w)^*.$$

Thus the frame operator weights the original column maps by $c_0(w)^{-1} > 0$. The period comparison below uses $J_w = {}_J W_w$. Let $R_{\mathrm{act}} = \sum_w r_w$ and $\rho = 1 - R_{\mathrm{act}}/R_{\mathrm{full}}$. By (57), $0 \le \rho \le 2f/b$.

For every block, $A_w^* A_v = J_w^* K_{wv} J_v$. In bases extending the ranges of the J_w, this compression retains a submatrix, and its Hilbert–Schmidt norm is bounded by the full norm. Consequently the same trace identity gives $\operatorname{tr}(Q_{\mathcal{H}}^2) \le \operatorname{tr}(Q_{\mathrm{full}}^2)$. Since $J_w^* J_w = I$, compression also bounds the sum of squared diagonal defects by $e_0 R_{\mathrm{full}}$. Taking traces of the compressed diagonal blocks and applying Hilbert–Schmidt Cauchy–Schwarz on their direct sum gives

$$t := \operatorname{tr} Q_{\mathcal{H}} \;\ge\; R_{\mathrm{act}} - \sqrt{e_0 R_{\mathrm{act}} R_{\mathrm{full}}} \;\ge\; (1-\rho-\delta_0) R_{\mathrm{full}}.$$

For all sufficiently large x we have $\sigma \le 1/20$, $\eta_0, \eta_1, \delta_0 \le \sigma/2$ and $\rho \le 1/4$, uniformly over the indicated translates. In this range the last bound is positive. Using $D \le (1+\eta_1) D_\star$, we obtain

$$\frac{D \operatorname{tr}(Q_{\mathcal{H}}^2)}{t^2} \le \frac{(1+\eta_1)\big[1+\eta_0+(1+\delta_0)^2\sigma\big]}{(1-\rho-\delta_0)^2}.$$

Write $q_0 = \sqrt{\rho}$. The numerator is at most $(1+\sigma/2)^2 + \sigma(1+\sigma/2)^3 \le 1+3\sigma$, and the denominator is at least $1-\sigma-2q_0^2$. In the given ranges,

$$\begin{aligned}(1+5\sigma+3q_0)(1-\sigma-2q_0^2) - (1+3\sigma)\\ = \sigma(1-5\sigma) + q_0\big(3(1-\sigma) - 2q_0(1+5\sigma+3q_0)\big) \;\ge\; 0.\end{aligned}$$

Thus $e = 2e_0$ and $\eta = 5\sigma + 3\sqrt{\rho}$ satisfy

$$\begin{aligned}\sum_w \|A_w^* A_w - I\|_{\mathrm{HS}}^2 \le e R_{\mathrm{act}}, \qquad 0 \le e \le \tfrac{1}{4},\\ \operatorname{tr}(Q_{\mathcal{H}}^2) \le (1+\eta)\frac{(\operatorname{tr} Q_{\mathcal{H}})^2}{D}, \qquad \eta \ge 0.\end{aligned} \tag{61}$$

The ramified rank bound $\rho \le 2f/b$, together with $f \le x^5$ and $b = \exp(x^\varepsilon)$, gives $\sqrt{\rho} \le \exp(-x^{\varepsilon/8})$ for sufficiently large x. With $E_x^0 = \exp(-x^{\varepsilon/8})$, the bounds $\bar e = 2(E_x^0)^2$ and $\bar\eta = 5x^{-C_H} + 3E_x^0$ work simultaneously at all indicated translates at scale x. Since C_H is fixed,

$$(\eta + 4\sqrt{e})x^{C_H} \le (\bar\eta + 4\sqrt{\bar e})x^{C_H} = 5 + (3+4\sqrt{2})E_x^0 x^{C_H} \le 6 \tag{62}$$

once x is sufficiently large.

3.6 The span of restricted columns

For a finite sequence of complex Hilbert space vectors $a_1, \dots, a_N$ with at least one nonzero vector, let d be the dimension of its span. The generalized Welch inequality for vectors of arbitrary norms states

$$\left(\sum_{i=1}^{N} \|a_i\|^2 \right)^2 \le d \sum_{i,j=1}^{N} |\langle a_i, a_j \rangle|^2.$$

[49, Theorem 1.3 and its proof, author version] For the images of orthonormal bases under the A_w, the established positivity of t ensures a nonzero vector. The inequality gives $t^2 \le d \operatorname{tr}(Q_{\mathcal{H}}^2)$ with $d = \operatorname{rank} Q_{\mathcal{H}}$. By (61), $d \ge D/(1+\eta)$. With $J_w = I$ the same argument gives $D/(1+5\sigma)$.

For $0 \le \alpha \le 1$, let $S_w \subseteq H_w^{\mathrm{act}}$ be complex subspaces with $\sum_w \operatorname{codim} S_w \le \alpha R_{\mathrm{act}}$. Put $M = \sum_w A_w S_w$, let P project orthogonally onto $M^{\perp}$, and let E_w project onto $S_w^{\perp}$. Since $PA_w = PA_wE_w$,

$$\begin{aligned} \operatorname{tr}(PQ_{\mathcal{H}}) &= \sum_w \|PA_wE_w\|_{\mathrm{HS}}^2 \le \sum_w \operatorname{tr}(E_w A_w^* A_w) \\ &\le \sum_w \operatorname{rank} E_w + \sqrt{\sum_w \operatorname{rank} E_w} \sqrt{\sum_w \|A_w^* A_w - I\|_{\mathrm{HS}}^2} \le (\alpha + \sqrt{e}) R_{\mathrm{act}}. \end{aligned}$$

From (61), $t \ge (1-\sqrt{e})R_{\mathrm{act}} \ge R_{\mathrm{act}}/2$. Set $z = \operatorname{tr}(PQ_{\mathcal{H}})/t$ and $p = \operatorname{rank}(P)/D$. Cauchy–Schwarz for the Hilbert–Schmidt pairing of $I - P$ with $Q_{\mathcal{H}}$ gives $(1-z)^2 \le (1-p)(1+\eta)$. Since $0 \le z \le 1$, $p \le \eta + 2z$. Therefore

$$\operatorname{codim}_{\mathcal{H}} M \le (\eta + 4\alpha + 4\sqrt{e}) D. \tag{63}$$

3.7 Projection to the common row dual

We now use $J_w = \jmath_w U_w$. Let $\Pi : \mathcal{H} \to H_{\mathrm{row}}$ be the orthogonal projection to $H_{\mathrm{row}} = s(VE_{\mathrm{row}})$ from (54). Since s is conjugate-linear and the inner product is conjugate-linear in its first argument, the row isometry gives a complex linear isomorphism $\mathcal{T} : H_{\mathrm{row}} \to E_{\mathrm{row}}^{\vee}$ by $\mathcal{T}(y)(F) = \langle s(VF), y \rangle_{\mathcal{H}}$. The common V gives the same map $\mathcal{T}\Pi$ for every word. By (36) the complexification of $B_{K,w}$ at ι_K is

$$B_{K,w,\mathbf{C}} = a_w \sqrt{c_0(w)}\, \mathcal{T}\Pi A_w.$$

In complexified row and column coordinates this is (54), with the common row change and the wordwise changes of basis specified there. Taking the conjugate embedding at the same infinite place conjugates all entries in the identity. Each nonzero scalar $a_w\sqrt{c_0(w)}$ preserves the complex image of the restricted map.

Write $C_K = (R \otimes K)^{\vee} = C \otimes K$ and $D' = \dim_K C_K = \dim H_{\mathrm{row}}$. For K-subspaces $S_w \subset N_w \otimes K$, take their scalar extensions $S_{w,\iota_K} \subset H_w^{\mathrm{act}}$ in the prescribed column models. Their complex codimensions are the original K-codimensions. Put $M = \sum_w A_w S_{w,\iota_K}$. There is a surjection $\mathcal{H}/M \to H_{\mathrm{row}}/\Pi M$. If $D' \ge (1-\delta)D > 0$ for $0 \le \delta < 1$ and $\sum_w \operatorname{codim}_K S_w \le \alpha R_{\mathrm{act}}$, then (63) and preservation of rank under field extension give

$$\frac{\operatorname{codim}_{C_K} \sum_w B_{K,w}(S_w)}{D'} \le \frac{\eta + 4\alpha + 4\sqrt{e}}{1-\delta}, \qquad 0 \le \alpha \le 1. \tag{64}$$

3.8 Counting the distinguished rows

At each g_i, let $\boldsymbol{P} = (P_z)_{z|g_i}$ be the tuple of local polynomial labels, with one factor for each of the f places of F above g_i. These label the finite-group types on the level $g_0 g_1$ reduction of the ambient row space. For $g = g_i$, denote by $z_P \in \mathbf{C}[\mathrm{GL}_r(\mathbf{F}_g)]$ the central idempotent of the sum of irreducibles with label P. The labels are the monic degree-r polynomials with nonzero constant term. The central projectors are orthogonal and sum to the identity. For a set S of label tuples at g_i let $e_{i,S}$ be the full row projector e multiplied by the central projector $\sum_{\boldsymbol{P} \in S} \bigotimes_{z|g_i} z_{P_z}$ in the finite quotient. All other local factors remain the

full row factors. The central projector commutes with e. Thus multiplication selects the allowed label spaces inside the range of e. In particular the resulting multiplicity counts the preimage of S in a basis of full rows chosen in the joint label spaces. Repeated labels contribute the dimension of their entire eigenspace. The finite projection in the distinguished row construction is the central idempotent for good labels at both primes. Its operator is the product of the two single-prime good-label projectors, since on each joint summand their actions are multiplication by the indicator functions of the two good conditions. In the common Petersson model its range is $H_{\mathrm{row}} = s(VE_{\mathrm{row}})$. The common row model in the joint label spaces. Full and distinguished row indices.

The joint good cut in this multiplicity space is cuspidal. Let $P = M_P U_P$ be a proper rational parabolic with first block size $0 < a < r$. For each good tuple some split place $z \mid g_0$ has a label for which no subset of the irreducible factor degrees sums to a. For $a > r/2$ use the complement $r - a$. The finite Green type at that place has no invariants under the unipotent radical of the $(a, r-a)$ parabolic. Indeed, Mackey's formula would put each cuspidal block entirely on one side of the cut. The radical of the actual local P_z contains this radical up to conjugacy in the full compact stabilizer K_z, since an integral basis adapted to its flag makes the parabolic standard.

For a smooth vector in this finite isotypic space, write f_P for its constant term. Fix all coordinates other than z. At an argument with z-coordinate $p_z \in P_z$, evaluation of f_P is invariant under right action on the vector by $U_{P,z} \cap K_z$. Indeed p_z normalizes $U_{P,z}$ and the constant term is left $U_P(\mathbf{A}_F)$- invariant. Averaging the finite K_z-type over this unipotent subgroup therefore gives zero for the evaluation. Iwasawa $G_z = P_z K_z$ gives vanishing for every z-coordinate. The real vectors in question are smooth because they are evaluations of bounded intertwiners on the smooth vector u, and the quotient in the unipotent integral is compact. Hence $f_P = 0$ for every such P. Summing over good tuples proves the assertion on the exact common cut.

For the local character projector a_z of the full row space at $z \mid g_i$, put $\omega_z(P) = (a_z z_P)[1]/a_z[1]$, where $[1]$ takes the identity coefficient. In the regular representation of a finite group G_f^{red}, the trace of left multiplication by a is $|G_f^{\mathrm{red}}|a[1]$. Thus ω_z is a probability law, as $a_z z_P$ are mutually orthogonal projections which sum to a_z. Haar normalization divides all identity coefficients of the finite group algebra by the measure of the principal congruence kernel. It cancels in the ratio, giving

$$M_i(S) = \frac{e_{i,S}(1)}{e(1)} = \sum_{P \in S} \prod_{z \mid g_i} \omega_z(P_z).$$

Each $e_{i,S}$ has the same principal congruence level as e in the sense of bi-invariance. It is self-adjoint, idempotent and of positive type. For example, a self-adjoint convolution idempotent a of compact support satisfies $a(1) = \|a\|_2^2$, and the kernel $a(x^{-1}y)$ is the Gram kernel of its right translates in L^2. The finite support of $e_{i,S}$ is integral and is unipotent modulo g_{1-i}, since the factor at this prime is unaltered. Thus the multiplicity estimate applies with g_{1-i} as the support prime and the same real data as for e. If $D_{i,S}$ is its multiplicity and $\eta_c = e^{-rf \log r/16}$, then

$$D_{i,S} \le (1 + \eta_c) D_\star M_i(S), \qquad D \ge (1 - \eta_c) D_\star.$$

This remains valid when $e_{i,S} = 0$.

The local character projectors. Write $s_0 = m/2 - 1$. The actual local factor a_z is the character projector of the upper unitriangular subgroup J defined by the following block order of lengths $s_0, k, 1, s_0$. Call the middle k indices B, and let t be the singleton immediately following B. Entries from B to column t are zero. The additive character $\psi_g \neq 1$ of $\mathbf{F}_g$ is applied to the sum of consecutive entries on the B-chain and on the complementary chain, which passes through t. The group has index g^k in the upper unitriangular group U_r. These are e_J and its ordinary Whittaker analogue e_U in the finite model. In the adapted split bases the row level uses e_J at every place over the two auxiliary primes. A change of basis conjugates e_J and leaves

its mass against the central z_P unchanged. The argument applies to every nontrivial residue-field additive character, including those obtained in the different local bases. Adapted split bases · Identity coefficients · Projectors at all places · Their common row record. The ordinary Whittaker law $\omega_U(P) = (e_U z_P)[1]/e_U[1]$ is a probability law by the same trace argument. We show that

$$\omega_z(P) \;\geq\; (1 - 1/g)^k \omega_U(P) \qquad \text{for every } P. \tag{65}$$

To see the scalar in (65), first move the left-hand s_0 indices past B, one at a time. In one step let a be the last index before B, and t the first of the remaining indices after it. The missing column from B is at t, and the complementary character uses the entry $a \to t$. Take X to be the root groups in the row $a \to B$, and Y those in the column $B \to t$. Remove X and the internal U_B from the current pattern to obtain the common subgroup N. The subgroups X, Y, U_B normalize N and preserve its character. In matrix multiplication the changes in character coordinates can only involve $a \to t$. Conjugation by X cannot change this entry because $B \to t$ is zero in N. Conjugation by Y cannot change it because $a \to B$ is zero there. Conjugation by U_B fixes all entries outside the rows and columns of B. For a representation V, on V^{N,ψ_g} the commutator of x in X and y in Y acts by the nondegenerate pairing $\psi_g(xy)$ of a row and a column. The averaging operators P_X, P_Y satisfy $P_X P_Y = |Y|^{-1} I$ on X-invariants and $P_Y P_X = |X|^{-1} I$ on Y-invariants, by character orthogonality. They commute with U_B, which normalizes X and Y. Hence replacing X by Y preserves the character multiplicity including the Whittaker condition in U_B. In the new order with a after B, this makes $B \to a$ the missing column. Repeating the exchange leaves B first and the complementary chain as a single tail.

For that final order, first take invariants of the cross-block columns other than the first tail column, then the tail's Whittaker eigenspace. The resulting space V_0 carries the group $P_k = \mathrm{GL}_k(\mathbf{F}_g) \ltimes \mathbf{F}_g^k$, where the translation group is the first cross-block column. Indeed the other cross columns are normal in the relevant block-upper subgroup and the tail unipotent group centralizes the first cross-block column modulo those columns. The rank of the new J-projector is the dimension of the U_k-Whittaker space in V_0. The full U_r-Whittaker rank is the dimension with the additional translation character $z \mapsto \psi_g(z_k)$. By Frobenius reciprocity this latter rank counts copies in V_0 of the irreducible mirabolic representation

$$T_0 = \mathrm{Ind}_{U_k \ltimes \mathbf{F}_g^k}^{P_k}(\psi_{U_k} \otimes \psi_g(z_k)).$$

[45, Proposition 1.2, p. 18] The restriction of T_0 to GL_k is $\mathrm{Ind}_{U_k}^{\mathrm{GL}_k} \psi_{U_k}$, as follows by using GL_k as representatives modulo translations in the induction model. Its U_k-Whittaker dimension is at least $(g-1)^k$. For d diagonal and w_0 the longest permutation, the double cosets $U_k d w_0 U_k$ have trivial stabilizer intersection because $w_0 U_k w_0^{-1}$ is lower unitriangular. Distinct d, d' give distinct cosets, since an upper triangular matrix with diagonal $d'^{-1}d$ can be lower unitriangular only when $d = d'$. Each coset thus supports a Whittaker function. It follows for every V that rank $e_J|_V \geq (g-1)^k$ rank $e_U|_V$. Apply this to the regular representation on the z_P-part and use $e_J[1]/e_U[1] = g^k$ to obtain (65).

The squarefree law. On a squarefree label $P = \prod_h P_h$, with distinct irreducible factors of degrees a_h, Green's character formula gives a unique irreducible in the label and degree $\prod_{j=1}^{r}(g^j-1)/\prod_h(g^{a_h}-1)$. It is the parabolic induction of the cuspidal representations of the distinct factors. The formulas apply with the label convention of the period spaces, since renaming cuspidal parameters within each degree preserves the squarefree formula. [44, Introduction, §7 Theorem 12 and §8 Theorems 13–14] Each cuspidal factor has a Whittaker line. [45, Theorem 0.5 and Theorems 2.3 and 3.3] The open parabolic–unipotent double coset in the induced model supplies the product of those lines, and the full Whittaker rank is at most one by Gel'fand's Theorem 0.5. Applying the regular representation trace calculation to this irreducible, and using $[\mathrm{GL}_r(\mathbf{F}_g) : U_r] = \prod_{j=1}^{r}(g^j - 1)$, gives

$$\omega_U(P) = \prod_h (g^{a_h} - 1)^{-1}.$$

There is a direct way to use these exact masses. For one place, draw a uniform permutation of r indices and, independently for each cycle of length a, draw a uniform element of $\mathbf{F}_{g^a}^\times$. Multiply the monic linear factors for all its a Frobenius conjugates, including repetitions when its degree is smaller than a. For prescribed squarefree P with n_a factors of degree a, the permutation type has probability $\prod_a (a^{n_a} n_a!)^{-1}$. Assign the specified factors to its cycles in $\prod_a n_a!$ ways and choose one of the a roots for each cycle. The resulting probability of P is exactly $\prod_h (g^{a_h} - 1)^{-1}$. Use independent draws at the f places. The product laws have the same probability for every tuple that is squarefree at all places, so they give the same probability to any condition implying squarefreeness at all places.

The bad-label mass. At each $g = g_i$ the good tuples are squarefree at all places, their distinct roots stay distinct after taking their prime-to-ℓ parts for all primes $\ell \notin \mathcal{B}(r, g)$, and no root agrees in prime-to-ℓ part with any of the m prescribed twisting values at its place or their inverses for any $\ell \neq g$. Here

$$\mathcal{B}(r, g) = \{g\} \cup \{\ell : \ell \text{ prime}, \ \ell^{Av_\ell(g^a-1)} > g^a \text{ for some } 1 \le a \le r\}.$$

The tuple must also have no common divisor degree in $[1, \lfloor r/2 \rfloor]$, and have a factor of prime degree p with $r/2 < p \le r - 3$ at some place. Good label conditions and both twist signs.

In the permutation model the probability of a repeated root at a place is $O(r^3/g)$. A cycle of length $a \ge 2$ lands in a proper subfield with probability at most $2ag^{-a/2}$. When all cycles draw elements of their full degrees, a repeat between two cycles requires equal lengths a, and then there are a choices for the second root out of $g^a - 1$. These bounds hold conditionally on the entire permutation. For distinct values u, v, equality of prime-to-ℓ parts means u/v has nontrivial ℓ-power order. For roots from two cycles of lengths a, a', the prime must divide $(g^a - 1)(g^{a'} - 1)$. With the first value fixed, the possible second values occupy at most one coset of the ℓ-primary subgroup of $\mathbf{F}_{g^{a'}}^\times$. For $\ell \notin \mathcal{B}(r, g)$ its size is at most $g^{a'/4}$. For roots in the same cycle of length a, their ratio has the form $u^{g^j - 1}$, $0 < j < a$. There are at most $g^{a/2} g^{a/4}$ choices for u with this ratio of ℓ-power order, since the kernel of this power map has order $g^{\gcd(a,j)} - 1$. Dividing by $g^a - 1$ and summing over at most $2r \log_2 g$ eligible primes and r^2 pairs gives $O(r^3 \log g / \sqrt{g})$ for these collisions.

Each specified twist value is in $\mathbf{F}_g^\times$. The equality test against a root from a cycle of length a has probability at most $(g^a - 1)_\ell / (g^a - 1)$, with $(n)_\ell$ the ℓ-primary factor. Since q, d are distinct primes and $qd \mid g - 1$, at least one of them divides the complementary factor. Thus the probability is at most $1/\min(q, d)$ for each $\ell \mid g^a - 1$. It suffices to consider these primes, because equality of the root and the twist value is already detected by any one of them. A union bound over roots, primes and the $2m$ values gives $O(r^2 m \log g / \min(q, d))$ at a place.

On squarefree outcomes, divisor degrees are the sizes of unions of cycles. For $1 \le a \le r/2$ let Z_a be the number of invariant a-subsets in a uniform permutation. Its expectation is 1. Size-biasing by Z_a is achieved by prescribing a uniform a-subset and drawing independent uniform permutations on it and its complement. A uniform permutation on a nonempty set has a fixed point with probability at least $1/2$, by the first two inclusion–exclusion terms for the events that individual points are fixed. With probability at least $1/4$ in the size-biased law there are fixed points on both sides, and swapping them gives a second invariant a-subset. Consequently $\Pr(Z_a = 0) = \mathbb{E}(Z_a - 1)_+ \ge \frac{1}{2}\mathbb{E}[Z_a \mathbf{1}_{Z_a \ge 2}] \ge 1/8$. Independence gives at most $r(7/8)^f$ for the obstruction of any degree. A prime cycle $r/2 < p \le r - 3$ has probability $1/p$, and these events are disjoint for different p. Their sum is at least $c/\log r$ for an absolute $c > 0$ and all large r, by the lower bound for $\pi(r) - \pi(r/2)$ with at most three primes removed. [72, Corollary 3, (3.8), p. 69] Failure at all places therefore costs at most $\exp(-cf/\log r)$.

For the actual product laws the comparison (65) gives at least $(1-1/g)^{kf}$ times the ordinary product-law mass of the good set. Let S_i^{bad} be the complementary label set. Bernoulli's inequality and the preceding bounds give

$$M_i(S_i^{\mathrm{bad}}) \ll \frac{kf}{g_i} + \frac{fr^3 \log g_i}{\sqrt{g_i}} + \frac{fr^2 m \log g_i}{\min(q,d)} + r(7/8)^f + \exp(-cf/\log r). \tag{66}$$

The comparison uses both internal Whittaker conditions of the actual finite projector.

The selected sizes give $f \geq x^2$, $m, f, r \leq x^6$, $g_i \geq \exp(8x^{\varepsilon_{\mathrm{row}}})$, $\log g_i \ll x^{\varepsilon}$, and $\min(q,d) \geq \exp(x^{\varepsilon_{\mathrm{row}}/6})$, where $\varepsilon_{\mathrm{row}} = \varepsilon/J > 0$ is fixed. Thus each $M_i(S_i^{\mathrm{bad}})$ is at most a polynomial in x times $\exp(-x^{\varepsilon_{\mathrm{row}}/6})$. The complement of the common good space lies in the sum of the two bad ranges. Counting both ranges by their dimensions, including all repetitions, gives

$$\frac{D-D'}{D} \leq \frac{1+\eta_c}{1-\eta_c}\big(M_0(S_0^{\mathrm{bad}}) + M_1(S_1^{\mathrm{bad}})\big).$$

Since $\eta_c \to 0$ uniformly and $x^{\varepsilon_{\mathrm{row}}/6} - x^{\varepsilon_{\mathrm{row}}/7}$ dominates every fixed multiple of $\log x$, this proves

$$D' = (1-\delta)D, \qquad 0 \leq \delta \leq \exp(-x^{\varepsilon_{\mathrm{row}}/7}) \leq \tfrac{1}{2}. \tag{67}$$

3.9 From ranks to the mean slope

For each word use $N'_w = I_w N_w$ and $C = R^\vee$ from (53). Multiplication by a nonzero fractional ideal I_w preserves the K-span of N_w. Thus N'_w has rank r_w and its K-space has the prescribed realization $\mathcal{M}_w$ at ι_K. Division in Section 4 gives $B_{K,w}(N'_w) \subseteq C$. The comparison below uses the squared-norm archimedean height bound h_w for this same map in the prescribed metrics on N'_w and C.

Use the squared-norm slope $\mu(L) = 2\widehat{\deg}_{\mathrm{n}}(L)/\operatorname{rk}_{\mathcal{O}_K} L$, where $\widehat{\deg}_{\mathrm{n}}$ is the Arakelov degree divided by $[K:\mathbf{Q}]$. In the Euclidean realization with one copy of $\mathbf{C}^{\operatorname{rk} L}$ for each complex place, the release doubles the slope from $-\log$ covolume per real dimension and adds $-\log 2 + \log|\operatorname{disc} K|/[K:\mathbf{Q}]$. Indeed, the reference covolume per K-coordinate is $2^{-[K:\mathbf{Q}]/2}\sqrt{|\operatorname{disc} K|}$. conversion from Euclidean covolume to the normalized slope Equip the K-spans of the Hermitian lattices with their given archimedean norms and their lattice unit balls at the finite places. For a nonzero K-linear map B between these Hermitian adelic bundles, write B_v for its scalar extension to the completion K_v. Its normalized adelic and squared-norm archimedean heights are

$$h_{\mathrm{ad}}(B) = \frac{1}{[K:\mathbf{Q}]}\sum_v n_v \log\|B_v\|_{\mathrm{op}}, \qquad h_\infty(B) = \frac{1}{n_\infty}\sum_{v\mid\infty} \log\big(\|B_v\|_{\mathrm{op}}^2\big).$$

Here $n_v = [K_v : \mathbf{Q}_p]$ for $v \mid p$ and $n_v = 2$ for the $n_\infty = [K:\mathbf{Q}]/2$ complex places. The absolute values at finite places extend the usual p-adic ones. [50, Definition 6.1, arXiv version] [51, Equation (20), p. 34] Rescaling an arithmetic map by $\lambda \in K^\times$ multiplies its image lattice by λ and adds $2\log|\mathrm{N}_{K/\mathbf{Q}}\lambda|/[K:\mathbf{Q}]$ to $h_\infty(B)$. For a unit of $\mathcal{O}_K$ both the lattice image and the archimedean height are preserved. The product formula gives $h_{\mathrm{ad}}(\lambda B) = h_{\mathrm{ad}}(B)$ for every $\lambda \in K^\times$.

For a Hermitian lattice L and any saturated sublattice M with the induced norms, $\widehat{\deg}_{\mathrm{n}}(L) = \widehat{\deg}_{\mathrm{n}}(M) + \widehat{\deg}_{\mathrm{n}}(L/M)$ with the quotient norms on L/M. The Harder–Narasimhan polygon of L is the upper concave hull of $(\operatorname{rk} M, 2\widehat{\deg}_{\mathrm{n}}(M))$ over these sublattices, including 0 and L. There is a unique saturated flag whose rank-degree pairs are exactly the vertices. The slopes of the successive quotient lattices are the segment slopes in decreasing order. [50, Propositions 4.22 and 5.13, arXiv version] Write $\lambda_j(L)$ for these slopes, repeated according to the ranks of the quotients. Their mean is $\mu(L)$ and their extrema are $\mu_{\min}(L)$ and $\mu_{\max}(L)$. The saturated step $L_{\geq a}$ retains the factors of slope at least a and has rank $\#\{j : \lambda_j(L) \geq a\}$.

For the usual slopes, which are half the squared-norm slopes, put $\mathcal{F}_u^L = L_{\geq 2u} \otimes_{\mathcal{O}_K} K$ in $L_K = L \otimes_{\mathcal{O}_K} K$ and define $\mathcal{F}_u^C$ in the same way. Chen's filtration shift gives, for every nonzero K-linear map $B : L_K \to C_K$ and every $u \in \mathbf{R}$,

$$B(\mathcal{F}_u^L) \subseteq \mathcal{F}_{u-h_{\mathrm{ad}}(B)}^C.$$

[51, Proposition 2.2.4, p. 42] If $B(L) \subseteq C$, the finite operator norms are at most 1, so $2h_{\mathrm{ad}}(B) \leq h_\infty(B)$. If $h_\infty(B) \leq h$, apply the inclusion at $u = (a+h)/2$ and intersect with C to obtain $B(L_{\geq a+h}) \subseteq C_{\geq a}$ for every $a \in \mathbf{R}$. The lattice inclusion also holds for the zero map.

Choose uniformly from the R_{act} donor slopes, with repetitions according to rank. The adjusted value is $X = \lambda_j(N'_w) - h_w$ when the chosen slope is the j-th of N'_w. Set

$$s = \mathbb{E}X = \frac{\sum_w r_w(\mu(N'_w) - h_w)}{R_{\mathrm{act}}}, \qquad U = \max X - s \geq 0.$$

For $\tau \geq 0$, take the steps $(N'_w)_{\geq s-\tau+h_w}$. The discarded rank fraction is $\alpha_\tau = \Pr(X < s - \tau)$. The retained periods land in $C_{\geq s-\tau}$. Their K-spans are subspaces of the same actual column models. By (64), the fraction of target slopes below $s-\tau$, denoted by β_τ, is at most $(\eta + 4\alpha_\tau + 4\sqrt{e})/(1-\delta)$. Choose $H_* \geq (s - \mu_{\min}(C))_+$. Integrating the two tails gives

$$s - \mu(C) \leq \frac{1}{D'} \sum_{j=1}^{D'} (s - \lambda_j(C))_+ = \int_0^{H_*} \beta_\tau \, d\tau,$$

$$\int_0^{H_*} \alpha_\tau \, d\tau = \mathbb{E} \min\{H_*, (s - X)_+\}.$$

Write $Z = s + U - X \geq 0$, with $\mathbb{E}Z = U$. If $H_* + U > 0$,

$$\min\{H_*, (Z - U)_+\} \leq \frac{H_* Z}{U + H_*}.$$

For $Z \leq U + H_*$, the right side bounds both 0 and $Z - U$. For $Z \geq U + H_*$ it bounds H_*. Thus

$$s - \mu(C) \leq \frac{4UH_*/(U + H_*) + (\eta + 4\sqrt{e})H_*}{1 - \delta}.$$

The fraction $UH_*/(U + H_*)$ is zero when either variable is zero, including their common zero. At $U = 0$ every adjusted donor slope equals s. At $H_* = 0$ every target slope is at least s.

In particular $UH_*/(U + H_*) \leq U$. For $U > 0$, one also has $UH_*/(U + H_*) \leq U \log(1 + H_*/U)$, by integrating $U/(U+\tau)$ on $[0, H_*]$. The form used in the arithmetic comparison uses this logarithmic bound and adds the nonnegative $\log|W|$.

$$s - \mu(C) \;\leq\; \log|W| + \frac{4U \log(1 + H_*/U) + (\eta + 4\sqrt{e})H_*}{1 - \delta}. \tag{68}$$

The U-term is zero at $U = 0$. The slopes of a Hermitian dual are the negatives of the original slopes in reverse order. [50, Properties 5.14(2), Hermitian case, arXiv version] This gives $s - \mu_{\min}(C) = s + \mu_{\max}(R)$ in Section 4. The height comparison there bounds this quantity by $x^{C_H} b$ with the same C_H as in the word count. Taking $H_* = \max(0, s - \mu_{\min}(C))$, the numerator satisfies $(\eta + 4\sqrt{e})H_* \leq 6b$ by (62). Since $1 - \delta \geq 1/2$ by (67), this term contributes at most $12b$. Using $UH_*/(U + H_*) \leq U$ also gives

$$s - \mu(C) \leq 8U + 12b. \tag{69}$$

4 Period lattices and centered prime-power sums

Fix a sufficiently large scale x at which the following comparisons hold. Choose a base index i_0, with CM field E, and an index $i(u)$ for each translate $u \in H$. The conditions at a common scale identify the rank parameters, auxiliary primes, root discriminant, and sampling sets T, T_2 under the specified identifications. The base index also seeds the common sampling degrees a_t for the family chosen in Section 6.1. The construction at this scale records the structural identifications. The two applications here also need the section-run and ramified donor identifications stated in Section 6.1. The weighted prime-power sums at the translates use the Artin map and the centering character of the base field. The stage weights $w_{u,j}$ are read at the corresponding index $i(u)$. At that index we use its Selmer estimate at its actual twist cutoff and its own period maps $B_{K,w}$, $w \in W(u)$, into the rational space underlying the common row dual $C = C(u)$. In what follows the period lattices and their coefficient field K refer to a fixed translate. All estimates are uniform across the translates.

Put $n = |H|$, $f = 2n$, $N_T = |T|$, $N_2 = |T_2|$, $L = x^{\varepsilon}$, $b = \exp(L)$, and $a_E = xb$. The selected family has $0 < \varepsilon < 1/52$. The size hypotheses for the lattice estimate include

$$x \geq 1, \qquad N_T \geq 1, \qquad m \leq x^{1/2+\varepsilon/2}, \qquad xm \leq k. \tag{70}$$

The construction supplies these hypotheses from the size window and the positive lower bound on the sampling density. Since $\varepsilon < 1$, we have $m \leq x^{1/2+\varepsilon/2} \leq x$, so $mb \leq a_E$. Section 5 uses the same size data. The scales for ordinary and exceptional finite-place errors are $B_0 = 2nmbx^{\varepsilon}$ and N_c. Fix nonnegative exponents e_D, e_L for the division and gain estimates, respectively.

4.1 From valuations to integral maps

The nonnegative weighted Selmer lengths in (6) have bounded finite subsums. Taking the supremum over those subsums gives a convergent sum $V \leq \sqrt{b}$ at the actual twist cutoff. For a finite place v of K over ℓ write $e_v = \operatorname{ord}_v(\ell)$ and $\lambda_v = \log \mathrm{N}\,\mathfrak{p}_v/[K : \mathbf{Q}]$.

At a coefficient place $v \mid \ell$ with nontrivial prescribed twist $\boldsymbol{\mu}(v)$, let $\mathcal{A}_{jv}$ index the local Hecke blocks of the twisted multiplicity space at stage j. For a block $\alpha \in \mathcal{A}_{jv}$, write O_α for its coefficient DVR, M_α for the integral lattice, and V_α for its generic fiber. Its split Hecke algebra has generic eigenspaces $V_{\alpha,\pi}$, indexed by the distinct eigencharacters π. Let V_α^+ be the sum over the distinguished good eigencharacters and P_α the projection along the other eigenspaces. The defect is

$$\delta_\alpha = \frac{\operatorname{length}_{O_\alpha}\bigl(P_\alpha M_\alpha/(M_\alpha \cap V_\alpha^+)\bigr)}{\operatorname{ord}_\alpha(\ell)},$$

$$\eta_{jv} = \sum_{\alpha \in \mathcal{A}_{jv}} \delta_\alpha, \qquad\qquad \eta_j = \sum_{\substack{v \in \operatorname{supp} \\ \boldsymbol{\mu}(v) \neq \boldsymbol{\mu}_0}} e_v \lambda_v \eta_{jv}.$$

The two lattices in the quotient span V_α^+, so its length is finite. The denominator puts the length in units with $v_\ell(\ell) = 1$. The definition agrees with the source's upper and lower lattices, since $P_\alpha M_\alpha = V_\alpha^+ \cap (M_\alpha + \ker P_\alpha)$. Local Hecke lattices and stage sum.

Type projectors on the localized Hecke lattice. Fix the selected index, nonbase tuple, stage and coefficient place, and the resulting local Hecke record M. At each place v of F over g_k, take the copy of $\mathrm{GL}_r(\mathbf{F}_{g_k})$ on the first block of the principal-level frame on the chosen split half. Let G be their product for $k = 0, 1$. The class-value model gives this action on the ambient principal-level forms and their integral value lattice. In fact an integral lift of a finite matrix translates the class representatives and transports the coefficient lattice in the same frame to the lattice of the translated class. This is the unipotent right action in the definition of the specialized row generators. In the ψ_0-eigenspace of the unipotent group

use the identical frame injection into this principal-level model. A central operator commutes with the ψ_0-projection and preserves the integral value lattice before passing to its level subspace. Its action on the level lattice is then obtained by taking the preimage under the frame injection. This argument also works for the image of the lattice under the projection. In particular no averaging denominator from the unipotent group enters the central operator. The action commutes with the normalized split Hecke operators, whose places are outside the two type primes. It therefore descends to M with its original residual Hecke projection. The class-value action and the integral lattices, Central action in the level frame, and The action on the localized lattice.

For a label s at a factor G_v of G, write $e_{v,s}$ for the sum of the central idempotents of all irreducible finite types of that label. For an irreducible type τ its idempotent is

$$e_\tau = \frac{\dim \tau}{|G_v|} \sum_{g \in G_v} \operatorname{tr} \tau(g^{-1})\, g.$$

The traces are sums of roots of unity of orders dividing $|G_v|$. The selected coefficient field contains the $|G|$-th roots of unity, so $|G| e_{\mathbf{s}}$ has integral coefficients for $e_{\mathbf{s}} = \prod_v e_{v,s_v}$. Let e_{good} be the sum of $e_{\mathbf{s}}$ over the good tuples of labels, with each tuple included once. The same integrality holds for $|G| e_{\mathrm{good}}$. The cyclotomic coefficient field and The two-prime central element.

Here is the local identification used to read these projectors on an actual automorphic fibre. At a chosen split half the block parahoric vector transforms on reduction under $\left(\begin{smallmatrix} I_r & * \\ 0 & D \end{smallmatrix}\right)$ by $\mu_v(\det D)$, with μ_v the indicated tame block character. This follows from the class-value transformation law at the same tuple and stage. If h is a lift in this subgroup, the lower-right blocks of $\overline{h}$ and $\overline{h}^{-1}$ are D and D^{-1}. Under the opposite split identification one takes the inverse transpose in the corresponding dual frame. Thus its determinant character is inverse. The congruence $\det((h^{-1})_{22}) \det(h_{22}) \equiv 1 \pmod{\mathfrak{p}_v}$ is enough, since the character on units is tame. This uses the determinant of the lower-right block in the parahoric vector, including for fibres without the prescribed global Speh block. The conjugate block determinant and The block character of the realized vector.

Suppose $e_{v,s}$ acts nontrivially on this vector, where s is monic, squarefree of degree r, has nonzero constant coefficient, and excludes both allowed scalar labels. The parahoric restriction gives a nonzero first-block type of label s against the second-block determinant character. To see its local consequence, use the normalized Jacquet module for block sizes (r, t). The depth-zero Jacquet comparison sends the indicated parahoric type to a $\mathrm{GL}_r(\mathcal{O}_v) \times \mathrm{GL}_t(\mathcal{O}_v)$ type, with the first principal-congruence kernel acting trivially and the second factor acting by the same determinant character. In a cell of the geometric lemma above a cuspidal support, the second compact type forces its t entries to be rank-one cuspidal support characters inertially equivalent to μ_v. Indeed the spherical-support criterion for GL_t, applied after twisting away μ_v, makes those characters unramified. If the full support has more such entries, one goes to the first block in that cell. The finite Harish–Chandra support then makes its scalar label a root of s. This is excluded. Thus there are exactly t entries of this inertia character. In the Bernstein–Zelevinsky classification [83, Sections 4 and 9] the two groups of cuspidal support types are on disjoint inertial lines. Their normalized product is irreducible, so the local representation is the product of a rank-r representation and a rank-t representation with only the scalar inertial line. The same Jacquet comparison puts the first-block type on the rank-r factor and gives a spherical vector on the rank-t factor after dividing by μ_v. In that comparison any other cell sending first-block cuspidal support to the second block violates the spherical-support criterion there.

Under local Langlands, depth-zero finite cuspidal support on the rank-r factor gives the Teichmüller lifts of the roots of s, as eigenvalues on the chosen tame generator. Each depth-zero cuspidal block of degree a contributes the a Frobenius conjugates of its finite-field parameter, with products for parabolic support. The full degree and nonzero constant coefficient account for all r eigenvalues in a tame inertia representation. Its monodromy vanishes because nilpotent monodromy preserves its one-dimensional inertia eigenspaces. The μ_v-twist of the spherical rank-t factor also has zero monodromy, and its inertial

polynomial is the scalar factor $(X-u)^t$ in finite-field labels, where u is the label of μ_v. In particular the full local polynomial is $s(X)(X-u)^t$. At the conjugate place the roots are inverted, with the generator transported by conjugation. These statements are on the finite vector used by the original group-algebra action. In particular, distinctness of the roots of s alone would give no monodromy assertion on the repeated scalar block. Jacquet cells and separation of the scalar support, Monodromy and the local polynomial on both blocks, and Application to the realized finite type.

At $\ell = g_k$ the local-global comparison is made on the same local representations of the regular algebraic cuspidal seeds, with their unramified normalization twists. It identifies the inertia semisimplifications and bounds the Galois monodromy by that of the automorphic parameter [122, Theorem 1.1]. In the spherical-after-twist scalar block, as well as in the squarefree first block, the automorphic monodromy is zero, so the Galois monodromy is zero. After the prescribed finite tame extension that kills the inertia type, these restrictions are crystalline.

For the first projection take X_{type} to be the set of Hecke fibres on which e_{good} has nonzero image. For a fibre carrying the prescribed rank-one Speh block, the automorphic transfer determines the remaining local types. The split Hecke polynomials give equal characteristic polynomials on a dense set of Frobenius classes in the CM field. The places with relative residue degree two have density zero. Continuity and Brauer–Nesbitt give equality of the semisimple Galois representations. The unramified comparison then gives the same isobaric base change by strong multiplicity one. For every occurrence in that fibre, the level-compatible first-block constituents have the tuple of semisimple labels obtained by removing the prescribed scalar block. Thus the central projector is the identity on the whole good fibre and zero on every fibre with the prescribed Speh block and a bad tuple. A fibre in X_{type} may have other occurrences or lack that global Speh block. Nevertheless there is an occurrence in that fibre and at every type place a compatible local finite type with the required labels. The preceding local calculation applies to it. Every other occurrence over the same Hecke point has the same isobaric transfer and hence the same local parameters at these split places. This proves the compatibility for every generic fibre retained in X_{type}. Action on the fibres carrying the prescribed block and The coarse support and all its retained fibres.

Project M to X_{type} in characteristic zero and use its image lattice. The good projection is unchanged. To recover the old intersection, multiply a preimage by $|G|e_{\mathrm{good}}$, which kills the discarded fibres and is $|G|$ on the good part. The normalized cost is at most $\sum_{\pi\ \mathrm{good}} \dim V_\pi\, v_\ell(|G|)$. At a type residue prime, for each good π also use $|G|e_{\mathbf{s}(\pi)}$ with its *specified* full tuple at the two type primes. Its support X_π contains π. All fibres in X_π have the same first-block labels at the active type prime, and any point there with the prescribed Speh block is good. Its action on π is $|G|$, giving the separate cost $\dim V_\pi\, v_\ell(|G|)$ in the pointwise congruence-ideal comparison. The local estimate uses only the type at the active prime; the scalar cost uses the full group G. The valuation bound in the type-prime comparison applies to this scalar. The coarse projection is charged across all primes by the separate allowance $\log|G| \le fr^2(\log g_0 + \log g_1)$. Recovering the original lattice defect and The individual tuple cut.

The first character detects the specified block. Fix a nonbase local record containing a good point, and take its common residual Hecke character. Write $n_V = r+t$, with $r>1$ and $t\ge 2$. The good point has the rank-r absolutely irreducible residual constituent and the t residual character factors of the prescribed Speh block. By the common characteristic polynomials in the integral Hecke algebra and Brauer–Nesbitt, every generic point in the record has residual Jordan–Hölder dimensions $r, 1, \dots, 1$, with t ones counted with multiplicity. This applies at every coefficient prime, including places where some of the characters have equal reductions.

For each unitary cuspidal seed of a Speh summand of length d, use its attached representation with initial arithmetic Frobenius exponent $u_0 = (1-n_V)/2 + (d-1)/2$. The d copies are its twists by ϵ_ℓ^{-k}, $0 \le k < d$. At split unramified places the unitary seed's Satake roots have absolute value one. Indeed the purity clause of BLGGT, Theorem 2.1.1(2), gives a common absolute value after an algebraic normalization, and the unitarity of the unitary seed's central character gives value one after undoing the norm shift. The

distinct infinitesimal entries of the full row give regularity of each seed after that shift. This is the same cuspidal attachment and normalization used for the full Hecke polynomial of the record.

If R_x contains χ_1, semisimplicity puts that line in one of these copies. The prescribed first character has Frobenius absolute value $q_w^{(1-n_V)/2+(t-1)/2}$. Comparison at any split unramified $w \nmid \ell$ gives $d-1-2k = t-1$. Thus $d \geq t$. The untwisted seed contains the line $\chi_1 \epsilon_\ell^k$, so its reduction has a rank-one Jordan–Hölder factor. All d copies contribute such a factor, whence $d \leq t$. We obtain $d = t$ and $k = 0$. Any other residual factor in the same seed would occur at least twice. The t one-dimensional factors are already accounted for, and the rank-r residual factor occurs once. Hence the seed has rank one. Equality of Frobenius characters after undoing the common norm shift identifies it with the specified rank-one automorphic character. The remaining residual multiset is the singleton $\{r\}$. Each remaining seed contributes at least one nonempty collection, repeated according to its Speh length, so there is one remaining seed of rank r and length one. Conversely the specified Speh block supplies χ_1. On the coarse type cut every point with this block is good. Therefore away from the good points the congruence relation module sees no first-character line. This gives the support used in the determinant comparison with the original Hecke algebra. [121, Theorem 2.1.1(2)–(4)] and The first character and the cuspidal complement.

In the projected Hecke algebra, let $\mathfrak{c}_{\alpha,\pi} \subset O_\alpha$ be the evaluations at a good π of the integral Hecke operators vanishing on the other, nongood summands. Order the good summands and filter the projected and intersected lattices by their successive partial sums. Apply an operator realizing a generator of $\mathfrak{c}_{\alpha,\pi}$ to a lift of an element of the π-graded piece of the projected lattice. The result is in the intersection lattice and its π-component is multiplied by this generator. Hence the normalized graded length is at most $r_{\alpha,\pi} v_\ell(\mathfrak{c}_{\alpha,\pi})$, where $r_{\alpha,\pi} = \dim V_{\alpha,\pi}$. The sum of these lengths bounds the defect on the type locus. Comparison after the type cut.

For the Galois comparison, let A be the integral Hecke algebra in one type block and $\mathscr{R}$ the A-span of the prescribed Galois matrices in the generic fibres. Their characteristic coefficients lie in A. This implies the same assertion for $M = \sum_i a_i r(g_i)$ in $\mathscr{R}$. In characteristic zero, expand $\log\det(1 - sM)$ as the formal trace series and group the words by their primitive cyclic period. A primitive word of length q has q rotations, so its powers sum to the logarithm of the factor $\det(1 - s^q M_{\mathrm{word}})$. The word matrix is a scalar times a Galois matrix, and only finitely many factors contribute to any coefficient. For a square array of blocks, use primitive closed walks among the blocks in the same argument. Each walk matrix is in $\mathscr{R}$. Thus determinants of square block arrays lie in A, also after adjoining scalar indeterminates.

The column relation module N in the comparison has a nonzero map to the generic Galois fibre at each nongood label. Its finite presentation gives a relation matrix P_N. At a chosen good label let B_i be the torsion module of its i-th diagonal constituent. The kernel of the i-th relation matrix on copies of B_i is $\mathrm{Hom}_{\mathscr{R}}(N, B_i)$. If its length is finite, Smith form gives a square minor after row compression whose determinant has valuation equal to that length. The evaluation of the order onto each constituent matrix ring is surjective by residual absolute irreducibility and Nakayama’s lemma. Lift each optimal compressor separately and take their linear combination with scalar variables. For each diagonal block the determinant polynomial after compression is divisible by the product of its Smith factors. At the variable corresponding to that block the normalized polynomial takes a unit value. The product of the normalized polynomials is nonzero modulo the uniformizer. It follows that the determinant of the compressed block array has a coefficient with valuation equal to the sum of the Smith lengths. All these coefficients vanish at every nongood label because of the nonzero generic kernel. Hence this coefficient lies in the Hecke congruence ideal of the chosen good label. The argument supplies the congruence length bound with coefficient one.

In applying this argument the character constituents can collide modulo the uniformizer. The combined character corner is selected by lifting its residual idempotent, distinguished from the residual irreducible constituent of rank at least two. If a is a lift, iteration of $a \mapsto 3a^2 - 2a^3$ squares the idempotence error up to an integral factor. The common residual characteristic polynomial makes the error nilpotent modulo the uniformizer in each fibre, so the iteration converges in the complete order. An inertia element

with cyclotomic value of infinite order separates the characters at the good point. The character corner polynomial can be made separable at the other generic labels by a sufficiently divisible perturbation that vanishes at the good point. The distinct absolutely irreducible generic constituents give full rational matrix factors, which contain diagonal matrices of simple spectrum on the entire corner. Adjoin the ordered roots of the resulting monic polynomial, removing each root before the next monic adjunction. The resulting algebra is free over A by successive division by monic polynomials. Its monomial basis is a basis of functions on the generic root orderings at each label after splitting. A determinant coefficient vanishing on all orderings over nongood labels therefore has basis coefficients vanishing on the original nongood labels. Evaluation at the chosen good ordering belongs to the original congruence ideal. Flat coefficient extension preserves this vanishing ideal and normalized lengths. Character corner and ordered roots and application to the column relations.

Relations in a realized Hecke block. Work in one nonbase local Hecke record. Use its normalized split Hecke algebra $A \subset O^X$ and actual lattice M generated by the localized integral class forms, with O finite over $\mathbf{Z}_\ell$. If necessary, extend this DVR to split the finitely many generic fibres. This preserves the normalized lengths of the Hecke lattice and its congruence ideals. The finite Hecke record is taken after specializing the common ray character. Its Galois representations are the r_x of dimension $d = r + t$, on the same eigencharacters $x \in X$. At a good eigencharacter π their constituents are $\rho_\pi, \chi_1, \dots, \chi_t$, with $\chi_c/\chi_1 = \epsilon^{1-c}$. Let Z denote the locus where χ_1 occurs. The good-type cut retains its prescribed good part of Z. The characteristic coefficients of $r_x(g)$, as functions of x, belong to A. This follows from the normalized split Hecke characteristic polynomial at Frobenius, Chebotarev over the CM field, and closure of A in the finite product of coefficient rings. CM primes of relative residue degree two have density zero and can be omitted from Chebotarev. The transport to the initial finite Hecke record uses the localized integral class forms themselves. The integral Hecke record and The characteristic coefficients and constituents.

Use the ordered-root extension A' of this algebra and the corresponding matrix order R' from the determinant comparison. Write $u(g) = r(g)\chi_1(g)^{-1}$ in R'. Let E be the scaled rank-one projector for the chosen ordered root in each generic fibre, with $E^2 = \delta E$ and nonzero $\delta(x')$ at every ordering x' over x. At the chosen good ordering its image is the first character. Let e denote the integral rank-one projector in the complementary rank-r corner. At the ordering over π it is a rank-one projector of the ρ_π-block. Write $\Delta = v_\ell(\delta(\pi'))$. The matrix

$$y(g) = (u(g) - 1)E$$

generates a left R'-module Y and satisfies $y(gh) = y(g) + u(g)y(h)$. Each of the R'-modules $Y, R'e, R'E$ is finite free as an O-module, since it is an O-submodule of the matrix order.

For a column place λ let $W = \rho_\pi \chi_1^{-1}$, with its recorded lattice T_ρ. Put $j_\lambda = \dim H^1_f(E_\lambda, W)$. The image of $H^1_f(E_\lambda, T_\rho)$ is a full lattice in this space. Choose integral finite cocycles $\kappa_{\lambda,a}$ representing a basis of that lattice. Evaluation of $R'e$ on a primitive vector of eT_ρ surjects onto T_ρ, since the evaluation of the order on the rank-r block is all of $\mathrm{End}_O(T_\rho)$. Evaluation of $R'E$ in the trivial u-constituent surjects onto $\delta(\pi')O$. These are surjections of integral lattices whose rational representations are subquotients of finite sums of the u_x. Apply the column lifting estimate with exponent e'_λ to the two surjections. It gives finite cocycles $K_{\lambda,a}$ in $R'e$ and $K_{\lambda,0}$ in $R'E$ evaluating, respectively, to $\ell^{e'_\lambda}\kappa_{\lambda,a}$ and $\ell^{e'_\lambda}\delta(\pi')\,\mathrm{ur}_\lambda$. Here ur_λ is the unramified cocycle which is 1 on Frobenius. A lift in cohomology gives an equality of cocycles after subtracting the coboundary of a lift of the evaluating vector. This uses the stated integral surjectivity on lattices. The two finite cocycle lifts.

Use the actual decomposition groups D_λ. The relation module N is the quotient of Y plus free left R'-generators n_λ, $s_{\lambda,a}$, $s_{\lambda,0}$ by

$$En_\lambda = 0, \qquad (1-e)s_{\lambda,a} = 0, \qquad (\delta - E)s_{\lambda,0} = 0,$$
$$y(g) + (u(g)-1)n_\lambda = \delta\left(\sum_a K_{\lambda,a}(g)s_{\lambda,a} + K_{\lambda,0}(g)s_{\lambda,0}\right), \qquad g \in D_\lambda.$$

The order is finite over O, so all these relations have a finite presentation. For a generic point x', the cochains obtained by evaluating the two source left ideals at vectors in the corresponding projector lines are rational finite cocycles. The finite local dimension of u_x is $j_\lambda + \dim H^0(E_\lambda, u_x)$. This is the Bloch–Kato dimension formula and the common labeled Hodge jumps, since the scalar constituents at π have jumps $0, \dots, t-1$ and the selected W has no local invariant vector. Thus, when $H^0(E_\lambda, u_x) = 0$, the $j_\lambda + 1$ evaluated cochains are dependent modulo coboundaries. When $H^0 \neq 0$, an invariant vector gives a nonzero coboundary relation with all scalar coefficients zero.

To see the nonvanishing needed for the determinant, at each column take such a nonzero relation

$$(u_x(g)-1)p = \sum_a b_a K_{\lambda,a}(g)v_e + b_0 K_{\lambda,0}(g)v_E,$$

where v_e, v_E span the indicated rank-one images in the generic fibre. If one relation has $Ep = 0$, map Y to zero, map this place's generators to $\delta(x')p$, $b_a v_e$, $b_0 v_E$, and map the other generators to zero. This is a nonzero map from N to that generic fibre. Otherwise rescale each relation so that $Ep = v_E$. Choose z with $Ez = v_E$ and map $a \in Y$ to az. Map the place's generators to $\delta(x')p - v_E$, $b_a v_e$, $b_0 v_E$. These assignments satisfy the relations, using $Ev_E = \delta(x')v_E$. At $x \notin Z$ the vector v_E is moved by some $u_x(g)$, so the map is nonzero on Y. The actual relation module and Its generic maps.

The localized cocycles of the relations. For a specialized torsion constituent $B = V/T$, a map $F : N \to B$ gives the cocycle $c_F(g) = F(y(g))$. This is continuous because a map from the finite O-module into the discrete torsion module factors through a quotient killed by a power of the uniformizer. The cocycle on a column place belongs to $H^1_f(E_\lambda, B)$ for $B_\rho = W/T_\rho$. Indeed the $K_{\lambda,a}$ evaluate to multiples of the chosen finite cocycles, and the column equation is an equality up to a coboundary. In the higher scalar blocks the local class is zero, since e and E act by zero. In the first scalar block the equation gives an unramified class, using $K_{\lambda,0}$.

For a row place, let L_λ be the O-span of the $y(g)$ for $g \in D_\lambda$. It is stable under the local action by the cocycle identity. The lattice $L_\lambda^+ = L_\lambda + OE$ is stable as well, and $L_\lambda^+/L_\lambda = O$ with trivial action. In fact evaluation on the first scalar block at π' kills L_λ and is nonzero on E. Both lattices are in finite direct sums of the actual u_x, by taking matrix columns in every fibre. Where these representations are crystalline, the extension by E shows that y is an integral finite cocycle: its rational extension is itself a subrepresentation of a crystalline representation. The same reasoning applies after any common extension over which the u_x are crystalline.

The common row lattice. For the ordinary row comparison the fixed dual lattice is the right ideal eR', with action on a matrix a given by $g \cdot a = \epsilon(g) a\, u(g)^{-1}$. It lies in a finite sum of the duals $u_x^\vee(1)$, by taking matrix rows in each fibre. At π' write $e_\pi = v_e\, v_e^*$, with $v_e^*(v_e) = 1$ over O. The map

$$p : eR' \longrightarrow T_\rho^\vee(1), \qquad a \longmapsto v_e^* a(\pi')\,\big|_W$$

is equivariant and surjective by integral matrix algebra surjectivity on the rank-r block. The row hypothesis bounds the colength l_λ of the image under p of the integral finite classes. This image is fixed for the whole relation module.

Indeed let $a(g)$ represent one such finite class on eR'. Regard R' and its submodule Y in the following cup calculation as modules with scalar action ϵ. The cochain $B(g) = -a(g)u(g)E$ satisfies

$$dB(g,h) = a(g)u(g)y(h).$$

This follows from $a(gh) = a(g) + \epsilon(g)a(h)u(g)^{-1}$. The right side is Y-valued. Its class is zero already in $H^2(E_\lambda, Y(1))$, since the map from this group to $H^2(E_\lambda, R'(1))$ is injective. In fact Smith form reduces the latter map to multiplication by nonzero scalars on copies of $H^2(E_\lambda, O(1)) \simeq O$. Apply $v_e^* \circ F$ to the resulting Y-valued coboundary. Its value in degree two is precisely the cup of $p(a)$ with c_F, by left R'-linearity of F. This works for every F with the same image of integral dual classes. Local Tate duality identifies its annihilator modulo the torsion finite condition with the dual of the common colength quotient. Hence the whole image has the row allowance l_λ. The integral right ideal, Its fixed target map, and Injectivity for the integral cup product.

The scalar row calculation similarly applies to $L_\lambda^+ \to O$, with the primitive vector E over 1. For target $(\operatorname{Frac} O/O)(-a)$ it puts $\ell^{c_B} c_F$ in the finite torsion condition. For $a > 0$ this condition is zero, and for $a = 0$ it is the unramified subgroup. Multiplication by ℓ^{c_B} and the local Euler characteristic bound the image modulo the desired condition by $c_B([E_\lambda : \mathbf{Q}_\ell] + 2)$ in normalized length. For the trivial target one uses divisibility of the unramified subgroup, so the ℓ^{c_B}-torsion of the quotient is an image of the ℓ^{c_B}-torsion of the whole group. The common row lattice and The scalar finite condition.

The normalized kernel and index count. For the rank-r and higher scalar blocks, suppose $c_F(g) = (u(g) - 1)v$. Choose an expression $E = \sum a_g r(g)$ in the order and put $s = \sum a_g \chi_1(g)$. Then $s(\pi') = \delta(\pi')$ by evaluation on the first scalar block, and $\sum a_g \chi_1(g) y(g) = (\delta - s)E$ belongs to Y. Left multiplication by E on this element is multiplication by δ, while E acts by zero on the target. Applying F proves $\delta(\pi')^2 v = 0$. The choice of the global coboundary vector costs at most $2\Delta \dim V$, and once chosen it determines $F|_Y$.

At a rank-r column, the torsion finite classes associated to the chosen integral basis make the a-th slot a rank-one torsion module killed by $\ell^{e'_\lambda}\delta$. The 0-slot is killed by δ. After these choices the remaining ambiguity of $F(n_\lambda)$ is an invariant vector in the torsion target. On a higher scalar block the a-slots are zero, leaving the 0-slot killed by δ and the corresponding invariant space. On the first scalar block $F|_Y = 0$ in the kernel. The n_λ-slot is killed by δ, the a-slots are zero, and the 0-slot is killed by $\ell^{e'_\lambda}\delta^2$, using an element on which $\mathrm{ur}_\lambda = 1$.

Ambient inertia at the fixed and moving auxiliary places. At every split auxiliary place in Q, or over the designated scalar-character prime $p_D(\ell)$, the nonbase level has a parahoric vector with the trivial compact type on the first block and the given tame determinant character on the second. The determinant character is the one for the selected tuple. It is nontrivial on inertia at the places of its moving conductor, and has order 15 at the scalar-character prime. The fixed copy of χ_1 has this inertia character by local class field theory: its norm and cyclotomic shifts are unramified there. At the other half, both readings use the inverse character from the block determinant comparison.

The two compact types separate the trivial inertial line from the scalar inertial line. In the parabolic Jacquet comparison the second block accounts for exactly t cuspidal characters on the scalar inertial line. Any further one on this line would occur in the spherical first block, which has only unramified support. The two blocks of support are disjoint, so their normalized product is irreducible. The indicated parahoric vector gives a spherical vector in the first factor and in the second after undoing the determinant twist. At these places ℓ differs from the residue characteristic. Applying local Langlands and the attached representation's away-from-ℓ compatibility gives zero monodromy and the inertia restriction

$$R_x|_{I_w} \simeq 1^{\oplus r} \oplus \chi_1|_{I_w}^{\oplus t} \quad \text{for every retained } x.$$

In particular $(R_x(\gamma) - 1)(R_x(\gamma) - \chi_1(\gamma)) = 0$ on the same matrix order for all $\gamma \in I_w$. At $p_D(\ell)$ one can use the order-three or order-five value, with order prime to ℓ. At a moving place no such unit conclusion is

used. Its cost is measured by the actual $\chi_1(\gamma) - 1$ for the selected tame generator. The selected parahoric vector, The vector on the actual tuple, and Both spherical blocks and the scalar inertia.

Consider next a type place w with $g_k \neq \ell$. The two-prime projector gives the full tame local shape and zero monodromy on each generic fibre. At this level the scalar block is the determinant-character line with its multiplicity t, so its inertia character is the recorded $\chi_1|_{I_w}$. The character of this fixed block is supplied by the local tuple even on fibres without the global rank-one Speh summand. Fix the actual tame generator γ_w of the recorded place and write $c_w = \chi_1(\gamma_w)$. In the common integral Hecke algebra A, divide the polynomial $\det(X - R(\gamma_w))$ by $(X - c_w)^t$ by monic division. The remainder vanishes in every generic fibre and is therefore zero in $A[X]$. Write P_w for the resulting monic degree-r quotient. The residual determinant on the record is constant. At a good point the label at this w excludes the reduction of c_w from the complementary rank-r inertia polynomial, even when ℓ is in the other type prime's bad set. Thus $P_w(c_w)$ is a unit in the common local Hecke algebra. The map $P_w(R_x(\gamma_w))$ vanishes on the complementary inertia eigenspaces. The other eigenspace has inertia action χ_1, so for every retained x and $\gamma \in I_w$ one has

$$P_w(R_x(\gamma_w))\,(R_x(\gamma) - \chi_1(\gamma)) = 0.$$

This identity holds in the integral matrix order because it holds on all its generic fibres. On the evaluating scalar block $P_w(R_x(\gamma_w))$ acts by a unit, as needed for the vanishing of the scalar inertia restriction in the congruence calculation. The same label exclusions give the residual invariant vanishings on W'_π and its Tate dual used by the finite Selmer condition. The type-prime cohomology above $\ell = g_k$ uses the characteristic-zero descent calculation instead. The integral polynomial on the full Hecke record and The realized fixed and moving local clauses.

Away from ℓ, S, Q the universal cocycle is zero on inertia. At the tame auxiliary places the inertia polynomial is $(r(g) - 1)(r(g) - \chi_1(g)) = 0$. On the scalar blocks the action of $u(g)$ on inertia is trivial, so it follows that $(\chi_1(g) - 1)c_F(g) = 0$. The designated tame generator determines the restricted class, giving cost $v_\ell(\zeta_\lambda - 1)$ at each place of Q. At the other fixed places one has either a unit $\zeta_\lambda - 1$ or the actual tame type polynomial which kills $r(g) - \chi_1(g)$ and has unit value at the first scalar inertia character. The restriction is then zero on the scalar blocks. The remaining global conditions are the selected relaxed finite Selmer group on B_ρ, the fine Selmer group on each higher Tate block, and the everywhere unramified group on the first scalar block. The latter has normalized length $v_\ell(h_E)$ by class field theory. An O-valued scalar Tate block has the same normalized fine Selmer length as the $\mathbf{Z}_\ell$-valued one, by taking a $\mathbf{Z}_\ell$-basis of O. The actual tame conditions and The negative Tate fine Selmer groups. Bloch–Kato, Proposition 3.8, Corollary 3.8.4, and Example 3.9.

Apply the determinant comparison to these kernels and images. Writing $c = \#\{\text{columns}\}$ and h_0 for the largest normalized length of a higher scalar invariant group at a column place, the resulting bound on $v_\ell(\mathfrak{c}_\pi)$ is at most

$$\begin{aligned}
&\sum_{\lambda\ \mathrm{row}} l_\lambda + \sum_{\lambda\ \mathrm{column}} \left[e'_\lambda + j_\lambda(e'_\lambda + \Delta) + 2\Delta + h^0_{\rho,\lambda}\right] \\
&\quad + s_{\pi,\ell} + \sum_{a=1}^{t-1} \mathrm{shaLen}_\ell(a) + v_\ell(h_E) + tc_B \sum_{\lambda\ \mathrm{row}} ([E_\lambda : \mathbf{Q}_\ell] + 2) \\
&\quad + 4d(c+1)\Delta + tch_0 + t\sum_{\lambda \in Q} v_\ell(\zeta_\lambda - 1).
\end{aligned}$$

For example the kernel count before enlarging the Δ term contributes $(c+2)(d-1)\Delta + 3c\Delta$ in addition to the $j_\lambda\Delta$ terms. The stated upper bound dominates this. Each global Selmer length has coefficient one at a good Hecke label. Returning to the integral row lattice multiplies by its actual multiplicity at that label, as in the preceding filtration of the eigenspace projections. The congruence length in the Hecke record.

The common jumps on the local coefficient lattices. Fix one localized record meeting the good locus and its coarse type cut. Write R_x for the rank-n_V representation at a retained generic point x, $R'_x = R_x\chi_1^{-1}$,

and $R_x^{\dagger} = (R'_x)^{\vee}(1)$. At a good π, put $W'_\pi = \rho_\pi \chi_1^{-1}$ and $T_{2,\pi} = (W'_\pi)^{\vee}(1)$. We take the restrictions via the decomposition map of the same record at $w \mid \ell$, with coefficient ring enlarged to split the embeddings of E_w if needed. Such an enlargement preserves the normalized length $\mathrm{length}_O / \mathrm{ord}_O(\ell)$, since tensoring a Smith diagonal with the larger DVR multiplies both numerator and denominator by the same ramification index.

For every x, realize its generic eigenspace on the compact unitary group with the recorded weight and level. Its isobaric base change and cuspidal attachments give R_x by strong multiplicity one and the normalized Hecke polynomials. The matching at unramified places above ℓ gives crystallinity when $\ell \notin \{g_0, g_1\}$. In this assertion the relevant unramified place is a place of E: the hyperspecial level and trivial local character at the conjugate halves give a spherical split component, and at an inert place the level is the stabilizer of the self-dual lattice and gives spherical base change. The varying conductors avoid ℓ, and the scalar-character prime has been chosen away from ℓ. The local fields above $2, q, d$ can ramify over $\mathbf{Q}_\ell$ and still have this spherical automorphic level. The polarized regular algebraic attachment gives crystallinity for each spherical cuspidal seed, then for their twists and sums. At g_0, g_1 the type projector gives zero automorphic monodromy on the entire local parameter. For $\ell = g_k$ the comparison [122, Theorem 1.1] applies as in the type cut. Its first assertion identifies the semisimplified Weil representation and bounds the monodromy partition by the automorphic one for each local cuspidal line. All the automorphic Jordan blocks have length one, so the potentially semistable attachment has $N = 0$ and is potentially crystalline, in even rank as well. Norm shifts and the Speh direct sums preserve this conclusion. For the type-prime lift itself use the per-good-point cut. For every retained x its roots at w lie in the common set consisting of the roots of s_w and the two allowed scalar labels; at the conjugate place take inverses. The label u_w of the recorded χ_1 belongs to that set. Division by u_w gives a Frobenius-stable set with 1 and at most $r + 2$ tame characters. These are the characteristic-zero descent characters on the period module. The Tate dual inverts this set and preserves the periods of character ratios. Crystallinity of all the spherical fibres, Compatibility on the full coarse cut, and One tame datum for the active type-prime cut.

Here is the weight comparison at each label. Compose a chosen embedding $\tau : E_w \hookrightarrow L_w$ with the placement of L_w in $\overline{\mathbf{Q}}_\ell$ and with the recorded complex identification. It induces the same place w on E and reads one element g of the host Galois group. Let ρ be the selected CM-type member of the pair $\{g, gc_\infty\}$. Write $s = e_{\rho,j-1}$, $v_a = v_{\rho,a}$ for $1 \le a \le k-1$, and $b_a = b_{\rho,a}$ for $1 \le a \le m$ in the fixed Harish–Chandra parameters. Put $c_0 = (t-1)/2$. The tied weight has infinitesimal-character list on the selected half

$$\mathcal{P} - \tfrac{1}{2}, \qquad \mathcal{P} = \{v_a\}_a \sqcup \{-b_a\}_a \sqcup \{-s + c_0 - k \,:\, 0 \le k < t\}.$$

The list at the conjugate half is its negative, by dualizing the highest weight, which reverses and negates its entries. The block character prior to its first norm shift has infinitesimal entry $-s - \frac{1}{2}$ and its negative at the two halves. Both statements use the actual level character: the diagonal character shifts every highest-weight entry by the same scalar, and the untwisted auxiliary character η has entry $1/2$ on the selected half. Dual highest weight and tied determinant shift and The ambient and block weights.

With arithmetic Frobenius the Hodge jump of a normalized cuspidal term with unitary infinitesimal entry z is $(n_V - 1)/2 - z$, when z includes its Speh shift. This follows from the labeled weights of the cuspidal attachment after undoing the common rank-dependent norm twist. Write σ for the sign of $\mathcal{P}$ in that Hodge formula. There is a common c_τ such that R_x has multiset $\sigma\mathcal{P} + c_\tau$ for *every* x. By the block-character formula and its first norm exponent $(1 - n_V)/2 + (t-1)/2$, the jump of χ_1 is $c_\tau - \sigma s - c_0$. Subtracting it therefore gives on R'_x

$$\{c_0 + \sigma(s + v_a)\}_{a=1}^{k-1} \sqcup \{c_0 + \sigma(s - b_a)\}_{a=1}^{m} \sqcup \{0, \dots, t-1\}. \tag{*}$$

The sign σ is common to all fibres at this embedding. For all embeddings it is determined by the sign character of the host and a single reversal fixed by the coefficient placement. Interlacing and the gaps make the entries distinct; in the middle list $(t+3)/2 \le |s - b_a| \le \eta_{\mathrm{sz}}$, and the outside list stays beyond

this gap. On W'_π delete the scalar list, by the good decomposition $\chi_c/\chi_1 = \epsilon_\ell^{1-c}$. In the dual replace each jump a by $-1-a$. In particular neither cuspidal window contains zero in its respective original or dual orientation. These calculations explain why the same $j_w = \dim H^1_f(E_w, W'_\pi)$ can be used for the finite cohomology dimensions on all generic fibres after adding the appropriate invariant dimension. The block list is nonnegative before dualizing. All fibers at every labeled coefficient embedding, The same shifts on the ambient and first character, and No zero on either cuspidal window.

Crystallinity, potential crystallinity over the common tame extension, and the labeled filtered dimensions pass to rational subquotients by the period comparison and filtered strictness. Thus every source and target of the integral row and column maps, and every graph or Tate-dual map at a type prime, has the indicated windows on its rational space. Their lattice maps are the actual integral surjections of the comparison. No integral splitting of a local crystalline representation is needed to pass the window hypotheses to those maps.

Signs at split, small and ramified primes. Let $\lambda \mid \ell$ be a host place and $n_\lambda = [E_\lambda : \mathbf{Q}_\ell]$. On an ordinary split prime the decomposition group is in the kernel of the CM sign character η_d. Hence σ is constant over the embeddings of a completion, and opposite at its conjugate completion. At the place with $\sigma = +1$ use R'_x and W'_π in the column comparison. At its conjugate use $R^\dagger_x$, $T_{2,\pi}$ in the row comparison. The outside jumps are positive in both comparisons. The Bloch–Kato finite dimension is at most $n_\lambda m$ on each cuspidal target, since only the middle list has negative entries and the cuspidal targets have no invariants. Their jumps and those of the full representations have absolute value at most $b + S_c + \eta_{\mathrm{sz}} + t + 2$. The local field is unramified at ordinary primes, including 3 outside the designated small set, where the level exponent is one. At the small primes the decomposition group still preserves the CM sign, by the prescribed splitting of the real character, and the completion is contained in the cyclotomic extension for the host conductor. At 2, $4 \mid Q_S$ and the depth is at least two. This gives the same paired declaration and dimensions for the deep cyclotomic estimates. The ordinary split declaration, The realized split modules, Three outside the small set, and The dyadic completion and level.

At $\ell = q, d$ the completion is $K_0(\sqrt{-\ell})$, with $K_0/\mathbf{Q}_\ell$ unramified. Here the other conductor factors give unramified extensions, and the ramified quadratic factor is the unique quadratic subfield of $\mathbf{Q}_\ell(\mu_\ell)$. Consequently the first cyclotomic extension has relative degree $(\ell-1)/2$, prime to ℓ. At q the whole decomposition group preserves the CM sign because $(-d/q) = 1$, and the paired split declaration applies with the same dimension count. At d a generator of inertia changes the sign. Over each unramified label there are two embeddings of the completion, whose outside jumps have opposite signs by $(*)$. Their two real labels can differ. Thus use the actual absolute endpoint of each cloud separately and the same upper bounds on both widths. After dualizing, the absolute endpoints are $s + \sigma(t+1)/2$ and the middle list lies in $[-\eta_{\mathrm{sz}} - (t+1)/2, \eta_{\mathrm{sz}} - (t+1)/2]$. After adding the scalar dual list this lies in $[-\eta, \eta]$ for $\eta = \lfloor \eta_{\mathrm{sz}} \rfloor + t + 2$. The cuspidal target contributes at most m entries in the central window per embedding. This supplies the two signed windows of the ramified comparison on the same crystalline row. The two embeddings here are paired by inertia, which is the pairing used for the two roots of the quadratic Eisenstein polynomial. The windows and the two local embeddings and The realized local representations at both primes.

At g_k the completions are all $\mathbf{Q}_{g_k}$. At a place with $\sigma = +1$, $(*)$ gives its direct window in $[-\eta', H]$ with $\eta' = \lfloor \eta_{\mathrm{sz}} \rfloor + t + 1$ and $H = \lfloor 3b \rfloor + 2$. The opposite sign uses the Tate dual, with the same window by the displayed bounds above. In each orientation at most $n_V + 1$ jumps are used after adjoining zero for a trivial quotient. The scalar and tame data therefore meet the jump and character cardinality requirements with $N = n_V + 1$. There are $2f$ host places for this comparison, and the separate type-prime estimate supplies its exponent on the indicated graph or dual surjection. The common sign and finite sets at type primes.

The two type primes. At $\ell = g_0, g_1$, perform the calculation on the type locus X_π selected by the tame projectors. All places above ℓ have degree one, and the type-prime size inequality gives $\ell - 1 \geq t$. The residual cyclotomic characters at a good point are distinct. The first residual scalar projector therefore

lifts inside the matrix order to an integral rank-one idempotent in every generic fibre. Indeed its lift at π has characteristic polynomial $X^{d-1}(X-1)$ modulo the uniformizer, common to every fibre of the local Hecke record. Iteration of $z \mapsto 3z^2 - 2z^3$ gives an idempotent in the complete order. Thus one can use $\delta = 1$ and the relation module Y restricted to X_π, with no columns. The congruence ideal computed in the projected Hecke algebra is exactly the ideal of evaluations at π of operators vanishing at $X_\pi \setminus Z$. The idempotent argument is also valid with the source's scaled projector and gives its bound with the additional allowance $4d\Delta$. Relations on the type locus and The projected congruence ideal.

A uniform lifting bound at a prime of tame type. This bound is an AI-proposed candidate. Let Σ be a finite Frobenius-stable set of tame inertia characters at p which contains the trivial character and has size at most N, with $N \geq 2$. Let q be the largest Frobenius period of a ratio of characters in Σ, with period one for the trivial ratio. Let $U \subset [-\eta, H] \cap \mathbf{Z}$ contain zero and have at most $N+1$ elements. Suppose

$$1 \leq \eta < H, \qquad 40(HN^2 + N^3) < p.$$

For an integral surjection of potentially crystalline lattices whose tame types and jumps lie in Σ and U, respectively, the cokernel on $H^1_f(\mathbf{Q}_p, -)$ is killed by $p^{7(Nq+\eta(N+q))}$. The bound is independent of the coefficient field and the tame splitting degree.

Take a common tame crystalline extension $K' = \mathbf{Q}_{p^a}(\varpi)$, with $\varpi^e = -p$, $e = p^a - 1$, where a is the least common multiple of the character periods. For a character of period d, its order divides $p^d - 1$. Thus all the types factor through the tame quotient of order e, and the representations become crystalline by descent on their potentially crystalline period modules. The unramified part contains μ_e, so the extension is Galois of degree ae with $p \nmid ae$. Each period is less than p. The choice $a = (p-1)!$ works as well. The two characters of a ratio belong either to the same cycle or to distinct cycles. The ratio period divides the length of the common cycle or the least common multiple of the two lengths. The trivial character occupies a separate cycle. Consequently $q \leq \lfloor N^2/4 \rfloor$.

Breuil–Kisin lattices and maps. After a finite free coefficient extension with field F and integers O, use the product over the unramified embeddings, with $A = O[[u]]$ on each factor. Put $v = u^e$, $E = v + p$, and let $\Delta = \mathrm{Gal}(K'/\mathbf{Q}_p)$. Frobenius sends u to u^p and shifts the factors. For the covariant Kisin module M, write $L = A\Phi(M)$. Let S be the Breuil divided-power ring of E and $\mathcal{L} = SL$. Let R be the rational analytic ring on $|u| \leq |p|^{1/e}$; M_R and L_R denote base change. We will use only the finite jets at the roots of E for calculations with principal parts. The comparison data give

$$E^H M \subset L \subset E^{-\eta} M, \quad \delta\mathcal{L} \subset \mathcal{L}, \quad E\delta M_R \subset M_R, \quad \delta\Phi = p\Phi\delta.$$

Here δ is the connection in the direction $u\,d/du$. It commutes with descent and has zero constant matrix coefficient in a Frobenius frame at $u = 0$. The signs are fixed by the rank-one condition $L = E^i M$ for jump i.

For these comparisons one may use the effective theory after the common Tate twist with Frobenius E^η, taking the dual to use Liu's contravariant convention. Liu's strong divisibility comparison gives stability under δ on the twisted Frobenius span. For untwisting, the Tate factor has connection coefficient in its Frobenius basis

$$-\eta \sum_{n \geq 1} u \frac{d}{du} \log\bigl(\varphi^n(E/p)\bigr).$$

This follows from its horizontal section $\prod_{n \geq 0} \varphi^n(E/p)^\eta$ in the original basis. Here $\varphi(E)/p = 1 + v^p/p$ is a unit in S, since $(v^p/p)^p \in pS$. The terms of the displayed derivative series are in S and tend to zero p-adically. This yields the integral connection on the signed Frobenius span. For $E\delta M_R$ work at a root of E in a filtration-adapted horizontal basis (d_i) of L_R. In the effective convention the recursive Breuil filtration has step r generated there by $E^{(r-i)_+} d_i$. For r at least the highest effective jump this step equals $E^r M_R$ after

completion at the root. Thus the original lattice is generated by $E^{-i}d_i$ at that root. This description gives the stated stability by Euler differentiation. Away from the roots the two analytic spans agree. The weight range is below $p-1$, regardless of ae. Liu, §2.2, Corollary 3.2.3 and Theorem 3.5.4.

For a short exact sequence of Galois lattices, the Kisin sequence $0 \to M_0 \to M_1 \to M_2$ is exact through the middle and has finite-length final cokernel. Apply the left exactness theorem for the covariant Kisin functor to the restrictions to the Kummer tower, using the common Tate twist if needed. Its final cokernel is killed by some power of (p, u). Coefficient endomorphisms commute with the functor, so the assertion holds with O-coefficients, and finite free change of coefficients preserves exactness. The semilinear tame actions commute with the maps. After inverting p the sequence splits as a sequence of modules since $M_2[1/p]$ is free. This gives exactness in the rational jet rings at all the roots of E. Čoupek–Gazaki–Marmora, Proposition 3.4.

Two Euler degree intervals. Put $\rho = 2(N-1)q$, $\beta = \rho + 1$, and $m_* = H + \eta + \beta$. Then $2m_* + 2 < p$. For the map $f : M_1 \to M_2$, let I_τ annihilate its cokernel in factor τ. The corresponding ideal on the Frobenius spans is $J_\tau = \varphi(I_{\tau-1})A$, by the decomposition of A into the p residue classes of degrees. The height inclusions imply $E^{H+\eta} I_\tau \subset J_\tau$. Let s_τ be least with $u^{s_\tau} \in I_\tau + (p^\beta)$. The binomial expansion $v^{m_*-1} \in (E^{H+\eta}, p^\beta)$ implies $p s_{\tau-1} \le s_\tau + e(m_* - 1)$. Taking the maximum gives $s_\tau \le e$ for all factors. One more application of the same expansion gives

$$v^{m_*} L_2 \subset f L_1 + p^\beta L_2.$$

Indeed $u^e \in I_\tau + (p^\beta)$ and the product of u^e with v^{m_*-1} belongs to $J_\tau + (p^\beta)$.

Write $\mathcal{L}_{i,\mathrm{hi}} = \mathcal{L}_i \cap u^{pe} F[[u]] L_i$. The coefficients of degrees below pe in S are integral because their allowed denominators $\lfloor d/e \rfloor!$ are units. Hence $\mathcal{L}_i = L_i + \mathcal{L}_{i,\mathrm{hi}}$ and the intersection with L_i is $v^p L_i$. The operator δ preserves the high module and acts by d on the graded piece of degree $d < pe$. In the inertia isotypic component c these degrees have the form $d(c, c') + je$, with $c' \in \Sigma$ and $d(c, c')$ the residue of $c - c'$ in $[0, e)$. Let P_1, P_2 be the monic polynomials of these degrees for the respective ranges $0 \le j < m_*$ and $m_* \le j < 2m_*$. Their Euler actions raise the corresponding vectors from degree 0 to em_*, and from degree em_* to $2em_*$, modulo the high module.

In a Lagrange denominator at a fixed root, roots of the same character differ by a p-adic unit. Each other character contributes at most one root with a difference divisible by p, since the whole range of j has length less than p. For such a difference D, the reduced denominator of D/e divides $p^h - 1$, where $h \le q$ is the ratio period. Since $|D/e| < 2m_* + 1 < p$, the numerator has magnitude less than p^{h+1} and $v_p(D) \le h$. The whole Lagrange denominator has valuation at most $(N-1)q$. Interpolating p^ρ on the first set of roots and zero on the second gives $p^\rho \in (P_1, P_2)$ over $\mathbf{Z}_p$.

For x in an inertia component of L_2, both $P_i(\delta) v^{m_*} x$ belong to

$$v^{m_*} f L_1 + p^\beta L_2 + \mathcal{L}_{2,\mathrm{hi}}.$$

For P_1 this follows from the preceding lift modulo p^β and the first degree interval on its source. For P_2 the result is in $v^{2m_*} L_2 + \mathcal{L}_{2,\mathrm{hi}}$, and the same lift applies again. The displayed sum is stable under δ. The Bézout identity puts $p^\rho v^{m_*} x$ in it. The high term lies in $v^p L_2$, and the term in $p^\beta L_2$ is divisible by v^{m_*} in that module by coefficient comparison. Cancel v^{m_*} and use the lift at $v^{p-m_*} L_2$. This yields $p^\rho L_2 \subset f L_1 + p^\beta L_2$. Nakayama gives $p^\rho L_2 \subset f L_1$. Faithful flatness of Frobenius pullback proves $p^\rho M_2 \subset f M_1$. Inertia and tame projections are integral since ae is a unit.

Principal parts of the Euler connection. Write $F = \mathrm{Frac}(O)$. At all roots of $E = u^e + p$ and all unramified embeddings, use the rational E-adic completions M_{an} and L_{an}. Put

$$Y = (E^{-\eta} M)^\Delta, \qquad Y' = vY,$$
$$Z = (E^{-\eta} M / M)^\Delta, \qquad J = (E^{-\eta} M_{\mathrm{an}} / M_{\mathrm{an}})^\Delta.$$

Brackets denote principal parts in J. Division by the distinguished monic E^{η} identifies A/E^{η} with the O-module of polynomials in u of degree $< e\eta$. Rational interpolation at its roots identifies the corresponding rational jets with $(A/E^{\eta})\otimes_O F$. In particular a vector of $M[1/E]$ regular in M_{an} belongs to M. Since $|\Delta| = ae$ is a unit, averaging shows that $[Y] = Z$, Z is finite free, and $J = Z[1/p]$. The module $Y' = vE^{-\eta}M^{\Delta}$ is p-adically complete.

The endomorphism $D[y] = [E\delta y]$ of J preserves Z. Indeed the connection preserves M_{an} after multiplication by E, and for $j < p$ every member of $\mathcal{L}$ has an integral representative in L modulo $E^{j}L_{\mathrm{an}}$. For the latter assertion expand each divided power $v^{n}/n!$ in E. Its coefficient at E^{k} for $k < j$, $k \le n$ is $(-p)^{n-k}/(k!(n-k)!) \in O$. Apply this calculation to the convergent divided-power expansions in an L-basis. For $y \in E^{-\eta}M$, put $h = H + \eta$. Then $E^{h}y \in L$ and

$$E^{2H}E\delta y = E^{H-\eta}\big(E\delta(E^{h}y) - h(\delta E)E^{h}y\big) \in \mathcal{L}.$$

Choose $\ell \in L$ representing this last vector modulo $E^{\,2H+\eta}L_{\mathrm{an}}$, which is possible since $2H + \eta < p$. Thus $E\delta y - E^{-2H}\ell \in E^{\eta}L_{\mathrm{an}} \subset M_{\mathrm{an}}$. Moreover $E^{\eta}E^{-2H}\ell \in M[1/E] \cap M_{\mathrm{an}} = M$, since $E\delta y \in E^{-\eta}M_{\mathrm{an}}$. This gives the required integral principal part.

Let $U \subset [-\eta, H]$ be the common set of jumps and put

$$B = \{i + h \,:\, i \in U,\ 1 \le h \le \eta\}, \qquad d = |B \setminus \{0\}| \le (N+1)\eta.$$

At a root use $t = E$, so $\delta t = e(t - p)$. The elementary divisors of M_{an} relative to L_{an} are t^{-i} for the labeled jumps i. The filtration on $t^{-\eta}M_{\mathrm{an}}/M_{\mathrm{an}}$ induced by $t^{s}L_{\mathrm{an}}$ therefore has a nonzero graded piece only for $s = -i - h$ with $1 \le h \le \eta$. For $\ell \in L_{\mathrm{an}}$,

$$(t\delta + sep)(t^{s}\ell) = t^{s+1}(\delta\ell + se\ell).$$

Hence the product of $D + sep$ over the distinct occupied degrees s annihilates the filtered space. It has simple roots, so D is semisimple on this space and on J, with spectrum in epB. In particular $D(J) \cap V = D(V)$ for every D-stable F-subspace $V \subset J$.

Let $W = Z \cap D(J)$, $b = (N+2)\eta$, and write

$$P(T) = \prod_{\beta \in B\setminus\{0\}} (T - ep\beta) = P(0) + TQ(T).$$

The factors use integers $0 < |\beta| < p$, hence $v_p(P(0)) = d$. For $w \in W$, the vector $-Q(D)w \in Z$ maps under D to $P(0)w$, because P annihilates $D(J)$. Also $p^{\eta}Z \subset [Y']$ by the identity

$$v\sum_{j=0}^{\eta-1} p^{\eta-1-j}E^{j} = E^{\eta} - p^{\eta}.$$

Dividing $P(0)$ by its unit factor gives $p^{b}W \subset D[Y']$.

The Frobenius series. For $y = y' + \Phi m$, $y' \in Y'$, $m \in M^{\Delta}$, write

$$\zeta(y) = -\sum_{n\ge 1}(p\Phi)^{n}\delta y, \qquad \mathrm{Mon}(y) = [E(\delta y - \zeta(y))], \qquad c_1 = (q+1)\eta.$$

In a horizontal frame δ is the Euler operator $u\partial_u$, so the formal series converges in positive u-degree. We shall use

$$p^{\eta}\zeta(y') \in \mathcal{L}, \qquad [E\zeta(y')] \in p^{\,p+1-2\eta}Z, \qquad p^{c_1}\zeta(\Phi m) \in \mathcal{L}.$$

Indeed $\varphi(E)/p = 1 + v^{p}/p$ is a unit of S, since $(v^{p}/p)^{p} \in pS$ by $v_p((p^2)!) = p + 1$. The inclusion $E^{\eta}L \subset M$ gives $p^{\eta}\Phi\mathcal{L} \subset \mathcal{L}$. For $y' = vE^{-\eta}x$, $x \in M^{\Delta}$, one has $\Phi y' \in p^{-\eta}v^{p}\mathcal{L}$. Using $\delta\mathcal{L} \subset \mathcal{L}$ and $\delta\Phi = p\Phi\delta$ gives for $n \ge 1$

$$(p\Phi)^{n}\delta y' = \delta\Phi^{n}y' \in p^{\,n-1-n\eta}v^{p^{n}}\mathcal{L}.$$

For $n \geq 2$ the monomial contributes at least $v_p((p^n)!) \geq p^{n-1} \geq (n-1)p$, which gives the first membership and p-adic convergence. Modulo E^η the expansion of $v^{p^n} = (E-p)^{p^n}$ is divisible by $p^{p^n-\eta+1}$. Represent the remaining $\mathcal{L}$-factor integrally modulo $E^\eta L_{\mathrm{an}}$. The discarded terms disappear in the principal part after multiplication by E, since $E^\eta L_{\mathrm{an}} \subset M_{\mathrm{an}}$. Thus $[E\delta\Phi^n y']$ is divisible in Z by $p^{p^n+n-(n+1)\eta}$. The minimum occurs at $n = 1$ and the exponents tend to infinity.

For the third membership take $w = \delta\Phi m \in \mathcal{L}$. It is invariant, and has zero constant coefficient in the formal horizontal frame. If $w \neq 0$, let $s = \min(e, \operatorname{ord}_u w)$ in the rational formal module at zero, taking the minimum over the components. If $s < e$, a nonzero leading coefficient has fiber inertia character $-s$ mod e when inertia acts on u by its tautological character. Its period $h \leq q$ gives $e \mid s(p^h - 1)$. Thus $sp^q \geq e$ in either case. In a Frobenius basis of L every coefficient f of w belongs to $S \cap u^s F[[u]]$. For $n \geq 1$, $\varphi^n(f)/u^{sp^n} \in S$. Indeed a term of degree $j \geq s$ moves to degree $(j-s)p^n$. For $j < ep$ its original factorial $\lfloor j/e \rfloor!$ is a unit, while for $j \geq ep$ the new degree is at least j. The new factorial therefore absorbs the denominator and preserves convergence of the divided-power coefficients. It follows that

$$(p\Phi)^n w \in p^{-n(\eta-1)} u^{sp^n} \mathcal{L}.$$

For $1 \leq n \leq q+1$ multiplication by p^{c_1} suffices. For $n = q+1+t$, $t \geq 1$, $sp^n \geq ep^{t+1}$, so the monomial contributes at least p^{p^t}. The net exponent after multiplying by p^{c_1} is at least $q+1+t+p^t-t\eta$, which is positive and tends to infinity since $p^t \geq tp$.

Correction in an exact triple. For these y, $\omega = \mathrm{Mon}(y)$ belongs to $D(J)$ and $p^{c_1}\omega \in Z$. The integral assertion follows from $y \in Y$, $DZ \subset Z$, and the existence of integral representatives for the coefficients of $p^{c_1}\zeta(y)$ modulo $E^\eta L_{\mathrm{an}}$. For the rational assertion solve $\delta s = \zeta(y)$ mod $t^\eta L_{\mathrm{an}}$ at each root. In an L_{an}-basis, $\delta = e(t-p)d/dt + C(t)$ with $C(t)$ regular. Set $s_0 = 0$. At degree j, obtain s_{j+1} by division by $-ep(j+1)$, for $0 \leq j < \eta$. Rational interpolation realizes the jets in L_{an}, and averaging gives an invariant solution. Now $s \in E^{-\eta} M_{\mathrm{an}}$ and $E(\delta s - \zeta(y)) \in M_{\mathrm{an}}$, so $D[y-s] = \omega$.

Set $k = p - 14H > b$. On Y' the difference $\mathrm{Mon}(y') - D[y']$ lies in $p^k W$, by the series estimate and p-saturation of W in Z. Since $p^b W \subset D[Y']$, successive approximation yields $p^b W \subset \mathrm{Mon}(Y')$. In detail, for a target $p^b w$ first choose y' with $D[y'] = p^b w$. Its error is in $p^k W$. If an error belongs to $p^{b+j(k-b)}W$, subtract a multiple $p^{j(k-b)}z$ with $z \in Y'$ whose leading term corrects it. The new error belongs to $p^{b+(j+1)(k-b)}W$. The corrections converge in Y' and the errors tend to zero in Z.

Finally let $0 \to M_0 \xrightarrow{i} M_1 \xrightarrow{g} M_2$ be the given exact triple with $p^\rho M_2 \subset gM_1$. It splits as a sequence of modules after inverting p, since the target is free. Thus $i : J_0 \to J_1$ is injective. For $w \in Z_1 \cap \ker g$ pick $y \in Y_1$ representing w. Then $gy \in M_2[1/E] \cap M_{2,\mathrm{an}} = M_2$. Lift $p^\rho gy$ to an invariant $x \in M_1$ by averaging. Write $E^\eta(p^\rho y - x) = im_0$. It follows that $p^\rho w = i[E^{-\eta}m_0] \in iZ_0$.

Given $y_2 \in Y_2'$, $m_2 \in M_2^\Delta$ with $\mathrm{Mon}(y_2 + \Phi m_2) = 0$, the two terms after multiplication by p^ρ lift to $y_1 \in Y_1'$ and Φm_1, $m_1 \in M_1^\Delta$. The error $\omega = \mathrm{Mon}(y_1 + \Phi m_1)$ belongs to $D_1(J_1)$, satisfies $g\omega = 0$, and $p^{c_1}\omega \in Z_1$. Use the preceding paragraph to write $p^{\rho+c_1}\omega = iw_0$ with $w_0 \in Z_0$. Semisimplicity gives $D_1(J_1) \cap iJ_0 = iD_0(J_0)$, hence $w_0 \in D_0(J_0)$. Choose $y_0 \in Y_0'$ with $\mathrm{Mon}(y_0) = p^b w_0$. Then $p^{\rho+c_1+b}(y_1 + \Phi m_1) - iy_0$ has zero monodromy and maps to $p^{2\rho+c_1+(N+2)\eta}(y_2 + \Phi m_2)$.

The target finite class. For $c \in H^1_f(\mathbf{Q}_p, T)$ the corresponding integral extension of the unit by T is crystalline over K', with its jumps and tame types in the same sets. Use the integral Kisin map to the unit to lift p^ρ to an invariant vector z. Then $y = \Phi z - z$ belongs to Y. The extension with Frobenius vector y is the pullback by p^ρ of the original étale extension, hence represents $p^\rho c$. The Frobenius–connection equation gives $\delta z = \zeta(y) - \delta y$ in positive u-degree. The series represents $\delta(\Phi z)$ in the analytic Frobenius span, and this vector lies in the rational kernel span. Since $E\delta z$ belongs to the analytic extension lattice and maps to zero in the unit, $\mathrm{Mon}(y) = 0$ by exactness at the roots. The correction from y to a vector of the required form uses $m = E^\eta y$. The cohomologous Frobenius representative of $p^{\rho+\eta}c$ is

$$p^\eta y + (\Phi - 1)m = y' + \Phi m, \qquad y' = (p^\eta - E^\eta)y \in Y', \quad m \in M^\Delta.$$

Its monodromy still vanishes because modifying the integral section by m changes the error by the class of $E\delta m$.

Realization of the corrected lift. The corrected vector $y = y' + \Phi m$ has zero monodromy. Multiplying first by p^{c_1}, $c_1 = (q+1)\eta$, puts $\zeta(y)$ in the integral Breuil lattice. With this scaled vector form $M_y = M \oplus Az$ with $\Phi_y z = z + y$ and z fixed under Δ. Its Frobenius span is $L_y = L \oplus A(z + y)$, with

$$E^{\eta} L_y \subset M_y, \qquad E^{H+\eta} M_y \subset L_y.$$

For the second inclusion use $E^{H+\eta} y \in E^H M \subset L$. The common effective Tate twist of this extension has Frobenius span $L_{\mathrm{tw}} = E^{\eta} L_y$ in the twisted basis. It lies in its original lattice M_{tw} and has height at most $r_{\mathrm{tw}} = H + 2\eta \le p - 2$.

Define $\delta_y z = \zeta(y) - \delta y$. The series equation for ζ gives $\delta_y \Phi_y = p \Phi_y \delta_y$. Since $\delta_y(\Phi_y z) = \zeta(y)$, the enlarged Frobenius span over S is stable. The equality $\mathrm{Mon}(y) = 0$ gives $E\delta_y z \in M_R$, so $E\delta_y$ also preserves $(M_y)_R$. The Tate factor preserves both conditions.

Apply the usual Kisin base change $S \otimes_{\varphi, A} M_{\mathrm{tw}}$ in height r_{tw}. Linearized Frobenius identifies it with $\mathcal{L}_{\mathrm{tw}} = SL_{\mathrm{tw}}$. Its filtration is the inverse image of $(\mathrm{Fil}^{r_{\mathrm{tw}}} S) M_{\mathrm{tw}}$. This is tested by rational jets through degree $r_{\mathrm{tw}} - 1$ at the roots of E. Thus the filtration is $\mathcal{L}_{\mathrm{tw}} \cap E^{r_{\mathrm{tw}}} (M_{\mathrm{tw}})_R$. The ordinary height-$r_{\mathrm{tw}}$ construction supplies $\varphi_{r_{\mathrm{tw}}} = p^{-r_{\mathrm{tw}}} \Phi_{\mathrm{tw}}$ on this filtration, with generating image and the required reductions modulo all p^n. The quotient by the filtration is p-torsion-free. Let $N_{\mathrm{Br}} = -\delta_{y,\mathrm{tw}}$. The Leibniz rule and preservation of $(M_{\mathrm{tw}})_R$ by $E\delta_{y,\mathrm{tw}}$ give $E N_{\mathrm{Br}}(\mathrm{Fil}^{r_{\mathrm{tw}}}) \subset \mathrm{Fil}^{r_{\mathrm{tw}}}$. Compatibility of Φ_{tw} and the connection gives

$$\varphi_{r_{\mathrm{tw}}}(E N_{\mathrm{Br}} x) = (\varphi(E)/p)\, N_{\mathrm{Br}} \varphi_{r_{\mathrm{tw}}}(x)$$

on that filtration. These are the conditions of the finite free strongly divisible category. Apply Liu's contravariant equivalence over $\mathbf{Z}_p$ and take the dual lattice to use the covariant convention. This yields a semistable lattice over K' with the commuting coefficient action. Its monodromy on the filtered isocrystal is obtained by evaluating the connection at $u = 0$. It vanishes there in the Frobenius frame, including the coefficient $\zeta(y)(0) = 0$ on the added vector. Thus the lattice is crystalline, and we undo the common twist. Liu, Proposition 2.3.4, Theorem 2.3.5 and Corollary 3.3.2.

The resulting lattice restricts on $G_{K'_\infty}$ to the étale realization of M_y, by the Kisin base change and the comparison of T_{st} and T_{cris}. Étale exactness gives the original integral kernel T and unit quotient. Integral full faithfulness of crystalline restriction then gives these same maps over K'. For tame descent, use $\varpi_{n+1}^{p} = \varpi_n$ with $\varpi_0 = \varpi$. Each $\gamma \in \Delta$ acts on the Kummer tower by $\varpi_n \mapsto (\gamma(\varpi)/\varpi)^{p^{-n}} \varpi_n$, using the unique p^n-th root in μ_e. The preimage of these actions in $G_{\mathbf{Q}_p}$ normalizes $G_{K'_\infty}$, maps onto Δ, and intersects $G_{K'}$ in $G_{K'_\infty}$. The semilinear action on M_y intertwines the corresponding actions on the étale restriction. By full faithfulness the intertwiners also intertwine the crystalline $G_{K'}$-lattice. Thus the two actions give an integral $G_{\mathbf{Q}_p}$-action on the exact sequence. Its class is finite by the injectivity of restriction on rational B_{cris}-cohomology. Comparisons of finite extension classes can be checked on $G_{K'_\infty}$ by full faithfulness over K', followed by restriction and corestriction of the p-adic unit degree ae. Liu, Lemma 3.3.4 and §3.5. Bhatt–Scholze, Corollaries 3.8 and 7.10.

The target representative, the monodromy correction, and realization cost respectively $\rho + \eta$, $2\rho + c_1 + (N+2)\eta$, and c_1. Thus the total exponent is

$$3\rho + 2c_1 + (N+3)\eta \le 7\big(Nq + \eta(N+q)\big).$$

This proves the local lifting bound on the given stable lattices.

Here the local type records, including the first-character twist, have trivial monodromy and a common finite set of tame inertia characters. Let a be the least common multiple of their Frobenius orbit lengths. Those lengths are bounded by the size of the common type set, which is less than ℓ. Consequently a is prime to ℓ, and the characters are killed by the common crystalline extension obtained by adjoining an

$(\ell^a - 1)$-st root of a uniformizer to the unramified extension of degree a of $\mathbf{Q}_\ell$. In particular y defines a rational finite class over $\mathbf{Q}_\ell$ itself. Its image in B_{cris} cohomology vanishes after restriction to this common extension, and restriction followed by corestriction multiplies by the nonzero extension degree in rational cohomology. With $f = F|_{L_\lambda}$ and $B = V/T$, use the stable graph lattice

$$D_f = \{(a, v) \in L_\lambda \times V \ : \ v + T = f(a)\}.$$

It has rational representation $(L_\lambda \otimes_O \operatorname{Frac} O) \oplus V$ and fits into $0 \to T \to D_f \to L_\lambda \to 0$. If the direct finite lift is available with exponent e_*, lift $\ell^{e_*} y$ along $D_f \to L_\lambda$ and project the lifted cocycle to V. It represents $\ell^{e_*} c_F$ in the torsion finite condition. If the dual finite lift is available instead, lift $\ell^{e_*} b$ along $D_f^\vee(1) \to T^\vee(1)$ to a cocycle $(b_0, \ell^{e_*} b)$. Pairing with $y(g)$ and reducing the defining congruence of $D_f^\vee(1)$ gives the cup of c_F with $\ell^{e_*} b$ as the reduction of minus the rational cup of y with b_0. The latter is zero by Bloch–Kato duality for finite classes. In both cases the finite lifting estimate is used over $\mathbf{Q}_\ell$, with the common tame datum. Both arguments bound the entire local image by the ℓ^{e_*}-torsion of the quotient by the finite condition, independently of F.

For a target of rank k that quotient has ℓ^{e_*}-torsion of normalized length at most $2ke_*$, by local Tate duality. In detail, the dual finite condition has free rank at most $k + h$, where $h = \dim H^0(E_\lambda, V^\vee(1))$, by the finite cohomology dimension formula in degree one. Its torsion has at most $k - h$ generators, since it is the quotient of torsion invariants by the h divisible invariant directions. Those directions come from the saturated sublattice in the invariant rational space. Thus the number of generators is at most $2k$.

The chosen type-prime inputs. Let $p = g_i$, $i = 0, 1$. At every $w \mid p$ the base field is $\mathbf{Q}_p$, and there are $2f$ such places. The tame type consists of the roots of the cuspidal type polynomial and the two allowed scalar labels, after division by the chosen first label. At the conjugate place use the inverse labels. The trivial character is included, and the type has at most $r + 2 \le n_V + 1$ characters. The type at each place means the characteristic-zero potentially crystalline descent type. Use

$$H = \lfloor 3b \rfloor + 2, \qquad N = n_V + 1, \qquad \eta' = \lfloor \eta_{\mathrm{sz}} \rfloor + t + 1.$$

The ambient first-character ratios are already in $[-\eta', H]$ on the up places. On the down places apply the Tate dual, sending a jump j to $-j-1$ and inverting the tame characters. This yields the same window and preserves the period of every ratio. The jump set, including the unit, has size at most $N + 1$. The scale relation $\eta_{\mathrm{sz}} + t + S_c + 2 \le b$ gives $1 \le \eta' < H$. For $40 \le C_T \le b$, $H \le 5b$, and $N \le 3kb$, the chosen prime size satisfies

$$C_T(HN^2 + N^3) \le 36(kb)^4 < (kb)^8 \le p.$$

Use the same signed data on the rational subquotients of the finite sums and on the graph lattices for the cocycle images. The local bound therefore applies with

$$e_{\max} = C_T\big(N\lfloor N^2/4 \rfloor + \eta'(N + \lfloor N^2/4 \rfloor)\big).$$

With $e_* = e_{\max}$ the graph comparison gives $2re_{\max}$ for B_ρ and $2e_{\max}$ for each scalar block. There are $2f$ places above either type prime, so the local contribution is at most $4f(r+t)e_{\max}$ per good label. This uses the same $s_{\pi,\ell}$, fine Selmer terms and tame budgets as above. The actual local lengths and labeled type locus · The common tame extension, and The type-prime exponent.

For each good π the integral type endomorphism on M is scalar $b_\pi \neq 0$ at π and zero off X_π. Its normalized scalar cost is at most $v_\ell(|G|)$. Compose it with a Hecke operator realizing an evaluation generator of the projected congruence ideal. The composition vanishes off the good part and acts by b_π times this generator at π. The endomorphism can have different X_π for different good points. Filtering the good projection and intersection lattices into their ordered generic eigenspaces makes the graded comparison separately for each π, with the endomorphism appropriate to that label. It costs $\dim V_\pi$ times the valuation of its scalar. The finite group order contributes at most $2fr^2$ per good label. The group has f factors

of $\mathrm{GL}_r(\mathbf{F}_{g_k})$ at each of the two type primes. In residue characteristic g_a the g_a-factor has order with valuation $r(r-1)/2$. The valuation of the other factor is at most $\sum_{j=1}^{r} 2j$, since $\log(g_{1-a}) \le 2\log(g_a)$. Their sum is at most $2r^2$. Integral type operators for individual labels · The type-prime multiplicity weight, and The order bound.

4.1.1 Local lattice and cohomology bounds

Write O for the local coefficient ring and ℓ for its residue characteristic. Write W'_π for the fixed good cuspidal lattice after removing the first character of the Galois block, and R'_z for the corresponding block representations. At a row place one needs a normalized length bound for the cokernel on integral finite cohomology for every lattice surjection onto $(W'_\pi)^\vee(1)$ whose generic source embeds in a finite sum of $(R'_z)^\vee(1)$. At a column place one needs an exponent of ℓ annihilating the cokernel for the prescribed subquotients of sums of R'_z. If this exponent is e_{col}, the column contribution is $(j_{\mathrm{col}}+1)e_{\mathrm{col}}+(j_{\mathrm{col}}+2)\Delta_\ell+h^0_{\mathrm{col}}$, where j_{col} is the dimension of rational finite cohomology, h^0_{col} is the normalized length of torsion invariants and Δ_ℓ is the valuation cost of separating the cyclotomic characters. The row and column contributions enter the global comparison of the torsion Hom lengths. Local declarations and Global comparison.

At an ordinary split prime $\ell \ge 5$ the column and conjugate row are paired. Their finite rational dimensions are at most $n_\lambda m$, where $n_\lambda = [E_\lambda : \mathbf{Q}_\ell]$. The large weights have the sign excluded from the finite dimension formula and at most m small jumps per label can contribute. All the relevant jumps have absolute value at most a common integer $H = \lfloor b + S_c + \eta_{\mathrm{sz}} + t + 2 \rfloor \ge 1$. The maximum-width unramified proof gives an annihilator exponent

$$2J(2d_0-1)+J(a_0)+2v_\ell((2H)!), \qquad d_0=\left\lceil \frac{2H}{\ell-1} \right\rceil, \quad a_0=\left\lfloor \frac{H}{\ell-1} \right\rfloor, \quad J(s)=s+v_\ell(s!).$$

The Wach quotient is killed by X^{d_0} in the tame parameter. Solving the horizontal coordinates successively costs $1+v_\ell(s)$ in degree s, so truncating at $2d_0-1$ gives $J(2d_0-1)$, once for the extension generator and once for its cochains. At the first cyclotomic divisor the operator $(\gamma-1)/X$ on the kernel quotient has diagonal entries $(1-\chi_{\mathrm{cyc}}(\gamma)^{-i})/\ell$, with $|i| \le 2H$. The correction polynomial removes the nonzero roots and its constant term has valuation $2v_\ell((2H)!)$. Elimination by $\gamma-1$ through degree a_0, then by the geometric inverse of $1-\varphi$ on X^{a_0+1}, gives the remaining $J(a_0)$. This is bounded by $20H/\ell+3$ by the factorial valuation bound, so the allowance $e_A = \lceil 23(2+2H/\ell+\log_\ell(2H+1)) \rceil$ suffices. The paired main length is at most $(2n_\lambda m+1)e_A$, with the invariant and Δ_ℓ terms charged separately. The split declarations above use these dimensions on the actual conjugate pair. The split orientation and Paired split allowance.

Scalar pushdown at the unramified places. For the ordinary unramified torsion term at $\ell > 2$, push a crystalline extension $0 \to Y \to U \to O \to 0$ through an equivariant map $Y \to (\mathrm{Frac}(O)/O)(-a)$, with $0 \le a < t$. The relevant jumps of Y have a common envelope $[-H, H]$, $H \ge 1$, and let J_Y be the set of distinct jumps in the convention where the Sen eigenvalue at level s is $(1+\ell^s)^w$ for $w \in J_Y$. Choose $u \ge 1$ least with $2H \le \ell^{u-1}(\ell-1)$, and put $q_0 = u+1$. The bound on the rationally trivial integral cocycles in $\mathcal{O}_{\mathbf{C}_\ell} \otimes Y(i)$ at level q_0 is

$$B_i = 2 + \max\left(1, \sum_{\substack{w \in J_Y \\ w+i \ne 0}} (q_0 + v_\ell(w+i))\right).$$

To see this, specialize the effective Wach frame of width $2H$ at the u-th cyclotomic point. The frame lies between the ordinary integral lattice and its multiple by ℓ, since the valuation of $(\zeta_{\ell^u}-1)^{2H}$ is at most one. The generator at level q_0 acts by a matrix $A \equiv I \pmod{\ell}$ times the cyclotomic action, because the first generator raises the parameter filtration modulo ℓ, and its ℓ^u-th power raises it by at least ℓ^u there. One power of ℓ moves an approximate invariant into frame coordinates and one uses the trace ideal of the

cyclotomic completion to move it to that completion modulo integral vectors. On each higher trace-zero layer the cyclic polynomial $\Phi_{\ell^r}(Z)$ annihilates the layer. Its evaluation at A^{-1} is ℓI modulo ℓ^2. Polynomial division therefore gives an inverse for $A\gamma-1$ with loss at most one. On the base layer apply the polynomial with the nontrivial Sen eigenvalues as roots. Its value at one has valuation $\sum_{w+i\neq 0}(q_0+v_\ell(w+i))$. These two corrections give B_i. [65, §§III.2–III.4] The integral graded steps.

The graded pieces of the integral period ring modulo the $(a+1)$-st power of its de Rham parameter ideal are $\mathcal{O}_{\mathbf{C}_\ell}(i)$ for $0\le i\le a$ at level q_0. Rational invariants lift across these quotients by the de Rham filtration basis. Thus subtracting an integral coboundary in each grade in turn bounds the total loss by $\sum_{i=0}^{a} B_i$. Transfer to the first cyclotomic level costs q_0-1. The original extension cocycle is rationally trivial in this ring by strict crystalline comparison, so applying a coefficient lift of the torsion map gives a scalar approximate coboundary after these losses. Local Tate duality allows a rational lift of the scalar pushdown class at cost $1+v_\ell(a+1)$. Its dual exponential has scalar valuation at least $-2a-v_\ell(a!)-2-a/(\ell-1)$ after the preceding period clearing. This follows by successively clearing the terms below degree a at losses $2+v_\ell(i-a)$. At degree a take normalized trace of the cyclotomic logarithm comparison. The integral dual exponential image contains ℓ times the integer ring at the first level. For this last bound the integral Iwasawa formula evaluates ∂^a of a $\psi=1$ series at the cyclotomic point. Prescribe the trace-zero parts by finite Fourier interpolation on roots of unity, and correct the value at one by adding the difference of two unit monomials whose exponents differ by ℓ. Interpolation costs one power and the correction costs $1+v_\ell(a)$ for $a>0$. The evaluation factor $\ell^{-a-1}/a!$ leaves loss at most one. At $a=0$ the constant is adjusted using the unramified Frobenius term of the $\psi=1$ equation, and the evaluation factor cancels the interpolation cost. Berger, Proposition II.5 and Theorems II.6 and II.8.

After clearing the dual exponential denominator one subtracts an integral class with the same dual exponential. The remainder lies in the rational finite subspace and changes no torsion class. Transfer from the first level has degree $\ell-1$, a unit. The resulting exponent is at most

$$(q_0-1)+\sum_{i=0}^{a} B_i+1+v_\ell(a+1)+\left\lceil 2a+v_\ell(a!)+2+\frac{a}{\ell-1}\right\rceil+1.$$

In the actual ordinary application $H\le h=\lfloor 2b\rfloor$, $a<t\le h$, and the union of distinct jump sets has cardinality at most $N_J=4n(r+t)$, where r is the cuspidal rank. With $v_0'=\lfloor\log_\ell(2h)\rfloor$, one has $q_0\le v_0'+3$ and the nonzero shifted jumps have valuation at most v_0'. Hence this exponent fits $c_B=3+v_0'+t(3+N_J(4+2v_0'))+4t+3$, the ordinary scalar allowance. This supplies the pushdown independently of the Frobenius-labelled row estimate below. The ordinary scalar allowance and The actual jump union.

At the designated quadratic ramified completions over q,d, restrict to $K_1=K_0(\mu_\ell)$ for the unramified part K_0. The restriction has index $(\ell-1)/2$, prime to ℓ. Cyclotomic prismatic evaluation gives a Wach module on $O_{K_0}[[v]]$ with $\xi=\varphi(v)/v$. The de Rham comparison gives its Frobenius bounds ξ^H and ξ^{-H} by testing at the embeddings of the first cyclotomic divisor. After shifting the module by v^H the Frobenius is integral and has height at most $2H$. In the Witt comparison with the Galois lattice, a comparison matrix U satisfies $\varphi(U)=UP$ with P integral. A negative minimum valuation of the entries modulo ℓ would be multiplied by ℓ on the left and could only increase on the right. Apply the same argument to successive Witt residues to get integrality. Applying it also to the shifted dual gives the reciprocal inclusion with factor v^{2H}. Specialization at the u-th root as above therefore gives the integral frame with loss at most one. The descent at level $u+1$ and the trace-zero comparison use the same unramified cyclotomic tower, so the scalar pushdown and torsion-invariant estimates apply with no additional exponent. Transfer back has unit index. At q,d the jump union has at most $2f(r+t)$ entries and $\lfloor\log_\ell(2h)\rfloor\le 5L$. The scalar exponent is consequently at most $60ft(r+t)L$, within the prescribed allowance. Gros–Le Stum–Quirós, Corollary 10.10, Bhatt–Scholze, Remark 5.4 and Theorem 5.6, and The scalar allowance at the two primes.

The small cyclotomic completions. Suppose the actual split row is over a completion contained in $\mathbf{Q}_p(\mu_{\mathrm{lcm}(p,N)})$, with $4\mid N$ at two. Put $s=\max(1,\mathrm{ord}_p N)$ and $N'=N/p^{\mathrm{ord}_p N}$. Over $\mathbf{Q}_p(\mu_{N'},\mu_{p^s})$ the

same prism uses $X = \varphi^{s-1}\nu$ and $Q = \varphi X/X$. Here $Q = p \bmod X$ and the cyclotomic action is trivial on the quotient prism modulo X. The height-$2H$ Witt sandwich after the common shift by X^H therefore gives $X^{2H}M_2 \subset f(M_1)$ for a surjection of crystalline lattices with jumps in $[-H, H]$. On degree i of the X-filtration the generator acts by $(1+p^s)^i$, and removal of degrees $1, \dots, a$ costs $J_s(a) = as + v_p(a!)$.

Lift the quotient generator and then the two cochains for its Frobenius and cyclotomic differences, each at cost $J_s(2H)$. On $Q^{-H}M/M$ the cyclotomic difference is annihilated by $Z\prod_{0<|i|\le 2H}(Z-((1+p^s)^i-1))$, by the filtered de Rham comparison at Q. The quotient is p-torsion-free, so this annihilation is integral. Applying the product without the initial Z to the rational error brings it into XM at cost $4Hs+2v_p((2H)!)$. Eliminate the first H degrees there at cost $J_s(H)$. On $X^{H+1}M$, inverting $1-\varphi$ by its geometric series completes the integral correction. The lift has exponent $9Hs+4v_p((2H)!)+v_p(H!)$ before transfer. Restriction and corestriction add at most $s-1+\lfloor \log_p N \rfloor$.

For the scalar estimates specialize the shifted frame at the least p^m-th root with $2H \le p^{m-1}(p-1)$. The depth of the frame is at most $s+m+\mathbf{1}_{p=2}$, squaring at two to gain the congruence modulo four. With $\lambda_N = s + \lfloor \log_p(2h) \rfloor + \lfloor \log_p N \rfloor + 1$ and N_J the prescribed union of jumps, the trace-zero inverse and the integral dual-exponential interpolation give invariant and scalar exponents at most $20(N_J+1)\lambda_N$ and $40(t+1)(N_J+1)\lambda_N$. The lift with $H \le h$ costs at most $32(2hs+1)+\lambda_N$ by Legendre's bound. The actual split sign and small-cloud dimension bound give at most $n_\lambda m$ finite directions, as in the ordinary split argument. This leaves the dyadic $O(fmb \log x)$ cost at the given conductor depth. These bounds use the selected crystalline realization and its split windows. Small-prime exponents, The dyadic range, and Use in the paired row.

At q use a paired declaration in the general comparison. The decomposition group preserves the CM type, so the large cloud has one sign at each place and the opposite sign at its conjugate. Assign the place with positive filtration jumps to the column. The dual twist at the paired row sends w to $-w-1$, so there too the large jumps are positive. The finite rational dimension on each side is at most $n_\lambda m$. Use the maximum-width lift with $s=1$, and transfer from the first cyclotomic level at index prime to q. Its exponent e_{lift} is $O(H)$ for $H \ll b$. The row length allowance is $n_\lambda m e_{\mathrm{lift}} + rB_0^{\mathrm{inv}}$, where B_0^{inv} is the invariant torsion exponent. The column uses $(n_\lambda m+1)e_{\mathrm{lift}} + (n_\lambda m+2)\Delta_q + h_\lambda^0$, with the scalar and invariant costs above. The negative Tate invariants vanish because $1 \le a < t < (q-1)/2$ and the residual inertia image has that order. Summing over places gives $O(fmb)$ plus a fixed polynomial in f, r, t times L, before weighting by $\log q$. General row and column comparison and The signed windows and dual twist.

Lifting at the mixed quadratic completions over d. Put $p = d$, $K = E_\lambda = K_0(\sqrt{-p})$ with K_0 unramified, and $n_\lambda = [K : \mathbf{Q}_p]$. The two embeddings over each unramified label have windows $[-\eta, \eta] \cup [b_+, b_+ + S]$ and $[-\eta, \eta] \cup [-b_- - S, -b_-]$, with $b_\pm > \eta$ and absolute jumps at most h. The crystalline map is a surjection of coefficient lattices. The target has rank r and at most m central jumps per embedding. Use a coefficient extension containing the images of $K_0(\mu_p)$. Both the normalized length and number of generators are preserved by this finite free extension.

On the cyclotomic prism use $\nu = [\varepsilon^{1/p}] - 1$. For the tame group $\Delta = \mathrm{Gal}(K_0(\mu_p)/K)$ of order $s = (p-1)/2$, average ν against the inverse Teichmüller character to obtain $u \equiv \nu \bmod \nu^2$ with $\delta(u) = \chi(\delta)u$. The invariant parameter is $X = u^s$. An invariant basis of the module is obtained by averaging lifts of a basis modulo ν, where the tame action is trivial. In the resulting $O[[X]]$ module the cyclotomic divisor is $E(X) = (X - a_+)(X - a_-)$. The two roots are conjugate quadratic uniformizers, with $v_p(a_+ - a_-) = 1/2$. Also $\varphi X/X$ is E^s times a unit. The cyclotomic generator γ with $\chi(\gamma) = 1+p$ acts trivially modulo X and acts on its degree-i graded piece by $(1+p)^{si}$. These identities use the filtered comparison at the two roots and the prime-to-p tame descent of the Wach module. Bhatt–Scholze, Proposition 5.3, Remark 5.4 and Theorem 5.6, and Gros–Le Stum–Quirós, Definition 10.8 and Corollary 10.10.

Write N_c for one unramified label of the module and $P_c = O[[X]]\varphi(N_{c-1})$. The jumps give $E^h N_c \subseteq P_c \subseteq E^{-h}N_c$. The same Witt sandwich as above at the first cyclotomic divisor gives an exponent $\lceil 2h/s \rceil$ in X on a surjection. Kernel exactness follows from crystalline comparison after inverting p and saturation

in the étale modules modulo p. Removing X degree by degree with $\gamma - 1$ costs $J(a) = a + v_p(a!)$ for degrees $1, \dots, a$, since s is prime to p. Set

$$\rho_{\mathrm{Wach}} = J(\lceil 2h/s \rceil), \qquad M = \lceil 4h/s \rceil, \qquad \sigma_{\mathrm{Wach}} = J(M-1).$$

Both costs are $O(1 + h/p)$.

For $W_c = E^{-h}P_c$, write $I_c^a = \{w \in W_c : (\gamma - 1)w \in X^a N_c\}$. A tuple in the I_c^M gives a finite cochain pair by inverting $\varphi - 1$ on the direct sum of the $X^M N_c$, using the convergent geometric series and $sM > h$. Conversely lift the quotient generator of a crystalline extension at cost ρ_{Wach} and remove the first $M-1$ terms in its cyclotomic difference at cost σ_{Wach}. The resulting classes come from the I_c^M. The same elimination takes I_c^1 to I_c^M modulo XN_c at cost σ_{Wach}. Its error is modulo $X^M N_c$, since $I_c^M \cap XN_c = X^M N_c$ by the first nonzero $\gamma - 1$ coefficient.

Compare W_c with constants $E^{-h}O[[X]] \otimes U$, $U = N_{c-1}/XN_{c-1}$. To construct the comparison and its inverse, lift the constant vectors to γ-fixed ones modulo X^M after inverting p. Each lift and its inverse cost at most σ_{Wach}, by the same degree equations and $J(i) + J(j) \le J(i+j)$. Applying Frobenius takes the truncation error into $\varphi(X)^M O[[X]] \subseteq X^M E^{4h} O[[X]]$. At either root a horizontal lift is the constant filtered vector to the required order. Indeed its γ-invariance forces all nonconstant coefficients in the logarithmic parameter to vanish in that range. Thus the two comparisons preserve the filtered sublattices and the I^a, and compose to $p^{2\sigma_{\mathrm{Wach}}}$ modulo X^M of the filtered lattice. The comparison in the constant model uses the conditions that the coefficient of degree j at each root belongs to Fil^{-j} at that root.

In the target U, split off a summand in $\mathrm{Fil}_+^{b_+}$. Its complement U_r can contain $\mathrm{Fil}_-^{-b_-+1}$ and have rank at most $2m$: combine a complement to the first saturated filtration with the second central subspace and saturate, then split off the remaining part of the first filtration. On an invariant residue modulo X of the filtered lattice, apply the polynomial in γ with roots $(1+p)^j$ for $j \in [-b_+ - S, -b_+ - 1] \cup [b_-, b_- + S - 1]$. Its value at one has valuation at most $\kappa = \lceil 2(S + S/(p-1) + \log_p h) \rceil$. Indeed in an interval of S integers at most $S/p^a + 1$ are divisible by p^a. The large-summand coordinates of the result have pole order at most b_+ at a_+ and are constant through degrees less than $B = b_- + S$ at a_-. The U_r remainder residues require at most $3 \operatorname{rk} U_r$ generators. The quotient modulo X of the filtered lattice is finite torsion-free in the constant model, by testing at $X = 0$ and the two roots after clearing E^h. At each of these three locations its invariant part in the remainder has at most $\operatorname{rk} U_r$ dimensions, since at a root only degree zero is invariant.

For a large-summand coordinate $G = A/E^h$ put $b_+ = b_1$. If $b_1 \le B + 2\eta$, divide A by $(X - a_+)^{h-b_1}$. Its coefficient in degree $h + k$ at a_- is $C_- \binom{b_1}{k} (a_- - a_+)^{b_1 - k}$ for $0 \le k < B$, where C_- is the constant of G there. With $k = \min(b_1, B-1)$, the root separation and the base-p carry bound on the binomial valuation show that $p^{y_c} C_-$ is integral for $y_c = \lceil \eta + 1/2 + \log_p h \rceil$. Subtract that constant. The remaining filtered vector can be lifted at cost $\rho_{\mathrm{Wach}} + y_c$; the constant lifts using the same Wach bound modulo X. If $b_1 > B + 2\eta$, lift the filtered test vector $(X - a_+)^{-b_1}(X - a_-)^B u_+$, for the given large basis vector u_+, after applying the horizontal comparison. On returning to the source constant model, multiplication by E^h and division by $(X - a_-)^{h-\eta}$ are integral. Extract the coefficient of degree $h - b_1 + B + \eta$ at a_+. It lies in $\mathrm{Fil}_+^{b_1 - B - \eta} = \mathrm{Fil}_+^{b_1}$ by the source gap. Its image is $p^{\rho_{\mathrm{Wach}} + 2\sigma_{\mathrm{Wach}}} u_+$. Multiplying it by G gives the lift in this case.

To count the classes, normalize the initial cochain at cost $\rho_{\mathrm{Wach}} + \sigma_{\mathrm{Wach}}$ and pass through the two horizontal comparisons at cost $2\sigma_{\mathrm{Wach}}$. The polynomial costs κ, after rescaling by an integral scalar so its value at one is p^κ. The large summands lift at cost $\rho_{\mathrm{Wach}} + \max(2\sigma_{\mathrm{Wach}}, y_c)$. The remaining classes come from the U_r residues. Returning their lifts and the source lifts from I^1 to I^M costs σ_{Wach} and makes the cochain comparison an equality of finite classes. Thus the cokernel D_F on integral finite cohomology, after multiplication by p^y, needs at most $6m$ generators per unramified label, with

$$y = 2\rho_{\mathrm{Wach}} + 4\sigma_{\mathrm{Wach}} + \kappa + \max(2\sigma_{\mathrm{Wach}}, y_c).$$

The maximum-width lift on the whole cokernel costs $O(h)$, by restriction to the first cyclotomic level and unit-index transfer. The target needs at most $(n_\lambda + 1)r$ generators by the finite dimension and tor-

sion bounds. Hence the normalized length of $D_F/p^y D_F$ is bounded by $y(n_\lambda + 1)r$ and that of $p^y D_F$ by $O(3n_\lambda m(h+1))$. This gives

$$\operatorname{len}_p D_F \ll (n_\lambda + 1)r(2h/p + 2S + 2\log_p h + 1 + \eta) + 3n_\lambda m(h+1).$$

The two-root comparison and Normalized row colength.

On the actual d-rows inertia switches the CM sign over each unramified label. The dual twist sends w to $-w-1$, so the deep bases become s_0-1 and s_0+1 for the respective original endpoints s_0. Thus $\eta = \lfloor \eta_{\mathrm{sz}} \rfloor + t$, $S = \lfloor S_c \rfloor$, and $h = \lfloor 2b \rfloor$ work. Removing the consecutive scalar weights leaves at most m central jumps with multiplicity one. Strictness preserves the windows on the prescribed subrepresentations. Since $\sum_{\lambda \mid d} n_\lambda = 2f$, the selected range $r/d \ll m$ and $r(S_c + \eta_{\mathrm{sz}} + t + \log b) \ll mb$ gives unweighted row cost $O(fmb)$. The scalar and inertia terms still use their separate bounds. The crystalline realization and labelled matching here are those of the actual type cut. The following estimate uses unramified completions and a different ordering of the labelled windows. The actual windows and target counts and The distinguished-prime size range.

The word at an inert completion. Let ℓ be an ordinary inert prime in the local defect table, and let $\lambda \mid \ell$ in E. The completion K_λ is unramified of degree $n_\lambda = \operatorname{ord}(\phi)$, where $\phi = \operatorname{Art}(\ell)$ inside the host group. Choose the coefficient extension used for this completion to contain all images of K_λ and the primitive ℓ-th root used in the alternating lift. The local group stays fixed under this extension. An O-basis of the new coefficient ring identifies the extended cochain complex with a finite direct sum of the old one as O-complexes, and the rational period comparison respects this decomposition. Hence both integral H^1_f and its cokernel commute with the coefficient extension. Lengths return to the original lattice in the same normalized units.

Label the embeddings by $\tau_{i+1} \circ \mathrm{Frob}_{K_\lambda} = \tau_i$. Indeed the map of the Witt scalars in the linearized Wach Frobenius is arithmetic Frobenius, and the target embedding of the coefficient factor following τ_i must precompose it to give τ_i. The corresponding complex labels are $g_i = g_\circ \phi^{-i}$, with $g_\circ$ the label of τ_0. This is a cyclic ordering of the embeddings of the actual completion, because the unramified decomposition group acts simply transitively on them. The CM sign of ϕ is -1, so the signs σ_i of the direct outside clouds alternate, and n_λ is even. For the dual words put

$$B_i = s_i + \sigma_i(t+1)/2, \qquad S = \lfloor S_c \rfloor, \qquad \eta = \lfloor \eta_{\mathrm{sz}} \rfloor + t + 2, \qquad s_i = e_{\rho(g_i), j-1}.$$

The B_i are integers. For $\sigma_i = +1$ the dual large interval is $[-B_i - S, -B_i]$, and for $\sigma_i = -1$ it is $[B_i, B_i + S]$. The other jumps are in $[-\eta, \eta]$, including the scalar jumps of an ambient dual. The fixed cuspidal dual $T_{2,\pi}$ has at most m central jumps at each label. The signed radius variables of the local comparison can thus be taken as $(M_i, h_i) = (\eta, B_i + S)$ on the up labels and $(B_i + S, \eta)$ on the down labels. These are common to all generic points in the coarse record and to all the subquotients used in the row surjection. The prime is at least eleven because the prescribed smaller odd primes split and the other small primes have their own branch. The inverse Frobenius order in the local coefficient factors, The word and its signed widths, and The natural word for these actual modules.

Set $\widehat{b} = \lceil b \rceil - (t+1)/2$. The structural design supplies the thick real labels and the three ranges for s_i. At a thick real label, $b \le s_i \le b + \eta_{\mathrm{sz}} + 1$, so $\widehat{b} \le B_i \le \widehat{b} + \eta$. Every other label has

$$w''(n) \le s_i \le w_B, \qquad w_B = \max\{w''(n) + \eta_{\mathrm{sz}},\ \max_{a \in T} b_a + b'/x + \eta_{\mathrm{sz}}\},$$

where the b_a in the maximum are the top-label parameters. All such integer s_i have $B_i \le w_P = w_B + (t+1)/2$, as charged in the row table. The chosen scale separates this upper bound from $\widehat{b}$. Thus the numerical thick indices in the local word are exactly the designated thick real labels. All jumps in the word have absolute value $\le h_B$ under the row-table convention. The local depth and width satisfy $S \le S_c$ and $\eta \le \eta_3 = \eta_{\mathrm{sz}} + t + 2$. The selected three real ranges and The thick labels after the sign shift.

The thick condition is constant on each conjugate pair and on the real, dyadic coordinate is membership in a translate $u_0 + T_2 \subset \mathbf{Z}/n_0\mathbf{Z}$. The design has $4L_w a < n_0$ for every $a \in T_2$, and excludes any relation $j_2(a_1 - a_0) = j_1(a_2 - a_0)$ on three different a_i with distinct positive $j_1, j_2 \le L_w$. If the dyadic part of ϕ has order $> 4L_w$, three thick positions $i_0 < i_1 < i_2$ within L_w steps would give that relation modulo n_0, with $j_u = i_u - i_0$, by eliminating the common orbit increment. Its difference has absolute value less than n_0 from the size condition. Thus the congruence is an equality, contradicting the raw design. This argument works at the cyclic boundary using integer lifts of the three positions. At $\ell < \ell_A$ put $j_A = \max(1, \lceil \log_{\ell-1} \ell_A \rceil)$. The selected inequalities give $2j_A + 2 \le L_w$, and the order bound gives $n_\lambda \ge j_A + 1$. Joining consecutive thick positions at forward distance at most j_A gives clusters of size at most two. Three joined positions would fit in $2j_A + 1$ consecutive slots. The clusters are separated by at least j_A small positions. The width inequalities of the design give $C(j_A + 1)(w + S + \eta + 1) \le \widehat{b}$ with the local small radius w used in the table. In fact $w + S + \eta + 1 \le 2b'$, $j_A + 1 \le L_w/2$, and the fixed multiple of $(L_w + 1)b'$ is absorbed by b at the common threshold. This supplies the long-order clustered word on this crystalline map. The modular relation in a thick interval, The actual depth and width allowance, and The application to the word.

For the other inert primes of dyadic order $> 4L_w$, the same count excludes three consecutive thick labels and yields at least one small label. Consider next dyadic order $d_H \le 4L_w$. Multiplying differences of two thick dyadic coordinates by d_H gives a multiple of n_0, and its absolute value is less than n_0. Hence they have the same dyadic coordinate. The embedding occupies that coordinate in the thick set at only one real label. Two successive thick positions at that label require $\phi = 1$ or $\phi = c_\infty$. The negative CM sign excludes the identity. Thus unless the orbit is $\{g, gc_\infty\}$ at one thick label, the word has a small label and no triple of thick labels. It alternates on both sides of the cyclic boundary. The small-radius inequalities for that branch are $20w, 20S, 20\eta \le \widehat{b}$ and $\eta < B_i$ at every label. They follow from $w \ll b'$, $\eta + S \ll w''(n) \le b'$, $s_i \ge w''(n)$, and the selected separation of b from b'. A small position in the band comparison has the same s_i with $|B_i - s_i| = (t + 1)/2$. This is the witness for the charged band in the row table.

If $\phi = c_\infty$ and the real label is thick, the two positions are thick with opposite signs. Let $M = \max_i B_i$ and $H = M + S + \eta$. The size conditions give $\widehat{b} \ge 100$ and $2\eta + S \le \widehat{b}/10$; thus $H \le 1.1\widehat{b}$ and $B_i \ge (M + S)/2$. For $\ell \ge 11$ use $(2\ell - 1)/(\ell - 1)^2 \le 1/4$ and $\log_\ell(2H) \le \widehat{b}/20$. The latter follows from $\log_{11}(2.2\widehat{b})/\widehat{b}$ decreasing for $\widehat{b} \ge 100$ and its value at 100. It follows that

$$2\left(H\frac{2\ell - 1}{(\ell - 1)^2} + 3 + \left(\log_\ell(2H) + 1\right)\right) + \eta + 1 \le 0.70\widehat{b} + 9 \le \widehat{b} \le B_i.$$

This gives the required two all-thick lower bounds with the actual signed radii. Together the two alternating cases supply the word data for the separate local lifting steps, while the scalar, invariant and fine Selmer terms use their own estimates. The short and conjugate orbit alternatives, The all-thick word, and The realized alternating comparison.

AI-proposed candidate (signed Frobenius-labelled crystalline lifting). Let $\ell \ge 11$ be prime and $K_\ell/\mathbf{Q}_\ell$ be unramified of degree n_ℓ. Take a finite coefficient field E'_ℓ splitting K_ℓ, with valuation ring O and valuation normalized by $v(\ell) = 1$. Label its embeddings of K_ℓ by $i \in \mathbf{Z}/n_\ell$ with $\tau_{i+1} \circ \sigma_{K_\ell} = \tau_i$, where σ_{K_ℓ} lifts residue Frobenius. Let $F_\ell : T_1 \twoheadrightarrow T_2$ be an equivariant surjection of finite free stable O-lattices in crystalline representations, and let $r_\ell = \mathrm{rk}_O\, T_2$. A labelled jump a means $\mathrm{gr}^a\, D_{\mathrm{dR}}(T_1[1/\ell])_{\tau_i} \ne 0$ in the decreasing filtration.

For every C above an absolute positive threshold, the estimate holds with the following data. Take nonnegative integers b_*, w_*, S_*, η_*, and $1 \le j_* \le n_\ell - 1$. At each label choose an integer $s_i \ge 0$ and a sign, independently of the other labels. Its jumps lie in $[-\eta_*, \eta_*]$ together with either $[s_i, s_i + S_*]$ or $[-s_i - S_*, -s_i]$, according to that sign. The multiplicities are arbitrary. Set $g_* = w_* + S_* + \eta_*$ and $I = \{i : b_* \le s_i \le b_* + \eta_*\}$. Require $s_i \le w_*$ outside I and $b_* \ge C(j_* + 1)(g_* + 1)$. For each $t_0 \in I$, either $t_0 - 1, \dots, t_0 - j_*$ are outside I, or there is $1 \le d'_* \le j_*$ with $t_0 - d'_* \in I$ and $t_0 - a \notin I$ for $1 \le a \le d'_* + j_*$, $a \ne d'_*$. All indices in this condition are cyclic. A pair within a cluster can be adjacent. The same jump windows contain those of the target by strictness.

Put $H_{\mathrm{loc}} = b_* + \eta_* + S_*$ when $I \neq \emptyset$ and $H_{\mathrm{loc}} = g_*$ otherwise. These bounds include the endpoints $s_i + S_*$ of all the prescribed outer windows, even when an outer window contains no jumps. Define $H^1_f(K_\ell, T)$ as the full inverse image in $H^1(K_\ell, T)$ of the rational Bloch–Kato finite subgroup defined by B_{cris}. It includes the integral torsion. Then the cokernel D_{F_ℓ} of the induced map on these groups has finite length and satisfies

$$\frac{\operatorname{length}_O D_{F_\ell}}{\operatorname{length}_O(O/\ell O)} \leq C(n_\ell + 1) r_\ell \left(\frac{H_{\mathrm{loc}}}{(\ell-1)^{j_*+1}} + (g_* + 1)(j_* + 1) + \log_\ell \max(1, H_{\mathrm{loc}}) \right). \tag{71}$$

The denominator is the absolute ramification index of E'_ℓ. Exact clustered estimate and Period and comparison data.

The cut exponent. In the proof write $p = \ell$, $G_0 = g_*$, $g = \max(1, G_0)$, $h = \max(1, H_{\mathrm{loc}})$, and $j = j_*$. Choose the minimal bounds $[-M_i, h_i]$ enclosing the jumps and zero. Their widths are at most $2h$ and at most $2G_0$ at thin labels, while a single signed outer window gives width at most $h + g$. Group successive large labels whose cyclic distance is at most j. The predecessor condition makes each group have at most two large labels and leaves at least j thin labels between groups. Cut at the last label of each group. If the next group's first label lies more than $2j + 1$ steps away, insert cuts at successive distances j. Thus all arcs have length from j to $a_* \leq 3j + 1$ and start with j thin labels. An all-thin cycle can be divided into lengths from j to $2j$. This construction also works when two large labels are adjacent.

The rational finite-group dimension is at most $n_\ell r_\ell + d_0$, where $d_0 = \dim(T_2[1/\ell])^{G_{K_\ell}}$. The torsion embeds in the quotient of $T_2[1/\ell]/(T_2[1/\ell])^{G_{K_\ell}}$ by the image of T_2, so it requires at most $r_\ell - d_0$ generators. This proves the factor $(n_\ell+1)r_\ell$ when a uniform annihilating exponent is converted into normalized length. For the full-parameter Wach modules, the factor at cut i is the factor of e_{i+1}, with e_i the idempotent of τ_i. Thus Frobenius maps the factor at cut $i - 1$ to that at i, and specialization by φ^{-s-1} at cyclotomic level s on the latter reads τ_{i-s}. Write P_i^{tar} for the target factor and P_i^{im} for the image of the source. Let s'_i be the least exponent of μ_{cyc} killing their quotient. For labelled bounds $[-M_i, h_i]$ put $H_0 = \max_i(M_i + h_i)$. A common rank-one crystalline twist by the conjugate Lubin–Tate characters shifts the i-th interval by M_i to $[0, M_i + h_i]$. Its Wach lattice is free of rank one, and tensoring preserves the cokernel exponent. The effective Witt sandwich of height H_0 gives $s'_i \leq H_0$. Use the same argument and the target widths for the map from a crystalline extension of T_2 to the trivial quotient. Kedlaya–Pottharst–Xiao, Example 6.2.6, and [65, §§II.1 and III.3]

If $c\mu_{\mathrm{cyc}}^{k+1} P_i^{\mathrm{tar}} \subset P_i^{\mathrm{im}}$, apply the Wach congruence $\gamma - 1$ on the image and divide by μ_{cyc}. Its leading term gives $c(\chi(\gamma)^{k+1} - 1)\mu_{\mathrm{cyc}}^k P_i^{\mathrm{tar}} \subset P_i^{\mathrm{im}}$, after subtracting the correction in degree $k + 1$. Choose $\chi(\gamma)$ generating $\mathbf{Z}_\ell^\times$. Descending from s'_i to power $a \leq s'_i$ costs at most

$$\sum_{\substack{a<k\leq s'_i \\ \ell-1\mid k}} (1 + v_\ell(k)) \leq (s'_i - a)\frac{\ell}{(\ell-1)^2} + c_H, \qquad c_H = \lceil 1 + \log_\ell(H_0 + 1) \rceil.$$

For the inequality count multiples of $(\ell-1)\ell^j$ in that interval, allowing one endpoint error at each power.

Let $q_{\mathrm{cyc}} = \varphi(\mu_{\mathrm{cyc}})/\mu_{\mathrm{cyc}}$ and $u'_i = \lceil s'_i/(\ell-1) \rceil$. Since $q_{\mathrm{cyc}} = \mu_{\mathrm{cyc}}^{\ell-1} + \ell B$ with B integral, the binomial expansion shows $\mu_{\mathrm{cyc}}^{u'_i} q_{\mathrm{cyc}}^{u'_i + c_H + 2} P_i^{\mathrm{tar}} \subset P_i^{\mathrm{im}}$. Indeed a term with exponent v on $\mu_{\mathrm{cyc}}^{\ell-1}$ has degree $a = (\ell-1)v + u'_i$. When $a < s'_i$, $(s'_i - a)\ell/(\ell-1)^2 \leq u'_i - v$, which is within its available ℓ-exponent after paying c_H. Otherwise the definition of s'_i applies. If the labelled jumps at i are in $[-M_i, h_i]$, the two Frobenius height inclusions with $q_{\mathrm{cyc}}^{M_i}$ and $q_{\mathrm{cyc}}^{h_i}$ now yield

$$s'_{i-1} \leq M_i + h_i + c_H + 2 + u'_i.$$

For this descent the height inclusions follow by testing poles at the first cyclotomic roots. Quotients by powers of their monic polynomial are coefficient-free, so the bounds descend integrally. The linearized Frobenius is injective and the parameter ring is free over its Frobenius image with basis $1, \mu_{\mathrm{cyc}}, \dots, \mu_{\mathrm{cyc}}^{\ell-1}$.

Thus membership of a Frobenius image in the span of the previous image descends to membership before Frobenius. Labelled Frobenius and the cut recurrence.

In cut notation the recurrence uses the width $M_{c+1} + h_{c+1}$ when descending to c. Iteration around the orbit, retaining H_0 as a second upper bound, gives the minimum of H_0 and the integer part of the geometric sum with ratio $1/(p-1)$ of the successive widths plus $c_H + 3$. At a cut the first j widths are at most $2G_0$. Inserting the resulting bound in the γ-descent at degree zero gives one common exponent r_0 for the two maps under consideration, with

$$r_0 \ll \frac{h}{(p-1)^{j+1}} + \frac{g}{p} + 1 + \log_p h.$$

This descends to the tame invariant modules by averaging a preimage over the tame group. The same exponent lifts constant residues modulo its parameter. The cut exponent and its constants.

Truncation on an arc. Use the tame invariant parameter

$$X = \frac{1}{p-1} \sum_{\delta \in \Delta} \delta(q_{\mathrm{cyc}}) - p, \qquad P_s = \varphi^s(X+p), \qquad R = O[[X]].$$

The tame average of μ_{cyc} against the inverse Teichmüller character is a parameter transforming by that character. Its $(p-1)$-st power generates the invariants, and X has that initial degree modulo p. The same averaging lifts an invariant Wach basis from modulo μ_{cyc}. The first cyclotomic roots are zeros of $X + p$, and monic division makes $X + p$ associated to q_{cyc} in the full parameter ring. Thus $\varphi^a X$ is associated to $X \prod_{s<a} P_s^{p-1}$. Each P_s has constant term p, is congruent to X^{p^s} mod p, and is associated to a monic distinguished polynomial of degree p^s. Use γ now for a generator of the pro-p cyclotomic group and put $u = \chi(\gamma)$, with $v_p(u-1) = 1$. It preserves the divisors and $\gamma(X)/X \equiv u^{p-1}$ mod X.

Set

$$M = \left\lfloor \frac{h+2g}{p-1} \right\rfloor + 1, \qquad \sigma = \left\lceil \left(3g + \frac{M}{p^{j-1}}\right) \frac{p}{(p-1)^2} + 1 \right\rceil, \qquad J_j = X^M \prod_{s<j} P_s^{3g}.$$

For $b \le pk/(p-1)$, $p^\sigma \varphi^j(X)^k \in (p^b, J_j)$. For $b > \sigma$ put $e = \lceil M/p^{j-1} \rceil$. The choice of σ gives $(p-1)k - 3g \ge b - \sigma + e$. Retain $\prod_{s<j} P_s^{3g}$ in the factorization of $\varphi^j(X)^k$, and expand the remaining power of $P_{j-1} \in (p, X^{p^{j-1}})$. Each term of its power $b - \sigma + e$ either has p-exponent at least $b - \sigma$ or X-exponent at least $p^{j-1}e \ge M$. This proves the congruence. Applying φ^{a-j} gives the same estimate modulo $J'_a = \varphi^{a-j}(J_j)$ on an arc of length a.

The horizontal formal frame for γ at the preceding cut is specified by its constant vector at $X = 0$. The equation in degree k divides by $u^{(p-1)k} - 1$. Thus the k-th coefficient of the frame and of its inverse has denominator at most $p^{k+v_p(k!)}$. The same estimate solves $(\gamma - 1)v = q$ for $q \in XN$ with zero prescribed constant. Since $k + v_p(k!) \le pk/(p-1)$, the congruence makes each transported frame and solution integral modulo its transported ideal after scaling by p^σ. The quotients by those ideals are coefficient-free by monic division, so these comparisons are compatible with the lattice maps. The truncation congruence.

Index the levels $s = 0, \dots, a-1$ backwards from the arc endpoint c. The rational completion at P_s uses the transported field at label τ_{c-s}. Its logarithmic parameter t_s is scaled by u under γ, and the coefficient action has order dividing p^s after tame invariants. The endpoint lattice $N_{i,s}$ for source or target i has coefficient of t_s^k in Fil^{-k} of the transported constant space. In the filtered de Rham comparison this gives the containing lattice of the Wach image, and Berger's elementary divisors give equality of colengths. Under coefficient splitting the inclusion and equality of the sum of colengths give equality at each factor. The horizontal frame transported by Frobenius is the constant crystalline frame there, since its positive parameter terms vanish at that level. [65, §§II.2–II.3 and III.2]

Round the jumps to 0 at thin levels, to $b_s, 0$ at up levels, and to $0, -b_s$ at down levels. The b_s line of the rounded filtration is the subspace Fil^{b_s} for an up level. The zero part at a down level is $\mathrm{Fil}^{-\eta_*}$. Denote

their saturated constant lattices by U_s. The rounded local lattice $B_{i,s}$ differs from $N_{i,s}$ by multiplication by P_s^g in either direction. Put $\Lambda = \prod_s P_s^{H_s}$ with $H_s = g + g_*, g + b_s + S_*, g + \eta_*$ at thin, up, and down levels respectively. The arc lattice W_i is $\Lambda^{-1}\,\mathrm{span}_R(\varphi^a N_{i,c-a})$. Its constant partner has the same Λ^{-1} factor on the free constant-lattice series.

The preceding truncation gives partners in both directions, each with cost p^σ and the same local errors of order at least $\mathrm{ord}_{P_s} J'_a - H_s$. At the first $a - j$ levels the first term is $(p-1)M$; at the last j thin levels it is $3g$. These orders are at least the lattices' required orders, namely η_* at an up level, $b_s + S_*$ at a down level, and g_* at a thin level. Thus every partner error belongs to $N_{i,s} \cap B_{i,s}$. An error on the constant side forces divisibility by $P_s^{H_s}$ before dividing the series by Λ. The partners commute with the lattice maps modulo these same errors, since maps between horizontal frames are constant in the initial parameter. Their constant map F_0 is precisely reduction of the Wach map at the preceding cut modulo X.

Write $\Theta = \prod_s P_s^g$. The linearized Frobenius matrices and inverses have denominators only at the first divisor, with pole orders given by the labelled jumps. Iterating gives exactly the height tests at the P_s. Weierstrass saturation shows that $\Theta^{-1}N_{i,c} \subset W_i$ consists of the elements whose localizations lie in $P_s^{-g}N_{i,s}$ for all s. On this submodule the map lifts p^{r_0} times any target by the cut bound. Membership in $XN_{i,c}$ is tested by membership in XW_i and local membership in all $N_{i,s}$, since X is a unit at the rational cyclotomic divisors. Finally $(\gamma - 1)W_i \subset XW_i$ by the Wach action on the previous cut and the constant terms of the divisor multipliers. Partner comparisons and level tests.

A residue fixed modulo $XN_{2,c}$ has localizations in constants plus $N_{2,s}$. Apply γ^{p^s} to each nonzero degree to see that the coefficient in its filtration quotient is zero. In degree zero average the finite coefficient action to take a constant. To pass from $P_s^{-g}B_{2,s}$ back to $B_{2,s}$, kill the nonzero degrees in the rounded boundary windows. For the reverse comparison use the union of the intervals $[-w-g, -w-1]$ for actual jumps w of the source. Each union has one central interval and at most two translated intervals of total width $O(g+1)$. For an arc of length a, the polynomial is the product of $Z^{p^a} - u^{kp^a}$ over the indicated nonzero degrees k, times $\sum_{j_0=0}^{p^a-1} Z^{j_0}$ to make degree zero constant. It kills the desired degrees since the coefficient action has order dividing $p^s < p^a$. Its valuation at one is $a + \sum_k (a + 1 + v_p(k))$. Counting multiples of successive powers of p in each interval gives a common bound $\kappa_{\mathrm{arc}} = O((g+1)(a_*+1) + \log_p(2h))$ for the two polynomials P, P'. The two polynomial windows.

Integral constants on the signed levels. Let L be one of the free constant lattices and write $L_R = L \otimes_O R$. In these coordinates, before division by Λ, the rounded condition $\mathcal{N}$ is membership at each s in

$$\begin{array}{ll} P_s^{H_s} L_R & \text{at a thin level,} \\ P_s^{H_s - b_s}(U_s)_R + P_s^{H_s} L_R & \text{at an up level,} \\ P_s^{H_s}(U_s)_R + P_s^{H_s + b_s} L_R & \text{at a down level.} \end{array}$$

The condition $\mathcal{I}$ allows subtraction of Λd_s with $d_s \in L[1/p]$ at each level. Divisibility over $R[1/p]$ by products of the distinct distinguished powers descends to R by monic division, also after projecting to L/U_s.

Without a down level the conditions coincide. With one down level and no up level, write $y = \Lambda z$. The class of z modulo $P_s^{b_s}$ in L/U_s is a constant, integral because 1 occurs in the integral remainder basis. Lift it to $\lambda \in L$. Subtracting $\Lambda\lambda$ puts y in $\mathcal{N}$. With two down levels $i = 0, 1$ put $g_i = P_i^{b_i}$ for their two powers, relabelling these two indices. The quotient $R/(g_0, g_1)$ is finite torsion over O. If it is killed minimally by ϖ^v for a uniformizer ϖ, the constant 1 has order ϖ^v. A largest Smith-length coordinate of 1 is primitive, so its constant submodule O/ϖ^v splits. The two integral constants in L/U_i therefore agree modulo $U_0 + U_1 + \varpi^v L$. Take lifts $\lambda, \lambda + \varpi^v\lambda'$ in L, respectively for the second and first constants. A scalar series B congruent to zero modulo g_1 and ϖ^v mod g_0 exists by the definition of v. Subtracting $\Lambda(\lambda + B\lambda')$ suffices. The coefficients of this scalar series are integral, so only the two ordinary vectors have to be lifted.

For an up level u and a down level d, put $g_u = P_u^{b_u}$, $g_d = P_d^{b_d}$, and write $y = (\Lambda/g_u)z$. Then $z \in (U_u)_R + g_u L_R$, and $q_d z \equiv g_u d_0 \bmod g_d$ over the fraction field, where $q_d : L \to L/U_d$. Write the remainder of g_u modulo g_d as $\varpi^v \alpha$ with α primitive. Applying an integral functional taking α to one shows that the remainder of $q_d z$ is αv_0 with $v_0 \in q_d(U_u) + \varpi^v(L/U_d)$. Consequently one can remove $\Lambda\lambda + (\Lambda/g_u)\alpha\lambda'$, where $\lambda \in L$ and $\lambda' \in U_u$. If $u < d$, then $b_u < p b_d$ gives $\deg(P_u^{b_u}) < \deg(P_d^{b_d})$. The monic remainder is primitive and one can take $\lambda' = 0$. The order $d < u$ can require that filtered constant. The integral constant conditions.

Put $A_0 = 2\sigma + r_0 + 2ga_*$. For $d < u$ we need $2b_d \le p b_u$ and $2A_0 < b_u$. To lift the filtered constant, let $v_2 \in U_{2,u}$ and take $y_2 = (\Lambda/g_u)g_d v_2$ in $\mathcal{N}_2$. Its arc partner belongs to $\Theta^{-1}N_{2,c}$ by the level tests. Lift p^{r_0} times that partner and return to the constant model. The two partner factors are $p^{2\sigma}$. Their local errors force divisibility by Λ on the constant side. Thus there is y_1 with $\Theta^2 y_1 \in \mathcal{N}_1$ and $F_0 y_1 \equiv p^{2\sigma+r_0} y_2 \bmod \Lambda$. Divide $\Theta^2 y_1$ by Λ/g_u and reduce modulo g_u. The residue belongs to the source up filtration and maps to Dv_2, with $D = p^{2\sigma+r_0}\Theta^2 g_d \bmod g_u$.

For its scalar content use Gauss valuation assigning X the value θ/p^u for $0 < \theta < 1$. The values of P_u and P_d are θ and θp^{d-u}, by the leading terms of their distinguished polynomials. Every P_s has value at most one. With $A' = 2\sigma + r_0 + 2ga$, choose $A'/(b_u - b_d/p) < \theta < 1$. This gives

$$v_\theta(p^{2\sigma+r_0}\Theta^2 g_d) \le A' + \theta b_d/p < \theta b_u = v_\theta(g_u).$$

If the remainder has content ϖ^v, the Gauss inequality for an expression in the ideal (ϖ^v, g_u) gives $v/v_\varpi(p) \le A' + \theta b_d/p$. Letting θ decrease to the lower endpoint bounds the content by $2A_0$ in normalized units. A linear functional on the integral remainder module taking its primitive part to one now gives $p^{2A_0}U_{2,u} \subset F_0(U_{1,u})$. This also lifts the required filtered constant when the coefficient valuation is ramified. Filtered constant and the weighted Gauss estimate.

Together the cases show that for any $y_2 \in \mathcal{I}_2$ there is $y_1 \in \mathcal{I}_1$ with $p^{2A_0}y_2 - F_0 y_1 \in \mathcal{N}_2$, using the same integral scalar series for the lifted constants. Start with an invariant target residue w mod $XN_{2,c}$. Apply the first polynomial $P(\gamma)$ so its localizations belong to constants plus $B_{2,s}$, and take its constant partner. Lift its constants as above and return to an arc partner z_1. The discrepancy $p^{2\sigma+2A_0}P(\gamma)w - fz_1$ belongs locally to $B_{2,s}$, by the two partner comparisons. Lift it in $\Theta^{-1}N_{1,c}$ with cost p^{r_0} and add the lift to $p^{r_0}z_1$. This gives an exact arc lift whose localizations are in constants plus $P_s^{-g}N_{1,s}$. The second polynomial $P'(\gamma)$ takes these into constants plus $N_{1,s}$. Its γ-difference thus belongs locally to $N_{1,s}$, and globally to XW_1, so it lies in $XN_{1,c}$. Modulo $XN_{2,c}$ the image of the resulting invariant residue is

$$p^{2\sigma+2A_0+r_0}P(1)P'(1)\,w.$$

The valuation cost is at most $2\sigma + 2A_0 + r_0 + 2\kappa_{\mathrm{arc}}$. The simultaneous arc lift.

The finite classes represented by the cut residues. Use the integral Herr cocycle convention $d_\gamma y = d_\varphi z$, with $(d_\varphi v)_i = \varphi(v_{i-1}) - v_i$ and $d_\gamma = \gamma - 1$. On the direct sum of the $X^M N_i$ the operator d_φ is invertible by its negative geometric series. One Frobenius step multiplies by $P_0^{(p-1)M}$ in the parameter part, which clears the possible pole of order h and leaves a factor in (p, X). Hence if $d_\gamma y$ is in this direct sum there is a unique second cochain there. Changing y by a member of the same direct sum changes the pair by a coboundary.

Such a pair is finite when the first cochain has only finitely many cyclotomic denominators. To verify this, write the associated integral étale extension with quotient vector e fixed by the tame group. If the pole order of the first cochain at P_s is at most m_s, choose nonnegative integers c_s eventually zero with $c_{s-1} - c_s \ge m_s$ for $s \ge 1$. Put $D = \mu_{\mathrm{cyc}}^k \prod_s P_s^{c_s}$ with $k \ge c_0 + m_0$, large enough for $\mu_{\mathrm{cyc}}^k N$ to be Frobenius-stable. In the enlarged full-parameter free lattice with added generator De, the quotient multiplier $\varphi(D)/D$ is integral and the off-diagonal entry $\varphi(D)y/\mu_{\mathrm{cyc}}^k$ is integral. The cyclotomic action has scalar χ^k modulo μ_{cyc} because the second cochain is in $\mu_{\mathrm{cyc}}N$. This lattice generates the étale module, since its cyclotomic

factors are units in the étale period ring. It lies in the positive period module by the Witt-coordinate minimum argument for an integral Frobenius matrix. The finite-height criterion with scalar reduced cyclotomic action makes the extension crystalline. Herr, the integral comparison, [65, §§I.2 and II.1] and The finite-class criterion.

At a cut c, take an invariant residue w with $q = d_\gamma w \in XN_c$. Solve $(\gamma-1)v = q$ in the horizontal formal frame with zero constant, and take an integral approximation $\tau(q)$ to $p^\sigma \varphi^j(v)$ modulo J_j in the Frobenius pull span. Its class modulo that ideal is unique by saturation and depends only on q mod $X^M N_c$. The first j steps are thin, so J_j times this span lies in $X^M N_{c+j}$. The first cochain supported at $c+j$ with value $p^\sigma \varphi^j(w) - \tau(q)$ has cyclotomic difference in $X^M N_{c+j}$. Solve for its second cochain as above. This defines a finite class $\mathcal{R}_c(w)$ functorially, since changing the integral approximation adds a coboundary. It kills XN_c, by taking $v = w$ in that case. The shifted cut residue.

Represent a target class ξ by a crystalline lattice extension of the trivial quotient. At each cut c choose x_c mapping to p^{r_0} in the quotient Wach lattice. Kernel exactness identifies $q_c = (\gamma-1)x_c$ with an element of $XN_{2,c}$. The vector $R_c = p^\sigma \varphi^j(x_c) - \tau(q_c)$ lifts $p^{r_0+\sigma}$ at $c+j$. Propagate each R_c by Frobenius to the next shifted cut. These quotient lifts give a pair representing $p^{r_0+\sigma}\xi$, with second cochain in $X^M N_2$. For the arc from c to $d = c+a$, put $w_d = \varphi^{a-j}(R_c) - p^\sigma x_d$. The quotient cancels. The height test and exactness of the Frobenius pull spans place w_d in the target arc lattice. Its γ-difference is congruent to $-p^\sigma q_d$ mod $X^M N_{2,d}$. The first difference of the propagated quotient lifts at $d+j$ is $\varphi^j w_d + \tau(q_d)$. By the functorial truncation of τ the first cochain of $\mathcal{R}_d(w_d)$ is p^σ times that difference modulo $X^M N_2$. Second cochains there are uniquely determined, so

$$\sum_d \mathcal{R}_d(w_d) = p^{r_0+2\sigma}\xi.$$

The simultaneous arc lifts and functoriality of $\mathcal{R}_d$ therefore lift this class after the common additional power $p^{2\sigma+2A_0+r_0+2\kappa_{\mathrm{arc}}}$. The resulting annihilating exponent for the cokernel is

$$e = 2r_0 + 4\sigma + 2A_0 + 2\kappa_{\mathrm{arc}}.$$

The common factor is used on every summand of the same class decomposition, so e has no dependence on the number of cuts. Decomposition and lifting of the cut classes.

The two mixed sign inequalities hold for the actual clustered windows. Indeed $b_u, b_d \in [b_*, b_* + \eta_*]$ and $b_* \geq C(j+1)(g_*+1)$ imply $b_u < p b_d$ and $2b_d \leq p b_u$. For $2A_0 < b_u$, use $j \geq 1$ and the width bound $H_0 \leq h + g$ in the explicit cut estimate and the definition of σ. The terms linear in h in $2A_0$ are at most $6hp/(p-1)^3 \leq 0.066h$. The rest is $O((g+1)(j+1)+\log_p(2h))$; the assumed separation and $h \leq b_* + \eta_* + S_*$ make it smaller than the remaining margin for large enough absolute C. Finally $a_* \leq 3j+1$ gives

$$e \ll h/(p-1)^{j+1} + (g_*+1)(j+1) + \log_p h.$$

When $H_{\mathrm{loc}} = 0$, the change to $h = 1$ is absorbed by $(g_*+1)(j+1)$. The finite-group generator bound at the start now proves (71). [64, Lemma 3.8.1 and Corollary 3.8.4]

The published integral Wach comparisons, average-width lattice classification and descriptions of individual finite cohomology groups give parts of this local analysis. [65, §§III.4, IV.1][67, Theorem B][68, Theorem 3.12 and Corollary 3.16][66, Remark 2.18 and Example 7.1]

For the actual inert row B the branch is $11 \leq \ell < \ell_A$, $\mathrm{ord}_H(\ell) > 4L_w$. At each host completion E_λ the labelled word is ordered by inverse Frobenius and must satisfy the size and cyclic predecessor conditions just stated. The selected separation parameter is $j_* = \max(1, \lceil \log_{\ell-1} \ell_A \rceil)$. Here the size parameters are $h_B = b + S_c + \eta_{\mathrm{sz}} + t + 2$ and $G_P = w_P + S_c + \eta_{\mathrm{sz}} + t + 2$, where S_c is the cloud width, η_{sz} is the weight displacement, t is the twist cutoff and w_P bounds the small step of the word. With $g_* \leq G_P$ and either $H_{\mathrm{loc}} \leq h_B$ or $H_{\mathrm{loc}} \leq g_*$, the row allowance after comparison is at most

$$(3C+1)(n_\lambda+1) r_\ell \left(\frac{h_B}{(\ell-1)^{j_*+1}} + (G_P+1)(j_*+1) + L \right), \qquad n_\lambda = [E_\lambda : \mathbf{Q}_\ell].$$

Indeed put $q_* = (\ell-1)^{j_*+1}$. For positive H_{loc}, split its logarithm into those of q_* and H_{loc}/q_*. Since $\log_\ell y \le y$ for $y > 0$ and $\ell \ge 11$, their sum is at most $j_* + 1 + H_{\mathrm{loc}}/q_*$. The same bound holds with the zero-height convention. Thus the bracket is bounded by $2h_B/q_* + (G_P+2)(j_*+1)$ in the first case, and by $2G_P + (G_P+2)(j_*+1)$ in the second. Both are at most three times the bracket in the row allowance. The rational allowance is chosen between $3C$ and $3C+1$ times the common quantity without that coefficient. This allowance is used on the row lattices entering the local-to-global comparison. The Hecke congruence comparison uses the relation module above and the cutoff for its non-Selmer part. Application to the actual word, Row B bound, and Inert branch selection.

The all-thick alternating word. Use the same unramified crystalline setup with $\ell \ge 11$. Here n is the even number of Frobenius labels, r the target rank, and m bounds the number of target jumps in the local central interval $[-\eta, \eta]$ at each label. The labelling may start on either sign. The two lattices have the prescribed jumps in the central interval and in the deep signed interval of absolute endpoints b_i, $b_i + S$. Suppose the signs alternate throughout the cyclic word and that all the deep windows have lower absolute endpoint $b_i \ge h/2$, with $h = \max_i(b_i + S)$. Put $B_i = b_i + S$, $H = h + \eta$, and

$$a = \left\lceil \frac{(2\ell-1)H}{(\ell-1)^2} + 4 + \log_\ell(2H) \right\rceil .$$

The prescribed size condition is $b_i \ge 2a_{\mathbf{R}} + \eta + 1$, where $a_{\mathbf{R}}$ is the expression inside the ceiling. Since b_i is integral, $b_i \ge 2a + \eta$. Use $Q = X + \ell$, $Q_1 = \varphi Q$, $L_{u,i} = N_{u,i}/XN_{u,i}$, and $P_{u,i} = \langle \varphi N_{u,i-1} \rangle$ for the source and target Wach modules. The cut estimate gives a Q^a-annihilator of their comparison and of the Wach comparison from a crystalline extension of the target to its trivial quotient. Indeed X^{d_0} and ℓ^{e_0} annihilate the comparison, with $d_0 = \lceil H/(\ell-1) \rceil$ and $e_0 = \lfloor H\ell/(\ell-1)^2 \rfloor + \lceil 1 + \log_\ell(H+1) \rceil$. Their sum is at most a; the binomial expansion gives Q^a.

Represent a free module of input classes by one crystalline extension with trivial quotient. At a down label take a tame-invariant étale section e, and Wach lifts of $\ell^N e$ and $Q^a e$ modulo the target. Write their differences from these two multiples as μ, ν. Then $\widetilde{\mu} = \mu - \ell^N Q^{-a}\nu$ belongs to $Q^{-a}N_2$ at that label. The operator $\delta = (\gamma - 1)/X$ acts on $Q^{-a}N_2/N_2$ with graded values $d_k = (\chi(\gamma)^k - 1)/(-\ell)$ in the degree k of the first-divisor germ. The degrees are in $[-a-\eta, \eta-1]$ and $[b_i - a, b_i + S - 1]$. The zero eigenspace has dimension at most m. Set $\Lambda = 1 + \log_\ell(2H)$. In the central interval the nonzero valuations of the d_k sum to at most its length divided by $\ell - 1$, by the factorial bound. For the displaced interval count multiples of each power of ℓ. The error is at most one per power through $\log_\ell(2H)$. The product over the nonzero d_k has constant coefficient of valuation at most $(2a+S+2\eta)/(\ell-1)+\Lambda-1$. Normalize it as $\ell^{s_1} + \delta U(\delta)$. Since $\delta\widetilde{\mu} \equiv -\ell^N(\gamma-1)(e + Q^{-a}\nu)/X \bmod N_2$, applying this polynomial and killing its zero eigenspace image replaces $\ell^{s_1} e$ by an integral étale section E with $\ell^N E$ Wach and $(\gamma-1)E \in XN_2$. The image condition has at most m generators. The latter containment cancels ℓ^N by saturation in an étale basis.

Set $M = \lfloor h/(\ell-1) \rfloor + 1$. At the next label, which is up, the relative negative height is at most η. Use the horizontal correction of $\ell^\sigma \varphi E$ modulo $X^M Q^\eta$, with $\sigma = \lceil (\eta+M)\ell/(\ell-1)^2 + 1 \rceil$. For the correction R_c on the up label immediately after a down label c, one has

$$\ell^N R_c \in P_{\mathrm{ext},c+1}, \qquad (\gamma-1)R_c \in X^M Q^\eta P_{2,c+1}.$$

The first containment saturates in the pulled étale basis. The second forces the regular first-divisor germ to have zero coefficients in degrees $1, \dots, \eta-1$. Membership in the extension Wach lattice therefore imposes only Fil^0 on the constant, of codimension at most m by strictness. At a down label c, the vector R_{c-2} is placed at the preceding up label $c-1$. The difference $w_c = \varphi R_{c-2} - \ell^\sigma E_c$ lies in $Q^{-\eta}P_{2,c}$ and has gamma difference in $XN_{2,c}$. The alternating section with entries $\ell^\sigma E_c, R_c$ has first cochain w_c, τ_c, where τ_c is the up correction. After multiplication by ℓ^σ, subtract the coboundary of the vector with down entries $-\ell^\sigma w_c$ and zero up entries. The resulting first cochain agrees modulo $X^M N_2$ with the shifted-cut cochains of

the w_c, and both second cochains belong to $X^M N_2$. At a down label the Frobenius error belongs there because $\varphi(X)^M P_{2,c} \subset X^M Q^{(\ell-1)M} P_{2,c} \subset X^M N_{2,c}$. Thus the shifted-cut cochains represent $\ell^{s_1+2\sigma}$ times the input. These first two steps impose n conditions with at most m generators.

For the remaining comparisons, one and two Frobenius steps transport a horizontal frame or its inverse modulo XQ^D at costs $s(D) = \lceil \ell D/(\ell-1)^2 \rceil$ and $v(D) = \lceil D/(\ell-1)^2 + 1 \rceil$, respectively. For degree k with coefficient denominator $J(k) \le \ell k/(\ell-1)$, the first uses $\varphi X \in XQ^{\ell-1}R^\times$. The second uses $\varphi^2 X \in X(QQ_1)^{\ell-1}R^\times$ and $Q_1 \in (\ell, Q^\ell)$. When $J(k) > v(D)$, $(\ell-1)k - J(k) + v(D) \ge \lceil D/\ell \rceil$, which is enough by the binomial expansion.

At a down label c, let d_0 be the degree-zero coefficient of w_c. Its other coefficients below degree b_c are in $U = \mathrm{Fil}^{-\eta} D_{2,c}$, since gamma acts by $\chi(\gamma)^k$ in degree k. Multiply w_c by $Q^\eta Q_1^{B_{c-1}}$ to enter the two-step pull span. Apply the inverse horizontal frame modulo $XQ^{b_c+\eta}$, project to the quotient by U, and divide by Q^η by distinguished division. The remaining scalar is $Q_1^{B_{c-1}} \bmod Q^{b_c}$, of content at most $B_{c-1} - b_c/\ell + 1 + \log_\ell(2H)$. Indeed expand $Q_1 = \ell\varepsilon + uQ^\ell$ with unit factors and take the term with index $\lfloor (b_c-1)/\ell \rfloor$; among terms of minimum coefficient valuation the lowest degree survives. A primitive coefficient functional gives

$$\ell^{e_0(c)} d_0 \in \varphi^2 L_{2,c-2} + U, \qquad e_0(c) \le v(b_c+\eta) + B_{c-1} - b_c/\ell + 1 + \log_\ell(2H).$$

At an up label i, take $v_i \in \varphi^2 L_{2,i-2}$ in Fil^{b_i}. Approximate its horizontal germ to order $2H$, with multiplier $\ell^{v(2H)}$, and multiply by $Q_1^{B_{i-1}} Q^{-b_i}$. The result x belongs to $N_{2,i}$ by the first-divisor tests. At zero its value is $\ell^{v(2H)+B_{i-1}-b_i} v_i$. Modulo $Q^a N_{2,i}$ the germ is on the single horizontal line with degrees $[-b_i, -b_i + a - 1]$, giving an annihilator for δ. Lift $Q^a x$ to $N_{1,i}$ and divide this lift by Q^a. The same class modulo $fN_{1,i}$ now has a second δ-annihilator from $Q^{-a}N_{1,i}/N_{1,i}$. Its degrees are in $[-b_i - S - a, -b_i - 1] \cup [-a-\eta, \eta - 1]$. These sets are disjoint from the first set. Lagrange interpolation and $v_\ell(d_k - d_j) = v_\ell(k-j)$ give an integral Bezout combination equal to ℓ^{E_0} with

$$E_0 \le (3a + S + 2\eta)/(\ell-1) + \Lambda - 1.$$

Indeed the interval containing the omitted diagonal factor uses a factorial bound on both sides of that factor. The other interval is displaced, with at most one extra multiple per power through $\log_\ell(2H)$. Thus $\ell^{E_0+v(2H)+B_{i-1}-b_i} v_i$ lifts in $L_{1,i}$.

For a down label let $J_i = \ell^{B_i} \varphi L_{2,i-1} \subset L_{2,i}$. Adapt an integral basis of $\varphi L_{2,i-1}$ to the saturated filtration intersections. For each central basis vector of jump j use the horizontal approximation modulo $XQ^{B_i+\eta}$ and multiply by Q^{-j}. This gives an integral Wach vector with constant $\ell^{s(B_i+\eta)-j}$ times the basis vector. Use ℓ^{B_i} times the other basis vectors. The relative determinant has Q-order $-\sum j$ and is a unit off Q, so these columns have colength at most $rS + ms(B_i+\eta) \le rS + 2mc_D$, where $c_D = (H+1)/(\ell-1) + 1$. Forcing a constant into J_i therefore costs one quotient of this length and one quotient with at most m generators. The same holds after projecting the two lattices through a rational quotient.

Apply this comparison at $c-2$ to the constant congruence for w_c modulo U. It gives a representative $\ell^{B_{c-2}} \varphi v_{c-1}$ with $v_{c-1} \in \varphi^2 L_{2,c-3}$. Impose the up filtration condition on v_{c-1}, apply the up lift, and transport one step modulo XQ^{B_c} at cost $s(B_c)$. The residue comparison exponent is

$$e_0(c) + E_0 + v(2H) + B_{c-2} - b_{c-1} + s(B_c) - B_{c-2}.$$

The term $B_{c-1} - b_{c-1} = S$ remains from the large parts after using the bound for $e_0(c)$. Force the remaining constant in U to lie in Fil^0. All other graded coefficients already pass Wach membership. This step uses one length quotient and three quotients with at most m generators.

Finally let $z \in N_{2,c}$ be the resulting integral discrepancy. Use the down comparison to write $z(0) = \ell^{B_c} \varphi z_1$. Then $\ell^{B_{c-1}} \varphi^{-1} z_1 \in L_{2,c-2}$ by the up height bound. Use the down comparison there to write this last vector as $\ell^{B_{c-2}} \varphi z_3$. Impose $\varphi^2 z_3 \in \mathrm{Fil}^{b_{c-1}}$. The up lift now lifts $\ell^{E_0+v(2H)+S} z_1$. Multiply its lifted Wach representative by Q^{B_c} after one Frobenius step to match the integral discrepancy modulo $XN_{2,c}$. This uses

two length quotients and three quotients with at most m generators. The shifted-cut map kills $XN_{2,c}$ and sends all source residues to integral finite classes.

The factor of H/ℓ^2 in the total exponent is bounded by

$$8\frac{(2\ell-1)\ell^2}{(\ell-1)^3}+2\frac{\ell^3}{(\ell-1)^3}+5\frac{\ell^2}{(\ell-1)^2}+\frac{\ell(2\ell-1)}{(\ell-1)^2}\le 31.35.$$

The four contributions are the degree-root polynomials, the first two transports, the two-step frame estimates, and the difference $s(B_c)-b_c/\ell$. The coefficients of S and $\eta/(\ell-1)$ are at most $2+3/(\ell-1)+\ell/(\ell-1)^2<4$ and $6+2\ell/(\ell-1)<10$. For the logarithmic allowance use $a\le H(2\ell-1)/(\ell-1)^2+4+\Lambda$. The cost of the three interval estimates is $3(\Lambda-1)$ in addition to their length terms. The down-content term adds Λ. The two σ's, three two-step ceilings and one one-step ceiling contribute at most $11+2\ell/(\ell-1)^2$ after extracting their terms in H,S,η. Together with the constant part of $8a/(\ell-1)$, this leaves at most

$$\left(4+\frac{8}{\ell-1}\right)\Lambda+8+\frac{2\ell}{(\ell-1)^2}+\frac{32}{\ell-1}\le 16.22\Lambda.$$

Thus $e_C\le 32H/\ell^2+4S+10\eta/(\ell-1)+25\Lambda$. After this exponent there are at most $4n$ generator quotients and $3n/2$ length quotients. The former have exponent bounded by the ordinary unramified exponent $e'=e_{\mathrm{PC}}$ on the cokernel. For this exponent use the actual source jumps $\mathcal{J}$ and define M_0 to be the largest negative endpoint and H' the maximum total width at a label. The correction at the first divisor uses the nonzero roots indexed by the negatives of

$$J=\{j+a\neq 0\,:\,j\in\mathcal{J},\ 1\le a\le M_0\}\cup\{1,\dots,M_0\}.$$

The cut comparison charges $e_{\mathrm{WA}}=\lfloor H'\ell/(\ell-1)^2+3+\log_\ell(H'+1)\rfloor$ twice. The product of the nonzero root factors costs $\sum_{j\in J}v_\ell(j)$, and the remaining gamma correction at zero costs $A+v_\ell(A!)$ for $A=\lfloor M_0/(\ell-1)\rfloor+1$. Thus the ordinary correction kills the cokernel by $\ell^{e_{\mathrm{PC}}}$ with e_{PC} the sum of these charges. The quotient by ℓ^{e_C} has at most $(n+1)r$ generators by the local finite cohomology generator bound. The normalized row length is therefore at most

$$(n+1)re_C+4nme'+\tfrac{3}{2}n(rS+2mc_D).$$

When $m=0$, the three length quotients are replaced by multipliers ℓ^S, since $Q^SN_{2,i}\subset Q^{B_i}P_{2,i}$ at a down label. This gives a lift at e_C+3S. The full-window input and normalized row cost. The strict up-label congruence and The jump-root correction.

In the chosen-family comparison the unramified degree is $n=n_\lambda$, the target rank is r, and the count of central jumps is the specified m. Put $M_\ell=1+h_B/\ell+\log_\ell h_B$. The local width used by the ordinary bound for e' is the maximum of the sums of the positive and negative endpoints at each label. It is at most H, since each labelled total width is $b_i+S+\eta$. The actual windows satisfy $H\le h_B+\eta_3$, $S\le S_c$, $\eta\le\eta_3\le h_B$ and $h_B\le 2b$. The ordinary unramified numerical bound is $e'\le 23(2+H/\ell+\log_\ell(H+1))\le 69M_\ell$. For the latter inequality use $H\le 2h_B$, $H+1\le 3h_B$, and $\ell\ge 11$. The same size bounds give $c_D\le 4M_\ell$. If flat_ℓ denotes the family allowance at least $(w_P+\eta_3)/\ell+S_c+\log b$, then

$$32H/\ell^2\le 32h_B/\ell^2+3\eta_3/\ell,\qquad 10\eta/(\ell-1)\le 11\eta_3/\ell,\qquad 25\Lambda\le 125\log b.$$

The last estimate follows from $2H\le 8b$, $\log b\ge 1$ and $\log\ell\ge 1$. Hence $e_C\le 32h_B/\ell^2+225\,\mathrm{flat}_\ell$. Substitution in the normalized row length uses $n_\lambda rS\le(n_\lambda+1)r\,\mathrm{flat}_\ell$. It fits the prescribed coefficients $368, 2082, 22365$ of $(n_\lambda+1)rh_B/\ell^2$, $(n_\lambda+1)r\,\mathrm{flat}_\ell$, and $n_\lambda mM_\ell$. The three row allowances, The ordinary exponent bound and Their actual labelled application.

Alternating words containing a small label. Use $p = \ell$ and the same labelled Wach convention as above. In the local calculation b is the large reference endpoint. Take sizes $b_i > \eta$ in $[b, b+\eta]$ at large labels and at most w at small labels. The common window at an up label is $[-\eta, \eta] \cup [b_i, b_i + S]$ and at a down label it is $[-\eta, \eta] \cup [-b_i - S, -b_i]$. Suppose that the signs alternate, there is at least one small label, and there are no three consecutive large labels. The target has at most m central jumps per label. Here $20\max(w, S, \eta) \le b$. Put $h = b + \eta + S$ if some label is large and $h = \max(1, w + S)$ otherwise, and set

$$G = w + S + \eta, \qquad \lambda_0 = 1 + \log_p(2h), \qquad u = h/p^2 + G/p + \lambda_0.$$

If all labels are small, the ordinary unramified exponent is at most $20h/p + 3 \le 23u$. If $u > b/160$ and there is a large label, $h \le 1.1b$ gives $20h/p + 3 \le 323u$. In either case use that exponent on the whole target.

In the other range, take each cut immediately before a small label. If an arc ending at c has a labels, write $s = 0, \dots, a-1$ backwards from c. Its small label is $a - 1$, and $1 \le a \le 3$. Use $Q = P_0 = X + p$, $P_s = \varphi^s Q$, and

$$M = \lfloor 2h/(p-1) \rfloor + 1, \qquad J = X^M P_0^{3G} R,$$
$$\Lambda = \prod_{s<a} P_s^{H_s}, \qquad H_{a-1} = G, \quad H_s = h \ (s < a-1).$$

After a steps the transported truncation ideal is $\varphi^{a-1}(J)$. It has order $(p-1)M > 2h$ at the large levels, order at least $3G$ at the small level, and order at least M at zero. These are the simultaneous error divisibilities used in the horizontal comparisons after dividing by Λ. The cost in either direction is $\sigma = \lceil (3G+M)p/(p-1)^2 + 1 \rceil \le 4u$. The uniform cut exponent is $r_0 \le 3u$: in the recursion the small succeeding width is at most G, and the full labelled width is at most $H = h + \eta \le 1.05h$. Thus r_0 is at most

$$\frac{p}{(p-1)^2}\left(G + c_* + 3 + \frac{H}{p-1}\right) + c_*, \qquad c_* = \lceil 1 + \log_p(H+1) \rceil \le 2\lambda_0.$$

This also applies to the map from a crystalline extension of the target to the unit. The shifted-cut classes therefore reassemble $p^{r_0+2\sigma}$ times the input as in the earlier cochain comparison.

At an arc ending at c let $L_2 = N_{2,c-a}/XN_{2,c-a}$. Use the constant space transported from this cut. Intersect the saturated deep sublattices at the up levels, and saturate the sum of the central sublattices at the down levels. The former has codimension at most the number of up levels times m. The latter has rank at most the number of down levels times m. Take an integral projection onto the latter lattice and intersect its kernel with the former. The resulting saturated submodule L_2^{good} has codimension at most am. Choose an integral complement containing the down central sublattices. This is possible by taking a complement first within the projection kernel. The projection to the good space preserves each filtered coefficient subspace outside the deep transition degrees, as does the complementary projection. At an up level the good space is in the deep filtration. At a down level it kills the central subspace, and every filtered subspace below the deep transition is central.

For the transition degrees k use the intervals $[-b_s - S, -b_s - 1]$ at an up level and $[b_s, b_s + S - 1]$ at a down level. Apply to the horizontal coefficients the polynomial

$$P(Z) = \left(\sum_{j=0}^{p^a-1} Z^j\right) \prod_k \left(Z^{p^a} - \chi(\gamma)^{kp^a}\right).$$

It kills the transition coefficients and averages the degree-zero coefficient on the finite coefficient orbits at all three possible divisors. In fact γ^{p^a} fixes the coefficient algebras and dilates the local parameter by $\chi(\gamma)^{p^a}$. Interval counting gives $v_p(P(1)) \le 13S + 3\lambda_0$, since the at most three intervals have length S and $v_p(1 - \chi(\gamma)^{kp^a}) = a + 1 + v_p(k)$. Divide $P(1)$ by its unit when counting this loss.

Take a horizontal partner of the original invariant residue, apply $P(\gamma)$, and use the two integral projections. Their partner residues sum to $p^{2\sigma}P(1)$ times the original residue modulo $XN_{2,c}$. In the good

component each scalar coordinate has its poles only at up levels, with order at most b_s. At a down level it equals a rational constant modulo $P_s^{b_s+S}$. For the other component, the rational invariant residue space injects into the sum of its constant space at zero and the degree-zero local quotients at the a divisors. Indeed in a nonzero degree k a power of gamma acts through the nontrivial scalar $\chi(\gamma)^{kp^a}$. Divisibility at the rational divisors together with vanishing at zero gives the kernel $XN_{2,c}$ of these tests. The integral residue module is finite and torsion-free over O, by distinguished division and the same membership tests. The residual component thus costs at most $(a+1)am$ generators. Summing over arcs costs at most $4nm$ generators, since their lengths sum to n.

For a one-letter arc an up input passes the target Wach tests and lifts at the cut. A down input splits into its integral constant and a part divisible by the down modulus, with each part lifting at the same exponent. For a two-letter arc let u_0, d_0 be the up and down levels, and write $A = P_{u_0}^{b_{u_0}}$, $B = P_{d_0}^{b_{d_0}+S}$. The input is $f \in A^{-1}R$ with $f \equiv z \bmod B$. If $\deg A < \deg B$ in the reduction, the remainder is primitive and z is integral. Split off the constant and lift the Wach remainder. If

$$p^{d_0}(b_{d_0} + S + \eta) < p^{u_0}(b_{u_0} - \eta),$$

use $(B/A)v$ as a target Wach vector. Lift its partner at the endpoint and use the inverse source frame. Divide the numerator by $P_{u_0}^{H_{u_0}-b_{u_0}-S} P_{d_0}^{H_{d_0}-\eta}$. In the remainder modulo $P_{u_0}^{b_{u_0}+S-\eta}$ the source coefficients lie in the deep up filtration. Their image is $p^{r_0+2\sigma} P_{u_0}^{S} P_{d_0}^{b_{d_0}+S+\eta} v$. The degree condition makes the scalar primitive, so a coefficient functional gives a filtered source lift of $p^{r_0+2\sigma}v$. Multiply by f and use the source horizontal partner. Its gamma difference has the allowed up pole and vanishes to the required order at the down divisor, and the image also matches at zero.

If neither degree comparison applies, let s be the small size. For large up and small down one has $p(s+S) \le b_{u_0} \le p(s+S+\eta)+\eta$. For small up and large down one has $b_{d_0} + S \le ps \le b_{d_0} + S + (p+1)\eta$. In both ranges the adjacent pair is in the band

$$|ps - b| \le 320\{p\eta + pS + b/p + s + p\log_p(2h)\}.$$

For the first range multiply $P_0^{b_{u_0}}$ by $P_0^{p\lceil b_{u_0}/p\rceil - b_{u_0}} (u_0')^{\lceil b_{u_0}/p\rceil}$, where $P_1 = p\varepsilon_0 + u_0' P_0^p$. Expand $(P_1 - p\varepsilon_0)^{\lceil b_{u_0}/p\rceil}$ in the remainder modulo $B = P_1^{s+S}$. For the second range expand P_1^s directly modulo $B = P_0^{b_{d_0}+S}$. The smallest surviving degree of minimum coefficient valuation gives respectively the bounds

$$\lceil b_{u_0}/p\rceil - (s+S-1) + \log_p(2h), \qquad s - \lfloor (b_{d_0}+S-1)/p\rfloor + \log_p(2h)$$

on the content of A mod B. Each is at most $\eta + 2u$. After rounding, make z integral and lift the constant and the Wach remainder. The two-letter loss relative to the common horizontal partner is at most $11u + \eta \mathbf{1}_{\text{band}}$ in all ranges.

The arc with two up levels. Number the three divisors backwards from the endpoint c. Let their sizes be x, y, s, with signs up, down, up, with $x, y \in [b, b+\eta]$ and $\eta < s \le w$. Put $B = y + S$, $h = b + \eta + S$, $H = h + \eta$, $\lambda = 1 + \log_p(2h)$, and $u_s = h/p^2 + (s + S + \eta)/p + \lambda$. Take $u = u_w \le b/160$, the other range being covered by the ordinary unramified estimate. Let L_a be the constant lattice at the initial cut $c-3$ for the source or target and use the common transported horizontal coordinates. The comparison map on these constants is $f_0 : L_1 \to L_2$. Take $v \in L_2$ in the deep filtration at both up levels. A scalar input is

$$f \in (P_0^x P_2^s)^{-1} R, \qquad f \equiv z \bmod P_1^B \quad (z \in F),$$

where the congruence is in the rational completion. Here $Q = X + p = P_0$ and $P_j = \varphi^j Q$. The horizontal partner of a coefficient series agrees through the local reductions and at $X = 0$ after the common factor p^σ. The partner is functorial modulo $XN_{a,c}$. A series which belongs to the local filtered lattices has a

Wach partner. An invariant residue requires only that the gamma difference belong to those lattices, since gamma acts trivially at zero. Such Wach vectors lift at the cut exponent $r_0 \le 3u$. Integral constants lift at the same exponent, by taking a preimage of the constant at the initial cut and applying the source frame.

First suppose $s > \alpha+\eta$, where $\alpha = \lceil \alpha_{\mu,c-2}/(p-1) \rceil$ and $\alpha_{\mu,c-2}$ is the least full-parameter Wach cokernel exponent at the first, small up label. Thus X^α kills the invariant Wach cokernel and $\alpha \le \lceil H/(p-1) \rceil$ by the common labelled twist. For $A \ge 1$ put $\rho(A) = \lceil pA/(p-1)^2 \rceil$. A one-step horizontal frame or its inverse can be made integral modulo XQ^A after multiplication by $p^{\rho(A)}$, since a term of degree k has denominator exponent at most $pk/(p-1)$ and its Frobenius transform is divisible by $X^k Q^{(p-1)k}$.

Use the target frame on v modulo $XQ^{s+\eta}$ and divide by Q^s to obtain a vector z_2 in the target Wach lattice at label $c-2$. Its leading germ lies in the allowed deep subspace, and the error has order at least η. Lift $X^\alpha z_2$ to a source Wach vector and multiply that lift by Q^{s+S}, putting it in the source Frobenius span. Apply the inverse source frame modulo XQ^{s+S}. In the remainder module modulo $Q^{s+S-\eta}$ the integral source coefficient series is in the source deep subspace. Indeed its image in the quotient by this subspace has order at least $s+S-\eta$ by the central weight bound. Its image under f_0 modulo the same power is

$$p^{\rho(s+\eta)+\rho(s+S)} X^\alpha Q^S v.$$

These congruences are integral by distinguished division on the saturated quotient of L_1 and on L_2. The scalar $X^\alpha Q^S$ mod $Q^{s+S-\eta}$ is primitive since it reduces to $X^{\alpha+S}$, with $\alpha+S < s+S-\eta$. An integral functional on the remainder module taking it to one gives $v_1 \in L_1$ in the small up filtration and $f_0(v_1) = p^k v$, where $k = \rho(s+\eta)+\rho(s+S) \le 3u_s$.

Write δ_D for the normalized content of $D = P_0^x P_2^s$ modulo P_1^B, and $d = \lfloor (s - B/p - \eta/p^2)_+ \rfloor$. The two scalar estimates are

$$\delta_D \le (s-B/p)_+ + 3u_s, \qquad P_2^s \in (p^d, P_1^B P_0^\eta).$$

For the first, write $P_{j+1} = p\varepsilon_j + u_j P_j^p$ with unit factors. Put $k_0 = \lceil x/p \rceil$. If $D \in (p^a, P_1^B)$ with $1 \le a \le B-k_0$, expand $P_1^{k_0+a-1}$ in $P_1 - p\varepsilon_0$ and $p\varepsilon_0$. Multiply D by $P_0^{pk_0-x} u_0^{k_0}$ and cancel the distinguished power $P_1^{k_0+a-1}$ modulo p^a to obtain $P_2^s \in (p^a, P_1^{B-k_0-a+1})$. In this latter quotient expand P_2^s and discard the terms with $pj \ge D' = B-k_0-a+1$. Divide the remaining sum by the least coefficient power p^{v_*}, where v_* is the minimum of $s-j+v_p\binom{s}{j}$. Reduce modulo the coefficient uniformizer. The smallest index j_* attaining the minimum contributes a nonzero leading coefficient at degree $p^2 j_* < pD'$, since $P_1 \bmod p = X^p$ and the two power series factors are units. Every other surviving term begins in a higher degree. Thus the remainder has content v_*. Using the term with index $\min(s, \lfloor (D'-1)/p \rfloor)$ gives $a \le \max(0, s-(B-k_0-a+1)/p+\lambda)$. If the index is s, the remainder is primitive and $a \ge 1$ is already impossible. Choosing the first integer greater than $((s-B/p)_+ + k_0/p+\lambda)/(1-1/p)$ contradicts this. It is at most $(s-B/p)_+ + 3u_s$ and satisfies $k_0 + a \le B$ because $k_0 \le b/8$, $s \le b/20$, and $u \le b/160$.

For the second estimate, expand P_2^s first in P_1^p. A term of index j with $pj < B$ is already divisible by p^d. For a term with $pj \ge B$, factor out P_1^B and expand P_1^{pj-B} in P_0^p. For $\eta \ge 1$ put $q = \lfloor (\eta-1)/p \rfloor$. The terms below P_0^η have total exponent at least $s-j+\max(0, pj-B-q) \ge s-(B+q)/p \ge s-B/p-\eta/p^2$. For $\eta = 0$, all terms with $pj \ge B$ vanish modulo P_1^B. This proves the second estimate in the integral ring.

The congruence for the input gives $v_p(z) \ge -\delta_D$. For $s > \alpha+\eta$ set $a = \lceil (\delta_D - d)_+ \rceil \le 6u_s$. Write $P_2^s = p^d V + P_1^B P_0^\eta W$ over R. Then $f_* = (p^{a+d} z)V/P_2^s$ has denominator only P_2^s and is congruent to $p^a z$ modulo both P_1^B and P_0^η. The series $(p^a f - f_*)v$ belongs to every target filtered lattice and lifts as a Wach vector after the usual cut exponent. The source horizontal partner of $f_* v_1$ is an invariant residue. At the small up level use the chosen deep source vector, at the thick up level use the order η of the gamma difference, and at the down level use its order B. Its image is p^k times the partner of $f_* v$. Both comparisons use exact evaluation at $X = 0$, so padding the two exponents and adding gives the required residue with additional loss $a + \max(k, r_0) \le 9u$.

If $s \le \alpha + \eta$, multiply by $p^{\lceil \delta_D \rceil}$ and split off the now integral down constant. The other summand passes the target Wach tests and both lift at the cut exponent. The only additional term is $(s - B/p)_+$. It is at most $\eta + 3u_s$ by the bound on α. When it is positive and $\eta > 0$, the same inequality before rounding gives

$$|ps - b| \le 3\{p\eta + pS + b/p + s + p\log_p(2h)\}.$$

Indeed $ps > b$ and $ps - b \le p\eta + b/(p-1) + p(2\eta + S)/(p-1) + p$, with $p \le p\eta$. This is the band condition at the adjacent small and large levels. For $\eta = 0$ the band term vanishes. The additional cost for every up, down, up arc is therefore bounded by $10u + \eta \mathbf{1}_{\mathrm{band}}$.

Three letters in the order down, up, down. Write the three levels backwards from the end c of the arc, so that level 2 is the first, small letter. Let $b_0, B \in [b, b+\eta]$ be the sizes at levels 0, 1, and let $\eta < s \le w$ be the size at level 2. Put

$$d = b_0 + S, \qquad y = s + S, \qquad h = b + \eta + S, \qquad G = w + S + \eta,$$
$$u = h/p^2 + G/p + 1 + \log_p(2h).$$

Here $p \ge 11$, $20\max(w, S, \eta) \le b$, and $u \le b/160$. The two down windows have deep jumps in $[-d, -b_0]$ and $[-y, -s]$, respectively. The up window has deep jumps in $[B, B+S]$. Each window also allows the central jumps in $[-\eta, \eta]$.

Let O be the coefficient discrete valuation ring, with fraction field F and valuation $v_p(p) = 1$. In the labelled tame invariant ring $R = O[[X]]$ put $P_0 = X + p$ and $P_j = \varphi^j(P_0)$. The P_j are pairwise coprime over $R[1/p]$, are associated to distinguished polynomials of degrees p^j, and satisfy $P_j(0) = p$. We use

$$\varphi(X) \in XP_0^{p-1}R^\times, \qquad P_1 = p\varepsilon_1 + u_1P_0^p, \qquad P_2 = p\varepsilon_2 + u_2P_1^p$$

with units ε_i, u_i. The first identity follows from the cyclotomic formula after taking tame invariants. It also gives the second by writing $X = P_0 - p$ in $P_1 = p + \varphi(X)$. The third is its Frobenius transform.

Use $N_i(j)$ for the source and target Wach lattices at label j. Write $\mathcal{F}$ for the lattice map and $F_0 : L_1 \to L_2$ for the map on the constant lattices $L_i = N_i(c-3)/XN_i(c-3)$. Let $v \in L_2$ belong to the transported Fil^B at level 1. The scalar input is

$$\begin{aligned} f \in P_1^{-B}R, \qquad & z_0, z_2 \in F, \\ P_1^B(f - z_0) \in P_0^dR[1/p], \qquad & P_1^B(f - z_2) \in P_2^yR[1/p]. \end{aligned} \tag{1}$$

Both congruences use the same f and v. We will produce an integral invariant source residue with target data $p^e fv$ at the three levels and at $X = 0$. Its exponent satisfies

$$e \ll u + \eta \mathbf{1}_{\mathrm{band}}. \tag{2}$$

Here the band is the one for a small letter adjacent to a large one, with inequality $|ps - b| \le C_b(p\eta + pS + b/p + s + p\log_p(2h))$ for $C_b \ge 2$. The construction will use the pair at levels 1, 2.

We recall the two comparisons used below. Put

$$M = \lfloor 2h/(p-1) \rfloor + 1, \qquad J = X^M P_0^{3G} R,$$
$$H_0 = H_1 = h, \qquad H_2 = G, \qquad \Lambda = \prod_{j=0}^{2} P_j^{H_j}.$$

The residues at c are taken in $W_i/XN_i(c)$, where $W_i = \Lambda^{-1}\varphi^{3*}N_i(c-3)$ and the superscript $*$ means the R-span of the indicated Frobenius image. After multiplication by p^σ, with $\sigma \ll u$, the horizontal comparison in either direction identifies a series in $\Lambda^{-1}(R \otimes_O L_i)$ and an element of W_i up to the error $\varphi^2(J)$ in their numerators. The resulting errors are integral in every rational local Wach lattice and vanish

at $X = 0$. The map on horizontal coordinates is F_0. In precise terms, $\varphi^2(J)$ has order $(p-1)M > 2h$ at P_0, P_1, order at least $3G$ at P_2, and order at least M at $X = 0$. Those orders suffice after division by Λ. This comparison also holds on the first two letters, with the undivided transport error $\varphi(J)$. Indeed the formal horizontal matrix and its inverse at the starting cut have their coefficient of degree j killed modulo integral coefficients by $p^{j+v_p(j!)}$. Integral truncation modulo J after one Frobenius, and then application of further Frobenius iterates, give the asserted precision and the same loss σ.

Use the decreasing filtration convention. A local coefficient of order k belongs to the Wach lattice when its vector lies in Fil^{-k}. Thus the conditions for membership in $N_i(c)$ are exactly the three rational local Wach conditions on W_i. The condition for membership in $XN_i(c)$ adds vanishing at $X = 0$. These statements follow by testing the divisibilities of the inverse linearized Frobenius matrix after clearing its cyclotomic denominators. The divisibilities descend integrally because all the quotients by powers of distinguished polynomials are O-free. Also $(\gamma - 1)W_i \subset XW_i$, since the Wach action is trivial modulo X and $\gamma(P_j)/P_j \equiv 1 \bmod X$. We can therefore test invariance of the residue by the three local conditions on its $\gamma - 1$ difference. At either end of the arc, the cut lifting bound supplies an integral Wach preimage after multiplication by p^{r_0}, with $r_0 \ll u$. On reduction at the starting cut it also gives $p^{r_0}L_2 \subset F_0L_1$.

The range $B \geq p(y+\eta)$. Let $U_i = L_i \cap \mathrm{Fil}^B$ at level 1. These are saturated, and the map preserves the filtrations. The series $P_0^d P_2^y P_1^{-B} v$ satisfies all three target Wach conditions. Its integral horizontal partner after p^σ therefore belongs to $N_2(c)$. Lift p^{r_0} times that partner at the cut and take its source horizontal partner after another p^σ. Write $A_1 \in R \otimes_O L_1$ for the resulting numerator. The local source Wach conditions imply divisibility of A_1 by

$$P_0^{H_0-\eta} P_2^{H_2-\eta} P_1^{H_1-B-S}.$$

Indeed the maximal pole orders at the down levels are η, and at the up level the maximum is $B+S$. These divisibilities can be imposed simultaneously by Weierstrass division and the disjoint rational divisors. Denote the divided numerator by A_1'. The quotient of the level 1 source filtration by $U_1[1/p]$ has pole order at most η; consequently the class of A_1' modulo $P_1^{B+S-\eta}$ lies in $U_1 \otimes_O R/(P_1^{B+S-\eta})$. After the same division, the target congruence is

$$F_0 A_1' \equiv p^{r_0+2\sigma} P_1^S P_0^{d+\eta} P_2^{y+\eta} v \pmod{P_1^{B+S-\eta}}. \tag{3}$$

The errors in the numerator comparison have order at least H_1 at P_1, which suffices for (3). The filtration statement and the congruence hold over O by the saturation of U_1 and of the distinguished power quotient.

For a nonzero element of a free O-module, its content is the valuation of the ideal generated by its coordinates. Set $m = B - \eta$ and $k = \lceil (d+\eta)/p \rceil$. We bound the content of $P_0^{d+\eta} P_2^{y+\eta}$ in $R/(P_1^m)$. Suppose this product belongs to (p^z, P_1^m) for an integer $z \geq 1$ with $k + z \leq m$. Multiplication by $P_0^{pk-d-\eta} u_1^k$ replaces the first factor by $(P_1 - p\varepsilon_1)^k$. In the binomial expansion of $P_1^{k+z-1} = ((P_1 - p\varepsilon_1) + p\varepsilon_1)^{k+z-1}$, each term is divisible by $(P_1 - p\varepsilon_1)^k$ or by p^z. Thus $P_1^{k+z-1} P_2^{y+\eta} \in (p^z, P_1^m)$. The distinguished polynomial P_1 is a nonzerodivisor modulo p^z. This follows by reducing modulo a coefficient uniformizer and then cancelling successively through its powers. We obtain

$$P_2^{y+\eta} \in (p^z, P_1^{m-k-z+1}).$$

In the expansion of $(p\varepsilon_2 + u_2 P_1^p)^{y+\eta}$ modulo P_1^t, the surviving terms of minimal p-valuation have distinct orders modulo the uniformizer after division by that power of p. The term with smallest order survives their sum. Taking $t = m - k - z + 1$ and the term indexed by $\min(y+\eta, \lfloor (t-1)/p \rfloor)$ bounds the content above by $\max(0, y + \eta - t/p + 1) + \log_p(2h)$. Here the binomial coefficient has valuation at most $\log_p(2h)$, by counting the carries. Using $B \geq p(y+\eta)$ gives

$$(1 - 1/p)z \leq (\eta + k)/p + 1 + \log_p(2h).$$

The first integer strictly above the resulting upper bound for z is at most $4u$, since $k \le (h+\eta)/p + 1$. It satisfies $k + z \le m$, as $k < b/8$, $u \le b/160$, and $m \ge 19b/20$. For this z the product has content $< z$. The factor P_1^S in (3) does not change the content when the modulus is increased to P_1^{m+S}, by cancellation modulo each power of the coefficient uniformizer.

Accordingly an O-linear functional on $R/(P_1^{m+S})$ sends the scalar factor $P_1^S P_0^{d+\eta} P_2^{y+\eta}$ to p^z. Applying the same functional to the source in (3) gives a vector $v_1 \in U_1$ with $F_0 v_1 = p^{r_0+2\sigma+z} v$. Use the single series $f v_1$ on the source. Its difference under $\gamma - 1$ lies in the up Wach lattice since its pole order is at most B. At both down levels its difference vanishes to the orders in (1), since the respective constants $z_0 v_1, z_2 v_1$ are fixed. The integral horizontal partner of $p^\sigma f v_1$ is therefore an invariant source residue whose image agrees with $p^{r_0+3\sigma+z} f v$ at all three levels and at $X = 0$.

The range $B < p(y+\eta)$. First make z_2 integral. If $B < py$, the scalar $P_1^B \bmod P_2^y$ is primitive, since its degree modulo the coefficient uniformizer is less than that of P_2^y. Then (1) gives $z_2 \in O$. In the range $py \le B < p(y+\eta)$, put $k_2 = \lceil B/p \rceil$. Multiplying P_1^B by $P_1^{pk_2 - B} u_2^{k_2}$ gives $(P_2 - p\varepsilon_2)^{k_2}$. The term of index $y-1$ and the minimum-valuation argument used in (3) bound its content modulo P_2^y by $k_2 - y + 1 + \log_p(2h) \le \eta + 1 + \log_p(2h)$. The original scalar has at most this content as well. This second range forces $-p(S+\eta) < ps - b \le \eta$, so the band applies. In either range we can thus write

$$\begin{gathered} p^{a_2} f = z_2' + f_2, \qquad z_2' \in O, \qquad f_2 = P_2^y P_1^{-B} g, \qquad g \in R, \\ a_2 \le 2 + \log_p(2h) + \eta \mathbf{1}_{\text{band}}. \end{gathered} \tag{4}$$

Here $z_2' = p^{a_2} z_2$, and the divisibility for f_2 in (4) is integral by saturation. Write $z_0' = p^{a_2} z_0 - z_2'$ for the level 0 constant of f_2.

The content bound at level 0 will pay for making a source vector horizontal there. The scalar P_2 lies in $(p, P_0^{p^2})$, so $P_2^y \bmod P_0^d$ has content at least $y - d/p^2$. On the other hand $P_1^B \bmod P_0^d$ has content at most $B - d/p + 1 + \log_p(2h)$, by the term of index $\lfloor (d-1)/p \rfloor$ in the expansion using $P_1 = p\varepsilon_1 + u_1 P_0^p$. This index is at most B since $d \le pB$. From $z_0' P_1^B \equiv P_2^y g \bmod P_0^d$ we obtain

$$v_p(z_0') \ge y - B + d/p - d/p^2 - 1 - \log_p(2h). \tag{5}$$

We take the valuation of zero to be $+\infty$.

Set $a = c-1$, which has the first two letters behind it, with the up letter at its level 0 and the small down letter at its level 1. Put $V_i^{(2)} = \varphi^{2*} N_i(c-3)$. The horizontal truncation at a gives $\lambda \in V_2^{(2)}$ representing $p^\sigma v$ transported from the starting cut, modulo $\varphi(J) V_2^{(2)}$ after inverting p. The error comparison here and below uses the formal horizontal jets at $X = 0$ and at the cyclotomic divisors. The truncation ideal is stable under γ; horizontality and saturation give $(\gamma - 1)\lambda \in \varphi(J) V_2^{(2)}$. The two local Wach tests now give

$$n_2 = P_1^y P_0^{-B} \lambda \in N_2(a).$$

At the up divisor P_0, the approximation to v comes from Fil^B and the error has order $(p-1)M - B \ge \eta$ after multiplication. At the small down divisor P_1 the factor P_1^y gives the required order.

We next lift $p^E n_2$ in the Wach lattice at a, with $E \ll u$. Let A_a, r_a be the bounds for the exponents of X, p, respectively, killing the cokernel of $N_1(a) \to N_2(a)$, and put $K = A_a + r_a$. Then $P_0^K = (X+p)^K$ kills the same cokernel. Put $H_* = h + \eta$. The integral Wach descent gives

$$A_a \le \lceil H_*/(p-1) \rceil, \qquad r_a \le H_* p/(p-1)^2 + \lceil 1 + \log_p(H_* + 1) \rceil.$$

In particular $K \le H_*(2p-1)/(p-1)^2 + 3 + \log_p(2h)$. Since $H_* \le 1.15b$ and $u \le b/160$, these inequalities give $K < b/4 + b/40$, $2K + \eta + 1 \le B$, and $B + K + \eta < 2h$.

Use the O-linear operator $D = (\gamma - 1)/X$ on the Wach lattices and their cokernel. For $\chi = \chi_{\mathrm{cyc}}(\gamma) \in 1 + p\mathbf{Z}_p$ with $v_p(\chi - 1) = 1$, write $c_m = (\chi^{-m} - 1)/(-p) \in \mathbf{Z}_p$. At the first cyclotomic divisor, $\gamma(P_0)/P_0 \equiv$

χ mod P_0 by its action on the local parameter. The ratio is 1 mod X. Applying φ to the latter congruence gives $\gamma(P_1)/P_1 \equiv 1 \bmod XP_0$. These congruences imply, for $0 \le j < K$ and $r \in R$,

$$(D - c_{B-j})(P_0^j r n_2) \in P_0^{j+1} R n_2 + P_0^K N_2(a). \tag{6}$$

In fact set $g_a = P_1^y P_0^{-B}$ and expand the left side by semilinearity. The term involving λ has scalar $P_0^j\big((\gamma(r)(\gamma(P_0)/P_0)^{j-B}(\gamma(P_1)/P_1)^y - r)/X - c_{B-j}r\big)$ times n_2. The parenthesis is integral by reduction at $X = 0$. It is divisible by P_0, since at $P_0 = 0$ the operator divides by $X = -p$ and its constant reduces to zero. The term from $(\gamma - 1)\lambda$, divided by P_0^K, has no pole at $X = 0$, has order at least $(p-1)M - B - K \ge \eta$ at P_0, and has order at least $3G + y$ at P_1. Hence it belongs to $N_2(a)$ by the two local tests. It follows from (6) that $\prod_{m=B-K+1}^{B}(D - c_m)$ kills $[n_2]$ in the cokernel.

A second annihilating polynomial comes from the source local filtration. If $K > 0$, take the product of $D - c_m$ over the distinct $m = w' + j$, where $1 \le j \le K$ and w' ranges over the source jumps at a. This polynomial maps $P_0^{-K}N_1(a)$ into $N_1(a)$. The operator D preserves $P_0^{-K}N_1(a)$ because the action is trivial modulo X on $N_1(a)$ and $\gamma(P_0)/P_0 = 1 \bmod X$. At P_0 the quotient is filtered by Laurent degree, with possible graded degrees $-w' - j$. On that grade D acts by $c_{w'+j}$: γ dilates the local parameter by χ, fixes the coefficient algebra at the first divisor, and $1/X$ starts with $-1/p$. The product kills the quotient by taking one factor for each graded degree. Membership descends integrally by Weierstrass division. Lift $P_0^K n_2$ in $N_1(a)$, divide its preimage by P_0^K, and apply this polynomial. Since the map commutes with D, the second polynomial also kills $[n_2]$.

The first index interval $[B - K + 1, B]$ is disjoint from the second set, contained in $[B + 1, B + S + K] \cup [1 - \eta, \eta + K]$, because $2K + \eta + 1 \le B$. Also $v_p(c_m - c_{m'}) = v_p(m - m')$ for distinct indices. For each index m of the first interval, the sum of these valuations over all the other indices in both polynomials is

$$\le \frac{3K + S + 2\eta}{p - 1} + 3\log_p(2h).$$

Indeed each of the three intervals contributes at most its cardinality divided by p^t, plus one, to the count of indices congruent to m mod p^t; no nonzero difference exceeds $2h$. Take the ceiling of the displayed bound for E. Modulo the monic first polynomial, use Lagrange interpolation at the c_m to express p^E as a multiple of the second polynomial. The coefficient of each Lagrange numerator is p^E divided by the product of all the corresponding differences $c_m - c_{m'}$, so it is integral by this estimate. The resulting polynomial congruence over O follows by monic division. Thus $p^E[n_2] = 0$, with $E \ll u$ since $K \ll h/p + 1 + \log_p(2h)$. If $K = 0$ the map at a is already onto and we use $E = 0$. Fix one $n_1 \in N_1(a)$ mapping to $p^E n_2$.

Let $U_i' = \varphi^* N_i(a)$ at c. On this span the only local difference from the endpoint Wach lattice is at P_0. In particular $P_0^d U_i' \subset N_i(c)$ and $U_i' \subset W_i$. Take the horizontal series at a whose zero coefficient is $\overline{n}_1 \in N_1(a)/XN_1(a)$. Truncate it through $\lfloor (d-1)/(p-1) \rfloor$, pull it forward by φ, and make it integral by multiplication by p^{σ_1}, where $\sigma_1 = \left\lceil \frac{p}{p-1} \left\lfloor \frac{d-1}{p-1} \right\rfloor \right\rceil$. The result $Y_1 \in U_1'$ is horizontal modulo P_0^d in the rational completion at P_0, since $\varphi(X)$ vanishes to order $p-1$ there. Its $\gamma - 1$ difference belongs to XU_1' and to $P_0^d U_1'$. After division by X, which is a unit at the rational divisor P_0, the second divisibility still holds integrally. Thus $(\gamma - 1)Y_1 \in XN_1(c)$.

The single target vector $V = f_2\varphi(\lambda) = g\varphi(n_2)$ lies in U_2'. At P_0 it is congruent modulo P_0^d to $p^\sigma z_0'$ times the transported v. Reducing n_2 at $X = 0$ at the intermediate label a gives $p^{\sigma+y-B}$ times the transported v, since $P_0(0) = P_1(0) = p$. Thus $\mathcal{F}(Y_1)$ is congruent at P_0 modulo P_0^d to $p^{\sigma+\sigma_1+E+y-B}$ times that same transported v. Put

$$\delta = \left\lceil \max(0, \sigma_1 + E + y - B - v_p(z_0')) \right\rceil, \qquad C = p^\delta z_0' \, p^{B-\sigma_1-E-y} \in O,$$

using $\delta = 0$, $C = 0$ if $z_0' = 0$. By (5) and $\sigma_1 \le pd/(p-1)^2 + 1$, we have $\delta \ll E + h/p^2 + 1 + \log_p(2h) \ll u$. The difference $p^\delta V - C\mathcal{F}(Y_1)$ is in $P_0^d U_2'$ by the congruence and saturation. Hence it belongs to $N_2(c)$. Choose its Wach preimage Z_1 after multiplication by p^{r_0} at c. The one source vector

$$Z_1 + p^{r_0} C Y_1$$

is invariant modulo $XN_1(c)$ and has exact image $p^{r_0+\delta}V$. The span U_1' already meets the source conditions at levels $1, 2$, including the degree-by-degree up filtration wherever a pole occurs. At level 0 the same vector has a horizontal constant modulo the Wach lattice. The target vector V represents $p^{\sigma} f_2 v$ since $\varphi(\lambda)$ represents the transported $p^{\sigma}v$ modulo $\varphi^2(J)$. Multiplication by f_2 can reduce the order of this error only at P_1, by at most B, which leaves order $> h$. This comparison includes the value at $X = 0$.

For the term $z_2'v$ of (4), lift $p^{r_0}v$ in L_1 and use its constant horizontal representative after the same loss p^{σ}. A horizontal constant has zero $\gamma - 1$ difference, so its integral partner gives an invariant residue at the three levels. Multiply this residue by p^{δ} and add it to the source vector above. By (4) the result realizes $p^{a_2+r_0+\delta+\sigma} f v$. In either range, if a target residue was already represented by fv, the simultaneous agreement at P_0, P_1, P_2 and $X = 0$ gives congruence to its prescribed multiple modulo $XN_2(c)$. This proves (2).

With the scalar comparisons in all four arc types, each good partner lifts after an additional exponent $18u + \eta \mathbf{1}_{\mathrm{band}}$. In the down, up, down comparison the two exponents before the common horizontal loss are $r_0 + 2\sigma + z \le 15u$ and $a_2 + r_0 + \delta \le 18u + \eta \mathbf{1}_{\mathrm{band}}$. For the latter use $a_2 \le 2u + \eta \mathbf{1}_{\mathrm{band}}$, $E \le 9u$ and $\delta \le 13u$. In fact the bound for K gives $3K/(p-1) \le 8.01h/p^2 + 0.9\lambda_0$, and the other terms of the interpolation bound give $2.2G/p + 4\lambda_0$. For δ, the content inequality at the thick down divisor leaves only the difference $pd/(p-1)^2 - d/p + d/p^2 \le 3.31h/p^2$, plus $O(\lambda_0)$; the roundings and the valuation term are at most $3\lambda_0$. The band constant 320 of the row table covers all the scalar branches. Combining the exponent with the cochain reassembly, the two horizontal projections, and the degree polynomial gives

$$e_B \le r_0 + 4\sigma + 13S + 3\lambda_0 + 18u + \eta \mathbf{1}_{\mathrm{band}} \le 40u + 13S + \eta \mathbf{1}_{\mathrm{band}}.$$

After this exponent the quotient is generated by at most $4nm$ classes. The generic unramified exponent e' kills this part as well. The remainder has normalized length at most $(n+1)re_B$, by the generator bound for finite cohomology. Consequently the normalized row length is at most $(n+1)re_B + 4nme'$.

For the actual family, write $\mathrm{flat}_\ell = (w_P + \eta_3)/\ell + S_c + L + \eta_3 \mathbf{1}_{\mathrm{band}}$ and $M_\ell = 1 + h_B/\ell + \log_\ell h_B$. The realized word has $h \le h_B \le 2b$, $w \le w_P$, $S \le S_c$, $\eta \le \eta_3$, and $n = n_\lambda$. Here $b = \exp L$ in the family parameters. For $L \ge 1$ and $\ell \ge 11$ one has $\lambda_0 \le 1 + \log_\ell(4 \exp L) \le 2L$. It follows that the largest coarse exponent $323u$ is bounded by $323h_B/\ell^2 + 646\,\mathrm{flat}_\ell$. The sharper arc exponent is bounded by $40h_B/\ell^2 + 80\,\mathrm{flat}_\ell$. The generic comparison gives $e' \le 20h_B/\ell + 3 \le 20M_\ell$. This fits the same three row coefficients $368, 2082, 22365$ as in the all-thick case. The actual band uses the local size of an adjacent small label, which differs from $s \in [w''(n), w_B]$ by at most $(t+1)/2$. The local large reference endpoint lies in $[b-(t+1)/2, b-(t+1)/2+1]$. These two translations give the family inequality

$$|\ell s - b| \le C\{\ell(\eta_3 + S_c + L) + b/\ell + s\}.$$

For $\ell \ge 2C$ the size inequality $C(\eta_3 + S_c + L) \le w''(n)/8$ gives $s \le 4b/\ell$, so the band charge multiplied by r is at most $16mb/\ell$. The range $\ell < 2C$ has bounded total logarithmic weight. For $\ell \ge \ell_A$ use the tail bound for $\sum \log \ell/\ell^2$ and $r/\ell_A \le 2m/L$. For $\ell < \ell_A$ of small host order, the allowed primes are congruent to ± 1 modulo $2^{e+2-\lfloor \log_2(4L_w) \rfloor}$; the sum of $r \log \ell/\ell^2$ over those two classes is $O(mL^2)$ with the actual modulus. Together with $\sum_{\lambda \mid \ell}(n_\lambda + 1) \le 4f$, the other terms sum within $O(fmbL^2)$ up to ℓ_*. The exceptional support above ℓ_* is governed by its separate polynomial size and cardinality bounds.

The cuspidal Selmer term at the chosen twist. For each good point π of a nonbase block α let $s_{\alpha,\pi}$ be the normalized length of the relaxed finite Selmer group of $W_\pi'[1/\ell]/W_\pi'$. The comparison needs

$$s_{\alpha,\pi} \le v_\ell + \ell^{-2}, \qquad t \sum_{z \in Q_\alpha} v_\ell(\zeta_z - 1) \le \ell^{-2}, \qquad v_\ell = \min_{g \in G_E} v_\ell(\epsilon(g)^{rt} - 1).$$

Here Q_α is the chosen tame conductor support, ζ_z is the first character on its inertia generator, r is the cuspidal rank, and ϵ is the cyclotomic character. These are the cuspidal and tame inputs at the same

character. The higher fine Selmer input above controls the separate negative Tate-character contributions. The two selected-twist bounds.

AI-proposed candidate (an integral Selmer bound for tame anticyclotomic twists). Fix the stage, its residue prime ℓ, and the finite list of good signatures. Write E/F for the CM host. The normalized cuspidal lattice $T_s = \rho_s \otimes \chi_s^{-1}$ of signature s has odd rank $r \geq 100$, with assigned sign $\sigma_s \in \{1, -1\}$ and fixed coefficient DVR O_s of fraction field K_s. Assume the integral polarization $(T_s^\vee(1))^c \simeq T_s \epsilon_\ell^t$ for positive odd t, where c is complex conjugation and ϵ_ℓ is the ℓ-adic cyclotomic character. The local realization above ℓ is de Rham and potentially crystalline, crystalline away from the type primes, for both ρ_s and T_s. The labelled jumps of T_s avoid $[-1, t]$. The dimensions of the two conjugate $D_{\mathrm{dR}}/\operatorname{Fil}^0$ quotients sum to r, so their total over all labels is $r[F : \mathbf{Q}]$. Use the integral finite-cocycle comparison, dimension formula and duality for these local realizations. The rational representations $T_s[1/\ell]$ and $T_s^\vee(1)[1/\ell]$ have no invariants at the finite decomposition groups. The polarization and Hodge gap and The local realization and coefficient hypotheses.

The residual signature matrices span the full matrix algebra. The untwisted ρ_s is unramified off the prescribed type primes and ℓ. At the type primes away from ℓ its reduction is tame, and some inertia matrix minus the scalar $\bar{\chi}_s$ on that element is invertible. Impose the resulting local invariant and cohomology conditions at the fixed relaxed places for the lattice and its Tate dual. At a designated twist prime away from ℓ there is one tame inertia element z, with $\epsilon_\ell(z) = 1$, acting on each T_s by a scalar d_s with $d_s^2 - 1$ a unit. In each finite quotient the order of z is prime to ℓ and the residue characteristic at that prime. Residual coefficient stability and The common tame element.

The moving pairs at auxiliary order ℓ^M are split conjugate degree-one places unramified over their rational primes, of norm 1 mod $2\ell^M$. They avoid the fixed excluded places, the ramification of the T_s, and any additional prescribed finite set. Each pair carries an anticyclotomic coordinate to $\mathbf{Z}/\ell^M\mathbf{Z}$ surjective on inertia at its two places, unramified away from the pair, and trivial at every decomposition group over ℓ. First-list Frobenius elements have the conjugacy class of z modulo the prescribed open normal subgroup. At the second-list prime indexed by i, the first-list coordinates have values δ_{ij}. The auxiliary prescription permits arbitrarily large M. The arithmetic statement requires one bounded-size array of such pairs for all signatures, with Selmer determinant polynomials on that array computing the normalized lengths after specialization. Each has a coefficient of bounded degree and valuation at most v_ℓ, uniformly in M. The fixed coefficient valuation is taken before adjoining the specializing roots. Admissible coordinates and the two Frobenius lists and The common determinant assertion.

For $\delta_A, \delta_D > 0$, the required choice is one finite ℓ-power character ψ with $\psi^c = \psi^{-1}$, trivial at the decomposition groups above ℓ, and ramified exactly at the chosen pairs Q. For each $(a, \ell) = 1$ use the twist $(T_s \otimes_{O_s} K_s'/O_s')(\psi^{a\sigma_s})$, where $K_s' = K_s(\operatorname{im}\psi)$ and O_s' is its valuation ring. The Selmer condition above ℓ is generated by the reductions of the prescribed integral finite cocycles multiplied by elements of K_s'/O_s'. The condition is unrestricted at the fixed relaxed primes and Q, and unramified elsewhere. Let $L_{a,s}$ be the length of this whole discrete Selmer group over O_s', divided by $v_{O_s'}(\ell)$. Then

$$\sum_s (L_{a,s} - v_\ell)_+ \leq \delta_A \quad ((a, \ell) = 1), \qquad t \sum_{v \in Q} v_\ell(\psi(\tau_v) - 1) \leq \delta_D. \tag{72}$$

Here τ_v maps to a generator of the inertia image and $v_\ell(\ell) = 1$. Finite length is part of the assertion. Finite cocycles and normalized length and The character and the two bounds.

Mazur and Rubin reduce residual relative Selmer groups for abelian varieties satisfying their hypotheses (H.1)–(H.7) at $\ell \geq 3$, in cyclic extensions of prescribed order with splitting and ramification conditions. Their local conditions are relative Kummer images, and their critical and silent primes use zero- and one-dimensional Frobenius coinvariants. Their simultaneous-abelian-variety discussion specifies an additional Frobenius separation condition. [91, Prop. 9.17, Thm. 9.21 and Rem. 10.4] Smith obtains distributions for favored admissible twists of a twistable module, for a fixed finite cyclic target and under the non-self-dual or alternating hypotheses of his theorem. His admissibility condition controls ramification

under projection to the quotient of order ℓ. The bound here concerns a joint integral determinant on moving anticyclotomic prime pairs and allows smaller inertia orders on individual pairs, including at $\ell = 2$. [92, Def. 2.2, Cases 2.12–2.13 and Thm. 2.14]

The following determinant and valuation arguments explain the arithmetic input and recover (72) from it. Taking determinants of a polarized comparison modulo a power of the coefficient uniformizer shows that the depth of a congruence between the twisted signature and its conjugate Tate dual is bounded by v_ℓ in normalized valuation. The length estimate requires the cohomological argument in addition to this congruence test. At $\delta_A = \delta_D = \ell^{-2}$, each summand in (72) gives the individual bound needed at the matched good point. The joint application to the actual lattices also needs the signature matching and the toric and lattice tests on the same array. The rank-one fine Selmer bounds enter separately.

In the classical host the prescribed ranks give odd rt, and $q, d > 4rt$. At an odd $\ell \neq q, d$ the host is unramified so the mod-ℓ cyclotomic image contains -1. At q, d it contains the subgroup of squares, of order $(\ell - 1)/2 > rt$. In either case some cyclotomic value to the rt-th power differs from one modulo ℓ, giving $v_\ell = 0$. At two the cyclotomic automorphism acting as inversion on the two-power roots fixes their maximal real subfield and can be chosen to fix both odd quadratic factors by the Chinese remainder theorem. Its value to the rt-th power is -1, giving $v_2 \leq 1$. Every two-adic unit is odd, so $v_2 = 1$. Cyclotomic defect on the host.

The signatures on the chosen type cut. The fixed relaxed locus consists of the places over g_0, g_1 and the designated scalar-character prime, with the places over ℓ removed. The scalar prime is p when $\ell \neq p$ and p' otherwise. On a good Hecke point the compact unitary base change gives a sum of conjugate-dual discrete general-linear representations, with the recorded split types and infinitesimal character. This uses the discrete-series and sign hypotheses of the compact case and $[F : \mathbf{Q}] \geq 2$. The discrete constituents are Speh representations of cuspidal ones. Remove the prescribed rank-one Speh block of length t. At either type prime the full degree of the remaining label forces each constituent to account for its whole rank in the label polynomial. If there are at least two cuspidal copies, one has rank $a \leq r/2$, giving a factor of degree a at every place over that prime. The good labels include a place excluding that degree. Thus the complement is one rank-r cuspidal representation with Speh length one. Labesse, Corollary 5.3, Mœglin–Waldspurger, the discrete-spectrum classification, and The finite labels on the complement.

The two algebraic normalizations of this regular odd-rank conjugate-dual cusp form give ρ_s and T_s. Use Galois attachment with labelled Hodge data, local compatibility at the finite places, and crystallinity at unramified data above ℓ. At $\ell = g_k$, the finite-type label gives regular semisimple inertia. Its eigenspaces are lines preserved by monodromy, hence monodromy is zero and the representation is potentially crystalline. Full degree means the depth-zero support accounts for the whole local representation, with the tame Teichmüller lifts of the label roots as inertia eigenvalues. These applications require the actual matching of the split type cut with the Hecke labels and the algebraic weight. Shin, Theorem 1.2 and Corollary 1.3, Barnet-Lamb–Gee–Geraghty–Taylor, Theorem 1.1, and Inertia and monodromy on the type.

Choose a type prime g whose bad residue-prime set omits ℓ. The two bad sets are disjoint. At its split places the reduced inertia eigenvalues remain distinct, and arithmetic Frobenius acts on their lines by the permutation of the g-power cycles, up to inversion of that permutation. A proper stable subspace of dimension a therefore supplies a factor of degree a and another of degree $r - a$ in all these local labels. Since the smaller degree is excluded at one place, this proves absolute residual irreducibility. At each $g_k \neq \ell$, including the other type prime, the separate label exclusion for the fixed first character gives zero residual inertia invariants on T_s and its Tate dual. The selected scalar prime has tame character image of order 15 and unramified ρ_s. Project an inertia lift in its procyclic closure to the pro-3 factor when $\ell \neq 3$ and to the pro-5 factor otherwise. Its scalar on all signatures is of order three or five, up to inversion, so $d_s^2 - 1$ is a unit. This also gives zero inertia invariants at every place over that scalar prime. For the finite coefficient modules the residual vanishing passes through their residue filtrations. Local duality and the Euler characteristic away from ℓ give acyclicity at the fixed relaxed places. The moving characters are

unramified there and preserve the inertia exclusions. The good labels, The scalar prime, and The fixed local vanishing.

For the second normalization the Hodge indices at conjugate labels satisfy $J_{\tau c} = t - 1 - J_\tau$. The prescribed weight window removes $[-1, t]$, so each pair has r negative jumps in total. The gap excludes jump zero both on T_s and on its Tate dual, and thus excludes rational invariants above ℓ by filtered strictness. Elsewhere outside the fixed locus, the Satake parameter of T_s is the unitary parameter times $N(v)^{(1-t)/2}$, using arithmetic Frobenius. Temperedness gives magnitude less than one after the shift. This excludes the eigenvalue one and the eigenvalue $N(v)$ corresponding to an invariant in the Tate dual. The same shift and conjugate-dual Satake identity give $(T_s^\vee(1))^c \simeq T_s \epsilon_\ell^t$ at the Frobenius classes. Chebotarev and residual absolute irreducibility give the integral polarization by the trace argument below. The jump window, The shifted dual operation, and Local hypotheses for the attached signatures.

To transport a bound from an untwisted signature, compare its twisted trace with the good cuspidal trace. The degree-one Frobenius identities give equality of traces on G_E by Chebotarev. The group-ring image of the signature lattice is the full integral matrix ring by residual absolute irreducibility and Nakayama's lemma. An element killing the other representation has trace zero against every group-ring multiple, so kills the signature as well by the perfect matrix trace pairing. Thus the other integral image surjects onto the full matrix ring. It has the same rank by its matrix realization, so this is an isomorphism. The images of the matrix units construct an integral change of basis. The actual Hecke ring O_α is finite free over the signature DVR O_0. A basis of this extension shows that continuous cochains and the rational finite conditions above ℓ commute with that base change, since the ray twist is trivial there. The finite condition on the torsion quotient is generated by reductions of these integral finite cochains. The conditions away from ℓ use the same relaxed set.

The ambient coefficient ring for the specialized roots may have infinite residue field. Its torsion quotient is still a union of finite modules over the original finite ray group ring over O_0. The universal Selmer matrix therefore computes the normalized length in that ring as well. The actual two signed readings have the same kernel on $O_0[G_E]$. Their images identify a single finite character order over O_0. Each coordinate root is the character value at an inertia element with that coordinate one and the others zero, hence belongs to the order. The specialized determinant is the same element of the order on both readings. The normalized valuations agree by uniqueness of extension on its fraction field. This transports the length bound to the actual Hecke ring using the same signed character. Signature matching and The local conditions and normalized length. The common character order.

The signature estimate uses the universal ray twists on the paired auxiliary primes. To construct their order-ℓ^M coordinates, invert the places over ℓ and choose invariant-factor ideal generators of the localized class group, with principal powers (α_i). Choose the ideals prime to the prime carrying the inertia scalar. Adjoin to E the $2\ell^M$-th roots of unity and the ℓ^M-th roots of the α_i, the localized units and their conjugates. The resulting Galois extension is unramified at that scalar prime. At a prime P split in it, the quotient of the residue multiplicative group by ℓ^M-th powers annihilates both the localized units and the principal-power relations. Sending the ideal generators to zero extends the quotient to the ray group, split at ℓ. Subtract the conjugate ray character to obtain a coordinate surjective on inertia at both P and P^c. The first Frobenius prescription is compatible with the extension since the scalar inertia acts trivially on it. For the second prescription, multiply the chosen Frobenius element by first inertia elements supported on individual coordinates. They preserve the used cocycle values by the local vanishing at the first primes. Chebotarev gives degree-one pairs with the prescribed used and unused coordinates, outside the finite exclusion set.

Over their finite group algebra the Selmer functor has a square presentation. Normalize a cocycle to have value zero at the scalar inertia, by subtracting the unique coboundary with its given value there. This is possible since the ray coordinates kill that element and $d_s - 1$ is a unit. The chosen representative commutes with changes of coefficient modules. The finite local cocycles form a saturated submodule of evaluations, so their conditions pull back through coefficient injections and lift along coefficient quotient

maps. Poitou–Tate and the Euler characteristic give equality of the product of the two residual obstruction orders with the residual Selmer order. The Hodge contribution in this calculation is fr by the conjugate pairing of jumps. The number of relations in a minimal presentation is therefore at most the number of generators. One pads by zero relations to make a square presentation. After specializing at a ray character, its determinant valuation equals the normalized length by Smith form. Residual obstruction bound and Square presentation.

The determinant calculation must produce, for all signatures, a coefficient of normalized valuation at most v_ℓ in bounded degree, uniformly in the large ray order. When $v_\ell = 0$, the two residual evaluation representations are nonisomorphic by the determinant test $\bar{\epsilon}^{rt} \neq 1$. Polarization of the Frobenius evaluations, using the inertia scalar with $d^2 - 1$ a unit, makes their paired forms span the bilinear forms on the residual Selmer spaces. A perfect combination is realized by the second prime pairs. If A is their residual inertia evaluation and U the scaled Frobenius evaluation, reciprocity shows that a leading vector of $\ker(A - ZU)$ lies in the radical of the paired form. This follows from its first two coefficients, since $Ax_0 = 0$ and Ux_0 lies in $\operatorname{im} A$. In the nonsingular case this excludes every nonzero kernel vector over the Laurent-series field. Specialize the first-order ray parameters to 1 on the selected second Frobenius elements and to the prescribed values proportional to Z on paired inertia. The tame relation puts the reduction of any residual kernel vector that lifts through second order in $\ker(A - ZU)$. The linear block on the residual kernels therefore has kernel dimension bounded by the radical dimension, giving a unit determinant coefficient in bounded degree in the nonsingular case. For the alternating case the bound is one and gives a unit adjugate coefficient. In residue characteristic two the polarizing operator uses the unit $d_s - d_s^{-1}$. The positive v_ℓ argument uses the additional prime step described below. Inertia and Frobenius evaluations.

For positive v_ℓ, work modulo $I = \mathfrak{m}_0(\ell^{v_\ell})$, with $e = v_{\mathfrak{m}_0}(\ell^{v_\ell})$. If w bounds the degree of a unit adjugate coefficient, put $h = (e + 1)w$. Specialize the used coordinates to scalar multiples of z in $A = (O_0'/IO_0')[z]/(z^{2h+2})$ over a sufficiently large residue field. A nonzero lowest-degree homogeneous part of the adjugate survives by choosing the scalar residues off its zero set. The resulting entry a_* divides z^h, since its lowest unit term has degree at most w and its lower terms are divisible by the uniformizer. If all determinant coefficients of smaller degree than $2h + 2$ belong to I, the determinant vanishes over A. The chosen column and row of the adjugate are then kernel vectors, and $\ell^{v_\ell} z^{h+1} a_* \neq 0$, since multiplying a divisor of z^h by that scalar cannot give zero in A.

Poitou–Tate on A and its additive dual gives a dual Selmer class not killed by $\ell^{v_\ell} z^h$. Indeed, if all were killed by that scalar, their strict conditions at the relaxed places would let them be represented in the submodule killed by $\ell^{v_\ell} z^{h+1}$. The coefficient filtrations preserve the orthogonal finite conditions. The relaxed local invariants are unchanged since z already kills them. The duality order equality would then force the old square presentation on the additive dual to have unchanged cokernel on quotienting by $\ell^{v_\ell} z^{h+1}$. Multiplication of its left kernel row would give $\ell^{v_\ell} z^{h+1} a_* = 0$. For the surviving dual class the paired cocycle evaluations are nonzero at some Frobenius. To prove this, polarize a hypothetical vanishing using the two distinct inertia scalars. For a product matrix U of primal and dual coordinate vectors this gives the trace comparison of the two group algebra maps f, f', with $f'(g) = \epsilon_\ell(g)^t f(g)$, against U and its conjugate counterpart. Both matrix images are full by residual irreducibility and Nakayama. Let $\mathfrak{a}$ annihilate every entry of U. Testing the trace comparison on left and right group-algebra lifts of matrix units shows that if $f'(b) = 0$, then each entry of $f(b)$ is in $\mathfrak{a}$. If $f'(b)$ has entries in $\mathfrak{a}$, subtract their linear combination of f'-lifts of matrix units and apply the same argument. Thus $f' \mapsto f$ gives an endomorphism of the full matrix algebra modulo $\mathfrak{a}$. It is conjugation, since it is an isomorphism over the residue field and its lifts of matrix units give a basis over the local ring. Equality of determinants then gives $(\epsilon_\ell(g)^{rt} - 1)U = 0$ for every g. Hence the coordinate products are killed by ℓ^{v_ℓ}. Reducing the universal primal class modulo the ideal generated by its evaluations on the selected common kernel kills all these evaluations. The scalar inertia comparison makes the reduced class a coboundary, so representability puts

the adjugate entry a_*, and hence z^h, in that ideal. The dual class would then be killed by $\ell^{v_\ell} z^h$, contrary to the choice. Duality and the adjugate and Frobenius pairing. The comparison modulo the annihilator.

Chebotarev realizes that pairing at a new pair with the prescribed unused ray coordinate. In the specialization over $A[X, Y]/(X^2, Y^{2r})$, the new inertia variable is X and the Frobenius variable is Y. The tame relations are $Xu_i = Yt_i$ after rescaling the inertia evaluations t_i by units. Here u_i are the Frobenius evaluations and $1 \le i \le 2r$ includes the conjugate place. Cup-product reciprocity gives a relation among the t_i modulo (X, Y) whose coefficients pair nontrivially with the old adjugate applied to the u_i. Lift this relation in place of the corresponding old row and append the $2r$ tame relations. In the new square determinant only terms using X once and Y to degree $2r - 1$ from the tame rows can survive. Their coefficient is the nonzero Frobenius pairing. Thus the original determinant has a coefficient outside I below degree $2h + 2 + 2r$. With the new inertia variable set to zero, a unit-leading linear Frobenius factor annihilates the $2r$ added generators in the kernel of the map to the old module. Multiplying minors by its $2r$-th power therefore preserves the earlier Fitting certificates with at most $2r$ extra degree. Induction over the fixed finite signature list proves bounded degree independently of the large ray order. The determinant on the new pair and Preserving the earlier signatures.

The integral universal tests in the fixed row space. At a fixed coefficient place, stage and side, let B be the reference lattice and let A be the integral operator projecting the prescribed tests into its span. Write $\Theta_Q(a)$ for the universal theta test at the same rational datum a, in the ray group algebra before adjoining specializing roots. The first requirement is a single finite c_0 such that all coefficients of $\ell^{c_0}\Theta_Q(a)$ belong to the fixed coefficient ball. This exponent is allowed to depend on the finite list of tests.

The normalized theta sections can be read on the level cover of modulus $M' = M_0 P_d$, where $M_0 = g_0 g_1$ and P_d is prime to ℓM_0, containing the scalar-character prime and the split rational primes below Q. The relevant ray fields have bounded local conductor at ℓ, with bound given by M_0. Local reciprocity puts all the corresponding completions in a single finite extension followed by unramified extensions. The complete coefficient field, after a finite enlargement, contains their coefficients because its residue field is algebraically closed. The fixed type and scalar-character values are included in this enlargement. The finite Fourier roots used in the marking have order dividing M'. The cover in the prescribed coefficient branch.

At a moving split prime of norm q the universal coefficient of a matrix of size $r \times s$ is the sum of projected determinants over complementary planes, recorded in the ray coordinates, divided by $q^{r(s-r)}$. This is a unit denominator at ℓ. Replacing marking functions by linear combinations of vacuum translates uses additive Fourier inversion on quotients whose orders divide a fixed power of M'. Reversing a marked half uses the finite Fourier transform on the paired position and momentum quotients. Its scalar at the moving primes is the product of powers of q and the normalized fixed splitting-character values at those primes. These values are units at ℓ: taking a power in a fixed ray class group reduces the algebraic infinity type to conjugates of a principal generator supported over that prime. The half changes at M_0 have finitely many fixed denominators. The universal determinant coefficient, The additive vacuum expansion, and The finite half changes.

For each fixed source class one may translate its representative into split support away from the level and ℓ. Decompose the finite adelic argument into a rational unitary element, a split-supported element and an element of a compact subgroup fixing the chosen translates. Weak approximation for the norm-one determinant torus and strong approximation for special unitary away from a fixed split place give this decomposition. At ℓ the split-supported element and the moving half changes are identity. Thus denominators of the rational element and its inverse there are uniformly bounded by the fixed source class and the fixed half changes. Substitution in the harmonic polynomial costs a fixed further power of ℓ. The split-supported translate.

Use the framed theta class identity of the base row with the lifted marking. Its terms are cubical isogeny pullbacks and finite theta-group translates, followed by the harmonic jet in the same invariant

differentials. The new torsion and isogeny degrees are prime to ℓ, so these steps extend on a good abelian scheme after an unramified extension. Prescribed theta-group lifts at the additional levels can be made integral by taking prime-to-ℓ roots of units. The harmonic contractions remove the change of Hodge splitting. In the short formula let $N = M'u$ with u prime to ℓ. The projection index and the number of torsion labels divide powers of N and M', respectively. The remaining derivative and trace-dual frame loss divides a power of $2\,\mathrm{disc}(E)N$ of fixed degree. The framed normalization is the fixed one from the base class. Thus a fixed power of ℓ bounds the finite sums on the cover, even when $q-1$ is divisible by arbitrarily high powers of ℓ. Mumford, Theorem 4, The framed short-series identity, and The fixed part of the short denominator.

The determinant coordinate is constant on components of the cover over the ray field. The shifted coefficient of index a is the sum of the unshifted coefficients of index $a-\delta$ restricted to components of determinant coordinate δ. This uses the $0,1$ component idempotents. There is no division by the order of the ray group. At each moving prime the determinant of a row translation shifts the local determinant label and cancels with the shift in the argument's ray coordinate. At the fixed type primes the local character equations are the original ones. Hence the shifted coefficient is invariant under the base level and satisfies the fixed type character in that base space. Descent can be tested on the two pullbacks to the cover fibre product, by the verified analytic invariance and the fixed algebraic reading. It requires no averaging over the deck group. The integral component fibres, The shifted coefficient and base invariance, and The fixed type subspace.

To bound coordinates in the fixed space, use finitely many algebraic CM points with framed evaluation injective on its cusp sections. Such points are dense on every component by the rational Cayley parametrization of the unitary group acting on a special point. Their abelian varieties have good reduction after one finite extension, and all additional prime-to-ℓ torsion markings then lift over unramified extensions. Invert a fixed nonsingular evaluation matrix. The uniform bound on the lifted evaluations gives a uniform bound in the fixed type-space basis. The fixed ball is a full lattice in this space. The following comparison fixes the normalization before varying the cover. Serre–Tate, Theorem 6(a), The full fixed type lattice, and The scalar comparison.

The common phase on the marked covers. Fix the stage, embedded base model, and side. Let n in this paragraph denote the number of definite-source columns. The base component constant is c_0^{fr}, and $\mu(h)$ is the nonzero adelic splitting factor. The base comparison of Section 2.6.7 uses one fixed source orthonormal lattice basis b. More precisely, in that basis the Gaussian times the product of normalized one-column theta readings has prefactor $\mathrm{ph}_0\, c_0^{\mathrm{fr}}(k_0(h))\mu(h)$, with $|\mathrm{ph}_0| = 1$. This is the comparison of the vacuum for rational arguments in the base frame. It uses the base component constants chosen with the fixed central level translation in the phase-one realization. Multiplying all the vacuum constants by the scalar of this translation changes only ph_0. The fixed base comparison and The choice of base constants.

The rational source change of basis b preserves the definite Hermitian form and the self-dual source lattice. It preserves the one-dimensional vacuum line, whose archimedean factor is invariant under the compact complex structure and whose finite factors are the normalized lattice vacua. Its action is a scalar u_b of modulus one. This operator commutes with the row action, so the scalar works for every adelic argument. With the convention $F_h(Xb) = u_b F_h(X)$ on a rational pure-real matrix X, compare the vacuum identity at X with the base comparison at Xb^*. For each component choose X with every column equal to a rational vector at which the normalized one-column theta function is nonzero. The rational vectors are dense in the real uniformizing space and the theta function has a nonzero constant Fourier coefficient in its normalized coordinates. The Gaussian and the product are identical in the two readings. Cancelling them gives

$$c_0^{\mathrm{vac}}(k_0(h)) = \mathrm{ph}\, c_0^{\mathrm{fr}}(k_0(h)), \qquad \mathrm{ph} = u_b \mathrm{ph}_0.$$

A fixed source change of basis and Comparison in all base components.

For a marked cover C, construct its vacuum constants c_C^{vac} with the same vacuum and the same splitting. For a fixed h take the rational row representatives γ_0, γ_C for its base and cover dictionary frames. Both place the same finite lattice attached to h, so $\delta = \gamma_0 \gamma_C^{-1}$ is a rational isometry carrying the cover lattice onto the base lattice. At each CM embedding the base frame matrix is δ times the cover frame matrix. In these frames the period map sends w at the cover to the same coordinates as δw at the base. The Hermitian form in the Gaussian also agrees, by the isometry. The two framed points have the same abelian variety, polarization and tangent frame after forgetting the extra marking and applying δ. Indeed the induced lattice identification carries the base torsion labels to the restricted cover labels on the period quotient. This is precisely the restriction to the original marking, with the base coefficient embedding unchanged. The rational lattice isometry and The common marked fibre.

Let t_h be the ratio of the cover theta-line generator to the base generator at those identified framed points. The same cubical line, with its rigidification, is pulled back when the extra marking is added. Choose transported symplectic coordinates on its fibre. A holomorphic reading of the base generator ϑ_0 then gives the reading $t_h \vartheta_0$ of the cover generator. This follows by identifying the analytic trivializations with the same normalized frame and applying linearity in the generator. The tangent frames agree, so for a pulled-back section of theta degree n and any dominant tangent weight a,

$$\mathrm{rd}_C(\pi_C^* f; h) = t_h^n \ \mathrm{rd}_0(f; h).$$

Rational equivariance of the vacuum gives

$$(\omega(1, \gamma_0 h) f_0)(\delta w) = (\omega(1, \gamma_C h) f_0)(w).$$

Take w with finite part zero and every period column equal to a z with $\vartheta_0(z) \neq 0$. The real period coordinates are onto. The Gaussian is the same and the product on the cover is $t_h^n \vartheta_0(z)^n$. Cancelling gives

$$c_C^{\mathrm{vac}}(k_C(h))\, t_h^n = c_0^{\mathrm{vac}}(k_0(h)).$$

Consequently $c_C^{\mathrm{vac}}/\mathrm{ph}$ realizes the same pullback comparison as the prescribed cover constants c_C^{fr}, on all pulled-back sections of the stated weights. The multiplier t_h is allowed to depend on the framed point and cancels in the comparison. Line and section covariance and Cancellation in the transported theta coordinates.

For uniqueness on a cover component met by the adelic dictionary, it suffices to use a section of tangent weight 0 and theta degree n with nonzero reading on the base component below it. The scalar-weight test exists on each base component. To see the nonvanishing used in its construction, let L be the self-dual period lattice with alternating form E_L, and let the prescribed half-ideal $\mathfrak{m}$ satisfy $\mathfrak{m}\overline{\mathfrak{m}} = (M)$, where $M \geq 2$. The form E_L/M is unimodular on $\mathfrak{m}L$, since the dual of $\mathfrak{m}L$ for E_L is $\overline{\mathfrak{m}}^{-1}L$. On the complex torus with period lattice $\mathfrak{m}L$ the base theta function is thus a section of a line of type $(M, \dots, M)$. Its translates by $M^{-1}\mathfrak{m}L/\mathfrak{m}L$ span the corresponding theta representation. This line is the M-th power of a principal line, after a possible change by a degree-zero line. The M-th power is generated at every point. In fact, one can tensor M translates of a principal theta section whose translation parameters sum to zero and for which none vanishes at the given point. The theorem of the square identifies their tensor product with that power. It follows that the theta function is nonzero at some division point of the prescribed half-lattice. Take the same marked point in every source column and take the harmonic polynomial 1. The torsion class formula for the fixed even theta degree n produces a scalar-weight section whose reading is this nonzero theta product times nonzero translation factors. Only a section of the automorphic bundle is needed in this test. Mumford, Section 1, The prescribed division points and The scalar-weight test section.

An upper-triangular change brings a frame with nonzero reading to a unitary frame and multiplies the reading by a nonzero weight character. Right translation of h at infinity reaches that frame up to the rational lattice stabilizer, whose multiplier is also nonzero. It preserves the finite dictionary component on both levels. Pull the test section back and evaluate the two pinnings there. Both readings include the

same nonzero $\mu(h)$ and positive Lie factor. Since the base reading is nonzero, the common cover reading is nonzero, and cancellation proves

$$c_C^{\mathrm{vac}}(k) = \mathrm{ph}\, c_C^{\mathrm{fr}}(k)$$

for every component, on every admissible marked cover. Indeed, for an arbitrary cover component take its marked finite lattice and any unitary frame of its prescribed signature. Realize the finite marked lattice by a finite adelic element, using the local isometries with the fixed marked PEL datum. Combine it with a real adelic element carrying the reference frame to the chosen frame. The resulting dictionary point has the chosen component, since the rational change between its lattice and the original one preserves the restricted marking and identifies the tangent frames. This is the same marked-lattice argument used for the base-to-cover frame comparison. Componentwise uniqueness · The common cover phase, and All cover components occur.

The fixed harmonic factor and the coefficient ball. Fix the coefficient field K_0 of the base row model. The harmonic polynomial map J is defined over K_0, using the fixed relabelling of the two signs and the extremal products of flag minors at the actual weights. The analytic map with the fixed unit extremal vector and orthonormal frames is a scalar ka times the map using J and the normalized vacuum. For this comparison the two maps have the same rational source action up to the oscillator scalar character. Their images intersect on the common extremal flag-minor polynomial. Both maps are injective on the irreducible source type, and that common vector makes their images equal. The induced automorphism of the source intertwines the two actions up to their scalar ratio. The trace of the action is nonzero on the rational elements near the identity, which generate the image densely by the rational Cayley parametrization at the definite real group. Taking traces there makes the ratio one, and Schur's lemma gives a scalar. In the normalization of the two extremal vectors it has absolute value one. The comparison extends by linearity to all coefficient vectors. In particular, the coordinates of an input with coefficients in a field extending K_0 are in the same field: choose a nonzero maximal minor of the fixed matrix of the polynomial coefficients of the source basis and solve for those coordinates. Neither J nor ka involves an auxiliary cover. The common normalized extremal vector and The fixed coefficient map.

The normalized short-series identity therefore uses the single factor $(\mathrm{ka}\,\mathrm{ph})^{-1}$ for all covers and coefficient vectors of the fixed stage weight. After clearing the fixed denominators of J, the cubical pullbacks and finite theta-group sums give the bounded cover readings explained above. The coefficientwise, determinant-shifted descent uses the original weight and type equations, so its output belongs to the fixed base cusp space. Let Y denote the faithful reading of the generic coefficient space for the side, with the same basis used for its integral ball. A side comparison scalar λ appears when converting the fixed analytic normalization to this reading. For each generic coefficient x, the descended bounded coefficient is an element x' of this space with $Y(x') = \lambda Y(x)$. If some $x \neq 0$, a nonzero coordinate in this shared basis gives λ in the fraction field. Clearing its single denominator scales the fixed ball by a fixed power of ℓ. If every x is zero the bound holds directly. In particular, one fixed scalar suffices for every split support and ray order, including the specializations used below. The normalization before choosing covers and The scalar comparison in the fixed ball.

The lattice determinant on the same ray coordinates. Fix one stage and side, its reference Hecke lattice B of rank n_B, and its projector A commuting with the normalized split Hecke operators. For the fixed test generators a_h, the columns $A\Theta_{\varnothing}(a_h)$ span the base test lattice in $B \otimes K$. Use the universal symbols $\Theta_Q(a_h)$ in this same coefficient model. Suppose their coefficients all enter the fixed ball after multiplication by ℓ^{c_B}, with c_B independent of the ray order and the selected pairs. Applying A carries that ball into B. Write C_Q for the matrix whose columns are $\ell^{c_B} A\Theta_Q(a_h)$ in a basis of B. The common coefficient model.

At an admissible split pair let v be the chosen half, $q = \mathrm{N}v$, and r the target rank. In the finite matrix chart with s columns, summing the group-ring determinant coefficients over the ray index gives the indicator of a full rank matrix modulo v. For rank r the complements of the row space are the graphs of maps from a fixed complement to the row space, giving $q^{r(s-r)}$, exactly the averaging denominator. At smaller rank no such quotient determinant is nonzero. The k-th split Hecke coset tests whether the image of the matrix lies in a subspace of codimension k. If the image has codimension d, there are $\begin{bmatrix} d \\ k \end{bmatrix}_q$ such subspaces. Thus the sum with coefficients $(-1)^k q^{k(k-1)/2}$ equals the primitive indicator, by the finite binomial identity $\sum_k (-1)^k q^{k(k-1)/2} \begin{bmatrix} d \\ k \end{bmatrix}_q = \prod_{i=0}^{d-1}(1-q^i)$. Dividing the Weil translation by its determinant character and modulus factor gives the corresponding normalized Hecke operator $T_{v,k}$ on the same row model. A translated summand with k-th determinant ideal shifts the residual ray indices by k times the Artin index of a uniformizer. Changing the summation variable inserts the arithmetic Frobenius character to the k-th power. Consequently deleting a pair with trivialized coordinate operates on the surviving series by

$$\mathcal{P}_v(U) = 1 + \sum_{k=1}^{r} (-1)^k q^{k(k-1)/2} T_{v,k} U^k, \qquad U = \psi'(\mathrm{Frob}_v),$$

where ψ' denotes the remaining coordinates, using the inverse if the chosen half is reversed. The identity holds in the group ring because characters separate its characteristic-zero scalar extension. The universal matrix coefficients, The Weil translation factor, and The distribution after deletion.

Delete the second pairs after trivializing their coordinates. Their Frobenius prescriptions identify the determinant factor D_M of $\det C_Q$ as the product of $\det \mathcal{P}_v((1+X_{L,a})^{e_a})$ with $e_a = 1$ or $\ell^M - 1$. The polynomial $\det \mathcal{P}_v(U)$ has degree at most rn_B and value one at zero. Its reduction has vanishing order at most rn_B at one. Composition with $(1 + X_{L,a})^{e_a}$ preserves this order because its linear term is a unit. The reduction of D_M therefore has a unit monomial with each exponent at most rn_B. For the remaining determinant g_M at zero, delete the first pairs as well. At each of these, the normalized Hecke characteristic polynomial is that of the signature $W_s = \eta_s \rho_s$, evaluated at inverse arithmetic Frobenius. The first-pair prescription can impose $\bigcap_s \ker \tilde{W}_s$ and its conjugate in the ray construction. The selected scalar then acts by $\delta^{\pm 1}$, with $\delta^2 - 1$ a unit, on every residual signature. The commuting residual Hecke algebra is block local after decomposing its finite module by maximal ideals. Every character of a nonzero block occurs on a socle vector and hence is one of the stage's residual signatures. Thus each $\mathcal{P}_v(1)$ has unit determinant. This gives $v_\ell(g_M(0)) = v_\ell(\det C_\emptyset)$. The residual Hecke normalization and The bounded order and base valuation.

By local Smith form the unscaled specialized columns have signed index $v_\ell(\det C_Q(\psi)) - c_B n_B$. Once they are in the prescribed twist lattice, a specialization with $v_\ell(\det C_Q(\psi)) \le v_\ell(\det C_\emptyset) + \epsilon$ therefore gives the index allowance ϵ. Uniform integral core comparison and The twisted lattice.

The specialized columns in the twisted Hecke lattice. Use the same fixed basis tests and rational coefficient reading as for the coefficient ball. Fix one selected ray character χ and write $\chi_b = \chi$ for $b = 1$ and $\chi_b = \chi^{-1}$ for $b = 0$. The Galois character on this side is the pullback of χ_b by the chosen ray coordinates. Its passage through the Artin map commutes with the inversion. Let O' be the coefficient valuation ring after adjoining the values of the chosen character. All split pairs in Q avoid the finite exceptional set and the level support of every chosen basis test. The same finite-support and compatibility conditions apply to every primitive power of χ. The two signs and Specialization of the common tests.

At the chosen half v of a split pair let k_v have cardinality q_v, let s be the source rank and r the row rank, and let $\kappa_v : k_v^\times \to \mathbf{Z}/\ell^M\mathbf{Z}$ be the local ray coordinate. Regard the r rows of y as elements of $k_v^{1\times s}$. For each subspace U of dimension $s - r$, take their images in $k_v^{1\times s}/U$ and evaluate their exterior product by the fixed nonzero determinant form there. Choose g_U invertible with $U g_U = 0^r \times k_v^{1\times(s-r)}$. The first r coordinates of $z g_U$ identify this quotient with $k_v^{1\times r}$. Thus the quotient determinant is c_U times the determinant of the first r columns of $y g_U$ for one $c_U \in k_v^\times$ independent of y. It is nonzero exactly when

the rows together with U span the full space, which says that U complements the row space when y has rank r. Put $\chi_{b,v} = \chi_b \circ \kappa_v$, with χ_b evaluated on the chosen coordinate alone, and extend the determinant test $\Phi_v(y) = \chi_{b,v}(\det(y_{\text{first}}))$ by zero when that determinant vanishes. The specialization of the universal coefficient on all y is

$$q_v^{-r(s-r)} \sum_U \chi_{b,v}(c_U)\, \Phi_v(y g_U).$$

At smaller rank every term is zero. At full rank the terms where U is not complementary to the row space are zero on both sides. Each coefficient of this sum belongs to O', since q_v is prime to ℓ. Products over the chosen halves commute with specialization by multiplicativity of χ_b. The shift by the determinant coordinate of the row argument g contributes exactly $\xi_b(\mathrm{detRoot}_W\, g)$, where ξ_b is the idele class character of the selected Galois character and the determinant root is a CM idele lifting the unitary determinant. Its ambiguity by a fixed-field idele disappears because ξ_b is trivial on such ideles. This is the row determinant twist in the analytic series. It follows by writing the shifted coefficient as the coefficient at $a - \delta(g)$ and changing the summation index by $\delta(g)$. Specialized determinant coefficients · The determinant shift, and The determinant-root convention.

Lift every g_U to the integral linear group at the split place and use its dual inverse at the conjugate half. The finite Schrödinger action identifies this change of matrix variables with right translation in the definite source, with the fixed splitting and the normalized Haar factor. The translated elements are supported on the chosen split pairs and are identity at the coefficient prime. The before- and after-translation source lattices therefore differ by denominators dividing a power of $\prod_{v\in Q} q_v$. The dual-Weyl coefficient lattice, formed by the prescribed tensor and dual operations on these lattices, has the same comparison after inverting that product. Thus the original harmonic coefficient of the chosen test remains integral at the translated finite argument. The linearity of the oscillator and the theta evaluation now gives an O'-linear combination of the legal twisted generators of the same datum. The character of the determinant on the row has already been incorporated in this statement. The finite source action and The translated coefficient lattices.

For completeness, a legal twisted generator is the image of an integral form in the prescribed twisted source lattice. Write Γ for the rational source group, J_χ for its finite-adelic level, and λ_χ for the level character with the same splitting and side. For a translated finite argument h and harmonic coefficient P, the corresponding source form is

$$F_{h,P}(x) = \sum_{\substack{\gamma\in\Gamma \\ h^{-1}\gamma^{-1}x\in J_\chi}} \lambda_\chi(h^{-1}\gamma^{-1}x)\, \gamma P.$$

This is a finite sum wherever it is supported, since the class stabilizers are finite. With the side's choice of the splitting the actual generator also carries the source scalar $\xi_b(\mathrm{detRoot}_V\, h)$. This scalar is a root of unity in O'. The level-character values are units in O'. At the old type primes the original fixed type values and splitting values are in the coefficient field, including the fixed roots of order dividing $g_i(g_i - 1)$ for each type prime g_i. At the extra primes the factors are the fixed side character and the selected finite-order character. Integrality follows from their monic root-of-unity equations and the integral closedness of O'. For every x, each term occurs in the dual-Weyl lattice at x by covariance under Γ and J_χ. Reindexing the sum gives the rational and level equivariance equations. Hence this is an element of the actual M^{tw} of integral source forms. Integral level-character values and The integral trace form.

In the normalized theta class sum of $F_{h,P}$, only the double class of h contributes. Each transporter contributes the same value by the theta kernel's dual covariance, and their number is the stabilizer cardinality. This cancels the inverse stabilizer cardinality already in the class sum, as an equality in characteristic zero. The resulting value is the original twisted generator, with the source scalar just described; that scalar follows by comparing the two Weil splittings through their determinant characters. The integral form $F_{h,P}$ itself has no stabilizer denominator. The shared normalized rational reading of the twisted class map is

faithful over the extended fraction field, so the equality of adelic functions gives equality of vectors in the very coefficient space in which the universal group-ring column was formed. An inadmissible universal datum is zero on the coefficient side. For an admissible datum the argument above shows that its unscaled specialization is in $\mathrm{thP}(M^{\mathrm{tw}})$. Cancellation in the theta class sum, The change of splitting, and Faithful comparison of the actual columns.

Write e_G for the prescribed finite-group projector and $A = |G|e_G$ for its integral multiple, as in the lattice determinant. The two signs are unramified at the fixed level in this comparison, so the specialized columns satisfy the original type equation. They therefore belong to the corresponding base-changed space $V_0 = e_G F^{(0)}$ after projection, and the projected legal generators lie in $e_G \,\mathrm{thP}(M^{\mathrm{tw}})$. The latter module is finitely generated over O'. The integral projected theta module that they generate is finitely generated as well, since it is a submodule over a DVR. The chosen specialized determinant is nonzero, so this module spans V_0. Its signed index relative to the original ball after base change is at most the valuation of the determinant of the unscaled specialized columns. Replacing e_G by A adds $\dim V_0\, v_\ell(|G|)$ to both the untwisted and twisted determinant valuations. Likewise the fixed coefficient ball denominator adds the same amount on both arrays of columns. These terms cancel in the difference. This gives the lattice comparison for the selected twist on the same ray array as the Selmer and toric estimates. The projected lattice and signed index.

For simultaneous specialization, fix the residue prime and the finite list of signatures, toric labels and lattice tests. Clear the fixed denominators in the latter tests. The signature and toric or lattice polynomials are over fixed discrete valuation rings, respectively O_0 and O, for all ray orders M. Their normalized valuations have a common discrete group $(1/e)\mathbf{Z}$, by taking a common multiple of the two ramification indices. The adjoining of torsion roots comes at evaluation. The variable list has fixed size as M varies.

For the toric or lattice order hypothesis, first set the second coordinates to zero. The resulting polynomial f_M is congruent to $D_M g_M$ modulo $((1+X_q)^{\ell^M}-1)_q$. The reduction of D_M has a unit monomial of degree at most a fixed N_*, by the product of bounded powers of the first coordinates times a series with nonzero constant term. Higher terms in the stated reduction congruence have degree at least ℓ^M. Also $v_\ell(g_M(0)) \le \tau$, the permitted untwisted valuation after clearing denominators. Round τ down to $(1/e)\mathbf{Z}$. For $a = 1/(e(N_*+1))$, the weighted Gauss valuation $\min_\alpha(v_\ell(c_\alpha) + a|\alpha|)$ is multiplicative. This follows by rescaling the variables by a uniformizer of valuation a and multiplying the nonzero reduced polynomials. Hence the weighted valuation of $D_M g_M$ is at most $\tau + aN_* < \tau + 1/e$. A monomial attaining it has coefficient of valuation at most τ and degree at most $\tau/a + N_*$. In any fixed degree $j \ge 1$ the cyclotomic relation has coefficient valuation at least $M - v_\ell(j)$. For sufficiently large M the same witness survives in f_M, also when regarded in the full variable list. The Selmer determinants already have bounded-degree witnesses with threshold v_ℓ.

Extract a subsequence along which every coefficient *valuation* of the finite list converges in $(1/e)\mathbf{Z}_{\ge 0} \cup \{\infty\}$. This set is compact with infinity, even for an infinite residue field. A bounded-degree witness persists, since only finitely many such monomials are possible. For each valuation array, take the coordinatewise minimal exponent vectors among the terms attaining its minimum. This set is finite by Dickson's lemma. Choose c so large that the weights $w_q = \ell^{c(q-1)}$ distinguish all vectors in each of these finite sets. Positive weights make the least weight among minimum-valuation terms occur at one vector. At a primitive ℓ^d-th root ζ, substitute $1 + X_q = \zeta^{w_q}$. The coordinate valuation is $w_q/\varphi(\ell^d)$. Choose d so large that the least term's weight contribution is smaller than both $1/e$ and the allowed error for every polynomial. No term with larger coefficient valuation can attain that minimum. For this fixed d, all sufficiently high degrees have valuation above the largest threshold plus one, uniformly in M, by nonnegativity of the coefficient valuations. In the remaining finitely many degrees finite valuations stabilize and those tending to infinity eventually exceed that same bound. Take such an $M \ge d$. The specialized polynomial has one least term, so its valuation is at most the threshold plus the allowed error. Every prime-to-ℓ power of ζ gives the same coordinate valuations. Each coordinate keeps its prescribed tame support, and its generator attains the least nontrivial valuation there. Increasing d before choosing M makes their sum, weighted by t, at

most ℓ^{-2}. This proves simultaneous specialization from the paired polynomial data, on the same actual twisting record for all the comparisons. The order hypothesis, The fixed rings and evaluations, and Use at the chosen character.

Multiplicities in the weighted defect. Fix j and a coefficient place v carrying a nonbase tuple. The bottom multiplicity space M_j is the direct sum of its actual finite isotypes $M_j[\sigma]$. For each good auxiliary Hecke point the full-image theta comparison identifies its eigenspace with the image of the corresponding carried bottom isotypes. The positive Petersson scalar on every nonzero bottom isotype makes the auxiliary theta map injective on their direct sum. Specifically the composite with its adjoint acts on each summand by the strictly positive scalar, so has zero kernel there. The same argument applies to the base tuple and to every used nonbase tuple. Their small-level images have the same rank $\dim M_j = D_0$. For the dimension of a good fibre one uses the surjectivity onto its full eigenspace in the theta comparison. The positive scalar alone proves injectivity from the bottom. The full-image comparison and sum of good dimensions and Injectivity and rank for both tuples.

The sum of multiplicities over the residual records at (j, v) counts each carried bottom isotype at most once. Indeed the rank-r spherical parameter at any admissible split place belongs to σ. It is thus independent of the record through which the isotype is read. The common tuple at (j, v) gives the same added length-t string and the same normalizing twist. The full normalized auxiliary polynomial is therefore the same for every such reading of the isotype. Passing to the coefficient ring for the fixed place v gives the same reduction. The coefficient of degree $n_V - k$ is the recorded Hecke generator times the unit $(-1)^k q_u^{k(k-1)/2}$, where q_u is the residue cardinality of the defining place $u \nmid \ell$. Thus all generator reductions agree and the residual records coincide. This uses the actual normalized Hecke action and the common tuple, including the same χ_1. Summing dimensions over disjoint collections of isotypes now gives

$$\sum_{\alpha \in \mathcal{A}_{jv}} \sum_{\pi \text{ good}} \dim V_{\alpha,\pi} \le D_0.$$

A record with no good fibre contributes zero. Uniqueness of the residual record and Summing the actual good multiplicities.

The lengths in the local theorem are normalized by $\mathrm{ord}_{O_\alpha}(\ell)$. To insert the estimate at a period coefficient place $v \mid \ell$, multiply by $e_v \lambda_v$, since $\mathrm{ord}_v(\ell) = e_v$. The sum of these place weights over $v \mid \ell$ is $\log \ell$. Thus the local per-point bounds enter the weighted stage sum with the common D_0, including the separate coarse-type and individual type-prime endomorphism costs. In the normalization of (91) this is the same η_j and D_0 that are used in the division comparison.

The large-prime windows on all embeddings. Fix the actual stage, Hecke record and good point π. Let $\ell > \ell_*$ be different from g_0, g_1. The threshold $\ell_* = mb/D_*$ tends to infinity, so $\ell \ge 5$. Write $c_0 = (t-1)/2$, $\eta = \eta_{\mathrm{sz}}$, $B = b/\ell_*$, and $s_* = \lfloor 2B \rfloor$. The exceptional set $\mathscr{E}$ consists of the primes $\ell > \ell_*$ dividing an integer in at least one interval

$$\begin{gathered}[\, W - \eta - 2s_* - 3,\ W + S_c + 2\eta + 2s_* + 4\,], \\ W \in \{b, w''(n)\} \cup \{b_a + \lfloor b'/x \rfloor : a \in T\}.\end{gathered} \tag{L2}$$

The three possible sizes W come from the thick, thin and top labels, respectively. At a labelled embedding write the signed jumps of $W'_\pi = \rho_\pi \chi_1^{-1}$ as the low set $\mathcal{L}_i$ and outer set $\mathcal{C}_i$ of the realized Hodge dictionary. The low jumps are $c_0 \pm (s - b_{\rho,a})$ with the one sign of that embedding, and satisfy

$$2B + t/2 + 3 \le |z - c_0| \le \eta \quad (z \in \mathcal{L}_i).$$

These inequalities are precisely the selected lower gap and upper closeness for the small block at $s = e_{\rho,j-1}$. The outer jumps are $c_0 \pm (s + v_{\rho,a})$ with the same sign. For some centre W in (L2) at this label, their signed

distances from c_0 are all in $[W, W + S_c + \eta + 1]$. Indeed $s \in [W, W + \eta + 1]$ on the three ranges and $0 \le v_{\rho,a} \le S_c$. Call the positive outer sign the lower-end case and the negative sign the upper-end case. The remaining t jumps of the full normalized representation $R_x\chi_1^{-1}$ are $[0, t-1]$, independent of x. The actual three centres and the low gap · The exceptional intervals.

At this scale all the centres lie between $2B + t/2 + 3$ and b. The lower ends of (L2) are positive, and the upper ends are $< 3b$. For the lower bound use $b'/b = xe^{-x^{\varepsilon/3}}$ together with $w''(n) = b'/(xn)$ and $b_a + \lfloor b'/x \rfloor \ge \lfloor b'/x \rfloor$. The latter two centres are larger than any fixed power of $U = x^{C+2}e^{x^{\varepsilon/4}}$ for large x, where B, η, S_c, t are bounded by fixed powers of U. Also $\ell_*^2 > 8b$, as follows from $\ell_* = mb/D_*$ and $D_* = O(U^7)$. Consequently each positive integer in (L2) has at most one prime divisor above ℓ_*. The count and size are therefore

$$|\mathscr{E}| \le (|T| + 2)(S_c + 3\eta + 4s_* + 8), \qquad \ell < 3b \quad (\ell \in \mathscr{E}). \tag{L3}$$

An interval of real length $H \ge 0$ has at most $H+1$ integers, which gives the constant 8. The conductor size condition $(Q_S q d)^{36} t^{150} \le b^{1/2}$ gives $Q_S q d \le e^{L/72} < b' < \ell_*$ for large x. In particular $E_\lambda/\mathbf{Q}_\ell$ is unramified for every $\lambda \mid \ell$ under consideration. Also $t + 1 < \ell$. Counting the positive windows · The centre sizes on the same family · The scale and thin centre · Conductor primes below the threshold.

For the actual local comparison, number the embeddings of the unramified E_λ by $\tau_{i+1} \circ \sigma_{E_\lambda} = \tau_i$, where σ_{E_λ} is arithmetic Frobenius. Set

$$a_i = \min(\mathcal{L}_i \cup \{0\}), \qquad b_i = \max(\mathcal{L}_i \cup \{t-1\}), \qquad \eta_i = b_i - a_i + 1.$$

In the lower-end case set $w_i = 1 - a_i$. In the upper-end case set $w_i = 1 - \min(\mathcal{L}_i \cup C_i \cup \{0\})$. In both set

$$h_i = 1 + w_i + \max(\mathcal{L}_i \cup C_i \cup \{t-1\}), \quad U_i = w_i + (\mathcal{L}_i \cup C_i), \quad \sigma_i = \left\lfloor \sum_{j \ge 1} \frac{h_{i+j}}{\ell^j} \right\rfloor.$$

The jump lists are simple by the prescribed separation. The shifted jumps lie in $[1, h_i - 1]$. The shifted scalar block $J_i = [w_i, w_i + t - 1]$ and shifted low jumps lie in $I_i = [1, \eta_i]$ in the lower-end case and $I_i = [h_i - \eta_i, h_i - 1]$ in the upper-end case. Here $h_i \le b + S_c + 2\eta + 3$ and $\eta_i \le 2\eta + 1$. For large x the first upper bound is at most $2b(\ell_* - 1)/\ell_*$, since the remaining parameters are bounded by fixed powers of U. The geometric sum gives $\sigma_i \le \max_j h_j/(\ell - 1) \le 2B$, hence $\sigma_i \le s_*$. Thus $\eta_i + 4\sigma_i + 6 < \ell$. The lower gap gives $|u - y| > \sigma_i + 2$ for $u \in U_i \cap I_i$, $y \in J_i$. It in fact holds for all the shifted nonscalar jumps.

For $u \in U_i \setminus I_i$ let $D_i(u)$ be u in the lower-end case and $h_i - u$ in the upper-end case. Then $D_i(u) \le W + S_c + 2\eta + 2$. More precisely, $D_i(u) - \eta_i$ equals $z - b_i$ in the lower-end case and $a_i - z$ in the upper-end case, where $z = u - w_i$ is an outer jump. Thus $D_i(u) - \eta_i \ge W - \eta$. If $D_i(u) \equiv v \bmod \ell$ for $-2\sigma_i - 2 \le v \le \eta_i + 2\sigma_i + 3$, then $D_i(u) - v$ is a multiple of ℓ in (L2). This is impossible for $\ell \notin \mathscr{E}$. This verifies the Frobenius carry, gap and residue conditions simultaneously at the actual ordered embeddings, with the same shift on the full representations and each scalar line. The signed local hypotheses · The shifts on the family.

Crystalline admissibility at the prescribed signed jumps. *AI-proposed candidate.* Let $K/\mathbf{Q}_\ell$ be unramified of degree n, with $\ell \ge 5$, and let O be the ring of integers of a finite extension $F/\mathbf{Q}_\ell$ containing the images of $K(\mu_\ell)$. Write ϖ for its uniformizer and k for its residue field. Number the embeddings of K into F by $\tau_{i+1} \circ \sigma_K = \tau_i$, cyclically modulo n, where σ_K is the unramified Frobenius. Use filtration jumps for the Hodge numbers, so the cyclotomic character ϵ has jump -1.

Take a crystalline representation with stable free lattice T_0 of rank r and jump multiset U_i at τ_i. Let $t \ge 1$ and take crystalline characters $\chi_{e+1} = \chi_1\epsilon^{-e}$, $0 \le e < t$, where χ_1 has jump w_i at τ_i. Put

$$B_i = [w_i, w_i + t - 1] \cap \mathbf{Z}, \qquad W_i = h_i - w_i,$$
$$\sigma_i = \left\lfloor \sum_{j \ge 1} \frac{h_{i+j}}{\ell^j} \right\rfloor, \qquad d_i = \#\{u \in U_i \,:\, u < w_i\}.$$

The integer W_i is a local width. The weights $U_i \sqcup B_i$ lie in $[1, h_i - 1]$. At an up label the edge interval is $[1, \eta_i]$, and at a down label it is $[h_i - \eta_i, h_i - 1]$. At every label B_i lies in the edge interval. Write U_i^{edge} and U_i^{out} for the weights of U_i inside and outside that interval, counted with multiplicity. Assume the position conditions

$$\eta_i + 4\sigma_i + 6 < \ell, \tag{73}$$

$$|u - \beta| > \sigma_i + 2 \qquad (u \in U_i^{\mathrm{edge}},\ \beta \in B_i), \tag{74}$$

$$u^{\mathrm{in}} \not\equiv y \pmod{\ell} \qquad (u \in U_i^{\mathrm{out}},\ -2\sigma_i - 2 \le y \le \eta_i + 2\sigma_i + 3), \tag{75}$$

where $u^{\mathrm{in}} = u$ at an up label and $u^{\mathrm{in}} = h_i - u$ at a down label. These are conditions on each labelled edge and its complement, with the intervening weights already included in the cyclic width σ_i.

Call a crystalline F-representation *full* when its jumps are $U_i \sqcup B_i$ and its characteristic polynomial at every $g \in G_K$ is congruent modulo ϖ to that of $T_0[1/\varpi] \oplus \bigoplus_{e=0}^{t-1} F(\chi_{e+1})$. The congruence is independent of the stable lattice used to compute it. A finite $O[G_K]$-module is admissible if it is a quotient of a stable lattice in a finite direct sum of individually full representations. Lattices in that sum may mix the summands.

For a finite O-module C with trivial G_K-action and a stable free lattice T in one of the specified slots, let $M = T(\chi_1^{-1})$. A cocycle z in $M \otimes_O C$ gives an extension $E_{z,a}$ of O/ϖ^a by $M \otimes_O C$ for each $a \ge 1$ with $\varpi^a C = 0$, using

$$g(m, s) = (gm + sz(g), s).$$

The admitted classes $A_C(M)$ are those for which $E_{z,a}(\chi_1)$ is admissible for all sufficiently large a. Write $H^1_f(K, M)$ for the preimage of the crystalline finite condition in $H^1(K, M[1/\varpi])$. The local comparison is

$$\begin{aligned} A_C(T_0(\chi_1^{-1})) &= \mathrm{im}\big(H^1_f(K, T_0(\chi_1^{-1})) \otimes_O C \longrightarrow H^1(K, T_0(\chi_1^{-1}) \otimes_O C)\big), \\ A_C(O(\epsilon^{-e})) &= 0 \qquad (1 \le e < t), \\ A_C(O) &= \ker\big(H^1(K, C) \longrightarrow H^1(I_K, C)\big). \end{aligned} \tag{76}$$

The scalar statements apply to every stable lattice in the character line, by an $O[G_K]$-linear choice of basis. In the first formula membership means that the cocycle is a finite sum of terms $c_j(g) \otimes x_j$, where the c_j are integral crystalline cocycles and $x_j \in C$, up to a coboundary with values in $M \otimes C$. A zero class in the second formula means an actual coboundary, with its element in $O(\epsilon^{-e}) \otimes C$. In the last formula it means that the restriction of the cocycle to inertia vanishes, since the twisted coefficient action is trivial. These are equalities at the given coefficient level.

The comparison keeps the endpoint lattices fixed and treats every finite coefficient module. The short-weight residue-field comparison of Atobe, Chida, Ibukiyama, Katsurada and Yamauchi allows lifts of the endpoints. The rank-one calculations of Chang and Diamond treat pairs of characters. [126, Lemma 7.6] [66, Theorems 7.8 and 7.12]

The residual bound. First count the admitted classes at $C = k$. Use $\bar{V} = T_0(\chi_1^{-1})/\varpi$ and $\omega = \bar{\epsilon}$. Write $\chi_e^{\mathrm{res}} = k(\omega^{-e})$ for the residual character. Here are the elementary consequences of (73)–(75) used in the count.

$$h_i + \sigma_i + 1 \le \ell(\sigma_{i-1} + 1), \tag{77}$$

$$U_i \text{ avoids } [w_i - \sigma_i - 2,\ w_i + t + 1 + \sigma_i] \pmod{\ell}. \tag{78}$$

At an up label U_i avoids $[-\sigma_i, 0]$ modulo ℓ. The outlying weights avoid $[0, \eta_i]$. At a down label U_i avoids $[h_i, h_i + \sigma_i + 1]$ modulo ℓ. The outlying weights avoid $[h_i - \eta_i, h_i]$. Moreover at a down label a member of $(U_i \sqcup B_i) - w_i$ congruent modulo ℓ to some $a \in [0, W_i]$ is equal to a. For (77) multiply the series before taking the floor at $i - 1$ by ℓ. The result is h_i plus the series at i. For the avoidance intervals, both the relevant intervals and the edge weights fit within an interval of length less than ℓ after using the indicated inward

coordinate. Congruence for the edge weights is therefore equality, excluded by the range or (74). All the required inward intervals for the other weights are contained in the interval of (75). The last assertion uses $[w_i, h_i]$ in the original coordinate at a down label.

Wach comparison at the first cyclotomic divisor. Let Δ be the tame subgroup of the cyclotomic group $\Gamma = \mathrm{Gal}(K(\mu_{\ell^\infty})/K)$. Take its pro-$\ell$ generator γ with $\epsilon(\gamma) = 1 + \ell$, and use the full cyclotomic parameter μ, with $\varphi(\mu) = (1+\mu)^\ell - 1$. The tame invariant ring after splitting the unramified labels is $O[[X]]$ at each label. One may take

$$X = \frac{1}{\ell - 1} \sum_{\delta \in \Delta} \delta \left(\frac{\varphi(\mu)}{\mu} \right) - \ell.$$

Indeed Δ linearizes on an eigenparameter of leading term μ, obtained by averaging $\epsilon(\delta)^{-1}\delta(\mu)$. The invariant ring is the series ring in the $(\ell - 1)$-st power of this parameter. The displayed X has zero constant term and leading term $\mu^{\ell-1}$ modulo ϖ, so it is an invariant generator. In the full ring $X + \ell$ is a unit times the first cyclotomic divisor $\varphi(\mu)/\mu$. Each summand vanishes at the primitive ℓ-th roots, and after dividing by that distinguished polynomial the resulting series has invertible constant term modulo ϖ.

Use the integral Wach lattice of a crystalline lattice in the full cyclotomic ring, allowing a Tate twist so that the usual effective construction applies. Berger's intersection formula $N(T) = N(T[1/\varpi]) \cap D(T)$ shows that $N(T)$ is ϖ-saturated in the integral étale module $D(T)$. It also justifies the use of the ramified coefficient ring O. Indeed the underlying $\mathbf{Z}_\ell$-Wach lattice is stable under the O-action by functoriality. After splitting the labels, its reduction modulo ϖ injects into the étale module modulo ϖ. Each component is a finite torsion-free $k[[\mu]]$-module of the expected rank as inverting μ recovers the étale module. Lifting a basis and applying Nakayama, with ϖ-torsion-freeness, gives freeness over $O[[\mu]]$ at each label. Averaging a basis modulo μ over Δ and lifting gives a Δ-fixed basis of the full-parameter module because Γ acts trivially modulo μ. Its invariant lattice is free over $O[[X]]$, and $(\gamma - 1)N_i \subset XN_i$. For a general Tate twist the formula $N(T(j)) = \mu^{-j} N(T) \otimes \epsilon^j$ preserves this congruence.

Index the invariant lattice at i by its first-divisor specialization. It is the component $e_{i+1}N^\Delta$ for the unramified idempotent e_{i+1} at τ_{i+1}, and the map φ^{-1} in the specialization reads the filtration at τ_i. Thus the Frobenius step into i is from $i-1$. To record the first divisor comparison, let J_i be the multiset of jumps of any such crystalline lattice. In the rational completion at the first divisor the Wach lattice is the filtered lattice with basis $z^{-u} d_u$, $u \in J_i$, for a Hodge-adapted constant basis d_u and coordinate $z = \log(1+\mu)$ near a primitive ℓ-th root. To see this label by label, use the constant frame of the crystalline periods and bracket the completed lattice between z^m and z^{-m} times the constant lattice for m large. On this finite window γ is diagonal by degree, with different eigenvalues $(1+\ell)^a$ in distinct degrees a. Projection onto those eigenspaces decomposes the completed lattice degree by degree. Stability under multiplication by z gives nested coefficient subspaces in the constant frame. An adapted basis gives generators $z^{b_j} d_j$, with $(1+\ell)^{b_j}$ the eigenvalues on its quotient by z. The natural $\theta \circ \varphi^{-1}$ specialization identifies that quotient, equivariantly, with the first labelled Sen space by Berger III.2.1–III.2.2. The eigenvalues there are $(1+\ell)^{-u}$, with multiplicities at τ_i. Thus $b_j = -u$ on each component. These comparisons respect the coefficient field and the unramified idempotents. Frobenius from the preceding component transports the horizontal frame at $\mu = 0$ to the first divisor, since $\varphi(\mu)$ vanishes there and the frame change at zero is invertible. Thus its linearized image in the same completion is the constant lattice.

In particular, for $-w_i \le u \le W_i$ in J_i, the bounds on the linearized image are

$$(X+\ell)^{W_i} N_i \;\subset\; \varphi_i^* N_{i-1} \;\subset\; (X+\ell)^{-w_i} N_i.$$

Indeed positive Wach Frobenius and its inverse have denominators only at the first cyclotomic divisor, and a Tate twist only multiplies by a power of it. The poles are bounded at that divisor by the comparison

with the constant frame. One may first remove all denominators by a power of $X+\ell$ and test the resulting divisibility of each integral matrix entry after inverting ℓ. Weierstrass division descends the test since $O[[X]]/(X+\ell)^j$ is O-free. Here and below we use Berger, *Limites de représentations cristallines*, II.1.1–II.1.3, II.2.1, II.3.1, III.2.1–III.2.2 and III.3.1.

There is also a truncated operator comparison. On N_i put $\nabla_i = (\gamma - 1)/X$, an O-linear operator preserving $(X+\ell)^s N_i$ for $s \geq 1$. Preservation follows rationally from the completed comparison, where X is invertible, and integrally from the same O-free quotients. In the rational first-divisor completion γ fixes the first roots and the constant basis, and sends z to $(1+\ell)z$. On the a-th layer of the line $z^{-u}d_u$, $0 \leq a < s$, the diagonal entry of ∇_i is

$$c(a-u) = \frac{1-(1+\ell)^{a-u}}{\ell} \in \mathbf{Z}_\ell.$$

The action on this truncated line is triangular and these diagonal entries within one line are distinct. Consequently the operator is killed by the polynomial with roots $c(a-u)$, using each distinct root once, even when different lines have the same root. The characteristic polynomial for $s = 1$ uses all multiplicities. These polynomial identities hold on the finite free integral truncation and can be reduced modulo ϖ. Their residual roots are $u - a$ in k. In particular at a down label, for $J_i = (U_i \sqcup B_i) - w_i$ and $s = W_i + 1$, a residual root zero arises only from the rational root zero. Thus there is a polynomial f_i with $f_i(0) \neq 0$ in k such that on the residual lattice

$$\nabla_i f_i(\nabla_i) N_i \subset X^{W_i+1} N_i. \tag{79}$$

In the same coordinates $\gamma X = X - X^2 + O(X^3)$ modulo ϖ. In the full parameter, $\gamma(\mu) = \mu + \mu^\ell + \mu^{\ell+1}$ in characteristic ℓ. The leading nonzero term of $\gamma X - X$ is $-\mu^{2(\ell-1)}$, and invariance under Δ gives the claim.

Models for residual subquotients. Twist the full representations by χ_1^{-1}. By the characteristic-polynomial congruence and Brauer–Nesbitt, reduction of any lattice in a sum of ν such representations has semisimplification with ν copies of the simple factors of $\bar{V} \oplus \bigoplus_{e=0}^{t-1} \chi_e^{\mathrm{res}}$. This holds also for a lattice mixing the generic summands, since its characteristic polynomial at each g is unchanged. For any subquotient of such a reduction, apply the exact étale cyclotomic module functor with k-coefficients and take Δ-invariants, an exact operation. Intersect the reduced Wach lattice with the étale subspace and take its image on the quotient. Call the resulting family of $k[[X]]$-lattices a model. Intersection and image commute with localizing in X, and all quotient lattices so induced are free because they embed in the corresponding étale spaces. The induced lattice inclusion is saturated for each short exact sequence of subquotients, so the induced lattice sequence remains exact modulo every X^j.

Write N_i for a model at i, $P_i = \varphi_i^* N_{i-1}$ for the $k[[X]]$-span of its Frobenius image, and D_i for the corresponding étale vector space over $k((X))$. The height bounds and congruence descend to the models:

$$X^{W_i} N_i \subset P_i \subset X^{-w_i} N_i, \qquad (\gamma_i - 1) N_i \subset X N_i. \tag{80}$$

For the intersection assertion on P_i, linearization by φ_i is flat on $k[[X]]$. The pull of the quotient lattice injects in the linearized quotient étale space, so the pull of the kernel is exactly the intersection of the pull with the kernel étale subspace. The same induced sequences show that (79) descends to the models at down labels. We may use one annihilating polynomial for all full bases, since their jump lists agree. The reduced ∇_i is k-linear, preserves $X^j N_i$ for all integers j, and its eigenvalues on $X^j N_i / X^{j+1} N_i$ are those on $N_i / X N_i$ shifted by $-j$. This follows from $\gamma X = X - X^2 + O(X^3)$ and the semilinearity. Write $s_{N_i}(a)$ for the algebraic multiplicity at $a \in k$ on $N_i / X N_i$. These multiplicities are additive in sequences with induced lattices. For a full reduction the multiplicities come from $J_i = (U_i \sqcup B_i) - w_i$, repeated for its number of generic summands, by the first-divisor calculation with $s = 1$.

There are standard models on the constituents of the split reference. The model N^0 of $\bar{V}$ has eigenvalue list $U_i - w_i$. On the character χ_e^{res} the standard lattice $N_{e,i}^0$ is the tame invariants of the full-parameter lattice

spanned by μ^e times an étale basis for $k(\omega^{-e})$. Its linearized Frobenius image equals $X^e N^0_{e,i}$, and it has the single eigenvalue e on its reduction modulo X. The full-parameter formula and the height bounds apply to this line also for $e = -1$.

Comparison of different models. For two families N, N' on the same étale object which satisfy (80), let $a_i \geq 0$ be the least integer with $X^{a_i} N_i \subset N'_i$. Their height bounds imply

$$X^{a_i+h_i} P_i \subset P'_i, \qquad \ell a_{i-1} \leq a_i + h_i.$$

For the second inequality use elementary divisors of the two lattices at $i-1$. After linearization their elementary divisor exponents are multiplied by ℓ, since $\varphi(X)$ has order ℓ. Iteration around the cycle gives $a_i \leq \sigma_i$. In particular

$$X^{\sigma_i} N'_i \subset N_i \subset X^{-\sigma_i} N'_i. \tag{81}$$

Each eigenvalue on N_i/XN_i differs from one on N'_i/XN'_i by some integer of absolute value at most σ_i. Indeed N_i/XN_i is a subquotient of the ∇_i-stable window $X^{-\sigma_i} N'_i / X^{\sigma_i+1} N'_i$, whose layers have the stated shifted eigenvalues.

Two models on a fixed residual representation have their sum as a model. Use the direct sum of their full sources, form the direct sum of the two subobjects, and compose the quotient map with addition. The bounds (81) and the noetherian property of $k[[X]]$ give a largest model on each residual representation that has a model. This assertion only uses the closure under finite sums: all models lie in one bounded family, so the sum of them is obtained after a finite sum.

The separation of the character block. The residual $\bar{V}$ has no subquotient χ_e^{res} for $-1 \leq e < t$. If it had one, the lattice induced from N^0 would have an eigenvalue in $U_i - w_i$. Comparison with $N^0_{e,i}$ places that eigenvalue in $[e - \sigma_i, e + \sigma_i]$ modulo ℓ, excluded by (78). More generally, on any subquotient of $\bar{V}$, every eigenvalue of any model belongs to $U_i - w_i$ modulo ℓ. It belongs to $(U_i \sqcup B_i) - w_i$ as an induced eigenvalue. Compare with the model induced from N^0. The eigenvalues on that induced model lie in $U_i - w_i$, so the comparison and (78) exclude all eigenvalues in $[0, t-1]$.

For $0 \leq e < t$, every model of χ_e^{res} is the standard one. To see this, take a Jordan–Hölder flag of the residual full source through the subquotient defining the given model. Induce lattices on all the factors. For each character factor of index e write its lattice at each i as $X^{a_i} N^0_{e,i}$. By (81), $|a_i| \leq \sigma_i$, and by the height bounds

$$-w_i \leq e + \ell a_{i-1} - a_i \leq W_i.$$

At an up label this implies $a_{i-1} \geq 0$, using $e + w_i \leq \eta_i$. At a down label it implies $a_{i-1} \leq 0$, using $W_i \leq \eta_i$. At a fixed i all the character factors therefore have the same weak sign for a_i. Their eigenvalues $e - a_i$ are in $[-\sigma_i, t-1+\sigma_i]$ and that interval injects into k. The $\bar{V}$-factors contribute no eigenvalues in this interval, by the preceding subquotient assertion and (78). Additivity of the eigenvalue list of the full source shows that the character list is exactly $[0, t-1]$, repeated ν times. The list of the original indices e has the same multiplicities by the characteristic-polynomial comparison. Summing the integer eigenvalues gives $\sum a_i = 0$ over the character factors at each label i. The common sign forces each exponent to vanish.

The eigenvalue counts on $\bar{V}$. Use the largest model N on $\bar{V}$, which contains N^0. The subquotient assertion gives

$$s_{N_i}(a) = 0 \qquad (0 \leq a \leq \sigma_i + 1). \tag{82}$$

The counts needed at the two edges are

$$\sum_{a=-w_i}^{-1} s_{N_i}(a) = d_i \quad \text{(up)}, \qquad \sum_{a=0}^{W_i} s_{N_i}(a) = r - d_i \quad \text{(down)}. \tag{83}$$

For the proof let $q_i(a)$ be the multiplicity on N_i/N_i^0. The two filtrations of N_i/XN_i^0, one through N_i^0 and the other through XN_i, give

$$s_{N_i}(a) + q_i(a+1) = s_{N_i^0}(a) + q_i(a).$$

Here multiplication by X shifts the eigenvalues by -1 from N_i/N_i^0 to XN_i/XN_i^0. This also follows from the semilinear formula for ∇_i by filtering the finite quotient by powers of X. More explicitly, under multiplication by X the conjugated operator on N_i/N_i^0 equals $(\gamma X/X)\nabla_i + (\gamma X/X - 1)/X$, which is $\nabla_i - 1$ on each layer of its X-filtration. By (81) the only possible eigenvalues supporting q_i are $u - w_i + j$, $u \in U_i$, $1 \le j \le \sigma_i$. Thus $q_i(0) = 0$ by (78). At an up label $q_i(-w_i) = 0$ by avoidance of $[-\sigma_i, 0]$. Sum the identity from $-w_i$ to -1. The same sum for N^0 is d_i because $[0, w_i - 1]$ has no congruences with the outlying weights and no wrap for the edge weights. At a down label $q_i(W_i + \sigma_i + 2) = 0$ by avoidance of $[h_i, h_i + \sigma_i + 1]$. Sum from 0 to $W_i + \sigma_i + 1$. The terms for N past W_i are zero by the subquotient assertion and that same avoidance. In the standard lattice the sum is $r - d_i$: the range $[w_i, h_i + \sigma_i + 1]$ has no outlying congruences and counts precisely the jumps $\ge w_i$ within the edge.

Admitted residual cocycles in Wach coordinates. Reducing the quotient slot of an admitted extension to k gives an extension $\bar{E}$ of the trivial residual character by the residual kernel. Twisting by χ_1^{-1}, the representation $\bar{E}$ is a quotient of a residual full source, because it is killed by ϖ. Take the corresponding model, with induced kernel N_i' and quotient the standard lattice of the trivial character by rank-one rigidity. Choose a lift v_i of its basis 1 in the model of $\bar{E}$. Put

$$b_i = \varphi_i v_{i-1} - v_i, \qquad g_i = (\gamma_i - 1)v_i.$$

The pair in the kernel is a cocycle in the étale Herr complex. Write $\Phi(x)_i = \varphi_i x_{i-1}$ and $(\gamma - 1)x_i = (\gamma_i - 1)x_i$. Herr's comparison identifies its class with the given residual class, and changing the lifts adds $((\Phi - 1)x, (\gamma - 1)x)$. In particular

$$(\gamma_i - 1)b_i = \varphi_i g_{i-1} - g_i, \qquad g_i \in XN_i'. \tag{84}$$

Use $P_i' = \varphi_i^* N_{i-1}'$. At an up label $b_i \in X^{-w_i}N_i'$ by the upper height bound. At a down label $b_i \in X^{-W_i}P_i'$: apply the lower height bound to $X^{W_i}v_i$ and subtract $X^{W_i}\varphi_i v_{i-1}$. The two sections have the same quotient, and the pull of the extension lattice has kernel P_i' by flatness. At a down label there is also $x_i \in N_i'$ with

$$g_i - (\gamma_i - 1)x_i \in X^{W_i+2}N_i'. \tag{85}$$

Indeed $f_i(\nabla_i)v_i/f_i(0)$ is another lift of 1 and (79) gives the claimed congruence. The difference of the two lifts is x_i. The kernel of the quotient lattice map is saturated, so the congruence holds in the kernel lattice.

Polar parts for the rank-r slot. Enlarge the induced kernel lattice to the largest model N of $\bar{V}$. All the conditions on (b_i, g_i) persist under the inclusion of lattices and their Frobenius images. By (82) the operator ∇_i is invertible on $N_i/X^{\sigma_i+2}N_i$. Choose $y_i \in N_i$ such that $\nabla_i y_i \equiv g_i/X \pmod{X^{\sigma_i+2}N_i}$. Subtract the coboundary of y, writing the new pair b_i', g_i'. Then $g_i' \in X^{\sigma_i+3}N_i$. The inequality (77) and the upper height bound show

$$\varphi_i g_{i-1}' \in X^{\max(\sigma_i+3, W_i+2)}N_i.$$

For the polar parts, i.e. the classes of b_i' in D_i/N_i, choose the representatives

$$c_i = b_i' \quad \text{(up)}, \qquad c_i = b_i - \varphi_i y_{i-1} + x_i \quad \text{(down)}.$$

Their defining conditions are

$$c_i \in X^{-w_i}N_i, \qquad \nabla_i c_i \in N_i \qquad \text{(up)}, \tag{86}$$

$$c_i \in X^{-W_i}P_i, \qquad \nabla_i c_i \in X^{W_i+1}N_i \qquad \text{(down)}. \tag{87}$$

In the up case N_i and P_i both lie in the indicated outer lattice, and (84) for the adjusted pair gives the ∇_i condition. In the down case the outer lattice contains P_i and N_i, and the numerator $(\gamma_i - 1)c_i$ equals $\varphi_i g'_{i-1} - (g_i - (\gamma_i - 1)x_i)$. Use (85) and the Frobenius estimate.

If all these polar parts of an adjusted pair vanish, then $b'_i \in N_i$ and (84) gives $\nabla_i b'_i \in X^{\sigma_i+2}N_i$. Invertibility on this truncation gives $b'_i \in X^{\sigma_i+2}N_i$. On the product of these deep lattices Φ maps each component into at least one additional X-layer, by (77), and an extra j layers in the source become ℓj extra layers in the target. The negative geometric series in Φ inverts $\Phi - 1$ there. Solve $b' = (\Phi - 1)m$. The rest $g' - (\gamma - 1)m$ is Φ-fixed by (84) and vanishes by contraction. The class is therefore a coboundary. The same argument applies to linear combinations of adjusted pairs, since the largest model and the indicated polar subspaces are common to all the pairs.

At an up label the solution space in D_i/N_i of (86) is the kernel on $X^{-w_i}N_i/N_i$. Its dimension is at most the multiplicity of zero, which is d_i by the layer shift and (83). We next bound the polar solution space at a down label by d_i.

The down-label calculation. Fix a down label i and put

$$L = X^{W_i+1}N_i, \qquad F_i = X^{-W_i}P_i.$$

Thus $L \subset N_i \subset F_i$, and ∇_i is a well-defined operator on D_i/L. We will prove

$$F_i \subset S + \nabla_i F_i + L, \qquad \dim_k S \le r, \tag{88}$$

where $S = X^{-1}\varphi_i(S_0)$ and S_0 is the k-span of a $k[[X]]$-basis of N_{i-1}.

For this, define ψ_i as the cyclotomic left inverse of φ_i. In the full parameter set $Z = 1 + \mu$ and write a vector uniquely as $\sum_{j=0}^{\ell-1} Z^j \varphi_i(u_j)$; set ψ_i equal to u_0. This uses the basis Z^j of the full scalar field over its Frobenius image. The Δ and γ actions commute with ψ_i. For each group element the remainders of the exponents of Z^j modulo ℓ permute the summands and preserve the $j = 0$ summand. Consequently ψ_i descends to D_i. We need

$$\psi_i(XF_i) \subset N_{i-1}, \qquad \gamma_i - 1 \,:\, F_i \cap \ker \psi_i \xrightarrow{\sim} XF_i \cap \ker \psi_i.$$

Choose a tame-invariant basis of the predecessor lattice and use its Frobenius image as a full-parameter pull basis. The change from μ^j to Z^j is triangular over k with diagonal one. Thus a full-parameter order bound ℓa on the pull coefficients is equivalent to order at least a on every u_j in the predecessor basis. On a $j \neq 0$ summand the difference $\gamma_i - 1$ in those predecessor coordinates is $Z^j\gamma_{i-1} - 1$, since $\gamma(Z^j) = Z^j\varphi(Z^j)$. The predecessor γ matrix is 1 mod X and the scalar action is $\mu \mapsto \mu + O(\mu^\ell)$. In order a this difference raises the order by exactly one, with leading multiplier j. Successive coefficient solution gives a bijection from full order ℓa to $\ell(a+1)$ on $\ker \psi_i$. For $2 - \ell \le a \le 0$, on tame invariants full order at least ℓa is equivalent to X-order at least a: the full orders are multiples of $\ell - 1$ and $(\ell-1)(a-1) < \ell a \le (\ell-1)a$. Use $a = -W_i$ and $a + 1$, both in that window by (73). By uniqueness the solution of the difference equation is tame invariant. For the first inclusion, ψ_i therefore sends XF_i to full order at least $1 - W_i$, which is greater than $-(\ell - 1)$. Its invariant output must have order at least zero. This proves both claims.

For $x \in F_i$, put $u = \psi_i(Xx) \in N_{i-1}$. The two claims give $Xx = \varphi_i u + (\gamma_i - 1)x'$ for $x' \in F_i$. Write $u = u_0 + Xu_1$ with $u_0 \in S_0$, and choose $z \in N_{i-1}$ so that $u_1 \equiv \nabla_{i-1} z \pmod{X^{\sigma_{i-1}+2}N_{i-1}}$. The height bound and (77) show

$$X^{-1}\varphi_i\big(X(u_1 - \nabla_{i-1}z)\big) \in L.$$

In fact the left side has order at least $\ell(\sigma_{i-1}+3) - w_i - 1$ in N_i, which is at least $W_i + 1$. Since φ_i commutes with γ, division of the expression for Xx by X now yields

$$x \equiv X^{-1}\varphi_i u_0 + \nabla_i(x' + \varphi_i z) \pmod{L}.$$

This proves (88), as $\varphi_i z \in P_i \subset F_i$.

The quotient F_i/L is finite-dimensional. Apply rank-nullity to its map by ∇_i into D_i/L. The inclusion (88) gives

$$\dim_k\big((F_i/L) \cap \ker(\nabla_i : D_i/L \to D_i/L)\big) \le r.$$

Inside N_i/L, the generalized zero eigenspace is killed by ∇_i by (79). Its dimension is $r - d_i$ by the layer shift and (83). Passing from D_i/L to D_i/N_i removes these $r - d_i$ kernel dimensions. The space of possible polar classes in (87) therefore has dimension at most d_i.

To finish the rank-r residual bound, choose any linearly independent admitted classes and an adjusted pair for each. If a linear combination of their polar vectors vanishes at every label, the same combination of their pairs is a coboundary by the contraction argument. Thus the vectors are independent in the sum of the polar spaces, and

$$\dim_k A_k(T_0(\chi_1^{-1})) \le \sum_i d_i. \tag{89}$$

The residual character slots. For $0 \le e < t$ use the standard model $N_{e,i}^0$ for the kernel in (84), as justified by rigidity. Suppress the character index and write $P_i = X^e N_i$. At an up label $b_i \in X^{-w_i} N_i$; at a down label $b_i \in X^{-(W_i - e)} N_i$. In both cases $\nabla_i b_i \in N_i$, since $g_i \in XN_i$ and $\varphi_i(XN_{i-1}) \subset X^\ell P_i \subset XN_i$. On a polar layer $X^{-j}N_i/X^{1-j}N_i$ the eigenvalue is $e + j$. For $j \ge 1$ in the two indicated windows one has $0 < e + j < \ell$, by (73) and B_i lying at the edge. Thus $b_i \in N_i$. For $e \ge 1$, Frobenius contracts the product of the N_i themselves because $X^e N_i \subset XN_i$. Solving for b and then using (84) makes the class a coboundary as before.

For $e = 0$ use the fixed standard basis of the unit at each label. Frobenius contracts XN_i, so remove the positive degree part of b_i by a coboundary. The resulting b_i are constants and $(\gamma_i - 1)b_i = 0$. The remaining $g_i \in XN_i$ are Φ-fixed and vanish. On the constant vectors over the cycle, the image of $\Phi - 1$ equals the kernel of their sum, by recursively solving the successive differences. The admissible residual classes have dimension at most one, including when ℓ divides n.

Admissibility and changes of finite coefficients. The finite admissible modules are closed under submodules, quotients, and sums. For a submodule take its preimage in the presenting stable lattice. For a quotient enlarge the kernel lattice. These lattices have finite index in a lattice in the same generic representation. For sums take the direct sum of the presenting lattices and full representations. In particular the standard constituent modules $(M \otimes_O C)(\chi_1)$ and $(O/\varpi^a)(\chi_1)$ are admissible, by the split reference, finite copies, and quotients.

For $\varpi^a C = 0$ the admissibility of $E_{z,a}(\chi_1)$ is independent of $a \ge 1$. To lower a within that range, take the quotient on the last coordinate. To enlarge a, take the pullback over O/ϖ^a of the extension and the larger standard last coordinate. This pullback is a submodule of the direct sum of two admissible modules after restoring the twist. An integral coboundary changes only the section, giving an isomorphic extension. For addition of cocycles take their two extensions at one level, pull back to the diagonal last coordinate, then push out by addition on the first coordinates. A scalar multiple is the pushout by scalar multiplication on the first coordinate. Thus $A_C(M)$ is an O-submodule of $H^1(K, M \otimes C)$.

For a coefficient map $C \to C'$ choose a common a. The new extension is a pushout on the first coordinate, which is a quotient of the direct sum of the old extension and the new standard first coordinate. Hence the cohomology map sends $A_C(M)$ into $A_{C'}(M)$. For an injection $C \hookrightarrow C'$, the extension over C is the submodule of the pushed-out extension with the first coordinate in $M \otimes C$. This uses the freeness of M. Thus a class over C is admitted whenever its image over C' is admitted. These assertions also apply when the pushed-out cocycle has been replaced by a cohomologous representative.

The ambient cohomology at finite level. For $M = T_0(\chi_1^{-1})$, the absence of a trivial subquotient of $\bar{V}$ gives $H^0(K, M/\varpi M) = 0$. The absence of $k(\omega)$ gives $H^2(K, M/\varpi M) = 0$ by local Tate duality. With

k-coefficients, the latter duality pairs with $\mathrm{Hom}_k(M/\varpi M, k(\omega))^{G_K}$, using the trace from k to $\mathbf{F}_\ell$ to identify the underlying finite dual. The same vanishings for $M = O(\epsilon^{-e})$, $1 \le e < t$, follow from $t < \ell - 1$ and the order $\ell - 1$ of ω.

Dévissage gives these vanishings for all $M/\varpi^a M$ in both cases. The long exact sequences of continuous cohomology give a finite free $H = H^1(K, M)$, and a natural isomorphism

$$H \otimes_O C \simeq H^1(K, M \otimes_O C) \tag{90}$$

for every finite C. Indeed, local cohomology of finite modules is finite. The transition maps in H^1 for the powers of ϖ are onto by H^2 vanishing, and the inverse limit computes the cohomology of M. The exact sequences for multiplication by ϖ^a show that H is ϖ-torsion-free, that $H^2(K, M) = 0$, and that $H/\varpi^a H$ is the finite-level H^1. Completeness and finite-dimensional reduction give finite generation. The natural map (90) is now an isomorphism on each cyclic finite module and hence on all finite modules. In the rank-r case H has rank nr by the local Euler characteristic. Since continuous cochains into $M[1/\varpi]$ have bounded image on compact source, localization identifies $H[1/\varpi]$ with the rational H^1.

For the trivial slot, local reciprocity gives

$$H^1(K, C) = \mathrm{Hom}_{\mathrm{cts}}(G_K^{\mathrm{ab}}, C) \simeq \mathrm{Hom}_{\mathrm{cts}}(\widehat{K^\times}^{(\ell)}, C).$$

Here $\widehat{K^\times}^{(\ell)} \simeq \mathbf{Z}_\ell^{n+1}$: one factor comes from the valuation and the other n from the logarithm on $1 + \ell O_K$. The finite residue unit group has prime-to-ℓ order. Thus (90) also holds for $M = O$ and the unramified line is a direct summand of $H^1(K, O)$. Its tensor with C is the unramified $H^1(K, C)$, with a class specified by its value on arithmetic Frobenius of the quotient by inertia. Coboundaries are zero in this twisted slot.

Crystalline and unramified classes are admitted. For $M = T_0(\chi_1^{-1})$, the integral Bloch–Kato subgroup $J = H^1_f(K, M)$ is saturated in H by definition. Its rank is $\sum_i d_i$ by Bloch–Kato's crystalline dimension formula [64, Corollary 3.8.4 and (3.8.5)] $\dim_F H^1_f(K, V) = \dim_F(D_{\mathrm{dR}}(V)/\mathrm{Fil}^0) + \dim_F H^0(K, V)$, for $V = M[1/\varpi]$. Here $H^0(K, V) = 0$ by the vanishing of residual invariants, and the jumps of V are $U_i - w_i$. An integral representative of a class in J defines an integral extension of the unit by M with crystalline generic fiber. After twisting back by χ_1, add the slots $\chi_2, \dots, \chi_t$ to obtain one full representation. Strictness of the crystalline filtered realization gives the exact jump multiset $U_i \sqcup B_i$. The triangular action gives the reference characteristic polynomial at every g, even before reduction. The original integral extension with the extra standard lattices realizes its reduction modulo ϖ^a as a quotient after killing the extra slots. Pushing out along maps $O/\varpi^a \to C$ and adding such classes shows $J \otimes_O C \subset A_C(M)$.

In the trivial slot a unit unramified extension at level O/ϖ^a is realized using two individually full representations. Use the split reference and another such representation with χ_1 replaced by $\chi_1\eta$, where η is unramified with $\eta(\mathrm{Frob}_K) = 1 + \varpi^b$ and $b \ge a$. This defines a continuous character by powers on $\mathbf{Z}_\ell$. An unramified character is crystalline of jump zero, and $\eta \equiv 1 \bmod \varpi$, so the second representation is full. After twisting by χ_1^{-1}, write e_1, e_2 for bases of the two distinguished lines. On the lattice with basis $e_1, (e_2 - e_1)/\varpi^b$, inertia acts trivially and arithmetic Frobenius acts by

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 + \varpi^b \end{pmatrix}.$$

This lattice is stable under the unramified group, since the matrix and its inverse are integral and the action is continuous. Reduction modulo ϖ^a gives the unramified extension with value one at Frobenius. Giving the remaining summands the same lattices in the source and the kernel realizes it as an admissible quotient after twisting back. The map $O/\varpi^a \to C$, $1 \mapsto x$, realizes the unramified class with arbitrary Frobenius value $x \in C$ by pushout. In the higher scalar slots the zero class is admitted by the standard constituent modules.

Exact comparison at each finite level. For each slot use the ambient module H in (90). In the rank-r slot take $J = H^1_f(K, M)$, in the higher scalar slots take $J = 0$, and in the trivial slot take the unramified line J. These are direct summands of H. In each case $J \otimes C \subset A_C(M)$, as just proved. For $C = k$ the residual bounds give $\dim_k A_k(M) \leq \operatorname{rank}_O J$, hence equality with $J \otimes k$.

Write $H = J \oplus J'$. If $x \in A_C(M) \cap (J' \otimes C)$, take a submodule $C_0 \simeq k$ when $C \neq 0$. By induction on length, the image of x in $J' \otimes (C/C_0)$ is zero. Thus x comes from $J' \otimes C_0$, and admissibility reflects the injection $C_0 \hookrightarrow C$. The residual equality forces $x = 0$. Subtracting the J-component of an admitted class is allowed, so $A_C(M) = J \otimes C$ inside the actual cohomology. This proves (76).

In particular the rank-r admitted module has O-length $(\sum_i d_i) \operatorname{length}_O C$, and its quotient in H^1 has length $(nr - \sum_i d_i) \operatorname{length}_O C$. The first scalar line contributes $\operatorname{length}_O C$ and the higher scalar slots contribute zero. Every comparison took place modulo powers of ϖ. The length normalized by $v_\ell(\ell) = 1$ is the O-length divided by $e(O/\mathbf{Z}_\ell)$, since $v_\ell(\varpi) = 1/e(O/\mathbf{Z}_\ell)$. The admitted classes are already in the prescribed finite, zero, and unramified subgroups, with zero additional loss.

If a finite free extension of coefficient DVRs is used to realize the labelled characters, the conclusions descend for classes over the original DVR whenever their extensions have been shown admissible. Continuous cochains and their cohomology commute with this finite free base change. For the finite condition use also its rational period kernel and the defining inverse image under localization. These comparisons commute with flat base change, as does the inertia kernel. Vanishing, and membership of the finite or unramified class in the corresponding subgroup, descend by faithful flatness. This assertion concerns the cohomology conclusions at the given coefficient module and includes their coboundary quotients.

The large-prime local conditions on the Hecke cocycles. Work over the complete split coefficient DVR O of the selected record. The full representations $u_x = R_x \chi_1^{-1}$ have the same residual characteristic polynomials for every Galois element, by the common residual Hecke record. At π their common residual semisimplification is the absolutely irreducible cuspidal rank-r factor and the t characters $\bar{\epsilon}^{1-c}$, $1 \leq c \leq t$. Because E_λ is unramified, the inertia cyclotomic image modulo ℓ contains a generator of $\mathbf{F}_\ell^\times$. The scalar characters are distinct, since $t + 1 < \ell$. In particular the trivial character occurs exactly once. Let $\mathscr{R}$ be the matrix order spanned by the $u(g)$ over the integral Hecke algebra A. Choose an integral idempotent $e_1 \in \mathscr{R}$ for the trivial residual factor. To see this, in the image algebra of the residual fibres take the central projector onto the simple trivial factor of its semisimple quotient. Lift it to the integral order. The idempotence error is nilpotent on every residue fibre. Iteration $a \mapsto 3a^2 - 2a^3$ converges in the order and gives the idempotent, since the error is squared up to an integral factor at every step. The order is a finite lattice in its rational span, so ambient convergence also gives convergence in the order. Along each residual composition series the rank of this idempotent is the sum of its ranks on the factors, namely one. The integral idempotent has the same rank in the generic fibre. At π the idempotent is the identity on the first scalar summand, by restriction to its saturated integral line and reduction. All the rational summands are stable under the order, so the rank-one idempotent is zero on the other summands. This gives the first-character projector with unit denominator. Unit separation of the character roots.

Write Y for the left order module generated by $y(g) = (u(g) - 1)e_1$ and $Y^+ = Y + Oe_1$. Evaluation on the trivial summand at π kills Y and is 1 on e_1. The cocycle identity makes Y^+ a stable lattice with quotient $Y^+/Y = O$ of trivial action. At a nongood fibre on the coarse cut, the first character is absent, while e_1 has rank one. Thus some $y(g)$ moves its image, and evaluation on a suitable column gives a nonzero map $Y \to V_x$. The determinant comparison therefore applies using Y as the relation module at these primes.

Let $F : Y \to B_s$ be an order-linear map to one torsion constituent at π, with lattice T_s, where $B_s = T_s[1/\ell]/T_s$. Then $z(g) = F(y(g))$ is a continuous cocycle, since the matrices $u(g)$ vary continuously and F has finite image. The torsion $F(Y)$ has bounded exponent since Y is finite over O. To verify the local admissibility of this cocycle, enlarge the coefficient DVR to O' to contain the embeddings at λ, write ϖ' for its uniformizer, and choose an integral crystalline character θ with labelled jumps w_i. Such characters

are products of the labelled Lubin–Tate characters of the unramified base, using the sign for the chosen jump convention. The shifted full representations are $u_x \otimes \theta$. They are crystalline with jumps $U_i \sqcup J_i$ and their characteristic polynomials agree modulo the new coefficient uniformizer with those at π, since they have the same twist. The constituent lines after the shift are $\theta\epsilon^{1-c}$ and have jumps $w_i + c - 1$.

Here is the integral quotient realizing the extension cocycle. Extend scalars to O' in Y^+, Y, T_s, B_s, F in the following formula. Identify $B_s[\varpi'^a]$ with $T_s/\varpi'^a T_s$ and choose a such that $\varpi'^a F(Y) = 0$. For every $a' \geq a$, the surjection

$$Y^+ \oplus T_s \longrightarrow (T_s/\varpi'^a T_s) \oplus O'/\varpi'^{a'}, \qquad (y + de_1, z_0) \longmapsto (F(y) + \overline{z_0},\ \bar{d})$$

has target action $(v, d) \mapsto (u_s(g)v + d\, z(g), d)$. The rational span of its source embeds in a finite direct sum of the full u_x, by using matrix columns for Y^+ and the fibre at π for T_s. These are globally semisimple representations, so take a stable rational complement and put any stable lattice in the complement. Extend the surjection by zero there. This realizes the target as a quotient of two full stable lattices. Now restrict to G_{E_λ} and tensor both lattices by θ. The target becomes the prescribed extension of $(O'/\varpi'^{a'})(\theta)$ by $T_s/\varpi'^a T_s$ shifted by θ. This argument uses the residual characteristic polynomial and the common Hodge list on every full summand. The common shift and full crystalline inputs · The full-lattice quotient for the cuspidal cocycle · The full-lattice quotient for scalar cocycles.

Apply the signed crystalline local comparison (76) with the shifted rank-r target and the shifted scalar lines. Untwisting by θ recovers the original normalized lattices in its three conclusions. It puts the cuspidal cocycle in the subgroup obtained from integral H^1_f by reduction, puts every higher scalar cocycle in the coboundaries, and puts the first scalar cocycle in the unramified subgroup. These implications descend to the original coefficient DVR. In fact the extension has an O-linear retraction taking 1 to 1, since O'/O is torsion-free and finite over the DVR. It gives a retraction on lattices, rational representations, their torsion quotients and cochains. It respects the finite condition by functoriality of the rational crystalline condition and maps the integral representatives to integral ones. It also respects coboundaries and restriction to inertia. Thus all local allowances at $\lambda \mid \ell$ are exactly zero. The three local conclusions in the actual record · No loss on descending coefficients.

The map from F to its global cohomology class is injective here. For a nontrivial target constituent, write $e_1 = \sum_g a_g u(g)$ with $a_g \in A$, and put $a = \sum_g a_g$. Then $a(\pi) = 1$ by evaluation on the first character, and $(1-a)e_1 = \sum_g a_g y(g)$ is in Y. If $z(g) = (u(g)-1)b_s$, evaluate F on this sum. It is zero because e_1 acts as zero on the target and $(1-a)e_1 = e_1(1-a)e_1$. It is also $(e_1 - a)b_s = -b_s$, hence the cocycle is zero. On the trivial target a coboundary is zero from the start. The generators $y(g)$ then give $F = 0$. Away from ℓ, the already verified tame relations impose the unramified conditions with the stated $t\sum_{z\in Q_\alpha} v_\ell(\zeta_z - 1)$ allowance. Thus the determinant argument yields

$$v_\ell(\mathfrak{c}_{\alpha,\pi}) \leq s_{\alpha,\pi} + V_\ell + v_\ell(h_E) + t \sum_{z\in Q_\alpha} v_\ell(\zeta_z - 1).$$

This is the large-prime zero of R_ℓ in (L1) for $\ell > \ell_*$ outside $\mathscr{E}$ and the type primes. The global Selmer, class-number, tame and type-projector terms are still measured at those places. The global comparison with unit projector · The local allowance vanishes.

Summing the defects over the actual coefficient places. Fix the member and stage, and use the common nonbase tuple of each coefficient place in all its good fibres. Write $n_V = r + t$, $L = x^\varepsilon$, $b = e^L$, let h_E be the class number, and put

$$V_\ell = \sum_{a=1}^{t-1} \mathrm{shaLen}_\ell(a), \qquad V = \sum_\ell (\log \ell) V_\ell.$$

The actual Speh length satisfies $t \leq \lfloor x^{400} \rfloor$, within the original fine Selmer range $\lfloor x^{401} \rfloor$. At the common tuple that bound gives $V \leq \sqrt{b}$. For a negative Tate character, choosing a $\mathbf{Z}_\ell$-basis of O_α makes the torsion

coefficient module and its localization map a direct sum of $[\operatorname{Frac}(O_\alpha) : \mathbf{Q}_\ell]$ copies of the $\mathbf{Z}_\ell$ version. An O_α-module of length one has $\mathbf{Z}_\ell$-length f_α, the residue field degree. Thus dividing the O_α-length of the extended group by $\operatorname{ord}_\alpha(\ell)$ cancels the degree $f_\alpha \operatorname{ord}_\alpha(\ell)$ from scalar extension. The same normalization gives $v_\ell(h_E) = \operatorname{ord}_\ell(h_E)$ for the everywhere unramified first scalar block. The allowed range of t · Fine Selmer at the actual family member.

Off g_0, g_1, write R_ℓ for the allowance for the local restrictions and the non-place error. The congruence comparison bounds the normalized valuation of the good-point ideal by

$$R_\ell + s_{\alpha,\pi} + V_\ell + v_\ell(h_E) + t \sum_{z \in Q_\alpha} v_\ell(\zeta_z - 1).$$

The two bounds on the selected twist give $s_{\alpha,\pi} + t \sum_z v_\ell(\zeta_z - 1) \le \nu_\ell + 2/\ell^2$. The cyclotomic calculation on the classical host gives $\nu_2 = 1$ and $\nu_\ell = 0$ for odd ℓ. The integral coarse type projector contributes at most $v_\ell(c_{\mathrm{row}})$ per good dimension, with $c_{\mathrm{row}} = \prod_{i=0}^{1} |\mathrm{GL}_r(\mathbf{F}_{g_i})|^f$. At g_0, g_1, use the individual type cuts and the tame potentially crystalline calculation instead, with the same global terms and $R_\ell \ll x^C \eta_{\mathrm{sz}}$, inclusive of the type endomorphism cost. Indeed its exponent is polynomial in $n_V + 1$ and linear in $1 + \lfloor \eta_{\mathrm{sz}} \rfloor + t$. The sum over host places, the coefficient-index cost, and the invariant terms are at most a fixed polynomial in the ranks times $1 + \eta_{\mathrm{sz}}$. The ranks are polynomial in x and $\eta_{\mathrm{sz}} \ge t \ge 1$. The global terms at a good point · Both bounds at the actual twist · The cyclotomic defect · The type-prime comparison.

At a single v, the sum of the good multiplicities over all records is at most D_0, by the good-fibre comparison. Then sum over the chosen nonbase $v \mid \ell$ with $e_v \lambda_v = [K_v : \mathbf{Q}_\ell] \log \ell / [K : \mathbf{Q}]$. These weights sum to at most $\log \ell$ by the local degree formula. Thus a collection $\mathscr{S}$ of rational primes contributes at most

$$D_0 \sum_{\ell \in \mathscr{S}} (\log \ell)\big(R_\ell + \nu_\ell + 2/\ell^2 + V_\ell + v_\ell(h_E) + v_\ell(c_{\mathrm{row}})\big). \tag{L1}$$

At the type primes the extra nonnegative $v_\ell(c_{\mathrm{row}})$ is harmless. This derivation preserves coefficient one on the higher fine Selmer sum after division by D_0. Weighting by the local degrees · Combining the prime contributions.

The ordinary primes below the large threshold. Use $A = \ell_A = kL/(md_2)$ and $\ell_* = mb/D_*$, where

$$D_* = nt^2\big(k + m(t+1) + t\big)L^2 + kn^2(S_c + t + L).$$

In particular $\ell_* L \le b$. The following rows are away from $2, 3$, Q_S, q, d, and g_0, g_1. With $M_\ell = 1 + h_B/\ell + \log_\ell h_B$, $h_B \le 2b$, the realized local estimates give $CfmM_\ell$ at split primes. For inert $\ell < A$ of host order $> 4L_w$ they give $CfrR_\ell^{\mathrm{long}}$. For other inert primes they give $Cfr(h_B/\ell^2 + \mathrm{flat}_\ell) + CfmM_\ell$, where

$$R_\ell^{\mathrm{long}} = \frac{h_B}{(\ell-1)^{j_\ell+1}} + (G_P + 1)(j_\ell + 1) + L,$$
$$\mathrm{flat}_\ell = \frac{w_P + \eta_3}{\ell} + S_c + L + \eta_3 \mathbf{1}_{\mathrm{band}}.$$

Here j_ℓ is the first integer with $(\ell-1)^{j_\ell} \ge A$. Include a common $Cfnt^2 n_V L^2$ in each row for the scalar pushdown, the column relations and the invariant terms. The split row uses the small count $m[E_\lambda : \mathbf{Q}_\ell]$ for the integral row target or column finite classes, with opposite signs at the two conjugate completions. The inert rows use the actual cyclic words of the family and its Tate dual. The degree and place counts are $\sum_{\lambda|\ell}[E_\lambda : \mathbf{Q}_\ell] = 2f$ and $\sum_{\lambda|\ell}([E_\lambda : \mathbf{Q}_\ell] + 1) \le 4f$. The ordinary envelopes · Their realization by local lifting · The congruence bound with the local errors.

For $2 \le \ell \le \ell_*$, use $\sum \log \ell \le \ell_* \log \ell_*$ and $\sum (\log \ell)/\ell \ll (1 + \log \ell_*)^2$. The common polynomial costs $O(fmbL)$, since $\ell_* nt^2 n_V L^2 \le mb$. The terms M_ℓ cost $O(fmbL^2)$, using also $\ell_* \log h_B \ll b$.

For the long-order row use $r/A \le 2m/L$, $G_P + 1 \ll b'$, $\log A \ll \log x$, and $rb'A(1 + \log x) \ll mb$. The latter follows from $b'/b = xe^{-x^{\varepsilon/3}}$ and the polynomial sizes of r, A. Hence $R_\ell^{\mathrm{long}} \ll b/(A(\ell-1)) + b'(1 +$

$\log A) + L$. The harmonic logarithmic sum up to A controls the first term and $A \log A$ controls the others. The resulting cost with the factor fr is $O(fmbL^2)$. For the h_B/ℓ^2 term at $\ell \geq A$, use $\sum_{u \geq A} (\log u)/u^2 \ll (1 + \log A)/A$.

For the small-order primes write $M_* = 2^{e+2-\lfloor \log_2(4L_w) \rfloor}$, where the dyadic real group has order 2^e. Modulo signs the odd classes modulo 2^{e+2} are generated by 5. A class 5^u of order at most $4L_w$ has $2^{e-\lfloor \log_2(4L_w) \rfloor} \mid u$. The identity $v_2(5^u - 1) = 2 + v_2(u)$ for $u \neq 0$ therefore puts the prime in the two classes ± 1 mod M_*. Here $M_* \geq R_0 + 1$, $M_* \leq 4n$, $k/m \leq R_0^2$, and $r \leq 2k$, with $R_0 = n/L$. By writing $\ell = uM_* \pm 1$ for $u \geq 1$, the sum of $(\log \ell)/\ell^2$ is $O((1 + \log M_*)/M_*^2)$. The small-order term thus costs $O(fmb(1 + \log x))$. The dyadic modulus and size bounds · The long-order comparison.

For the flat term, $w_P + \eta_3 \leq 9b'$, $rb' \ll mb$, and $\ell_* r(S_c + L) \leq 2mb$. The actual neighboring-small-letter condition gives $r\eta_3 \mathbf{1}_{\text{band}} \leq 16mb/\ell$ when $\ell > 2C$. It uses $|\ell s - b| \leq C(\ell\eta_3 + \ell S_c + b/\ell + s + \ell L)$ and $C(\eta_3 + S_c + L) \leq w''(n)/8$ as in the signed word calculation. The finitely many $\ell \leq 2C$ cost $O(fr\eta_3) = O(fmb)$. These inequalities and the same elementary sums give $O(fmbL^2)$ for the flat term. Summing the flat and band terms.

Exceptional, small and designated primes. For $\ell > \ell_*$ away from the type primes, the local allowance is zero off the exceptional set $\mathscr{E}$ of the exact local restriction calculation. Take $U = x^{C+2} \exp(x^{\varepsilon/4})$ with fixed sufficiently large C. The family bounds give $m, k, r, f, n, t, L, S_c, \eta_{\text{sz}} = O(U)$ and $D_* = O(U^7)$ by its displayed formula. The exceptional windows give $|\mathscr{E}| = O(U^8)$, using $|T| \leq 2n$ and $b/\ell_* = D_*/m$. For $\ell \in \mathscr{E}$, the ordinary split or remaining inert envelope costs $O(U^9)$ before weighting. The long row is absent since $\ell > \ell_* > A$. We have $b/\ell \leq D_*/m$, $(w_P + \eta_3)/\ell \leq 9b/\ell$, $\log_\ell h_B = O(L)$, and $fr = O(U^2)$. By the neighboring-letter bound just proved, $r\eta_3 \mathbf{1}_{\text{band}} \leq 16mb/\ell \leq 16D_*$. The common polynomial cost is $O(U^7)$. Also $\ell \leq 3b$, so the weighted sum is $O(U^{18}) = o(b)$, because $18(C+2)\log x + 18x^{\varepsilon/4} = o(L)$. The high exceptional envelope · Its weighted sum.

At a divisor of Q_S and at the separate unramified split branch at 3, take the host conductor $N = Q_S qd$, height $H = \lfloor 2b \rfloor$, and $s_N = \max(1, v_\ell(N))$. The cyclotomic lifting exponent is $32(2Hs_N + 1) + \Lambda_N$, with $\Lambda_N = s_N + \lfloor \log_\ell(2H) \rfloor + \lfloor \log_\ell N \rfloor + 1$. The scalar and invariant exponents are bounded by polynomials in the ranks and t times Λ_N. Here $s_N = O(1 + \log x)$ by $Q_S \leq n^2$, including the dyadic completion, and $s_N = 1$ at unramified 3. The conductor size gives $\Lambda_N = O(L)$, and the split dimensions use m per embedding. Hence the fixed-prime cost is $O(fmb \log x + x^C L^2)$. When $3 \mid Q_S$ use the small-set branch. Since $\sum_{\ell \mid Q_S \text{ or } \ell = 3} \log \ell \leq \log(3Q_S)$, the weighted cost is $O(fmbL^2)$. Cyclotomic exponents · The dyadic data · Unramified three · Weighting the small set.

At $\ell = q, d$ the quadratic ramified calculation gives

$$R_\ell \ll f\{r(b/\ell + S_c + t + \eta_{\text{sz}} + L) + mb\} + (frt)^C L.$$

The signs separate by completion at q. At d one uses the mixed-sign estimate in the same completion. The conductor choices give $\ell > 4tr + 1$, so $r/\ell \leq 1/4$. Also $Q_S qd \leq e^{L/72}$ by $(Q_S qd)^{36} t^{150} \leq b^{1/2}$. Thus $\log \ell \leq L$ at both primes. All the terms other than $frb/\ell + fmb$ are $o(b)$ by the bounds defining U. Their weighted total is $O(fmbL^2)$. At g_0, g_1 the selected prime interval gives $\log g_i = O(L)$, so the separate type allowances $O(x^C \eta_{\text{sz}})$ cost $o(b)$. The quadratic-prime allowance · The rank compared with the quadratic primes · The two logarithmic prime bounds · Weighting the type allowance.

The weighted stage bound. The terms $2(\log \ell)/\ell^2$ and $(\log 2)v_2$ cost a bounded constant. Prime factorization sums the class number and projector terms to $\log h_E + \log c_{\text{row}}$. The group order gives $\log c_{\text{row}} \leq fr^2(\log g_0 + \log g_1)$. Minkowski gives an integral ideal of norm at most $(4/\pi)^f |\operatorname{disc} E|^{1/2}$ in every class. The count of these ideals is at most the square of this bound times $\zeta_E(2) \leq \zeta(2)^{2f}$. Thus $\log h_E \ll f + \log|\operatorname{disc} E|$. By the abelian conductor-discriminant formula, $\log|\operatorname{disc} E| \leq 2f \log(Q_S qd)$. Both logarithmic costs are therefore $o(b)$.

Apply these estimates to (L1) on the rational primes in the chosen support. Using the nonnegative fine Selmer sum V gives one constant C_N and one threshold for the family with

$$\eta_j \le D_0\big[C_N(fmbL^2 + (\log x)^2 N_c) + V\big].$$

The term $(\log x)^2 N_c$ is a nonnegative enlargement of the bound just proved. Since $B_0 = fmbL$, $\sqrt{b} \le B_0$, and $(\log x)^2 \le L$ for large x, this yields

$$\eta_j \le D_0(|C_N| + 1)x^{\epsilon}(B_0 + N_c). \tag{91}$$

All estimates use the stage and cutoff of the actual tuple, so summing the normalized bound over the stages costs at most a further factor m. The weighted estimate before amplification · Amplification at the common cutoff.

This weighted bound is applied below through the integral class reading on the actual auxiliary Gram. The equality of the small-level ranks and the enumeration of the local factors use Section 6.1. The absolute covolume and the transport indices use the geometric section lattices in the same fitted coordinates.

4.1.2 The common fitted row map

Use the class-coordinate inclusion and the prescribed rational basis of the adjoint image at P_0 in (41). In these coordinates $P_{\boldsymbol{\mu}}^{\mathrm{fit}}$ is its projection $P_{\boldsymbol{\mu}}$ restricted to the P_0-character space. The cut is orthogonal for the class Gram, and the constant $1/|P_0|$ cancels in the matrix formula for the projection. After using the images of the base complementary basis as row coordinates, $D_{\boldsymbol{\mu}}^{\mathrm{fit}}$ is $D_{\boldsymbol{\mu}}$, by uniqueness on those basis vectors. The termwise identity (15) uses the original class values of y before projection and the prescribed coefficient substitution. Reading its terms on each vector in the fitted basis gives the matrix of the rational evaluation $V_{\boldsymbol{\mu}}^{\mathrm{fit}} y$. Thus the normalized principal tests in (16) use the same arguments, including the level character on the initial row. The phase and square normalized by the base tuple give

$$r_{\boldsymbol{\mu}} = (e_{\boldsymbol{\mu}}/e_{\boldsymbol{\mu}_0})q_{\boldsymbol{\mu}}, \qquad \iota_K(q_{\boldsymbol{\mu}}) = c_{\boldsymbol{\mu}}^{\square}/c_{\boldsymbol{\mu}_0}^{\square}.$$

Here $q_{\boldsymbol{\mu}} \in K^{\times}$ follows from the square ratios in (17). Positivity of the base rank and separation establish $\kappa \neq 0$ in Section 2. A coefficient at the base is in the kernel of $V_{\boldsymbol{\mu}_0}^{\mathrm{fit}}$ exactly when all its tests vanish. By separation and injectivity of the row realization, this is exactly the kernel of $P_{\boldsymbol{\mu}_0}^{\mathrm{fit}}$. Both kernels are over K. The base quotient map is surjective and $D_{\boldsymbol{\mu}_0}^{\mathrm{fit}}$ is an isomorphism. Consequently their composition has a right inverse over K. Applying that inverse to the fitted base jet matrix gives the unique matrix Ψ of the normalized combined test in the common row coordinates. At each tuple the same test identity gives

$$\Psi D_{\boldsymbol{\mu}}^{\mathrm{fit}} P_{\boldsymbol{\mu}}^{\mathrm{fit}} = r_{\boldsymbol{\mu}}^{-1} V_{\boldsymbol{\mu}}^{\mathrm{fit}}.$$

If $M_{\boldsymbol{\mu}}^{(0)}$ is the lattice of integral good top class values at P_0, the local base lattice is

$$R_v^{(0)} = D_{\boldsymbol{\mu}(v)}^{\mathrm{fit}} P_{\boldsymbol{\mu}(v)}^{\mathrm{fit}} (M_{\boldsymbol{\mu}(v)}^{(0)})_v. \tag{92}$$

This is the projected image of integral vectors before projection, as in (52). The calculation there applies at the base character level since both projections are restrictions of the principal map. Localization commutes with the matrix image. Hence the local generators at every selected tuple have the same integral initial rows used in the finite kernel calculation. Common row realization. Descent of the common test and kernel calculation.

Write $R = C^{\vee}$ for the row lattice. At a finite place v of the coefficient field K, use the auxiliary twisting tuple $\boldsymbol{\mu}(v)$ of §2. Its top pairing on a test row is the image of $\zeta_{w,\boldsymbol{\mu}(v)} B_{K,w}$ on the corresponding descended row, in the chosen complex embedding. The scalar is a unit at v, so the tested valuation equals that of $B_{K,w}$. The descended integral test rows generate R_v over $\mathcal{O}_{K,v}$. The ultrametric inequality extends the

valuation bound to their span and to the localized source. Let g_{wv} be the sum over stages of the local gain plus the reference exponent minus the exponent for w, with valuations normalized by $v_p(p) = 1$. Gains at the discarded primes are set to zero. The local division estimate applies on the actual source models, subject to the ramified donor identification in Section 6.1. On that source it gives

$$\operatorname{ord}_v B_{K,w}(\phi)(F) \geq \lfloor e_v g_{wv} \rfloor \qquad (\phi \in (N_w)_v,\ F \in R_v). \tag{93}$$

Here N_w is the source before division and $\operatorname{ord}_v(0) = +\infty$. Multiplying the resulting inclusion in $\mathfrak{p}_v^{\lfloor e_v g_{wv} \rfloor} C_v$ by the inverse local ideal and intersecting over the finite places gives

$$I_w = \prod_v \mathfrak{p}_v^{-\lfloor e_v g_{wv} \rfloor}, \qquad N_w' = I_w N_w, \qquad B_{K,w}(N_w') \subseteq C.$$

The support is contained in the places of the local comparison, including those over the level and exceptional primes. The divided source N_w' in (53) inherits the K-space and metrics of N_w, with the prescribed realization in (68).

For a positive-rank Hermitian $\mathcal{O}_K$-lattice J, use the determinant norm at each place. Its unit ball at a finite place is $\det J \otimes_{\mathcal{O}_K} \mathcal{O}_{K_v}$. In the squared-norm convention of §3, the normalized degree and slopes are

$$2\widehat{\deg}_{\mathrm{n}}(J) = -\sum_v \kappa_v \log \|a\|_v^2 \qquad (0 \neq a \in \det(J_K)),$$

$$\mu(J) = \frac{2\widehat{\deg}_{\mathrm{n}}(J)}{\operatorname{rk} J}, \qquad \mu_{\max}(J) = \max_{0 \neq V \subseteq J_K} \mu(J \cap V).$$

Here $J_K = J \otimes_{\mathcal{O}_K} K$, V runs over K-subspaces with the induced metrics, and $\kappa_v = [K_v : \mathbf{Q}_{v_{\mathbf{Q}}}]/[K : \mathbf{Q}]$ for the place $v_{\mathbf{Q}}$ below v. The absolute values extend the usual ones on $\mathbf{Q}$. The product formula makes the degree independent of a. [52, (36) and § 2.4] The determinant of N_w' is $I_w^{\operatorname{rk} N_w} \det(N_w)$ with the original norms at infinity, so the degree formula gives

$$\mu(N_w') - \mu(N_w) = -\frac{2 \log \mathrm{N}(I_w)}{[K : \mathbf{Q}]} = 2 \sum_v \lambda_v \lfloor e_v g_{wv} \rfloor.$$

Tensoring by I_w pairs saturated sublattices over every K-subspace, so the maximal slope has the same shift. Rounding costs at most $2 \sum_{g_{wv} \neq 0} \lambda_v$, and $\sum_{v|p} 2 e_v \lambda_v = 2 \log p$.

4.2 Valuation gains and Frobenius classes

At an unramified prime with contributing Frobenius class, the mixed factor in the adapted Cartier coordinates at odd Frobenius distance h is $\exp(-p^{-e_h} X^{p^h} Y)$, with $e_h = 1 + (p^h - p)/(p^2 - 1)$. The derivative of order $(p^h r, r)$ at the origin is $(-p^{-e_h})^r (p^h r)!$. Legendre's formula gives $v_p(n!) = \sum_{i \geq 1} \lfloor n/p^i \rfloor = (n - s_p(n))/(p-1)$, where $s_p(n)$ is the base-p digit sum. Since $s_p(p^h r) = s_p(r)$, the valuation of the jet is

$$r\left((p^h + 1)\frac{p}{p^2 - 1} - 1\right) - \frac{s_p(r)}{p - 1}.$$

The mixed factor lowers the baseline $r(p^h + 1)p/(p^2 - 1)$ by r before the digit correction. Let $Z = (z_{ab})$ be an $r_0 \times r_0$ matrix of independent variables over a field of characteristic zero, where $r_0 \geq 1$, and put $\partial_Z = (\partial/\partial z_{ab})$. For every integer $s \geq 1$ the ordinary Cayley identity gives

$$\det(\partial_Z)(\det Z)^s = \left(\prod_{i=0}^{r_0 - 1} (s + i)\right)(\det Z)^{s-1}.$$

[58, Theorem 2.1, (2.1)] Iteration gives the constant $C_{r_0}(s) = \prod_{i=0}^{r_0 - 1} (s + i)!/i!$ for $\det(\partial_Z)^s (\det Z)^s$.

AI-proposed candidate (a short interval for perturbed determinant jets). Let $p \geq 5$ be prime, let $h \geq 1$ be odd, and let K_0 be a characteristic-zero field with an additive nonarchimedean valuation $v : K_0 \to \mathbf{R} \cup \{+\infty\}$ normalized by $v(p) = 1$. Write $q_h = p^h$, $\rho_p = 1/(p^2-1)$, $\gamma_p = p\rho_p$, and $c_h = -p^{-e_h}$ with $e_h = 1 + (q_h - p)\rho_p$. Use two independent $r_0 \times r_0$ matrices X, Y of variables, where $n_* \geq 1$ and $0 \leq r_0 \leq n_*$ are integers. For $U = \sum_{A,B} U_{A,B} X^A Y^B \in K_0[[X,Y]]$, let $|A|, |B|$ denote total degrees. Put $r_1 = \rho_p + 1/(2q_h)$, $r_2 = \rho_p + 1/2$, and $\delta = 1/(8q_h)$. The coefficient hypotheses are

$$v(U_{A,B}) + (r_1 - \delta)|A| + (r_2 - \delta)|B| \geq \frac{1}{p-1} \qquad \text{for all } A, B,$$
$$U_{A,B} \neq 0 \implies |B| \geq 1.$$

They give a well-defined formal exponential in the jet

$$\Lambda_y = \det(\partial_X)^{q_h y} \det(\partial_Y)^y \exp\left(c_h \sum_{a,b} X_{ab}^{q_h} Y_{ab} + U(X,Y) \right)\Bigg|_{X=Y=0}.$$

These are ordinary derivatives, so $\partial_X^A \partial_Y^B$ at the origin extracts $A!B!$ times the coefficient of $X^A Y^B$. For $r_0 \geq 1$ and every nonnegative integer Y_0, there is an integer y with

$$\begin{aligned} Y_0 \leq y < Y_0 + S_0, \qquad S_0 &= \left\lceil 64 r_0^2 \big(h + 3 + 3\log_p(r_0+1)\big) \right\rceil, \\ v(\Lambda_y) &\leq r_0 y\big((q_h+1)\gamma_p - 1\big) + n_*(\log_p n_* + 1). \end{aligned} \tag{94}$$

For $r_0 = 0$ one has $\Lambda_y = 1$ for every $y \geq 0$. Statement in the release.

To see the uniformity in Y_0, write $D = \det X$ and use the factorial pairing $\langle P, Q\rangle = P(\partial_X)Q(0)$. With $X^{[q_h]} = (X_{ab}^{q_h})$ and an independent matrix Z, put $\Psi(X,Z) = \det(X^{[q_h]} - p^{e_h} Z) - D^{q_h}$. Let Q_s be the part of degree $r_0 q_h s$ in X of $\Psi(X, \partial_Y)^s \exp(U)$ evaluated at $Y = 0$. Conjugating the Y-derivatives by the mixed exponential gives $c_h^{r_0 y}(D^{q_h} + \Psi(X, \partial_Y))^y$. Adjointness of multiplication by D and differentiation by $\det(\partial_X)$, followed by Cayley's identity, gives

$$\frac{\Lambda_y}{c_h^{r_0 y} C_{r_0}(q_h y)} = \sum_{s=0}^{y} \binom{y}{s} a_s, \qquad a_s = \frac{\langle D^{q_h s}, Q_s \rangle}{C_{r_0}(q_h s)}, \qquad a_0 = 1.$$

Binomial expansion.

A term of Z-degree k_0 in Ψ^s has X-degree $q_h(r_0 s - k_0)$, and its coefficient is divisible by $p^{e_h k_0 + (s-k_0)_+}$. Each Z contributes p^{e_h}, and each factor free of Z is divisible by p by $\det(X^{[q_h]}) \equiv D^{q_h} \pmod p$. The weighted valuation of every coefficient of $\exp(U)$ is nonnegative by the coefficient bound on U and $v_p(j!) \leq j/(p-1)$. To form Q_s, a term with Z-degree k_0 extracts bidegree $(q_h k_0, k_0)$ and introduces a right factorial of valuation at least $k_0/(p-1) - r_0^2(\log_p(r_0 s) + 1)$ for $s \geq 1$. Since $q_h(r_1 - \delta) + (r_2 - \delta) = e_h + 1/(p-1) - (1 + q_h^{-1})/8$, the valuation of a coefficient of Q_s is at least $s/8 - r_0^2(\log_p(r_0 s) + 1)$. Here $(s-k_0)_+ + k_0(1 + q_h^{-1})/8 \geq s/8$. The digit-sum formula for $C_{r_0}(m_0)$ and for the left factorials in the pairing gives

$$\begin{aligned} v_p(C_{r_0}(m_0)) &\leq \frac{r_0 m_0}{p-1} + r_0(\log_p r_0 + 1), \\ v(a_s) &\geq \frac{s}{8} - r_0^2\big(3\log_p r_0 + 2\log_p s + h + 3\big) \qquad (s \geq 1). \end{aligned}$$

Coefficient bounds and factorial quotient estimate.

This last lower bound is positive for $s \geq S_0$. For the first S_0 binomial polynomials, successive row differences give determinant 1 for their evaluation matrix on any S_0 consecutive nonnegative integers. If all the normalized values in the window had positive valuation, the truncated sums would have positive valuation as well. The integral inverse of the evaluation matrix would give $v(a_0) > 0$. Thus one normalized value has valuation at most zero. The bound on $C_{r_0}(q_h y)$ and $q_h/(p-1) - e_h = (q_h+1)\gamma_p - 1$ give (94). Window argument.

Cayley's identity supplies the determinant factorial. For one-variable exponentials under truncated Dwork congruences, Krattenthaler and Müller obtain a congruence at step p^{ℓ} for normalized ordinary derivatives, giving a sharp valuation at multiples of p^{ℓ} under an additional unit condition. Here the bound on the Newton coefficients uses the weighted matrix perturbation and gives an interval of length S_0 in y. [55, Theorem 2, (2.4)–(2.6)]

The local row-depth proof uses (94) as an upper valuation bound at an available order in each odd block. The division estimate (93) supplies a lower valuation bound for every integral evaluation of the period maps. Row-block test.

In the continuous calculation on an even Frobenius cycle with degrees D_s, an allocation $z_{s,h} \geq 0$ of mixed terms at odd distances h satisfies $\sum_h p^h z_{s,h} + \sum_h z_{s-h,h} \leq D_s$ at s. Its saving relative to the baseline is $\sum_{s,h} z_{s,h}$. On a two-cycle with degrees A, B, the constraints $pz + z' \leq A$ and $z + pz' \leq B$ give $z + z' \leq \min\{(A+B)/(p+1), A, B\}$. Equality holds at the intersection if $A \leq pB$ and $B \leq pA$, and on the corresponding axis otherwise. The remaining cost is

$$\frac{A+B}{p^2-1} + \frac{(A-pB)_+ + (B-pA)_+}{p+1}.$$

On a longer cycle the source degree a_t supplies at most a_t/p^h mixed factors at distance h in this continuous calculation. Subtracting the degree at the receiving position and taking the positive part gives a term of the form $(a_t/p^h - w)_+$. The prescribed first gain sums such positive parts at selected odd distances with shifted source and receiving degrees and a local error. On selected orbits the final prescription takes a maximum with a second gain expression. The integral orbit plans account for competing uses of a source degree, rounding to usable orders and the contributions of the remaining directions. At an auxiliary prime $g_i = \pi_i \overline{\pi_i}$, the local jet identities and integrality condition specify the effect of isogeny pullback and translation on the harmonic jet. The upstairs jet agrees with the translated downstairs jet times a root of unity and $\prod_\tau \tau(\overline{\pi_i})^{d_\tau}$. The product times the downstairs jet is integral for the prescribed integral sections and tests when $1 + \sum_\tau d_\tau < g_i$, which bounds the downstairs denominator.

Use the Artin coordinates (3) on the Galois group Γ of the base field. Set $P = \{(0, i, s) : i \in \mathbf{Z}/2\mathbf{Z},\ s \in H\}$. For $t \in T$ the translation $p_t(g) = g + c_\infty + (0,0,t) = g + (1,1,t)$ maps $\Gamma \setminus P$ bijectively to P. Write $F_p + I_p$ for the arithmetic Frobenius coset at p, where I_p is inertia. For $a > 0$ and weights $w_r > 0$ on P define

$$A(a,t,w) = \sum_{g \notin P} \sum_{\substack{p \text{ prime} \\ h \geq 1}} \frac{2\log p}{|I_p|} \mathbf{1}_{\{g - hF_p \in I_p\}} (a/p^h - w_{p_t(g)})_+. \tag{95}$$

At an unramified prime with $F_p \notin P$, only odd powers contribute and each has weight $2\log p$. The sum vanishes at unramified primes with $F_p \in P$ and at primes dividing $Q_S q$, where the entire Frobenius coset is in P. At d the inertia is $I_d = \{0, (1,0,0)\}$, so each coset $hF_d + I_d$ with $h \geq 1$ has exactly one element outside P. The analytic sum includes a term of weight $\log d$ at every positive power. These support assertions use the base field's Frobenius and inertia properties. The contribution at d is bounded by a geometric series. The sample-size condition $|a_t/b - x| \leq 1$ from Section 5 gives $a_t \leq 3a_E/2$ at $x \geq 2$. The positive parts are at most a_t/d^h, so the sum over all samples and stages at d costs at most $2mN_T a_E \log d/d$, using $d \geq 4$. Geometric bound at d.

The arithmetic comparison uses the same upper limit

$$N = \lceil 2a_E \rceil. \tag{96}$$

For every u, j, $t \in T$ and $r \in P$ we have $0 < a_t \leq 2a_E$ and $w_{u,j}(r) \geq 1$. Thus $a_t/w_{u,j}(r) \leq N$ and every term with $p^h > N$ vanishes, including those with $p = d$. The finite sum used in both the arithmetic comparison and the Fourier expansion therefore equals (95). The finite character expansion uses this support condition.

AI-proposed candidate (Frobenius-weighted divisibility with uniform loss). Fix the target, precision and comparison constants. At every sufficiently large scale with the tied data, integral models, and local, rank, level and size hypotheses above, the prescribed exponents g_{wv} satisfy (93) for all integral pairs, places and words at every tied translate u. On the same data write g_{jv} for the retained stage gains, τ_v for the local loss from restricting the sources, and ι_v for the local index allowance. Their weighted sums are

$$\mathcal{G}(u) = 2\sum_v e_v\lambda_v\Big(\sum_{j=1}^{m} g_{jv} - \tau_v\Big), \qquad \mathcal{I}(u) = \sum_v e_v\lambda_v\iota_v.$$

The restriction loss is charged once per place. The gain and index comparison gives a constant C_{loss}, uniform in x and u, with

$$\max\left\{\sum_{t\in T}\sum_{j=1}^{m} A(a_t, t, w_{u,j}) - \mathcal{G}(u),\ \mathcal{I}(u)\right\} \le C_{\mathrm{loss}} N_T m b x^{e_L\varepsilon}. \tag{97}$$

At the unramified places served by the local comparison, its local inequality bounds the sum of the odd-power ramps for the served samples, after local errors, by $\sum_j g_{jv} - \tau_v$. The aggregate bound pays for unserved samples, discarded primes and the geometric term at d. The rounding of the local exponents and the index and level costs enter the slope error $R_{\mathrm{loc}}(w)$ below. Bost obtains improved local divisibility for commuting integral vector fields over an absolutely unramified local field when their reductions have vanishing p-th powers. His filtered slope inequalities use the heights of evaluation maps. The local target here consists of the prescribed Frobenius-weighted gains on the integral period maps, with the joint loss bounds (97). [54, Props. 3.9(2), 4.5–4.6]

4.3 Degrees and slopes of the source lattices

The reference lattice M_w is the definite column lattice at the initial stage. Write Δ_{cl} for the auxiliary weight size Sc and $a_{w,t}^{\mathrm{top}}$ for the top column weight coordinates.

Evaluation on the definite classes. Put $T = T_w$ and let L_c be the rank-k Hermitian lattice of a decorated class c at the initial stage. A finite adelic class representative acts by a unitary isometry on the prescribed local lattices. They are self-dual away from T, and at T the discriminant module of the prescribed lattice is $\mathcal{O}_E/T\mathcal{O}_E$. Thus $L_c \subset L_c^{\#}$ and $[L_c^{\#} : L_c] = T^{2f}$, where # denotes Hermitian duality and $[E : \mathbf{Q}] = 2f$. Discriminant module on the class. Its positive definite trace form makes its automorphism group finite. At a place of the type prime g_0, an element stabilizing the initial decoration reduces to the upper unitriangular subgroup on the whole rank k. The prime-to-g_0 part of such an element is trivial. Indeed reduction of roots of unity of that order is injective and all its eigenvalues reduce to one. An element of order g_0 acting on the underlying rational vector space would require at least $g_0 - 1$ dimensions. Here $g_0 - 1 > 2fk$, hence the stabilizer is trivial. The initial decoration at the type prime.

Evaluation on double-coset representatives consequently gives the full sum of coefficient spaces, one copy per class. The inverse extends a value by the real type and the prescribed level character; independence of the extension follows from the stabilizer calculation. At a finite place, changing an integral basis of L_c by a matrix in $\mathrm{GL}_k(\mathcal{O}_v)$ acts by the integral Weyl representation, including its determinant character. A local basis with column matrix Y for L_c has dual basis $(\bar{Y}\,{}^t H)^{-1}$ for the Hermitian pairing of matrix H. In the fitted coordinates, let M_b be the output frame and ρ_τ the indicated conjugate embedding. The coefficient change is $g_c = M_b\rho_\tau(H)^{-1}$ and the local lattice frame is $\gamma_c = M_b\rho_\tau(Y)$. Then $(\gamma_c^{-1}g_c)^t$ at the matching complex embedding is $(\bar{Y}\,{}^t H)^{-1}$, by Hermitian symmetry. The induced action includes $\det^{s_\tau}$. This gives the equivalence between the prescribed integrality tests in the polynomial and dual frames. It identifies the local lattice with the tensor of the integral Weyl lattices and determinant characters. Over each complex embedding use a unitary frame for H. The initial column norm is then the orthogonal sum

of the normalized Fischer norms, with stabilizer weight one at each class. This identifies the reference lattice at its finite and infinite places with

$$\bigoplus_c \bigotimes_\tau \left(\mathrm{W}_{\nu_\tau}(L_{c,\tau}) \otimes (\det L_{c,\tau})^{s_\tau}\right), \qquad s_\tau = \lambda_{\tau,k}, \quad \nu_{\tau,i} = \lambda_{\tau,i} - s_\tau.$$

Here $L_{c,\tau}$ is the extension of L_c to the coefficient field through τ. The last exponent can be negative. Comparison of frames and Class evaluation and metrics.

Two bounds for the integral Weyl lattice. Use the squared-norm convention for degrees. For the standard rank k lattice write $d = \sum_i \nu_i$. Divide the Fischer inner product by the squared length of the product p_ν of leading minors, so this integral vector has norm one. This is the normalization of the class metric. Apply to p_ν the generic unipotent matrix opposite the stabilizer of its highest weight line, with entries z_{ij} above or below the diagonal. The coefficients as polynomials in z belong to the integral Weyl lattice. Indeed, triangular elimination factors this matrix into elementary unipotents over $\mathbf{Z}[z]$, whose coefficients act by the integral divided row operators. These vectors span the rational representation, since multiplying the opposite unipotent subgroup by the stabilizing Borel gives the open cell and the highest weight vector generates the representation. On the unit polydisc the matrix has operator norm at most k. Its action on degree d polynomials in the Fischer norm has norm at most k^d, by the symmetric tensor realization. The same bound holds for each coefficient by Cauchy's formula. Choosing independent coefficients and applying the determinant inequality gives the standard mean slope $\mu_{\mathrm{std}}(\nu) \geq -2d \log k$.

We also need a bound for the norms of the integral linear forms that read off polynomial coefficients. If X^α is a monomial, its column degrees are ν_j; the multinomial inequality gives $\prod_j \nu_j!/\alpha! \leq k^d$. In the ordinary unnormalized Fischer norm $\|p_\nu\|_F^2 = \sum_\alpha \alpha! |p_{\nu,\alpha}|^2$. The sum of the absolute coefficients of a minor of order i is at most $i! \leq k^i$. It follows that $\|p_\nu\|_F^2 \leq k^{2d} \prod_j \nu_j!$. Consequently reading any coefficient has squared norm at most k^{3d} on the normalized representation.

For the trace metric on L_c every nonzero vector has squared length at least $2f$, by the arithmetic–geometric mean inequality for its totally positive integral norm. In real dimension $2fk$ the Euclidean dual therefore has $2fk$ independent vectors of length $O(k\sqrt{f})$, by Banaszczyk's transference bound [53, Theorem (2.2)]. Here one may use $\lambda_{\mathrm{last}}(\Gamma) \leq 2\mu_{\mathrm{cov}}(\Gamma)$: if the lattice vectors of length at most $2\mu_{\mathrm{cov}} + \delta$ spanned a proper subspace, a point at distance $\mu_{\mathrm{cov}} + \delta/2$ along a unit normal would have no lattice point within covering distance. The trace dual is $\mathfrak{D}_{E/\mathbf{Q}}^{-1} L_c^\#$. Minkowski's convex body theorem on a product of complex discs gives a nonzero element of $\mathfrak{D}_{E/\mathbf{Q}}$ of absolute value $O(\mathrm{rd}(E)^{3/2})$ at every embedding: the covolume of the different in $\mathbf{C}^f$ is $2^{-f}|\Delta_E|^{3/2}$. Multiplication and selection of k linearly independent vectors over E give an integral map $A : L_c \to \mathcal{O}_E^k$ of full rank whose operator norms at infinity are bounded by $B = Ck^{3/2}\sqrt{f}\ \mathrm{rd}(E)^{3/2}$. The coordinates of A are the Hermitian linear forms of these vectors.

Apply the polynomial action of A and read its coefficients. This action is integral at every finite place: Smith decomposition uses two invertible integral changes of frame and a diagonal matrix with integral entries, and the untwisted Weyl weights are nonnegative. These integral forms span the dual over the field, since A has full rank. Their norms on $\mathrm{W}_\nu(L_c)$ are at most $B^d k^{3d/2}$. Take the tensor products of these forms at the chosen embeddings. On any rational subspace one can choose as many of them as its dimension with independent restrictions. The resulting integral determinant and the product formula bound the slope on that subspace by $\sum_\tau d_\tau(2\log B + 3\log k)$. Tensoring by a metrized line adds its degree to every slope. In an orthogonal sum the maximal slope is at most the maximum over the summands: intersect a subspace with the first summand and project onto the rest; the quotient norm dominates the projected norm.

Since the Hermitian dual has opposite degree and its index over L_c is T^{2f}, the squared normalized degree of $\det L_c$ is $-\log T$. Normalized degree is preserved by base change. If the polynomial represen-

tation has rank R, its determinant character is $(\det)^{dR/k}$: restrict to the scalar matrices to determine the exponent. Comparing determinant lines in the local and unitary frames therefore gives, with $d_\tau = \sum_i \nu_{\tau,i}$,

$$\mu(M_w) \geq -\frac{\log T}{k} \sum_{\tau,i} \lambda_{\tau,i} - 2\log k \sum_\tau d_\tau,$$
$$\mu_{\max}(M_w) \leq -\log T \sum_\tau s_\tau + (2\log B + 3\log k) \sum_\tau d_\tau.$$

At a sampled embedding the top entry is $a^{\mathrm{top}}_{w,t}$, the last is between $-2b'$ and $-b'$, and the intervening entries are between zero and Δ_{cl}. At the other embeddings all entries are nonnegative and $O(\Delta_{\mathrm{cl}})$. Consequently $\sum d_\tau \ll N_T a_E + fk(b' + \Delta_{\mathrm{cl}})$, $-\sum s_\tau \leq 2N_T b'$, and $\sum_{\tau,i} \lambda_{\tau,i} \leq \sum_t a^{\mathrm{top}}_{w,t} + O(fk\Delta_{\mathrm{cl}})$. Since $k \geq x$ and $f \leq 2x^3$, $\log B \ll \log(kv_0)$ with $v_0 = \mathrm{rd}(E)/2$. This proves the two reference bounds in the actual class norms.

In the actual size range $a^{\mathrm{top}}_{w,t} \leq 2a_E$, the two column estimates give

$$\begin{aligned}\mu(M_w) \geq\ & -\sum_t \frac{a^{\mathrm{top}}_{w,t}}{k} \log T_w \\ & - O\big(N_T a_E \log k + fk(b' + \Delta_{\mathrm{cl}}) \log k + f\Delta_{\mathrm{cl}} \log T_w + fk^2 L\big), \\ \mu_{\max}(M_w) \ll\ & N_T a_E \log(kv_0) + f(b' + \Delta_{\mathrm{cl}})\big(k \log(kv_0) + \log T_w\big) + fk^2 L.\end{aligned}$$

Here $v_0 = \mathrm{rd}(E)/2 \geq 1$ and $L_a = \log(a_E \,\mathrm{rd}(E))$. For large x, $f \leq 2x^3$, $k \leq x^5$, and $X' = \exp(x^{\varepsilon/3})$ is at least x^{14}. Since $\Delta_{\mathrm{cl}} \leq b' = a_E/X'$ and $X' \leq a_E$, it follows that $fk(b' + \Delta_{\mathrm{cl}}) \leq N_T a_E$ and $fk^2 \leq N_T a_E$. Also $\log T_w \ll mL$, $\log k \leq 5L_a$, and $\log(kv_0) \leq 6L_a$. The first term of the lower bound is therefore $O(N_T a_E L_a m/k)$, and all other displayed costs are $O(N_T a_E L_a)$ since $m \leq k$. Consequently, for a uniform constant C_{sl},

$$\mu(M_w) \geq -C_{\mathrm{sl}} N_T a_E L_a (1 + m/k), \qquad \mu_{\max}(M_w) \leq C_{\mathrm{sl}} N_T a_E L_a.$$

Column estimates with their level dependence and Specialization on the family.

The division comparison also uses cusp lattices on the indefinite side of the theta tower. We first state the finite-degree comparison and then apply it to those lattices.

AI-proposed candidate (finite-degree bounds for adjoint section lattices). On a smooth projective component Z of flag dimension D_{flag}, the geometric input uses a totally complex coefficient field, a base-point-free birational moving system on L, and two nef and big line bundles L_A, L_2. The lattices are in the prescribed adjoint section spaces for $K_Z \otimes L_A \otimes L^j$. The comparison also uses the adjoint twist $K_Z \otimes L_2 \otimes L^{j+D_{\mathrm{flag}}}$. For $j < \ell$, multiplication by an integral moving section s takes the degree-j lattice into the next and satisfies $\|sf\|_{j+1,\sigma} \leq \sqrt{\beta}\,\|s\|_\sigma \|f\|_{j,\sigma}$ in the chosen metrics at each embedding σ. The lattices at degrees $0, D_{\mathrm{flag}}, \ell$ contain bases of their rational spaces with prescribed norm bounds. Geometric hypotheses.

Let d_Z be the degree of the moving system, r_Z the cusp-form lattice rank and ℓ the truncation degree, with $D_{\mathrm{flag}} \geq 1$ and $\ell \geq 3D_{\mathrm{flag}} + 5$. The comparison of degrees and ranks gives

$$(1 - z) r_Z \leq q_Z := \binom{\ell - 1}{D_{\mathrm{flag}}} d_Z \leq r_Z, \qquad z = 7\big(D^2_{\mathrm{flag}}/\ell + D_{\mathrm{flag}}\rho\big) \leq \tfrac{1}{2},$$

where $0 \leq \rho < 1$ and $0 < d_* \leq d_Z \leq d_*/(1 - \rho)$. Let h_Z be the Chow height in the moving frame divided by $(D_{\mathrm{flag}} + 1)d_Z$. Under the cover's quantitative bounds for the frames, multiplication, and residue terms, the finite-degree estimate bounds the normalized Arakelov degree of the target cusp lattice J_Z from below at the specified ℓ.

$$\widehat{\deg}_{\mathrm{n}}(J_Z) \geq q_Z\big[\ell h_Z - 3(D_{\mathrm{flag}} + 1)|h_Z| - A_{\mathrm{HS}} \log(2 + |h_Z|) - R_{\mathrm{HS}}\big]. \tag{98}$$

Here $A_{\mathrm{HS}} = 26(26D^2_{\mathrm{flag}} + 2z\ell)$ and $R_{\mathrm{HS}} = 4000b + 2c_{\mathrm{cmp}}$, where $r_Z c_{\mathrm{cmp}}$ bounds the degree loss when the modified cusp lattice is compared with the target lattice.

David and Philippon bound the arithmetic Hilbert function of a homogeneous coordinate quotient by using minors divisible by a Chow form. For an image of dimension D_{flag}, degree d_Z, and ambient projective dimension N_{amb}, Chen's slope formulation uses the quotient of the ℓ-th symmetric power of the ambient Hermitian bundle by the restriction kernel, and gives a bound for $\ell \geq (N_{\mathrm{amb}} - D_{\mathrm{flag}})(d_Z - 1) + 1$. The adjoint comparison controls the share q_Z/r_Z and the metric losses in the actual section lattices. Its hypotheses allow ℓ below this sufficient degree condition. [56, Prop. 4.10] [57, Thm. 4.8]

For a smooth projective variety X over an algebraically closed field of characteristic zero and a nef and big line bundle L, Kawamata–Viehweg vanishing gives

$$H^q(X, K_X \otimes L) = 0 \qquad (q > 0),$$

where K_X is the canonical bundle. For nef L, bigness is equivalent to positive top self-intersection. With zero boundary, the vanishing of $H^i(X, L^{-1})$ for $i < \dim X$ gives the displayed form by Serre duality. [59, Lemma 5.8(a) and Corollary 5.12(c), with zero divisor and Serre duality] The section-space calculation uses this on the geometric generic fibres of the smooth models, with the adjoint twists whose remaining factors are nef and big in the release's comparison.

The release's package for the component estimates includes an auxiliary L^2 extension estimate. For its compact case let (X, ω) be a compact Kähler manifold of dimension $D \geq 1$. It is weakly pseudoconvex, with a constant plurisubharmonic exhaustion. Let L and $\mathcal{B}$ be holomorphic line bundles with smooth Hermitian metrics. Suppose holomorphic sections $s_1, \dots, s_D$ of L have a finite common zero set S, with invertible derivative of $s = (s_i)_i$ at each zero, and $\sum_i |s_i|_L^2 \leq 1$ on X. Put $\Psi = D \log(\sum_i |s_i|_L^2)$, with value $-\infty$ at S. The transverse coordinates give neat analytic and log-canonical singularities along the reduced set S. Assume that for some $\delta > 0$,

$$i\Theta_{\mathcal{B}} + \alpha i\partial\bar\partial\Psi \;\geq\; 0 \qquad \text{as currents for every } \alpha \in [1, 1+\delta].$$

Write $dV_\omega = \omega^D/D!$ and use the norm $|\cdot|_{\omega,\mathcal{B}}$ on $K_X \otimes \mathcal{B}$ induced by ω and the metric on $\mathcal{B}$. The residue measure ν_Ψ on S assigns to $P \in S$ the mass

$$\nu_\Psi(\{P\}) = \limsup_{t \to -\infty} \int_{\{t < \Psi < t+1\} \cap U_P} e^{-\Psi}\, dV_\omega,$$

where U_P is a small coordinate neighborhood with no other zero in its closure. Transversality makes each mass finite. Demailly's extension theorem gives, for every $f \in H^0(S, (K_X \otimes \mathcal{B})|_S)$, a holomorphic section $F \in H^0(X, K_X \otimes \mathcal{B})$ with $F|_S = f$ and

$$\int_{X \setminus S} \frac{|F|_{\omega,\mathcal{B}}^2 e^{-\Psi}}{1 + \delta^2 \Psi^2}\, dV_\omega \;\leq\; \frac{34}{\delta} \int_S |f|_{\omega,\mathcal{B}}^2\, d\nu_\Psi.$$

[60, (2.4)–(2.5) and (2.7)–(2.8)] The product $|f|_{\omega,\mathcal{B}}^2\, d\nu_\Psi$ at each zero is independent of ω. [60, Remark (2.9)(a)] The release's chart version obtains the Kähler form from a positively curved auxiliary line bundle and imposes the curvature inequalities off S. The local plurisubharmonic potentials are bounded above near S, so the current inequality extends across S. The chart estimate uses an upper bound for the residue integral in terms of the metric weights and the derivatives of the s_i at their common zeros.

For the component degree bound, the residue filtration uses the algebraic vanishing above and the regular sequence cutting out the zero scheme to lift its sections. Those lifts and the short vectors give a nonzero polynomial determinant divisible by the Chow form raised to the power $\sum_{0 \leq a < \ell - D_{\mathrm{flag}}} \binom{a + D_{\mathrm{flag}} - 1}{D_{\mathrm{flag}} - 1} (\ell - a)$. Integral multiplication and the norm bounds control the determinant in the prescribed lattice. Determinant and norm comparison.

The component estimate controls the lattice of cusp forms on the indefinite side of the theta tower. The division theorem uses its contribution through a stage covolume bound. Put

$$E_{\mathrm{cov}} = B_0 \log\log x, \qquad M_j = (p_j - q_j) \sum_{\rho \in P} w_{u,j}(\rho) \log \Omega, \qquad (p_j, q_j) = (k + j - 1, m - j).$$

Here $\log \Omega$ is the period normalization computed below. We work at sufficiently large x, so $E_{\mathrm{cov}} \geq b$. The signature (p_j, q_j) belongs to the auxiliary model used to compare stage covolumes to the linked division lattice. If J_j^{row} denotes the stage lattice in its prescribed comparison frame and d_{row} the row rank used in the division theorem, the stage bound is

$$\frac{\mathrm{lcv}(J_j^{\mathrm{row}})}{d_{\mathrm{row}}} \leq M_j + C_3 E_{\mathrm{cov}}. \tag{99}$$

The log covolume here uses the frame of the indefinite row lattice. Stage covolume and main term and error scale.

4.3.1 The cover and its moving systems

For $j < m$ both entries of this auxiliary signature are positive. At the given bundle point take a deep cover of its unitary moduli space, and use its smooth flag compactification on each geometric component Z. The cover existence proof chooses an auxiliary split level M_{cov} with the required coprimality and passes to a coefficient field allowing descent of the components and the specified integral models. Good reduction and flatness are used for the section lattices. The degree comparison uses the top Chern number of the flag line bundle, proportionality on the components, and multiplicity bounds at the modified weight. Choice of cover, integral models, flat degree, and component degree comparison.

In this construction $r = p_j + q_j$, and the flag dimension $D = D_{\mathrm{flag}} = f\binom{r}{2}$ is the sum of the base dimension $f p_j q_j$ and the fibre dimension $f(\binom{p_j}{2}+\binom{q_j}{2})$. Let P_{cov} be the parameter of the cover construction, and let κ be the sum of absolute values of the base weight coordinates at this stage. The quotient-weight construction uses

$$\ell = \left\lceil P_{\mathrm{cov}}^2 D^2 k/m \right\rceil .$$

Its integral weight path has a moving weight u_1 and auxiliary weight u_2. On a smooth toroidal flag component X_0 write $H(a)$ for the canonical line of weight a and B for the reduced boundary, so cusp sections have line $H(a)(-B)$. The canonical weight κ_{can} has line $K_{X_0}(B)$. Let $v = u_W - c_0 h$ be the lowered weight, with $h = 0$ on the first signature block and $h = -1$ on the last. Divide each successive gap of $v - \kappa_{\mathrm{can}}$ by $n = \ell + D + 1$. With quotient and remainder arrays u_1 and R_1 having coordinate zero at the anchor, put $u_2 = u_1 + R_1$. The weight $a_s = \kappa_{\mathrm{can}} + u_2 + (D + s)u_1$ at $s = \ell$ differs from v by a constant weight at each embedding. This is a flat determinant character on the component. Indeed at principal level $M \geq 3$ an arithmetic determinant ζ is a unit with all conjugates of absolute value one and $\zeta \equiv 1 \pmod{M}$. If $\zeta \neq 1$, its difference from one would have absolute norm at least $M^{[E:\mathbb{Q}]}$ and at most $2^{[E:\mathbb{Q}]}$. The determinant section extends across the unipotent boundary, so the two cusp spaces have the same rank r_Z. Degree and quotient weight, the weight path, and the flat twist.

For generation on the open flag stratum U, let A_{flag} have all successive gaps one. The base logarithmic canonical line has weight $(q_j, \dots, q_j, -p_j, \dots, -p_j)$, by taking the determinant of the cotangent representation. This differs from rh by a trivial determinant character at the deep level. Thus a positive power of $H(h)$ pulls back an ample line from the minimal compactification. The bundle $A_{\mathrm{flag}} - H(h)$, in additive notation, is nef by the gap criterion and relatively ample on flag fibres. If a d-dimensional irreducible subvariety Y meets U and its image on the minimal base has dimension e, its mixed intersection with e factors of $H(h)$ and $d-e$ factors of $A_{\mathrm{flag}} - H(h)$ is positive by the projection formula and the degree on the generic fibre. The other mixed intersections are nonnegative, so $(A_{\mathrm{flag}}^d \cdot Y) \geq 1$. Mumford, Proposition 3.4.

Set $C_{\mathrm{flag}} = D(D + 1) + 1$. The selected gaps of u_i are at least $(D + 3)(\binom{D+1}{2} + 4r + D + 3)$, and the greatest canonical gap is $2r - 2$. Thus $H_i := H(u_i - \kappa_{\mathrm{can}})$ satisfies $H_i - C_{\mathrm{flag}} A_{\mathrm{flag}}$ nef. By bigness write $A_{\mathrm{flag}} \sim_{\mathbb{Q}} A + \Delta_0$ with A ample and Δ_0 effective. Choose rational $0 < \eta < 1$ small enough for $(X_0, \eta\Delta_0)$ to be Kawamata log terminal. Then $H_i - \eta\Delta_0$ is ample and its d-th intersection root on each Y meeting U is at least $C_{\mathrm{flag}} - \eta > D(D + 1)$. Fujino's effective freeness and point separation apply to $K_{X_0} + H_i$ with the constants $c(d) = D(D+1)$, because both $\sum_{d=1}^{D} d/c(d)$ and $\sum_{d=1}^{D} 2^{1/d} d/c(d)$ are at most one. The complete

cusp systems are generated over U and separate points there, hence are birational in characteristic zero. The chosen weight gaps and Fujino, Theorems 1.1 and 1.2.

Resolving the base ideals of the complete cusp systems as sections of their canonical lines by a common modification $\pi : Z \to X_0$ gives effective divisors E_1, E_2 disjoint from the pullback of U. The two moving bundles are

$$L = \pi^* H(u_1)(-E_1), \qquad L_2 = \pi^* H(u_2)(-E_2).$$

The birational moving systems give L_2 nef and big. For $L_A = \pi^* H(u_2 + Du_1)$, nefness comes from nonnegative weight gaps and bigness comes from the ample part of a positive power and the effective remainder. The canonical section of $\mathcal{O}_Z(E_2 + DE_1)$ is the required comparison section of $L_A \otimes (L_2 \otimes L^D)^{-1}$ on U. The canonical-weight identification realizes the cusp spaces along the path between global and open sections of $K_Z \otimes L_A \otimes L^a$, with equality with the global sections at $a = 0$. Adjoint generation on the open stratum. Moving and comparison bundles. Section comparison on the modification.

Degree and rank on each component. The moving cusp system is birational, so $d_Z = c_1(L)^D$ is its image degree. Intersecting with general hyperplanes measures the same degree in the short frame because that frame spans the complete rational cusp space. Put

$$V(a) = \prod_{\tau} \prod_{i<l} (a_{\tau,i} - a_{\tau,l}), \qquad T_Z(a) = c_1(H(a))^D.$$

Pushing the flag intersection down to the toroidal base expresses it as a top degree Chern polynomial in the canonical automorphic bundles. Mumford proportionality then computes it on the compact dual, where the flag over the Grassmannian is the complete flag. The leading coefficient of the Weyl dimension polynomial at weights sa, multiplied by $D!$, gives $T_Z(a) = c_Z V(a)$. Here $c_Z > 0$ for regularly decreasing a and is proportional to the component volume. The selected components have equal volume. Their prescribed trace self-dual lattices in the same Hermitian space are locally isometric. The host extension over its real subfield is unramified at finite places. At a nonsplit place first rescale the common pairing by its fractional trace-dual ideal, so that trace self-duality gives a unimodular Hermitian form. A nonsingular Hermitian space over the quadratic residue field has a vector of nonzero norm, since varying a scalar in the polarization identity uses a nondegenerate trace pairing even in characteristic two. Lift the vector, normalize its unit norm by the surjective unramified unit norm, and split off its orthogonal complement. Induction identifies the two lattices. At split places the factors are paired duals, so a basis and its dual give the isometry. The isometries conjugate the congruence subgroups in the special unitary group and preserve volume, since the adjoint determinant of unitary conjugation is one. The special unitary group is simply connected and absolutely almost simple with noncompact real factors in this signature. Strong approximation gives the real covolume as its fixed adelic quotient volume divided by the finite compact volume. Thus $c_Z = c_*$ on the selected components. Mumford, Theorem 3.2, Platonov, Theorem A, and the prescribed component lattices.

The fixed part of the moving system is controlled by the nonzero theta carrying section θ_Z^{tr} of weight $8(u_1 - u_1')$, where $u_1' = u_1 - 3rh$. Its squared Bergman comparison (113) is for the complete cusp space at u_1 in the same canonical frames. The invariant metrics on the canonical automorphic extensions have polylogarithmic bounds for their Gram entries and inverse determinants. These bounds pass to the flag quotient metrics, determinants, duals and tensor powers. The algebraic and analytic canonical extensions identify the frames in the pointwise estimate. Mumford, Main Theorem 3.1 and §1 and Lan, Theorem 5.2.12.

Take a complex transverse disc at a general point of a prime divisor Γ above the boundary of Z. The boundary branch equations below pull back to powers of the disc parameter z times units. Thus a nonzero local frame of each canonical weight line and its reciprocal have norms bounded by powers of $1 + |\log |z||$, including at a divisor of the principalization. If $e_\Gamma = \operatorname{ord}_\Gamma E_1$, the coefficients of each member of a Petersson orthonormal basis of cusp sections are divisible by z^{e_Γ} in the canonical frame. Their squared norms

sum to $\mathscr{B}_Z$, so $\mathscr{B}_Z$ is bounded by a constant times $|z|^{2e_\Gamma}$ times a power of the logarithm. By (113), the coefficient of the carrying section in its canonical frame is bounded by $|z|^{8e_\Gamma}$ times a power of the logarithm and a fixed constant. Its meromorphic germ therefore has order at least $8e_\Gamma$. This also excludes a pole when $e_\Gamma = 0$. The section is regular on the interior, so $\operatorname{div}(\pi^*\theta_Z^{\mathrm{tr}}) - 8E_1$ is effective, of class $8(L - \pi^*H(u_1'))$. The result holds over the coefficient field by base change. Both $H(u_1')$ and $H(u_1)$ are nef. For this criterion, decompose a weight of nonnegative gaps into elementary flag quotient determinants and constant weights. The Hodge eigenbundles are nef by semipositivity for their logarithmic extensions with unipotent monodromy. The determinant and quotient construction gives nefness of the elementary lines, and the flat lines are numerically trivial. The prescribed cross-block gaps of u_1 exceed $3r$, so the criterion applies to u_1' too. Intersect the two effective differences with the mixed products of the corresponding nef lines and sum. This gives

$$0 < d_* := c_* V(u_1') = T_Z(u_1') \le d_Z \le T_Z(u_1) = \frac{d_*}{1-\rho}, \qquad 1 - \rho = \frac{V(u_1')}{V(u_1)}.$$

Fujino–Fujisawa, Theorem 1.3, the elementary nef lines.

Let S be the transverse intersection of D general moving sections $s_1, \dots, s_D$, and choose s_0 nonvanishing on S. The intersection consists of d_Z points in the open where the moving map is an isomorphism. Koszul and Kawamata–Viehweg vanishing give surjectivity of restriction of $H^0(Z, K_Z \otimes L_A \otimes L^D)$ to S, since the higher cohomology of the twisted Koszul terms $K_Z \otimes L_A \otimes L^{D-i}$ vanishes for every $0 \le i \le D$. Choose lifts of the basis vectors at the d_Z points and multiply them by $s_0^{\ell - D - |\alpha|} \prod_i s_i^{\alpha_i}$ for $|\alpha| \le \ell - D$. They are independent by reading the first nonzero homogeneous term of any relation in the local coordinates s_i/s_0 at each point. The resulting adjoint sections are cusp sections at a_ℓ on X_0. The canonical weight formula shows this since the relative canonical divisor of the modification is exceptional and the effective fixed divisor imposes vanishing. Thus $r_Z \ge d_Z \binom{\ell}{D} \ge q_Z = d_Z \binom{\ell-1}{D}$.

For the upper bound put $n' = \ell + D + 3$ and $a = \lfloor n'/(1-\rho) \rfloor$. The bundles $H(v)$ and $H(n'u_1 - v)$ are nef. Indeed, if u_1 has quotient gap b_i and remainder $c_i < n$, the second weight has gap $2b_i - c_i$ minus the canonical gap. The selected gap bounds are $b_i \ge n + 1 + 2r$ within each block and $b_i \ge \ell + c_0$ between the blocks, with $\ell \ge D + 2r$. These suffice for nonnegativity, since the canonical gaps are -2 and $2r - 2$, respectively. Blow up S and write E_S for its exceptional divisor and b_S for the blow-up map. The sections s_i generate the nef line $N = b_S^* L(-E_S)$. If a nonzero section of $H(v)$ vanishes to order at least $a + 1$ at every point of S, intersect the effective class $b_S^* \pi^* H(v) - (a+1)E_S$ with $c_1(N)^{D-1}$. The restriction of N to each exceptional divisor is $\mathcal{O}_{\mathbf{P}^{D-1}}(1)$, and pullback classes restrict trivially. By the two nef comparisons,

$$(a+1)d_Z \le c_1(\pi^*H(v))c_1(L)^{D-1} \le n' T_Z(u_1) \le n' d_Z/(1-\rho).$$

This contradicts the choice of a. Restriction to the jets through order a is injective and $r_Z \le d_Z \binom{a+D}{D}$. For $0 \le i < D$, the ratio $(\ell - 1 - i)/(a + D - i)$ differs from one by at most $4D/\ell + \rho$, since $a(1-\rho) \le \ell + D + 3$ and $D \ge 2$. Multiplying the factors gives $q_Z \ge (1-z) r_Z$ for $z = 7(D^2/\ell + D\rho)$, as required for the finite-degree theorem. The selected gap bounds and the rank-share normalization.

Under the evaluation and integrality comparisons on the cover, finitely many detecting CM points embed the section space in a direct sum of coefficient extensions. Clearing their fixed evaluation denominators puts the integral section lattice inside the corresponding rings of integers. This gives finite generation. For fullness, the weight comparison gives a nonzero scalar c and a positive integer N such that cf^N is integral for a rational section f. Multiply c by an integer if required so that $v(c) \ge 0$ at the testing valuations. Then $N(v(c) + v(f)) \ge v(c) + Nv(f) \ge 0$, so cf is integral. Finite generation and fullness.

The theta span at the principal cover. Write $X_c = frL$. For a theta auxiliary rank $n_{\mathrm{aux}} = 8(fr + 1)$, use the identity refinement of the chosen principal level M_{cov}. The half-level marking group for this refinement is trivial. At a prime of the split half ideal $\mathfrak{m}$ the marked Schrödinger packet has position support

$\left(\begin{smallmatrix}1_r\\0\end{smallmatrix}\right) + \mathfrak{m}\,\mathrm{Mat}_{n_{\mathrm{aux}},r}(\mathcal{O}_{E_v})$. It is the Heisenberg translate of the unit self-dual vacuum of the corresponding lattice. Translation forces constancy on its cosets and modulation forces this support. For two supports to meet after a rational unitary change, their top blocks must be integral and both congruent to 1_r, forcing the level congruence on the change. For a nonzero rational isotropic line, rescale a nonzero dual coordinate to be primitive at one half-level prime. A modulation there trivial on the position lattice and nontrivial on a marked top row acts by a nontrivial scalar on the packet. Additive approximation along the line gives the same action on it. Hence the packet has zero coinvariant along that isotropic line. The marked vacuum, The split support test, and Isotropic coinvariants.

For two trace-self-dual lattices with common-index $m_{\mathrm{int}} = [L_1 : L_1 \cap L_2]$, the absolute overlap of the marked auxiliary vacua is at most $m_{\mathrm{int}}^{-n_{\mathrm{aux}}/2}$. The projection onto the first vacuum is an average over the quotient of its tensor lattice by the intersection. The translated terms acting on the second vacuum are orthogonal because their differences pair nontrivially with its self-dual lattice. The quotient has order $m_{\mathrm{int}}^{n_{\mathrm{aux}}}$. There are at most $m_{\mathrm{int}}^{d_{\mathrm{lat}}}$ possibilities for the intersection, where $d_{\mathrm{lat}} = 2fr$, by counting successive prime-index sublattices. Each compatible second lattice is a splitting over the intersection in its discriminant module. The splittings differ by homomorphisms between order-m_{int} groups with at most d_{lat} generators, giving at most $m_{\mathrm{int}}^{d_{\mathrm{lat}}}$ possibilities. In the mass-weighted overlap operator C, the diagonal entry is the common positive total mass and the relative off-diagonal row sum is at most $\sum_{m_{\mathrm{int}}\geq 2} m_{\mathrm{int}}^{2d_{\mathrm{lat}}-n_{\mathrm{aux}}/2} < 1$. Thus C^{-1} is the convergent Neumann series on bounded families and preserves the prescribed covariance. Vacuum overlap and The lattice row sum.

Double the kernel, using the conjugate polarization for the second theta function. The definite auxiliary space has Witt index zero, and $n_{\mathrm{aux}} > 4r$. The absolutely convergent unitary Siegel–Weil formula identifies the mass-weighted products of theta kernels with the Eisenstein sum on the doubled space. Its finite Schwartz value at the origin after partial Fourier transform is the overlap pairing. The singular orbits vanish by the isotropic coinvariant calculation. The graph orbits therefore give the periodization of C times the real differentiated Gaussian kernel $K(x, y)$. Yamana, Theorem 2 and p. 246, and Graph terms after differentiation.

On the slab of frames with Gram matrices in a fixed compact range about the signature matrices, the real kernel is integrable in y. In Cartan coordinates the determinant estimate decays by $\prod_i(\cosh t_i)^{-n_{\mathrm{aux}}}$; the density grows at most as $\exp(2r\sum_i t_i)$. At an orthonormal frame its diagonal is positive, by the sum of the paired Fischer norms. Integration against $K(x, y)$ depends only on the character projection under the diagonal compact torus at x. After removal of the Borel character, diagonal conjugation gives strictly positive degrees on the lower unitriangular chart, so a holomorphic torus invariant there is constant. Thus the projected function is determined by its value at x. This gives reproduction by a positive constant, the squared integral norm of the kernel divided by its diagonal value. Extend a bounded cusp section over its marked lattice orbits with the given covariance and apply the inverse C^{-1}. Integrating the graph identity against this bounded family reproduces the section. Absolute integrability from the row sum and the real kernel bound permits unfolding, which expresses the result in the finite span of the theta series and their highest-weight harmonic derivatives. Slab integrability and Reproduction and the unfolded span.

Integral short generators from the theta span. The harmonic tensors in this proof admit a small integral spanning set. For a unimodular auxiliary lattice Λ its trace-dual vectors have Euclidean minimum at least $\sqrt{[E : \mathbf{Q}]}/\operatorname{rd}(E)$, by the norm of the different and arithmetic–geometric mean. Transference gives a rational basis of Λ with polynomially bounded lengths times $\operatorname{rd}(E)$. Take upper vectors u_i, dual candidates v_i', and lower vectors w_j. Clearing the Gram inverse by its adjugate replaces the lower vectors by integral ones orthogonal to the upper vectors. The products of leading minors in these two blocks are harmonic. They span the highest-weight harmonic space because the full-rank orthogonal pairs form a dense orbit in the variety $XY = 0$. Its defining ideal has codimension p_jq_j and is generically reduced, as seen by solving for Y when X has full rank. For $\operatorname{rk} X = a < p_j$ the dimension drops by $(p_j - a)(n_{\mathrm{aux}} - q_j - a) > 0$, so complete intersection gives reducedness. Thus evaluation on the dense orbit separates the

apolar harmonics. At the standard pair the highest-weight evaluation is represented by the product of leading minors, by the diagonal character on the Gauss cell. Taking its unitary orbit therefore spans the space. The parametrization by u_i, v_i', w_j has degree at most $(2r+1)d_a$, where d_a is the harmonic degree. Interpolation on the integral grid from zero to that degree gives a spanning family. The determinant expansion and the basis lengths bound its log coefficient ℓ^1-norm by $O(rLd_a)$. [53, Theorems 2.1–2.2] The harmonic family and interpolation.

For the rational realization, the odd pullback $[N]^*$ of the rigidified theta line has symmetric finite-order theta-group lifts. On the isotropic N^2-torsion kernel the lifts are multiplicative by the Weil commutator identity. Averaging the translated sections and applying the harmonic jet is algebraic and descends by faithful flatness. A change of symplectic complement introduces a quadratic exponential paired across the two blocks; the associated contraction kills the harmonic polynomial, so the jet is independent of the complement. Rational theta values.

The marked phase of the stable packet. For the stable span use the identity refinement of the prescribed principal level M. Its comparison section is one. The full marked group of the universal abelian variety at this level still has order M^{fr}. Fix a marked component with self-dual period lattice L, and put $L_0 = \mathcal{O}_E^{n_{\mathrm{aux}}}$. Write $\Gamma_0 = L_0 \otimes_{\mathcal{O}_E} L$, $\Gamma_\Lambda = \Lambda \otimes_{\mathcal{O}_E} L$, and $\Gamma_\Lambda^\sharp = \Lambda \otimes_{\mathcal{O}_E} L^\sharp$, where $L^\sharp$ is the half-level shift in the stable packet. If the chosen half ideal is $\mathfrak{m}$ in the Schrödinger convention, this shift is $L^\sharp = \mathfrak{m}\bar{\mathfrak{m}}^{-1}L$. In the Heisenberg representation use the semicharacter

$$\epsilon_w(y) = \exp\bigl(\pi i E_n(y, w_0 y)\bigr),$$

where E_n is the alternating polarization pairing on the product. The action on the vacuum of a self-dual lattice Γ is fixed by $\epsilon_w(y)W(y)$ for $y \in \Gamma$. In particular the unmodified Weyl operator $W(2y)$ fixes it, since w_0 preserves Γ and $E_n(y, w_0 y)$ is integral. The semicharacter and vacuum and The packet in the shifted vacuum.

For the inner average, multiplication by N identifies the isotropic group of translates with $\Gamma_\Lambda/N\Gamma_0$. The defining conditions are that every N-multiple of a row in the dual of Λ kills the translate, so they impose precisely $Nx \in \Gamma_\Lambda$ for a torsion representative x. The two inclusions $NL_0 \subseteq \Lambda \subseteq N^{-1}L_0$ give the asserted quotient. Its order is odd because N is odd and $\Lambda = L_0$ at two. We can thus choose representatives $y = Nx$ in $2\Gamma_\Lambda$. In the pulled-back frame the normalized translate is the operator $W(y)$ in the period coordinates before dilation. The identity $N^2 E_n(x, x') = E_n(Nx, Nx')$ preserves the Weyl commutator. The even representatives give symmetric lifts with order dividing N^2, since the even period relation makes their N^2-th powers trivial on the corresponding theta space. The uniqueness of the odd-order symmetric lift identifies them with the lift in the algebraic rule. Torsion representatives for the projection.

The projection of the fixed unit vacuum of Γ_0 onto Γ_Λ is the orbit average over $\Gamma_\Lambda/(\Gamma_\Lambda \cap \Gamma_0)$. The orbit terms are orthogonal. If two representatives differ by a vector outside Γ_0, self-duality supplies a member of Γ_0 with nontrivial commutator against the difference, so the two vectors have distinct eigencharacters for that operator. The order of the quotient is J^r for $J = [\Lambda : \Lambda \cap L_0]$, by local freeness of L in rank r. The projection therefore has norm $J^{-r/2}$. Each projected term occurs $m_N = [\Lambda \cap L_0 : NL_0]^r$ times in the inner sum. Dividing by m_N gives its projector with factor J^r. If d_a is the harmonic degree, the invariant derivatives from dilation give N^{d_a}, which balances N^{-d_a} in the rule.

For the marking sum let $Q = \mathrm{Lev}_M^{n_{\mathrm{aux}}}$ be the product of the designated half-level marking groups. A marked point q has an $[N]$-lift supported on the level half. Before dilation this gives an even representative Y_q in $\Gamma_\Lambda^\sharp$. These representatives map bijectively to the quotient of $\Gamma_\Lambda^\sharp$ by its intersection with Γ_Λ. At primes of M, $\Lambda = L_0$, and the quotient is the half-level module of the marking. Away from M the two lattices coincide. The Chinese remainder theorem gives the global representatives, even because the quotient has odd order. Changing a lift at level M changes its dilation by the intersection and leaves the action on the projected vacuum unchanged. For a dual marked representative $u \in 2\Gamma_\Lambda$, the finite Fourier character is

$\exp(-2\pi i E_n(u, Y_q))$, by the specified symplectic pairing of the two halves. For the projected vector v one has $W(u)v = v$. Thus the commutator identity gives

$$\frac{1}{|Q|}\sum_q e^{-2\pi i E_n(u,Y_q)} W(Y_q)v = W(u)\,\mathrm{pr}_{\Gamma_\Lambda^\sharp}\, v.$$

The sign here uses $W(u)W(Y_q) = e^{-2\pi i E_n(u,Y_q)} W(Y_q)W(u)$. Applying this identity to the Fourier transform of the marked class indicator gives exactly the prescribed sum of translates in the analytic packet. For the identity refinement the only phase from identifying the marked class with its dual representative is fixed by the component marking and is independent of Λ, N. Indeed changing the representative changes the alternating pairing by an integer. The Fourier argument uses the full half-level group even in this identity refinement. The finite commutator calculation, The cosets of the selected half, and The marked Fourier coefficient.

The level shift from L to $L^\sharp$ is at the primes of M, while Λ differs from L_0 away from M. The two vacuum projections act in different factors of the restricted tensor product. Hence projecting first to Γ_Λ and then to $\Gamma_\Lambda^\sharp$ has the same marked level factor as when $\Lambda = L_0$. The orthogonal orbit sum gives the norm of the overlap with the chosen unit vacuum for the shifted standard lattice as the inverse square root of the relative level index. Its phase is fixed once for the component. Away from the level the shifted analytic packet uses the unit vector obtained by normalizing the projected standard vacuum. Passing to that unit vector leaves exactly the positive scalar $J^{-r/2}$ from those primes. Together with J^r from the inner average, the auxiliary index factor for the raw rule is $J^{r/2}$. The unscaled short rule has the further algebraic factor β_Λ^{-1}. The projections at the two disjoint sets of primes.

There is also a scalar comparing the principal theta generator in the marked abelian fibre with the periodization of the unit standard vacuum. This is common to every auxiliary lattice and all harmonic polynomials at that frame. Its modulus is constant on the unitary frames of the component. For example, in symplectic period coordinates with period matrix Ω the normalized torus mean of the classical theta norm squared is $2^{-g/2}(\det \Im\Omega)^{-1/2}$, where $g = fr$. Integration over the real period coordinates kills unequal Fourier indices and the remaining integral unfolds to the ordinary Gaussian integral. If h changes the abelian coordinates to the normalized period coordinates, the metaplectic transformation of the product of n_{aux} theta factors includes $(\det h)^{n_{\mathrm{aux}}/2}$. In the unitary frame the determinant of the fixed positive polarization form is $|\det h|^2/\det \Im\Omega$, independent of the frame. Taking the ratio therefore gives the constant modulus, with the symplectic basis and its unit vacuum held fixed on the component. Mumford, *Tata Lectures on Theta I*, pp. 188–189, and The norm of the principal comparison.

Both readings are holomorphic with the same right triangular weight. Orthonormalizing by that group extends the common modulus of the ratio to all frames in the component. On an open set where one reading is nonzero, the ratio is holomorphic with constant modulus, hence constant. The equality for every other reading on that set then extends to the connected component. In the factor $C_Z u_{\Lambda,N}$ for the fixed marked frame, C_Z contains the common principal and level factors. The nonzero $u_{\Lambda,N}$ contains the index and, for the unscaled rule, β_Λ^{-1}. This verifies the separation required by the simultaneous stable span. The supremum estimates for the raw rule use its Gaussian normalization and the same index $J^{r/2}$. The identity on the component.

The marked harmonic value in the Raynaud chart. Keep the rigidified Poincaré biextension of the integral chart and the trace-one element w_0, with $w_0 + \bar{w}_0 = 1$. The theta line is obtained on the diagonal $(x, w_0 x)$. To describe the connection, write $h(x, y)$ for the additive gauge of a rigidified Poincaré frame on the formal group. It vanishes on both axes. The two additivity defects of h are logarithms of the integral biextension multipliers. Likewise $h(y, x) - h(x, y)$ and $h(ax, y) - h(x, \bar{a}y)$ are logarithms of the integral

symmetry and adjoint multipliers. In the frame induced on the theta line, the level-N harmonic value of a germ F on n copies is

$$\mathcal{D}_{N,P}(F) = P(D)\big(F\exp(NH)\big)\Big|_0, \qquad H(x_1,\dots,x_n) = \sum_b h(x_b, w_0 x_b).$$

Here D consists of the translation-invariant derivatives in the prescribed framed tangent directions, acting on their indicated copies. Gauge and connection.

A change of integral Poincaré trivialization changes the germ and the exponential by inverse factors. With a fixed frame, the difference of two gauges is a symmetric, CM-adjoint biadditive form. In the formal logarithm coordinates of the characteristic-zero group this is bilinear. Adjointness implies that the two entries pair only conjugate embeddings. On the diagonal $(x, w_0 x)$, symmetry and $w_0 + \bar{w}_0 = 1$ give half its diagonal quadratic form. For each pair of tangent coordinates the resulting contraction sums over the copies. If Q_2 is this quadratic form, the Leibniz rule identifies multiplication by $\exp(NQ_2)$ on germs with $\exp(NQ_2(\partial/\partial X))$ on their differential polynomials in X. Harmonicity kills all its positive-degree terms. This proves independence of the rational gauge in the same framed tangent directions.

For an integral $\mathcal{O}_E$-matrix U on the copies with $U^*U = L\,1$, pullback identifies the product theta line with its L-th power by the rigidified cubical isomorphism. In the product frame its multiplier is the integral unit

$$W_U = \exp(H \circ U - LH), \qquad W_U(0) = 1. \tag{100}$$

Indeed the two additions in the biextension express the formula for $H \circ U$, modulo logarithms of integral units, as the sum of the terms $h(x_i, \bar{U}_{bi} w_0 U_{bj} x_j)$. Using the adjoint multipliers and summing over b leaves $\sum_i h(x_i, Lw_0 x_i)$. Repeated addition in the second argument leaves LH. Each of these operations uses the actual rigidified multiplication or adjoint map of the frame. The same calculation in the gauge trivialization gives its multiplier (100). The normalized derivative therefore satisfies

$$\mathcal{D}_{kL,P}\big(W_U^k (F \circ U)\big) = \mathcal{D}_{k,U_*P}(F),$$

where U_*P is substitution of the differential of U on the tangent directions in P. The equality is the chain rule for invariant derivatives, since $W_U^k \exp(kLH) = \exp(kH \circ U)$. The single-copy multiplier and The product chain rule.

The Raynaud expansion restricts a theta section to the same formal neighbourhood of the identity. The normalized line isomorphism restricts there to the one given by the rigidified Poincaré frame. This follows by pulling back the Poincaré bundle and its biextension maps under the uniformization, then restricting to each finite thickening of the identity. Its product restriction is the tensor product of the restrictions on the copies. Thus the coefficients of the expansion through degree $\deg P$ in the induced frame are the coefficients used by the rational connection. One can use another normalized symmetric lift of a finite torsion translation in the expansion. Comparing the two lifts over the generic fibre gives a scalar, and finite order makes the scalar a root of unity. These scalar factors have valuation zero. In particular, if $N = \alpha\bar{\alpha} \mid M_0$ uses the prescribed half and $x_b \in A[\bar{\alpha}]$, put $\eta = (M_0/N)\bar{\alpha}$. The pullback by η of the translate of $[\alpha]^*\theta_1$ at x_b is the normalized level-M_0^2 translate of $[M_0]^*\theta_1$, up to a root of unity. A lift x_b' can be taken in the marked half of $A[M_0^2]$. At a prime dividing N, multiplication by η maps the free half-torsion of level prime squared onto its level-prime subgroup. At the other primes take the lift zero. The level point is $M_0 x_b' = \alpha x_b$. For harmonic embedding degrees d_τ with $d = \sum_\tau d_\tau$, write $C_\alpha = \prod_\tau \tau(\alpha)^{d_\tau}$ in the coefficient field. The pullback chain rule multiplies the harmonic value by C_η, and the marked rule uses M_0^{-d} times that value. Since $C_\eta C_\alpha = M_0^d$, the local harmonic value equals C_α times the corresponding marked rule up to a unit. For $N = M_0$ every marked point occurs, since α is invertible on the complementary half. The marked evaluation and The marked chart and product restrictions.

Finite loss in the canonical connection. Let R be the complete valuation ring of a finite uniformizing extension, with all the eigenvalues in its fraction field. For an integral unit u, every nonzero invariant derivative of $\log u$ is integral, by iterating $D \log u = u^{-1} Du$. Differentiate an additivity defect of h with at least one derivative on each of the two variables being added. Only the term on their sum survives. At the identity its derivatives are the corresponding invariant derivatives with the two sets of directions combined. Thus all derivatives of h at the identity other than mixed bidegree $(1,1)$ are integral. Derivatives along only one argument vanish by rigidification.

For the mixed bilinear part b_h and integral eigenvectors u_τ, v_σ, adjointness implies

$$\big(\tau(a) - \sigma(\bar{a})\big)\, b_h(u_\tau, v_\sigma) = D_{u_\tau} D_{v_\sigma} \log(\mathrm{ad}_a)(0,0) \in R.$$

If $\tau \neq \bar{\sigma}$, one can choose a whose eigenvalue difference has valuation at most $\delta_p = v(p) v_p(|\Delta_E|)/[E : \mathbf{Q}]$, where p is the residue characteristic and the chosen host is Galois over $\mathbf{Q}$. Indeed if the embeddings have distinct residue characters, the difference can be a unit. Otherwise use a local uniformizer in the totally ramified extension of the maximal unramified subfield. The derivative of its minimal polynomial is the product of its differences with the other conjugates and generates the different. Integral global elements approximate the uniformizer to arbitrary precision. For $\tau = \bar{\sigma}$, the symmetry multiplier implies that the two orders of the mixed pairing differ by an integral pairing. Subtract a symmetric, CM-adjoint biadditive form in logarithm coordinates agreeing with one order in each conjugate pair. This leaves integral pairings in the other order and preserves the harmonic value by the gauge calculation above. This choice uses no division by two in R.

In the modified gauge, every possibly nonintegral factor in the repeated Leibniz rule for derivatives of $F \exp(kH)$ uses two derivatives and costs at most δ_p. Indeed the higher derivatives of the biadditive correction in invariant directions vanish. The repeated rule for differentiating the exponential has integral coefficients and constant exponential factor one at the origin. For integral F and an integral harmonic polynomial of total degree d,

$$v\big(\mathcal{D}_{k,P}(F)\big) \geq -d\delta_p/2.$$

This is uniform in k and in the integral chart. At either type prime $\delta_p = 0$. In the stable short span the bound $-d\, v(|\Delta_E|)$ suffices. This argument gives the uniform conductor bound. The smaller interval of gap-factor valuations at the non-type primes enters separately through the formal least depth of the chosen polynomial. The conductor units and the normalized loss.

Write $M_{\mathrm{cov}} = M$ and $S_\theta = d_a + frn_{\mathrm{aux}}/2$. The marking module has order M^{fr}, and its Fourier point indicators have unit coefficients.

The primary part of the short theta average. Fix a marked point in the outer finite Fourier sum. Let $L_0 = \mathcal{O}_E^{n_{\mathrm{aux}}}$, $L_N = \Lambda \cap L_0$, $r = p_j + q_j$, and $m_N = [L_N : NL_0]^r$, where the index is the additive index. In particular the denominator m_N uses the intersection lattice. The translating isotropic group $H_{\Lambda,N}$ has order $[\Lambda : NL_0]^r$. For its translates F_x of the product $[N]^* \theta_1^{\boxtimes n_{\mathrm{aux}}}$, with the marking lift included, the inner average is

$$A = N^{-d_a} m_N^{-1} \sum_{x \in H_{\Lambda,N}} \mathcal{D}_{N^2,P}(F_x).$$

At a trait in residue characteristic g, split the semistable chart and the relevant torsion after a finite extension, retaining the original valuation through ramification normalization. The two lattice indices.

Suppose $g \nmid M$. The marking lift then has order prime to g. Choose an E-unitary matrix γ carrying Λ to L_0 at every place above g and 2. The local isometries exist for these unimodular Hermitian lattices: at split places choose a basis and its paired dual, and at unramified nonsplit places split off vectors of unit norm and use surjectivity of the unramified unit norm. The relative CM extension of the chosen host is unramified at finite places. The indicated lattice conditions are open. The Cayley chart of the unitary group and weak approximation on its skew-Hermitian matrices give the rational choice of γ. Choose an

integral matrix G with columns in L_N and a basis there at these same places, by the Chinese remainder theorem and Nakayama's lemma. Choose integers k_1, k_2 prime to $2g$ so that

$$Q = k_1 N G^{-1}, \qquad P_1 = k_2 \gamma G$$

are integral. This is possible because $NL_0 \subseteq L_N$ and $\gamma L_N \subseteq L_0$ locally at the prescribed places. Put $k = k_1 k_2$ and $U = P_1 Q = kN\gamma$. Thus $U^* U = (kN)^2\, 1$. The subgroup $H = H_{\Lambda,N}[g^\infty]$ is the g-primary kernel of U: the condition defining $H_{\Lambda,N}$ is to be killed by all N-multiples of rows in the dual of Λ, and locally those rows are spanned by the rows of γ. Hence the same kernel occurs after the unit dilation k. Set $K = \ker(Q)[g^\infty] \subset H$. The equation $GQ = k_1 N\, 1$ implies that K is killed by N. The Tate lattice has rank r over the local $\mathcal{O}_E$-factors, so the elementary divisors of Q give

$$|K| = g^{v_g(m_N)}.$$

Indeed the cokernel of Q at g has the same elementary divisors as L_N / NL_0. The differential of P_1 on the chart has determinant

$$\delta_g = \prod_\tau \tau(\det P_1)^{m_\tau},$$

because the action on the m_τ tangent copies at τ is the scalar extension of the same matrix. At g, $\det P_1$ generates $\det L_N / \det \Lambda$ since $\gamma\Lambda = L_0$ locally and k_2 is a unit. The adapted matrices and The local index and determinant.

Fix a prime-to-g coset of the group of translates and include its contribution in a base section. Pull that section back by $[k]$ so that its level is $(kN)^2$. Multiplication by k permutes the primary translates, and changes the normalized derivative by a unit. The K-translations fix the product section. They fix $[N]^* \theta_1^{\boxtimes n_{\mathrm{aux}}}$ by its pullback linearization, and commute with the prime-to-g translations since the Weil commutator between the two primary parts is trivial. On the formal group the common cubical unit W_U of (100) conjugates the U-kernel action on germs to ordinary translation. In fact the pullback frame of $U^* \mathscr{L}^{\boxtimes n_{\mathrm{aux}}}$ has the canonical kernel linearization, and the rigidified cubical isomorphism transports it to the symmetric lifts. Two such lifts over a kernel point differ by a scalar with square one, by symmetry, and with odd order, by the finite-order condition. The scalar is one. This argument also explains why the relevant unit is the same for all points of the formal kernel.

Let H° and K° be the parts of H and K in the formal completion at the identity. The primary sum reduces to the sum of H°-orbits for representatives of $H/(H^\circ + K)$, each with multiplicity $[K : K^\circ]$. A normalized translate in each representative has integral germ: pullback preserves the weighted Fourier lattice, the finite-order translation preserves that lattice, and all the coefficient orders in this lattice are nonnegative in the Raynaud identity chart. Translation of a monomial into the local torus coordinates $t_i - 1$ uses integral binomial coefficients. The growing quadratic Fourier orders give convergence of the coefficients there. The same statements hold at residue primes dividing kN, since preservation of the weighted lattice uses the quadratic cancellation and finite order, without dividing by the group order.

For one representative, divide the integral germ by W_U. It is invariant under K° and descends by Q to an integral germ f. To see the integrality, the maps P_1, Q on the formal group are finite flat: their kernels are contained in multiplication kernels of the finite-height formal group, and the pullbacks of the parameters generate a primary ideal on the special fibre. Thus the completed coordinate rings give finite local maps of equal-dimensional regular rings, flat by the Cohen–Macaulay criterion. After splitting the generic kernels, their fibres are the translation orbits. An invariant belongs to the fraction field of the target and is integral over its normal coordinate ring. In particular it descends over the same valuation ring. Now the sum over H° is $|K^\circ|$ times the pullback of the trace of f by P_1. On the Raynaud extension P_1 induces isogenies on the abelian and toric parts, since an integral multiple of its inverse acts on both parts. The resulting map of the extensions is finite flat, as can be checked on the torus torsors after base change along the abelian isogeny. It is a local complete intersection between smooth group schemes. The

trace in its relative dualizing module is the Jacobian times a generator. In translation-invariant volume forms the Jacobian is δ_g, so in local coordinates it is δ_g times a unit. Completion at the identity on the target and passage to the identity factor on the source give the formal trace. Thus the trace belongs to δ_g times the target power series ring. Restoring the multiplier and summing the representatives shows that their contribution has formal germ

$$|K|\,\delta_g\,W_U\,(f_0\circ U)$$

for an integral germ f_0. Stacks Project, Lemma 49.12.4 and Lemma 49.12.6, and Trace of the adapted formal kernel.

The invariant chain rule shifts the polynomial to $(kN)^{-d_a}U_*P=\gamma_*P$. This is harmonic because γ is unitary, and integral in the local standard column lattice because $\gamma\Lambda=L_0$ there. The conductor estimate for f_0 gives loss at most $d_a v(|\Delta_E|)$. Since $|K|$ is the full g-part of m_N, the normalized sum satisfies

$$v(A)\geq v(\delta_g)-d_a v(|\Delta_E|).$$

This holds on every prime-to-g coset, hence on their sum, with the same determinant ideal. In particular the scalar in the inverse reflex ideal clears this determinant loss in the same framed trait. The normalized adapted average.

At $g\mid M$, $\Lambda=L_0$ locally, so the whole primary subgroup is killed by N and has order $g^{v_g(m_N)}$. The same pullback linearization fixes it, and its cardinality cancels the primary denominator. The marking lift in this half-level is sent by $[N]$ to a point of order dividing g^e on its primary part, with $g^e\|M$. If $g^b\|N$, factor $[g^{b-e}]$ out of the pullback. The primary order of the new lift divides g^{2e}, and the remaining level factor at g is g^e. Even if the raw marking translation is performed after a prime-to-g translate, its normalized pullback agrees by the coprime Weil commutator. The chain rule contributes a factor $g^{(b-e)d_a}$, cancelling that part of N^{-d_a}. Thus multiplication by g^{ed_a} covers the remaining scalar cost, in addition to the conductor bound. Both cases use the same residue-characteristic valuation and the same polynomials as the outer Fourier sum. The level-primary average and The residual marking exponent.

The adapted determinant factor belongs to the ideal $I=\prod_\tau(\tau(\mathfrak{a})\mathcal{O}_{K_{\mathrm{sec}}})^{m_\tau}$, where $\mathfrak{a}$ is the determinant ideal for $\Lambda\cap\mathcal{O}_E^{n_{\mathrm{aux}}}\subseteq\Lambda$, and m_τ are the Hodge multiplicities in the coefficient field K_{sec} of the cover. For the additive index J of this inclusion, $\mathrm{N}I=J^{r[K_{\mathrm{sec}}:\mathbb{Q}]/2}$. Minkowski's theorem on a product of complex discs gives $s\in I^{-1}$ with $\max_\sigma|\sigma(s)|\leq \mathrm{rd}(K_{\mathrm{sec}})^{1/2}J^{-r/2}$. It is integral over M and cancels the adapted determinant loss away from M. Thus one common clearing factor for the scaled section is

$$Z_M=M^{frn_{\mathrm{aux}}+d_a}|\Delta_E|^{d_a},\qquad \log Z_M\ll(frL)^2S_\theta.$$

The level primes, The adapted determinant, and The reflex ideal scalar.

In the frame where the supremum is approximated, the stable theta estimate uses a Gaussian lattice minimum at least e^{-40L}. Cauchy bounds for the degree-d_a derivatives cost $O(Ld_a)$; the Gaussian lattice sum costs $O(frn_{\mathrm{aux}}L)$ in log norm by counting lattice points in shells of that minimum. Its raw index factor $J^{r/2}$ cancels against s. Together with the harmonic coefficients and Z_M this gives log norm $O(X_c^2S_\theta)$ for the short integral spans. For weight a the size is $S_\theta=O(r|a|_1+(fr)^2)$, with $|a|_1=\sum_{\tau,i}|a_{\tau,i}|$. The theta key and flat twist carry this span to weight a. Their log norm cost is $O((fr)^2+rL(|a|_1+f^2r)+X_c^2)$ by the flat determinant bound $9rL$. The resulting log norms are at most $C(X_c^3|a|_1+X_c^5)$. The Gaussian and index factors and The weight change.

Choose a field basis from these vectors and multiply it by a short rational basis of field integers. This gives a rational basis in the lattice. Reduce an integral $\mathbf{Z}$-basis of the lattice into the parallelotope of the short rational basis by taking fractional parts of the rational coordinates. The short basis and reduced vectors generate the lattice. They number at most twice the integral rank, with lengths at most the integral rank times the length bound for the short rational basis. Integral generators from a short span. At each embedding, a bound on the sum of squared pointwise evaluations of an orthonormal basis of the moving space gives the multiplication bound. Indeed, Cauchy–Schwarz bounds a section's pointwise norm by its

norm in the moving space times the square root of that sum. Integration against the squared norm of the second factor gives the needed operator estimate. Integral sections multiply integrally in the prescribed integral models. The prescribed Bergman estimate is $\exp(4000X_c^3)$ for the weights in use. Its integral in the probability Petersson measure bounds the corresponding rational dimensions, since that integral is the trace of orthogonal projection. The moving and auxiliary weights have size $|u_i|_1 \le X_c^2(\kappa/\ell+1)$ for $i = 1, 2$, and the adjoint weights a_s have $|a_s|_1 \le X_c^2\kappa$. At $s = 0$ the formula for a_0 gives $|a_0|_1 \le (D+2)X_c^2(\kappa/\ell+1)$. Since $P_{\mathrm{cov}} \ge X_c^{12}$, the short spans and the Bergman bound give the following numerical bounds at large x. The moving frame has n_Z vectors of norm at most Λ. A short rational basis can be taken from each spanning family in degrees $0, D, \ell$. Write Λ_D, Λ_ℓ for the bounds in the last two degrees and J_{res} for the residue estimate in the component theorem. The construction gives

$$\log n_Z,\ \log\beta,\ \log d_Z \le P_{\mathrm{cov}}, \qquad \log\Lambda \le P_{\mathrm{cov}}(\kappa/\ell+1),$$
$$\log\Lambda_D,\ \log\Lambda_\ell \le P_{\mathrm{cov}}\kappa, \qquad J_{\mathrm{res}} \le 4\big(DP_{\mathrm{cov}}(\kappa/\ell+1) + \ell P_{\mathrm{cov}}\big) + 2DP_{\mathrm{cov}}\kappa/\ell.$$

The degree-zero short bound is absorbed in J_{res} in the numerical form of the component estimate. The comparison of the modified cusp lattice to the target lattice costs at most $r_Z c_{\mathrm{cmp}}$ in normalized degree, with $c_{\mathrm{cmp}} \le C_p(frL)^2$. The size estimates for this stage give

$$L \le frL, \qquad 1 \le (frL)^{12} \le P_{\mathrm{cov}}, \qquad \ell P_{\mathrm{cov}} \le \sqrt{b}, \qquad zP_{\mathrm{cov}}\kappa \le b, \qquad z \le \tfrac{1}{2}.$$

In fact $\ell \ge xD^2$ because $k \ge xm$, so $\ell \ge 3D+5$ beyond a fixed threshold. The quantities P_{cov}, k, D have polynomial bounds in x, so $\ell P_{\mathrm{cov}} \le \sqrt{b}$ follows too. The degree Vandermonde ratio $1-\rho$ for the two moving weights has one factor $1-3r/g$ per affected cross-block gap $g > 3r$. The quotient weights bound these gaps below by $H_{\mathrm{sep}}/(2(\ell+D+1))$, with base gap H_{sep}. For pairs starting in the distinguished block the gap is at least $\sqrt{b}/(4\ell)$. Thus

$$\rho \le 3fr^3\left(\frac{2(\ell+D+1)}{H_{\mathrm{sep}}} + \frac{4\ell}{\sqrt{b}}\right)$$

by $1-\prod_i(1-a_i) \le \sum_i a_i$ for $a_i \in [0,1]$. With $\kappa \le 4frb$, the three terms of $zP_{\mathrm{cov}}\kappa/b$ are bounded by $56fm/P_{\mathrm{cov}}$, $210r^4/(N_b m^2 k)$, and $336f^2r^4D/(P_{\mathrm{cov}}^2 T_0)$. Here T_0 is the order bound, N_b is the terminal column rank, and the weight path has $N_W \ge fr(r+1)$ steps. The inequalities follow from $\ell \le T_0$, $H_{\mathrm{sep}} \ge T_0 N_b f^2 m^2 k D P_{\mathrm{cov}}$, $T_0 \ge 2500r(D+1)$, $N_b \ge mT_0N_W/2$, $P_{\mathrm{cov}} \ge (2fm)^4 N_W$, and $P_{\mathrm{cov}}^3 T_0^2 \le \sqrt{b}$. Each term is less than $1/6$ in this range, giving $zP_{\mathrm{cov}}\kappa \le b/2$. Since $P_{\mathrm{cov}}\kappa \ge b$, we get $z \le 1/2$. Cross-block ratios. Size inequalities. Bergman and moving-frame bounds, degree-zero short vectors, short bases and residue costs, and size and flat-twist comparisons.

In particular $c_{\mathrm{cmp}} \ll b$. Also $D^2 \le \ell/14$ and $2z\ell \le \ell$, so $A_{\mathrm{HS}} \le 100\ell$. Since $\log\ell \le P_{\mathrm{cov}}/2$, these estimates give

$$A_{\mathrm{HS}} \log\max(A_{\mathrm{HS}}, 1) + R_{\mathrm{HS}} \ll b.$$

4.3.2 The component height comparison

The independent height comparison has three nonnegative allowances c_s, c_c, X and comparison heights s_Z. They satisfy

$$t_Z := h_Z + c_s \ge 0, \qquad (1-\rho)\max(0, s_Z + c_s) \le t_Z + c_c, \qquad \mathrm{mean}_Z\, s_Z \ge -X,$$
$$\ell X,\ \ell c_c,\ \big(6D + 10 + \ell(z+2\rho)\big)c_s \ll E_{\mathrm{cov}}.$$

These are the height inputs at the same moving weight and with the same component metrics. They use the arithmetic self-intersection comparison and short integral sections on the cover. For the last allowance the lower bound on the moving Chow height uses $c_s \le P_{\mathrm{cov}}(\kappa/\ell+1)$. Since $(6D+10)/\ell + z + 2\rho \le$

$(1 + 18/(7D))z < 4z$, that allowance is at most $4zP_{\mathrm{cov}}\kappa + 4z\ell P_{\mathrm{cov}} = O(b)$. Height inputs and their construction.

For the estimate of X, consider the root model at the smaller weight $u' = u'_1$. We use a gap $g > 0$ between each of its last q_j slots and every other slot at that label, including the other last-block slots. Write H_{sep} for the lower gap parameter of the undistributed weight. In the quotient construction $n = \ell + D + 1 \le 5\ell/4$, $H_{\mathrm{sep}} \ge 4(c_0 + 2r + n)$, and $H_{\mathrm{sep}} \ge 24r\ell$. The base gap is at least H_{sep} across the two blocks and at least $H_{\mathrm{sep}} - 1$ within the last block. Thus the floored quotient across the blocks is at least $(H_{\mathrm{sep}} - c_0 - 2r + 2 - n)/n \ge 3H_{\mathrm{sep}}/(4n)$. The internal last-block quotient is at least $(H_{\mathrm{sep}} + 1 - n)/n \ge 3H_{\mathrm{sep}}/(4n)$. Passing to u' subtracts $3r$ from the cross-block gap and leaves the internal quotients unchanged. Both gaps are therefore at least $g = H_{\mathrm{sep}}/(4\ell) > 0$. Summing consecutive gaps gives the same lower bound for all pairs involving a last-block slot. The first-block quotients are positive by the within-block bounds above, so u' is regular. The quotient and its gap hypotheses and The two kinds of adjacent gaps.

Use the integer cutoff $B_* = \lfloor 4fr^2(1 + \|u'\|_\infty/g) \rfloor$. Then $B_* \ge 4fr^2 \ge 128$ and $B_* + 1 > 4fr^2(1 + \|u'\|_\infty/g)$. The exceptional support at the deep level consists of the primes dividing $2\,\mathrm{disc}(E)$ or the forced part of the level M'. The host conductor has support in 2 and the two host quadratic primes q, d. The level contributes at most 15 further primes. Thus the exceptional set has at most 18 members and so at most B_*. Let δ'_Z be the normalized geometric degree on the automorphic component. Let S_Z be its limiting arithmetic self-intersection per $(D+1)\delta'_Z$, with the root extension degree divided out. The gap condition at every small slot, The host conductor, The chosen exceptional support, and The integer cutoff.

At the selected finite primes the integral Hodge–Tate comparison identifies the reduced root flags with flags in the trivial Tate bundle. The dual prefix determinants are generated by this bundle. A dominant root weight is their product with nonnegative exponents and a power of the trivial determinant, hence is nef on those fibres. The reduced Hodge–Tate flags.

For the finite denominator at a selected prime l, on a common framed toroidal model take the Weil characters at order l^4. Their cotangent image lies in the integral cotangent module Ω and contains $\kappa\Omega$, with $\kappa = \zeta_l - 1$ for odd l and $\kappa = 4$ at 2. For a finite flat Hopf algebra with augmentation ideal I over a valuation ring, take the lattice generated by its characters, I^2 and l^4 times the whole algebra. A functional y in its dual lattice has bounded powers because it takes integral values on the generic Cartier dual. Its coproduct differs from $y \otimes 1 + 1 \otimes y$ by an integral functional. For $v(c) > v(l)/(l-1)$, $F(Z) = \exp(cZy)$ converges in the lattice norm. Its coproduct is $(F(Z) \otimes 1)(1 \otimes F(Z))$ times an integral exponential of Gauss norm one. The product has norm $\|F\|^2$, since at twice the largest index attaining the maximum there is exactly one term of maximum norm. The coproduct and the integral factor give $\|F\|^2 \le \|F\|$, so the linear coefficient cy is integral. Dualizing and differentiating, and then taking the valuation limit, proves the cotangent bound for the finite Weil group with conormal $\Omega/l^4\Omega$. At a ramified coefficient label the trace dual identifies the eigen-submodule with the different times its eigen-quotient. This gives the different in the dual signed root block. The projected Hodge–Tate maps compose to zero and their reduced ranks add to r, which gives the complementary blocks after flattening the finite image lattices. Applying this calculation uses the prescribed finite Weil differential on the toroidal traits. Kajiwara–Kato–Nakayama, Section 9 and Lan, Proposition 3.3.1.5.

At a remaining prime $p_0 > B_*$, take a normalized curve in the reduced fibre and split the labels over its residue field. With the logarithmic extensions from the semiabelian model, the Hodge and conjugate sequences on the curve have the forms

$$0 \longrightarrow \mathcal{L}^\vee \longrightarrow \mathcal{H} \longrightarrow \mathcal{L} \longrightarrow 0,$$
$$0 \longrightarrow F_C^*\mathcal{L} \longrightarrow \mathcal{H} \longrightarrow F_C^*\mathcal{L}^\vee \longrightarrow 0.$$

Here $\mathcal{L}$ is the Lie bundle, F_C is Frobenius, and duality conjugates the CM label. The root lines on this curve compare with the induced Lie flags, dualizing and reversing the second flag. If d_τ is the degree of the Lie eigenspace at τ, the de Rham degree at τ is $d_\tau - d_{\bar\tau}$. The second sequence says it multiplies by p_0

along a Frobenius step. It is therefore zero on every Frobenius orbit, giving $d_\tau = d_{\bar\tau}$. The Lie flags on the curve and The logarithmic filtration comparison.

Write P_{Lie} for the largest degree of a saturated Lie subbundle, including the zero bundle, and $H_{\mathrm{Lie}} = -\deg \mathcal{L}$. Pull back a maximizing subbundle by Frobenius and place it in $\mathcal{H}$ via the second sequence. Project to $\mathcal{L}$ in the first sequence. The saturated image has degree at most P_{Lie}, and the saturated kernel in $\mathcal{L}^\vee$ has degree at most $H_{\mathrm{Lie}} + P_{\mathrm{Lie}}$ by dualizing its quotient of $\mathcal{L}$. Thus

$$(p_0 - 2)P_{\mathrm{Lie}} \le H_{\mathrm{Lie}}.$$

For a dominant weight μ of least cross-block gap $g_\mu > 0$, summation by parts on the two Lie flags bounds the negative of its root degree at a given label by $\Delta_\tau(d_\tau - P_{\mathrm{Lie}}) + (\mu_{\tau,1} - \mu_{\tau,r})P_{\mathrm{Lie}}$, where $\Delta_\tau = \mu_{\tau,p_j} - \mu_{\tau,p_j+1}$. Each proper prefix degree is at most P_{Lie}, and the two full block degrees agree. Since $d_\tau - P_{\mathrm{Lie}} \le 0$, replace Δ_τ by g_μ in this upper bound. The first-block degrees sum to $-H_{\mathrm{Lie}}/2$. Summing and using the Frobenius bound gives root degree at least $g_\mu H_{\mathrm{Lie}}/2 - 2f\|\mu\|_\infty P_{\mathrm{Lie}}$. This is nonnegative when $(p_0 - 2)g_\mu \ge 4f\|\mu\|_\infty$. Subdegrees from the two filtrations and The flag degree inequality.

The definition of B_* gives the strict inequality at u'. It persists under the small weight perturbation in the approximation, and adding its relatively nef projective polarization preserves the finite inequalities. This verifies relative nefness on the actual cutoff root flags constructed above.

The count in the original Hodge supremum norm. Fix the cutoff and a component. In additive notation let L_λ be its signed root line, where $\lambda = u'$, with the prescribed finite model and Hodge weight $\phi_\lambda = -\log\|\cdot\|_\lambda$. The log norm of an original section is tested on the whole abelian interior. Work on a common normal projective graph with smooth generic fibre on which the finite lines pull back and the generic boundary is a divisor. A section of a pullback of L_λ^n pushes down to an original section by normality, and its supremum on the interior is the original supremum. The degree-one projection formula applies to the arithmetic intersections.

The generic L_λ is nef. Its proper-prefix coefficients are positive and the scalar prefix regularizations lose curvature tending to zero. The last determinant factor is flat. Its top geometric degree is positive by the degree formula, so L_λ is big. Choose an ample arithmetic line $\bar A$ on this graph with strictly positive smooth metric. By bigness, for some fixed positive integer k the generic line $kL_\lambda - A - D_\partial$ has a nonzero section, where D_∂ is the reduced boundary. Put $\nu = k\lambda$. If indices for rational lines need to be cleared, take the section in a fixed integral multiple $q_0(L_\nu - A - D_\partial)$. It provides a nonzero section e of $q_0(L_\nu - A)$ vanishing on every boundary divisor. Take it over the component field. Multiplication by a single nonzero integer clears its finitely many vertical poles. Put $R = 2 - \log\|s_\partial\|^2 \ge 2$, using a smooth metric on the boundary line. In the fixed smooth reference metric the boundary vanishing gives

$$\log\|e\|_{\mathrm{ref}} \le -cR + C_0$$

with $c > 0$. All constants in this multiplication may depend on the fixed cover and cutoff.

For $s > 0$ small rational the line $L_s = (1 - ks)L_\lambda + sA$ is relatively nef on the actual finite model by the cutoff inequality already proved. Use the scalar prefix regularizations at $t_s = \exp(-1/s^2)$ in $L_{\lambda - s\nu}$ and the smooth metric on sA. In writing a root weight as a combination of dual prefix determinants, every proper-prefix coefficient is an adjacent gap and stays positive. The last coefficient is signed. Use the smooth flat archimedean metric on the full signed determinant for this last factor. It is the original metric there and its finite line stays fixed. At infinity the proper prefix curvature errors are $O(s^2|\log s|)$ against a fixed Kähler form and are absorbed by $s\bar A$. Thus the resulting smooth weight ϕ_s is strictly positive at infinity, so its generic line is ample. The normalized geometric degrees and arithmetic intersections of $\overline{L}_s$ tend to those of L_λ with the original Hodge metric, by the mixed-metric calculation.

The pointwise scalar bound needed for section norms is the global bound

$$0 \le \phi_{i,t} - \phi_i \le Ct^{1/3} + C\frac{1 + \log\log(1/t)}{\log(1/t)}R$$

for a proper prefix i. Equivalence of the boundary size functions on the common chart only changes C. The fixed smooth reference differs from ϕ_v by $O(1 + \log R)$. In a common local tensor frame use the identity $L_\lambda = L_s + s(L_v - A)$. The log norm cost per unit of weight of multiplication by e^{s/q_0} is bounded by

$$\frac{s}{q_0} \log \|e\|_{\mathrm{ref}} + Cs(1 + \log R) + C\big(e^{-1/(3s^2)} + s^2(1 + |\log s|)R\big).$$

Here the smooth sA terms cancel, and the last term is the change of the proper-prefix metric. The first term is at most $-csR/q_0 + C_0 s/q_0$. For s sufficiently small this absorbs the $s^2(1 + |\log s|)R$ term. The quantity $C(1+\log R)-cR/(2q_0)$ is bounded above for $R \geq 2$. The total cost is therefore at most $C's$ uniformly on the interior. The global pointwise regularization estimate and Boundary vanishing dominates the pointwise error.

For clarity, realize the multiplication at integral weights. Take $s = (m + m_0)^{-1}$ and $Q_s = 8q_0(m + m_0)$, increasing the fixed q_0 if needed to clear all indices. Then $Q_s L_\lambda = Q_s L_s + 8q_0(L_v - A)$. At degree n in the line $Q_s L_s$, multiply by e^{8n}. This is injective and preserves finite integrality. Its output has the original weight $nQ_s\lambda$. The pointwise estimate gives original log supremum at most $nQ_s t$ whenever the input has log supremum at most $nQ_s(t - C's)$. Integral multiplication and the weight identity.

Apply the arithmetic Hilbert–Samuel formula for $\widehat{\chi}_{\mathrm{sup}}$ to the smooth relatively semipositive line $Q_s\overline{L}_s$, shifted by $Q_s(t-C's)$ in log weight. The theorem uses generic ampleness and finite relative nefness, which hold for fixed s; the shift leaves them and the curvature unchanged. The small-section count is bounded below by this Euler term up to $O(\mathrm{rank}\log(2 + \mathrm{rank}))$, by Minkowski's theorem. That error has lower order than n^{D+1}. The geometric-degree term in the metric shift is exactly $t - C's$ times the normalized geometric degree. Divide by $(nQ_s)^{D+1}$ with the original field and geometric-cover normalizations and take the limsup over these multiples. The injection transfers this bound to the original Hodge norm. Finally let $s \downarrow 0$ at the fixed cutoff, using the convergence of the mixed intersections. The resulting lower bound on the original counting rate is $\delta'_Z \max(0, S_Z + t)$, with the positive part also following from the zero section. The arithmetic Hilbert–Samuel formula in the form stated by Yuan, Section 1.3 and The normalized limit at the original weight.

For the trace, let $b_1, \dots, b_\Delta$ be a basis of the root function field over the target flag function field, where Δ is the generic root-cover degree in the common flag reading. Choose a nonzero section a at one fixed weight k_0u', with $8 \mid k_0$, vanishing on the boundary and clearing the poles of all the b_i. Such a section exists by bigness. One can make all ab_i integral on the finite model with a single scalar. Their log suprema are bounded because their boundary vanishing absorbs the log-log Hodge weights. In the descended generic line frame take the field trace of fab_i. These Δ traces together determine f, by the nondegenerate trace pairing and $a \neq 0$. They lie at weight $(n + k_0)u'$. Their suprema cost at most the fixed factor $\Delta \max_i \|ab_i\|_{\mathrm{sup}}$, using the same original metric in the conjugate readings over the target point. At a horizontal boundary divisor the products vanish in all branches over it, so their trace also has positive order there in the descended canonical frame. At a finite trait the trace of elements integral in all branches is integral in the target valuation ring. Apply this after the scalar needed for the stated comparison of the integral root and automorphic frames. The same argument applies to the second change to the mixed frames, giving the two occurrences of $e_2(p)$. Rounding the scalar exponents costs $O(1)$ at this fixed cutoff. In the normalized counting rate the shift k_0 and fixed supremum factor disappear, and the number Δ of trace coordinates cancels the generic root-cover degree. Fixed-weight basis and injectivity of the trace tuple and Norm, finite costs and count comparison.

For $p \leq B_*$ and the primes dividing $2\,\mathrm{disc}(E)M_{\mathrm{cov}}$, the comparison costs at most twice

$$e_2(p) = \frac{1}{p-1} + \frac{40\,\mathrm{ord}_p\,|\mathrm{disc}(E)|}{[E : \mathbf{Q}]} + \mathbf{1}_{p=2}$$

per unit of absolute weight at p. The trace count divides by the root extension degree, which cancels the number of coordinates of the trace tuple. Thus the mixed integral lattice has a lower count with slope

$s_Z = S_Z - 2|u'|_1 \sum_p e_2(p) \log p$, with this prime support. The first summand contributes $O(\log B_*)$ by the dyadic Chebyshev bound below B_*. Above B_* there are at most B_* forced primes and $\log p/(p-1)$ is decreasing, giving the same bound there. The discriminant term contributes at most $40 \log \mathrm{rd}(E)$. Finite costs and trace comparison.

The normalized mean of the root intersections. This AI-proposed candidate concerns the actual nonterminal root models and their prescribed metrics and boundary extension. The abelian CM host E has degree $2f$ with $f \geq 2$. Its type extends one embedding of the host imaginary quadratic field, and its signature is (p_j, q_j) at every label, with $q_j \geq 1$ and $p_j + q_j = r \geq 3q_j + 1$. The underlying level is good with auxiliary parameter $A = 1$ and level integer M'. The exceptional primes, those dividing $2|\operatorname{disc} E|M'$, number at most $4fr^2$. The selected primes consist of these together with the primes at most B_*. The compatible root models and their component identifications are fixed for the given PEL datum before the weight and cutoff are tested. At the tested weight $u = u'$ the regularity, all small-slot gaps, and the strict inequality for $B_* + 1$ are those established above.

Let $V_Z > 0$ be the geometric volume factor of a component and $I_{B_*,Z}(\ell_u^{D+1})$ its arithmetic intersection with the field and generic root-extension degrees divided out. Put $\ell_u = -\sum_{\tau,i} u_{\tau,i} x_{\tau,i}$, $G(u) = \deg_{\mathrm{Fl}}(\ell_u^D) > 0$, and

$$\Lambda_{B_*} = \log(B_* \,\mathrm{rd}(E)) + |\Re(L'/L)(1, \chi_{d_0})|,$$

where χ_{d_0} is the prescribed host quadratic character. The proposed bound is, with an absolute C,

$$\frac{\sum_Z I_{B_*,Z}(\ell_u^{D+1})}{(D+1)G(u)\sum_Z V_Z} \geq -C\left(\Lambda_{B_*} \sum_{\tau,i} |u_{\tau,i}| + q_j \sum_\tau (\max_i u_{\tau,i} - \min_i u_{\tau,i})\right). \tag{101}$$

The geometric degree on each component is $V_Z G(u)$, and equality of the V_Z's converts this volume-weighted quantity into the arithmetic mean. The finite-place trace cost above then gives the bound for $\mathrm{mean}_Z\, s_Z$. Root models and the geometric degree, The arithmetic mean inequality, and The host character.

Mumford's proportionality accounts for the geometric volume factors. Tamvakis computes the compact-dual arithmetic flag intersections and the Maillot–Rössler eigenspace formula applies on abelian schemes with the coefficient discriminant inverted. Results on Hodge volumes for signature $(n-1, 1)$ give arithmetic intersections of one canonical determinant line. The bound here concerns all the graded lines of the actual finite root model at the indicated cutoff. Mumford, Theorem 3.2, Tamvakis, Theorem 4, Maillot–Rössler, Theorem 7.26, and Bruinier–Howard, Theorem 8.6.1.

The proof uses the change on adjoining one good odd prime, its summable flag residues, two calibrating intersections, and the balanced limit on each fixed nonterminal cover. Throughout this proof write $p = p_j$, $q = q_j$, $N = r$, and $D = f\binom{N}{2}$.

Change at a newly included good prime. For an odd prime l unramified in the host and off the level, take the finite framed torsion flag on the common model. On the abelian locus its successive quotients are finite flat residue-field lines. Let $X_{\tau i}$ be a fundamental-character line of one such quotient, with its generic frame. The defining equations $x_{\tau i}^l = a_{\tau i} x_{\varphi\tau,i}$ give its cotangent space at $\varphi\tau$ as $X_{\varphi\tau,i}/X_{\tau i}^{\otimes l}$, where φ is residue Frobenius. The conormal sequences of the successive isogenies add to that of $[l]$. The total cotangent class is therefore $[\omega_{\varphi\tau}/l\omega_{\varphi\tau}]$.

Pair the k-th relative Newton class of each generically framed $X_{\tau i}$ with an old flag class Q of degree $D+1-k$ and sum with sign $(-1)^k$, calling the result t_τ. On the right the relative Newton class of V/lV is $k[\mathrm{div}(l)]p_{k-1}(V)$. Indeed splitting into lines reduces to the difference of their k-th powers after one scalar divisor twist. Further powers vanish on the principal vertical divisor. Taking $V = \omega_{\varphi\tau}$ and dualizing to Lie gives $l^k t_\tau = t_{\varphi\tau} + L_{\varphi\tau}$, where $L_\sigma = k \log l\, D_{\mathrm{fl}}(p_{\mathrm{Lie},\sigma,k-1} Q)$ in volume-normalized units. Iteration around the finite orbit gives $t_\tau = \sum_{m\geq 1} l^{-km} L_{\varphi^m\tau}$. Raynaud, Theorem 1.4.1 and Corollary 1.5.1.

We identify these lines with the prescribed finite root lattices. Let $n = l^a$, $a \geq 4$, be the finite marking. At a point of the marked normal model, henselize the local ring. The ordinary n-torsion of the canonical semiabelian group is quasi-finite and flat. Its open factor at the identity of the closed fibre is finite flat. It is a subgroup, since multiplication and inversion preserve the identity factor on the closed fibre and the sources of these maps on the finite factors are proper. The generic framed Weil characters restrict to points of its finite Cartier dual and extend over the normal base by finiteness. A local lift of each character to the semiabelian group has differential determined modulo $n\Omega$. Indeed two interpolants agree on the torsion germ by generic density, and the conormal map there is multiplication by n. Faithful flatness descends the germ from the henselization. Thus the image generated by these differentials and $n\Omega$ pulls back to the same image on the common normal graph. This applies to both character maps, including at exceptional divisors over the toroidal strata. Finite factors over Henselian local rings and The prescribed Weil differential.

For the scalar comparison a Kummer cover realizes each required log torsion group as a classical finite group. The period-lattice comparison uses $w = \zeta_l - 1$, so $(w^{l-1}) = (l)$, and the intermediate lattice $T \subset H \subset w^{-1}T$. Its Hodge quotient on T is the framed Weil image Ω^{HT}. Perfect duality identifies the kernel on T with the conjugate image dual W. The kernel on H is $w^{-1}\,\mathrm{Lie}$. The relative Newton number of H is t_τ, since its coordinate graded lines have ideals $X_{\tau i}^{-1}$. Hence the comparison sequence is

$$0 \longrightarrow w^{-1}\,\mathrm{Lie}\,/W \longrightarrow H/T \longrightarrow \Omega/\Omega^{\mathrm{HT}} \longrightarrow 0.$$

This sequence descends to one finite normalization. Locally present the right quotient by $\Omega^{\mathrm{HT}} \hookrightarrow \Omega$ and express its extension as a pushout of that injection. The two comparisons of the middle terms and the identities on overlaps are equations modulo w^2. The completed absolute integral closure and its uncompleted version agree at this precision. The finitely many equations therefore hold on a finite normal graph, where the presenting injection stays injective by its generic isomorphism. This gives additivity for the supported Newton classes on the finite model. Both generic root readings use the same quotient metrics. Inoue, Lemma 2.13, Lau, Theorem 9.8, and Howard, Theorem 2.2.2.

Additivity and the scalar comparison from w give the old-minus-new correction for $p_{\tau,k}Q$,

$$k \log l \left(\sum_{m \geq 1} l^{-km} D_{\mathrm{fl}}(p_{\mathrm{Lie},\varphi^m \tau, k-1} Q) - \frac{D_{\mathrm{fl}}(p_{\tau,\leq p_j,k-1} Q)}{l-1} \right).$$

In using the old Q, the other term is the change in Q multiplied by the new $p_{\tau,k}$. The change in Q is supported on the l-fibre. On the common graph $w\Omega \subset \Omega^{\mathrm{HT}}$ implies $n\Omega \subset l\Omega^{\mathrm{HT}}$, so the framed differentials give a surjection $T/lT \to \Omega^{\mathrm{HT}}/l\Omega^{\mathrm{HT}}$. The conjugate surjection gives the injection from W/lW. At the generic valuation of every reduced special component their composite vanishes by the Hodge–Tate complex. Local freeness and complementary ranks then give exactness on that component. The new blocks therefore make a flag in a trivial bundle on each reduced component, so its positive-degree ambient Newton classes vanish. This proves the jump for every complementary Q, with generic degrees divided out on projection from the common graph. For $k = 1$ and $D_{\mathrm{fl}}(Q) = 1$, the correction is zero if Frobenius preserves the type. If it exchanges sides, the ranks alternate $q_j, p_j, \ldots$ and the correction is $(q_j - p_j) l \log l/(l^2 - 1)$. The comparison on the torsion flags, The two-level formula, and Its characteristic number.

The flat calibration. The full determinant at a label has generic line $\det \mathrm{Lie}_\tau \otimes (\det \mathrm{Lie}_{\bar{\tau}})^{-1}$. The signed ordering of the two blocks gives this determinant. The determinant of the polarized Hodge metric is constant in a flat frame, and its arithmetic monodromy is finite since all eigenvalues are integral units of modulus one at every embedding. A suitable power of the flat frame trivializes the generic line, extending over the unipotent boundary. On a selected reduced fibre its determinant is that of the trivial ambient flag. Off the cutoff, the two filtrations on vertical curves make its degree multiply by l along the finite Frobenius orbit, and hence zero. Thus the divisor of the powered flat frame has degree zero on every

vertical curve. Zariski's lemma on an ample surface section makes it a rational multiple of each full fibre. The normalized point height of this determinant is constant on a connected component.

For its value choose a split CM point on the component. The normalized Hermitian lattice is locally unimodular: split factors have dual bases, and in an unramified quadratic factor one splits off unit lines using the surjective norm on units, also at residue characteristic two. The real signatures and finite local isometries identify the rational space by Landherr's theorem. Strong approximation for its noncompact special unitary group reduces the adelic lattice transport to $\operatorname{diag}(\det g, 1, \dots, 1)$. The reference lattice splits into p_j positive and q_j negative CM axes, with the first axis changed by a norm-one fractional ideal. This modification does not change its normalized degree. The saturated tangent axis for an ideal I is $I\mathfrak{D}_{E/K_d}$ times the elliptic tangent axis by trace duality. The normalization of (160) therefore gives ordinary determinant degree $(p_j - q_j)\log\Omega/2$ at this point. If X_B is the finite comparison cost per rank, the cutoff value differs by at most rX_B, with $X_B \ll \Lambda_{B_*}$ from the prime sum. This gives the full determinant calibration and the corresponding upper and lower point comparisons. The same calculation on the polarized abelian cusp quotient applies when its small tangent block is nonzero. Shimura, Section 5, Yuan–Zhang, Section 2.

The omitted good primes. In the local cutoff identity the summable correction is a sum of terms $k\log p_0 \sum_{m_0\ge 1} p_0^{-km_0}$ times flag residues with $k \ge 2$. Those residues can be bounded uniformly in the Frobenius labels. Write $G_0 = \sum_{\alpha,k} p_{\alpha,k} Q_{\alpha,k}$ for a homogeneous flag polynomial of degree $D+1$ decomposed in the power-sum generators, and let Z_0 be a block power sum of degree $k-1$ at any label. For generic rank-r bundles, the flag pushforward π_* and evaluation on compact flags D_{fl} give

$$D_{\mathrm{fl}}(Z_0 Q_{\tau,k}) = [P_{\tau,k}]\pi_*(Z_0 G_0),$$

where the bracket extracts the coefficient of the single ambient power-sum generator $P_{\tau,k}$. This follows from the projection formula. At $k \ge 2$ replacing the large block by the negative of the small block leaves the coefficient unchanged, since the difference contributes only a product of positive-degree symmetric generators. Root signs can be absorbed in the absolute value.

For $G_0 = F_u$, a single-label flag pushforward of the exponential gives $\det(e^{-u_i y_l})/\Delta(y)$ up to the orientation sign, where y_l are the ambient roots. Cauchy–Binet expands this as the weight Vandermonde divided by $\prod_{j=0}^{r-1} j!$, times $\sum_\lambda c_\lambda s_\lambda(-u)s_\lambda(y)$, with $c_\lambda = \prod_i (r-i)!/(r-i+\lambda_i)!$. Taking the coefficient of the power sum of degree k in $s_\lambda(y)$ uses only hooks of size k, with magnitude $1/k$ by the symmetric-group character formula. Multiplication by the small-block power sum corresponds to summing the $(k-1)$-st weight derivatives in the q_j small slots. If the minimum gap from such a slot to all others is g, the a-th derivative of the Vandermonde costs at most $(r/g)^a$ relative to its value by the product rule. With U_τ the spread at label τ, the b-th derivative of a Schur factor costs at most $k^b U_\tau^{k-b} s_\lambda(1^r)$. Here one can first translate the weights at each label, which affects only products with an ambient degree-one generator in the pushforward. The hook formulas give $c_\lambda s_\lambda(1^r) = f^\lambda/k!$ and $\sum_{\lambda\ \mathrm{hook}} f^\lambda = 2^{k-1}$. Summing the derivatives with $a+b = k-1$ bounds the flag residue by

$$\frac{k}{(D+1)G(u)}|D_{\mathrm{fl}}(Z_0 Q_{\tau,k})| \le q_j U_\tau \frac{2^{k-1}}{k!}\big(r(1+2\|u\|_\infty/g)\big)^{k-1}.$$

The dyadic Chebyshev estimate bounds $\sum_{p_0>B_*} \log p_0 \sum_{m_0\ge1} p_0^{-km_0}$ by $O(B_*^{1-k})$, uniformly for $k \ge 2$. The chosen B_* is at least $2fr^2(1+\|u\|_\infty/g)$, which bounds the total omitted contribution by $O((D+1)G(u)q_j\,\mathrm{spread}(u))$, the scale needed in the normalized mean. This derivation uses the local change of arithmetic model from the cutoff identity. Darondeau–Pragacz, Theorem 1.1.

A cusp estimate for the Hodge calibration. For the framed period lattice let $\lambda(y)$ be its first minimum, $V_0 = 2f\,\mathrm{rd}(E)$, and $r_0 = V_0^{-1/2}$. There are $2f$ linearly independent integers of E of house at most V_0. Indeed

the dual of the integer lattice in the trace metric is the conjugate codifferent, whose nonzero elements have absolute norm at least $|\operatorname{disc}(E)|^{-1}$. The arithmetic–geometric mean bounds its first minimum below by $\sqrt{2f}/\operatorname{rd}(E)$. Euclidean transference gives the claimed house bound. For periods v, w shorter than r_0, the integral Riemann pairing of v with aw, for each such integer a, has absolute value less than one. It vanishes, so those periods span a totally isotropic E-space. Its real dimension is at most $s = 2fq_j$. The index of a shortest vector in its saturated E-line is at most $(V_0\sqrt{2f})^{2f}$, by taking its products with the independent integers and comparing the covolume of the saturated lattice with the cube from Minkowski's theorem. Banaszczyk, Theorem 2.1.

Let $n_y(t)$ count nonzero isotropic vectors of at most this index and length less than t. For $t \le r_0$ it is zero unless $\lambda(y) < t$, and then it is between one and $(3t/\lambda(y))^s$ by packing balls of radius $\lambda(y)/2$. Unfolding the integral over the finite-volume component gives $\int n_y(t)\,d\nu = ct^a$, with $a = 2f(r-1)$ and invariant probability measure ν. For this, reduction theory gives finitely many orbits of rational isotropic lines. On each line there are finitely many sublattices of bounded index. Their generators form unit orbits, and the line stabilizer contains a congruence subgroup of these units, by extending their action over a rational hyperbolic plane. Hence only finitely many vector orbits occur. Each vector stabilizer has finite arithmetic volume, being an extension of a smaller unitary group by a unipotent group with no rational characters. The unfolded cone measure at each real place is $\delta(h(v,v))\,dv$, of homogeneity $2r-2$. The resulting product gives the claimed power of t, with zero constant if no isotropic vectors occur. Borel, Proposition 15.6 and Borel–Harish-Chandra, Theorem 9.4.

Put $Z_0 = \log^+(r_0/\lambda(y))$ and $A = \int n_y(r_0)\,d\nu$. The lower count gives $\nu(Z_0 > u) \le \min(1, Ae^{-au})$ for $u \ge 0$. Integrating this tail and using the upper count gives

$$A \le 3^s \mathbf{E} e^{sZ_0} \le 3^s \frac{a}{a-s} \max(1, A)^{s/a}, \qquad \int Z_0\,d\nu \le \frac{\log^+ A + 1}{a} \le 3,$$

since $a \ge 3s$. For a saturated isotropic lattice I of E-rank $j \le q_j$, define $n(I)$ as the product of the small-side Gram determinants divided by the saturation index. The real covolume is $\operatorname{rd}(E)^{fj} n(I)$ by the trace determinant and agreement of the two Gram forms on an isotropic space. Minkowski bounds it below by $(\lambda(y)/\sqrt{2fj})^{2fj}$. Thus the cusp correction $\max(0, \sup_I(-\log n(I)))$ has mean at most $fq_j(6 + \log V_0 + \log(2fq_j \operatorname{rd}(E)))$. The normalized top power of the invariant Hodge form is precisely ν, so the same estimate controls its cusp integral in the Hodge calibration.

A tangent slope estimate for the Hodge calibration. For a principally polarized abelian variety of dimension g with stable Faltings height h_F, let $k_0 = \lceil 2 + g + h_F^+ \rceil$ and take the cubist model of the $n_0 = 3k_0$-th power of its polarization over a finite extension. The integral section bundle with its Haar L^2 norms is semistable of slope $\mu_{n_0} = -h_F/2 + (g/4)\log(n_0/(2\pi^2))$. In the tangent lattice choose a saturated line V over any finite extension and filter by the order of vanishing along its formal exponential at zero. The quotient of the first k_0+1 jets has rank k_0+1. Indeed in the very ample cube of the polarization choose one section nonzero at zero and one vanishing there with nonzero derivative on V. Their products of total degree k_0 have all the required orders. The order-i quotient maps integrally into $(V^\vee)^{\otimes i}$, since the ordinary derivative uses the integral invariant derivation and lower derivatives vanish. The cubist rigidification is used at zero. [89, Prop. 2.8, Thm. 2.10, Props. 2.11 and 2.13] Gaudron, §2.3.2.

At an embedding σ, the order-zero norm bound is $\|s\|_{\sup} \le C_\sigma \|s\|_2$, with $\overline{\log C_\sigma} = O(g\log(2+g+h_F^+))$ by the matrix lemma for period minima. On the exponential of a unit vector of V, the cubist frame multiplies this upper bound by $e^{\pi n_0 |z|^2/2}$. Apply Cauchy's derivative estimate at radius $\sqrt{i/(\pi n_0)}$ and use $i! \le i^i$. The derivative norm on the quotient is at most $C_\sigma(\pi e n_0 k_0)^{i/2}$. Summing the degree inequalities for the jet quotients and using semistability gives

$$(k_0+1)\mu_{n_0} \le -\tfrac{k_0(k_0+1)}{2}\widehat{\deg} V + (k_0+1)\overline{\log C_\sigma} + \tfrac{k_0(k_0+1)}{4}\log(\pi e n_0 k_0).$$

Division gives $\widehat{\deg} V = O(\log(2+g+h_F^+))$ as an upper bound. The comparison of absolute minimum and maximal slope adds at most the harmonic number $H_{g-1} = O(\log(2+g))$. For a smooth integral

group model the Néron mapping property injects its tangent lattice into the stable Néron tangent lattice. Saturation can only increase sublattice degrees, and h_F is at most minus the degree of the original tangent lattice in the Riemann norm. These give the same logarithmic bound for the tangent spaces used in the abelian quotient of the Hodge test. Gaudron, Theorem 3.2 and §4.1 and Ballaÿ, Corollary 9.3.

From point heights to the Hodge calibration. Let M_{Hod} be the small determinant line on the base of the flag component. On an algebraic point with toric rank a over E, the abelian quotient has small eigenspace rank $f(q_j - a)$. Write $x_{\mathrm{sm}}, x_{\mathrm{lg}}$ for the ordinary degrees of the orthogonal small and large saturated eigenspaces of its tangent lattice. Their sum is at most minus the tangent height h_G.

Here the determinant comparison is applied to the actual polarized abelian quotient. If I is its saturated isotropic period lattice of E-rank a, its full Tate lattice is $I^{\perp}/I$. Over each local factor of $\mathcal{O}_E$ the inclusion of I is a direct summand. Trace-perfectness of the original pairing makes restriction onto its inverse-different-valued dual surjective. The annihilator calculation then shows that the induced pairing on $I^{\perp}/I$ is trace-perfect, with the same polarizing ideal and host square root. For example, an integral linear form on $I^{\perp}/I$ extends from $I^{\perp}$ to the whole local lattice since the quotient is projective. Its pairing vector annihilates I, and the radical on $I^{\perp}$ is exactly I. At each real label, removing the paired isotropic planes gives ranks $(p_j - a, q_j - a)$. At every point-field embedding these ranks use the two labels selected by the image of the same host square root. The normalized Lie metric on the quotient is the one from this principal polarization.

If $q_j > a$, put $r' = r - 2a$; then $r' \geq 3(q_j - a) + 1$. After finite extension mark the quotient with a good auxiliary level of parameter $A = 1$ for this lower PEL datum, for instance using the original auxiliary marking modulus on a Tate basis of the quotient. Its full torsion basis is defined over a finite extension. This uses the self-dual quotient lattice and a new marking on it. At a finite valuation of the point field the torus extends as a closed split subtorus after passing to a splitting field, with locally free tangent quotient. The semiabelian quotient is the extension of this marked lower abelian variety on the trait, by full faithfulness for semiabelian models. The saturated ordinary eigenlattices used here are those of that quotient. In particular their formation and their normalized degrees commute with finite extension of the point field. Apply the flat calibration on the lower component. Its normalized cutoff determinant is constant, its CM frame gives $\frac{1}{2}((p_j-a)-(q_j-a))\log\Omega$ at each label, and the cutoff-to-ordinary comparison costs at most $r'X_B$ both at the CM point and at the point being evaluated. The same B and exceptional prime allowance suffice: the host, different, and auxiliary level modulus are unchanged, and $r' \leq r$. Consequently at each label

$$\left| \widehat{\deg}_n \operatorname{Lie}_\tau - \widehat{\deg}_n \operatorname{Lie}_{\bar\tau} - \tfrac{1}{2}(p_j - q_j)\log\Omega \right| \leq 2r'X_B.$$

Summing these comparisons and using $|\log\Omega| + X_B \ll \Lambda_B$ gives $x_{\mathrm{sm}} - x_{\mathrm{lg}} \ll fr\Lambda_B$. If $q_j = a$, the small tangent space of the abelian quotient has degree zero, and the argument needs only the torus calculation. The same comparison with $a = 0$ applies at interior points. The actual abelian quotient and paired eigenlattices and Flat comparison in the signature of the tested polarized abelian variety. Lan, Propositions 3.3.1.5 and 3.3.1.7 and Anantharaman, Chapter IV, Theorem 4.C.

If h_G exceeds this upper bound, addition gives $x_{\mathrm{sm}} \leq 0$. Otherwise the logarithmic tangent bound gives $x_{\mathrm{sm}} \ll f(q_j - a)\Lambda_B$, since $g \leq fr$ and $\log(2 + fr + h_G^+) \ll \Lambda_B$.

For the boundary point bound, let D_I be the index of the $\mathcal{O}_E$-lattice generated by a chosen E-basis of I. If $\mathfrak{j}$ is the coefficient ideal of its top exterior power in that basis, $\mathrm{N}\mathfrak{j} = D_I^{-1}$. The saturated torus tangent determinant at label τ is $\tau(\mathfrak{j}\,\mathfrak{D}^a_{E/\mathbf{Q}})$ in the projected cocharacter basis. This follows from intersection with the axis in the trace-dual basis, and gives finite normalized degree $\log D_I/[E : \mathbf{Q}] - a\log\mathrm{rd}(E)$ per small label. The winding period $2\pi i$ and tangent metric πH give the factor $(4\pi)^{-a}$ times the projected period Gram determinant. Use a torus frame followed by lifts of a frame of the abelian quotient. The determinant identity for the map to $I^\vee$ cancels the projected Gram determinant in $\phi - \frac{1}{2}\log n(I)$. This

weight extends smoothly in that local frame as the smaller Hodge determinant weight plus its constant cusp factor. Here $n(I)$ is the product of the Gram determinants divided by D_I. The archimedean constant in the weight is $\frac{1}{2}\log D_I - \frac{1}{2} f a \log(4\pi)$. Summing the finite index terms over the f small labels cancels the archimedean index term in the degree. The remaining torus allowance is $-f a \log \mathrm{rd}(E) + \frac{1}{2} f a \log(4\pi)$; the cutoff comparison costs at most $f q_j X_B$ on the full small line. Together with the abelian estimate this gives $-C f q_j \Lambda_B$ as a lower bound for the corrected point height in the actual cutoff line.

On the domain the corrected weight is $\phi + U/2$, with $\phi = -\log \|\cdot\|$ the small determinant weight and $U = \max(0, \sup_I(-\log n(I)))$. For a fixed isotropic plane I, pairing with it gives a surjection from the negative q_j-plane to $I^\vee$, with kernel the negative plane in $I^\perp/I$. The determinant of the quotient metric is the determinant of the input metric divided by the kernel's determinant. The Gram factor in $n(I)$ cancels the quotient term in $\phi - \frac{1}{2}\log n(I)$, up to a pluriharmonic log determinant. The resulting small determinant weight on the smaller domain is plurisubharmonic. The maximum is locally finite on the open domain, since positive corrections require bounded covolume in the integral exterior-power lattice. Thus the corrected metric is semipositive there.

Write $\psi = \phi + U/2$. On a smooth toroidal period chart the real part of a Siegel period matrix Z is bounded in a fundamental strip and $c\,1 \le Y = \operatorname{Im} Z \le CR\,1$, where $R = 2 - \log \|s_\partial\|^2$ for the local boundary section. The period norm is given, up to the fixed factor, by

$$(m + \operatorname{Re} Z\, n)^{\mathrm{t}} Y^{-1} (m + \operatorname{Re} Z\, n) + n^{\mathrm{t}} Y n.$$

Its matrix and inverse have size bounded by a power of R; this holds on every exterior power of the fixed integral period lattice as well. In the Gram covolume defining $n(I)$, the division by the index is the passage to the saturated integral lattice, so the bound applies uniformly to all the isotropic sublattices in that chart. Thus ψ grows at most as $C(1 + \log R)$ from above. In a polydisc with punctured coordinates q_k, add $\eta \sum_k \log|q_k|$ for each $\eta > 0$. This gives a plurisubharmonic weight bounded above at the divisors. After removal of the punctures, the maximum principle on a smaller polydisc bounds it by its values on the distinguished boundary, uniformly as $\eta \downarrow 0$. It follows that ψ extends as a plurisubharmonic weight locally bounded above. At a boundary point it dominates the smooth corrected weight for that point's actual cusp, so it is locally bounded below and its value there satisfies the point-height comparison just proved.

The finite small dual determinant M_{Hod} is relatively nef on the cutoff model. On a selected reduced fibre it is a quotient determinant of the framed trivial bundle. On a curve in any other fibre the Hodge and conjugate filtrations give equality of the degrees of the paired Lie blocks and $-\deg \mathrm{Lie} \ge 0$. The small dual determinant there has degree $-\frac{1}{2}\deg \mathrm{Lie} \ge 0$. On the generic compactification twice M_{Hod} is the full Hodge determinant plus a sum of flat signed determinants. The flat determinants are torsion and the full Hodge determinant is, up to a power, pulled back from the ample line on the minimal compactification. Thus the generic line is semiample and big. Pass to a smooth generic model and take its graph over the finite model if necessary. Pullback preserves the finite nefness and the point bound, and the degree in the projection formula cancels in the normalized intersections.

Choose an ample auxiliary arithmetic line $\bar{A}$ with a strictly positive smooth metric and nonnegative point heights. For fixed positive rational ϵ, regularize the bounded plurisubharmonic weight ψ by smooth weights $\psi_j \downarrow \psi$ on each complex fibre, averaged over complex conjugation. They have curvature at least $-\delta_j c_1(\bar{A})$, with $\delta_j \to 0$. In Demailly's regularization the decreasing continuous upper bounds for the Lelong numbers converge uniformly to zero for a bounded weight. Hence for large j the line with weight ψ_j plus $\epsilon\bar{A}$ is relatively nef and semipositive. The lower point-height bound survives since $\psi_j \ge \psi$. Shifting its weight by $C f q_j \Lambda_B$ makes it arithmetically nef. Its top intersection is nonnegative, and undoing the shift gives a normalized lower bound $-C f q_j \Lambda_B$. For fixed ϵ, monotone convergence of the bounded mixed Bedford–Taylor energies gives the bound at ψ, and then the intersection polynomial in ϵ gives the bound at $\epsilon = 0$. Demailly, Theorem 1.1 and Moriwaki, Corollary 5.5.

For the uncorrected weight ϕ the mixed-intersection limit in the compatible toroidal metrics equals its finite-energy intersection. Indeed the curvature has full mass by goodness of the canonical Hodge metric, and the log-log potential is integrable against the product Poincaré majorant. The mixed tail estimate also controls the weighted truncations in the energy formula. Concavity of that energy bounds the increase in normalized intersection from ϕ to ψ by

$$\frac{1}{2[K_{\mathrm{Hod}} : \mathbf{Q}]\ \deg(M_{\mathrm{Hod}}^{d_{\mathrm{b}}})} \sum_{\sigma} \int U_{\sigma}\ c_1(M_{\mathrm{Hod},\phi,\sigma})^{d_{\mathrm{b}}},$$

where d_{b} is the dimension of the base and K_{Hod} is its cutoff coefficient field. This uses the weight convention $-\log\|\cdot\|$, the same as for the normalized point degree. The curvature measures after division by the displayed geometric degree are the original invariant probability measures. The cusp mean bounds the increase by $O(f q_j \Lambda_B)$. Pushing the top product of within-block root differences down the flag fibres multiplies both the arithmetic intersection and its normalizing degree by the same positive geometric factor. This gives $P_B(H) \geq -C_{\mathrm{Hod}} f q_j h\, \Lambda_B$ in the normalization of the small Hodge wall, with h the geometric factor used there. Berman–Boucksom–Guedj–Zeriahi, Propositions 2.1, 2.4 and 2.8 and Berman–Freixas i Montplet, (4.4).

Point comparison on the complete reference flag. Fix one nonterminal cover and one component. In this paragraph all constants may depend on this choice. Use the ordinary semiabelian Lie lattices, their saturated eigenbundles on a common normal modification, and a smooth metric with orthogonal eigenspaces. Let $H \geq 0$ be a smooth ample height on its complete flag. For a point z over $\overline{\mathbf{Q}}$, write V for the pullback of the full integral Lie lattice, with this smooth metric. Then

$$\widehat{\mu}_{\max}(V) \leq C \log(2 + H(z)). \tag{1}$$

All degrees of point lattices are divided by the point-field degree.

Here is a proof that includes points on the semiabelian boundary. Stratify the fixed generic fibre by torus rank. On each irreducible stratum the maximal torus is a torus scheme. Pass to a finite extension splitting its generic fibre, and normalize its projective arithmetic closure in that extension. After shrinking the stratum the torus is split there. Its closure on the normal arithmetic scheme is a closed torus; the generic isomorphism with a split torus and its inverse extend. Noetherian induction on the omitted locus gives finitely many proper models covering all strata in this way. On each such model the differential of the closed torus gives a fibrewise injection. The Gram matrix in its cocharacter basis and its inverse are bounded for the pulled-back smooth metric by compactness. Thus the torus Lie lattice K at any of these points has bounded maximal slope and bounded absolute degree. For the split Euclidean lattice the subdegree assertion follows directly by taking an exterior product and applying the product formula to a nonzero coordinate. These assertions persist under extensions of the point field. Lan, Propositions 3.3.1.5 and 3.3.1.7, and Corollary 3.3.1.10.

For the metric on the abelian quotient, choose a symplectic basis adapted to a cusp of the underlying principally polarized Siegel space and restrict its chart to the PEL domain. In the normalized Lie coordinates the tangent norm matrix is πY^{-1}, where $Y = Y_0 + \sum_j y_j A_j$, $y_j = -\log|q_j|/(2\pi)$, and the A_j are positive semidefinite real monodromy matrices supported on the toric coordinates. The bounded matrix Y_0 is uniformly positive on $K_0 = \bigcap_j \ker A_j$, the abelian directions at the chart centre. A face with vanishing coordinates indexed by I has toric subspace $J_I = \sum_{i\in I} \operatorname{im} A_i$. In a basis adapted to the saturated period sublattice in J_I the quotient map discards its toric coordinates. Its period matrix is the remaining principal block: taking periods modulo the torus discards the same rows and columns. The induced polarization is principal because the quotient of a unimodular symplectic lattice by a primitive hyperbolic summand is unimodular. Equivalently one can compress $Y_0 + \sum_{j\notin I} y_j A_j$ to $J_I^{\perp}$ in a fixed Euclidean frame. On this complement the common kernel of the remaining compressed matrices is K_0, by semidefiniteness.

On its orthogonal complement their sum is positive definite. After shrinking the chart all finite y_j are sufficiently large: in the decomposition by K_0 the lower diagonal block tends uniformly upward, the off-diagonal block stays bounded, and the Schur complement is bounded below by half the lower bound on $Y_0|_{K_0}$. This also treats a zero-dimensional block by omission. There are finitely many faces on a chart. The inverse residual period matrix therefore gives, uniformly, an abelian quotient tangent norm at most C times the smooth norm on V. The normalized Lie coordinates extend across the boundary using dw in the additive directions and $d\log u/(2\pi i)$ in the multiplicative directions. Lan's canonical extension and its restriction to the boundary chart identify these with the algebraic semiabelian Lie and cotangent bundles. The use of the underlying Siegel space works after a finite compatible refinement and for every complex embedding of the fixed cover. Lan, *Toroidal compactifications of PEL-type Kuga families*, Proposition 6.9(1)–(4) and Lan, analytic comparison, Theorem 5.2.12 and Lemma 5.2.16.

Over a finite extension choose the stable Néron identity model of the abelian quotient A_z. The generic quotient homomorphism to A_z extends to it from the smooth semiabelian model. Its differential $f : V \to V_A$ has full generic rank. The closed torus Lie subbundle is saturated and has the generic kernel rank, so $\ker f = K$. For any saturated $W \subset V$ the determinant inequality for the map from its quotient by $W \cap K$ to $\operatorname{sat} f(W)$ gives

$$\widehat{\deg} W \le \widehat{\deg}(W \cap K) + \widehat{\deg}(\operatorname{sat} f(W)) + C \operatorname{rk} f(W).$$

The map is integral at every finite place and has the uniform operator norm just proved. For $W = V$ this implies $h_F(A_z) = -\widehat{\deg} V_A \le -\widehat{\deg} V + C$, up to the fixed normalization of the tangent metric. The ample height comparison for the fixed smooth line $\det V$ bounds the last expression by $C(1 + H(z))$. Apply the logarithmic maximal-slope bound for a principally polarized abelian tangent lattice to V_A in the preceding inequality for every W. This proves (1). If A_z has dimension zero the torus estimate gives (1) directly. The reference tangent slope on the complete model.

The determinant of the flat τ-space is $D_\tau = \det V_\tau \otimes (\det V_{\bar\tau})^{-1}$ in the homological Hodge sequence. Its monodromy determinant belongs to $E^1 \cap \widehat{\mathcal{O}}_E^\times = \mu(E)$, where $E^1 = \{a \in E^\times : a\bar a = 1\}$ is viewed diagonally in the finite ideles. Indeed the monodromy preserves the integral PEL lattice and the absolute values of a norm-one CM unit in every complex embedding are one. A fixed power of the determinant thus has a parallel generator on the connected component. This generator extends across the toroidal boundary without a zero or pole: the determinant residue of the unipotent local monodromy is zero and the canonical extended Hodge sequence identifies its determinant with D_τ. The holomorphic trivialization on the projective complex component gives an algebraic trivialization after a number field extension by cohomology and base change. Its divisor on the fixed arithmetic model is vertical and bounded above and below by a sum of whole fibres. The smooth model metric is bounded in this generator. Consequently $|\widehat{\deg} D_\tau(z)| \le C$. For the finite reference eigenspaces one may pass to a common graph to flatten them. The ordinary saturated point eigenspaces differ from the resulting bundles only by bounded factors at the fixed ramified primes: the integral multiple $\mathfrak{d}_\tau e_\tau$ of the coefficient idempotent sandwiches the projection and the intersection. Equivalently, uniform scalar inclusions of the corresponding locally free lattices on the graph can be pulled back at every point.

In the large block, the degree of a root prefix is the degree of a sublattice of V. In the small block it is the degree of the complementary prefix before reversing and dualizing, plus the degree of D_τ. All proper prefix degrees are therefore at most $C\log(2 + H(z))$ by (1). The final prefix is D_τ. Summation by parts in $\ell_\mu = -\sum \mu_{\tau i} X_{\tau i}$ now gives, uniformly for bounded dominant μ,

$$h_{0,\mu}(z) \ge -C(1 + \log(2 + H(z))), \qquad |h_{0,\mu}(z)| \le C(1 + H(z)). \tag{2}$$

The absolute bound follows from the ordinary height comparison for the finite list of smooth root lines on this fixed model.

Uniform finite costs at the wall. For a cutoff T include the fixed primes dividing $2\operatorname{disc}(E)M'$ even when they exceed T. Choose one common normal projective graph carrying the reference and actual

flag bundles and all needed torsion markings. The common generic identification is the identification of canonical flags on the complete characteristic-zero toroidal model. An inclusion between two root lines with a scalar multiplier can be tested at the height-one valuations of this graph. It then extends everywhere as a section of their invertible Hom bundle, by normality. Pulling the section back along a horizontal point preserves inclusion. This gives comparison at a boundary point even when that point never enters the abelian open. Every point of the original model lifts to the proper dominating graph after a finite field extension, which preserves its normalized heights.

At a newly marked unramified odd prime l, the finite Weil differentials of the l^4-marking and $l^4\Omega$ give the finite image Ω', also on the toroidal valuations. The established inclusions $(\zeta_l-1)\Omega \subset \Omega' \subset \Omega$ commute with the unramified eigenprojections. Intersecting two lattices in this sandwich with a common rational prefix, and taking the image in the rational graded line, preserves each scalar inclusion. Dualizing reverses the two inclusions. Thus each signed graded root costs at most $\log l/(l-1)$ in absolute normalized point degree. This argument includes exceptional divisors of the graph. At each fixed bad prime the printed trace different in the large block and the dyadic bound $4\Omega \subset \Omega'$ give a fixed scalar in both directions. One can also fix the proper marked graph at this prime first. Its Cartier comparison divisor is vertical there and has bounded coefficients relative to the whole fibre. Further markings preserve these bounds: the same generic l^4-frame already exhausts the Weil characters, their restrictions to the finite flat factor at the identity extend uniquely, and their differentials with $l^4\Omega$ have the same image after pullback to every height-one trait. At unmodified primes the two ordinary lattices agree. It follows from $\sum_{l\le T} \log l/l = O(\log T)$ that (2) loses at most $C(1+\log T)$ for these finite changes. The sum estimate follows since $\lfloor T/l \rfloor \ge T/(2l)$ and $\sum_{l\le T}\lfloor T/l\rfloor \log l \le \log(T!)$. Order, signs and quotient of the prescribed graded lattices.

The prescribed mixed Hodge intersections. On one fixed nonterminal cover use E_i, F_j for the increasing steps of the large and small Lie flags. The signed roots at τ are E_i/E_{i-1} for $1 \le i \le p$ and $(F_{q-a+1}/F_{q-a})^\vee$ in slot $p+a$. In a flag-adapted frame put G_i equal to the leading i by i Gram matrix and $\det G_0 = 1$. The first block has graded norms $(\det G_i/\det G_{i-1})^{1/2}$; the other block has the dual quotient norms in reverse order. The dual prefix determinants P_i are therefore the quotient determinants $\det E_i^\vee$ for $i \le p$ and $P_N \otimes \det F_{q-a}^\vee$ for $i = p+a$. The full P_N has the flat determinant metric, with the fixed scale of the given polarization. These are the metrics specified on the open abelian locus. In the released origin the logarithm of the singular metric of each P_i is the negative sum of the logarithms of its first i root norms; the same identity holds for the smooth approximants. Thus controlling all prefixes controls the actual signed roots by successive differences. The polarized Lie norm, The signed graded lines, and The fixed singular and regular prefix ties.

Use the fixed smooth toroidal flag and, when two compactifications are involved, resolve their fixed joint graph. The generic big-cell coordinates of the flag over the moduli base belong to the fixed field in the origin, by the transported flag steps. The generic flag condition makes the index of their field in the cutoff field equal to Fl. deg. The fixed field has the same index by the origin degree identity. These subfields therefore coincide. The resolution and its ample positive line can be fixed before varying the cutoff. The big-cell flag degree and The fixed-field inclusion and its degree. On an adapted Siegel chart the cotangent norm is $\pi^{-1}Y$, where $Y = Y_0 + \sum y_j A_j$, $y_j = -\log\|s_j\|^2$, with constant positive semidefinite monodromy matrices A_j in normalized period coordinates. Choose projective real-analytic norms for the boundary lines. Then Y_0 is real analytic and bounded. These are the normalized additive and multiplicative frames of the algebraic semiabelian Lie bundle by the canonical Hodge sequence and toroidal comparison. The coefficient of y_j is an intrinsic tensor near its divisor: on a simple point, a real analytic tensor U for which $U\log|q_j|^2$ is real analytic must vanish. Otherwise its first nonzero transverse Taylor term along a ray gives a nonzero power times a logarithm. Analytic continuation gives equality of the tensors on chart overlaps at crossings. The same argument preserves CM orthogonality. Lan, *Toroidal compactifications of PEL-type Kuga families*, Proposition 6.9 and analytic comparison, Theorem 5.2.12 and Lemma 5.2.16.

Choose the boundary norms and the collars in conjugate pairs, so that the regularized metrics are arithmetic Hermitian metrics. On smaller collars replace y_j by $y_{j,t} = -\log(\|s_j\|^2 + t^2)$, interpolating smoothly on the fixed outer collars and taking all $y_{j,t} \geq 2$. The resulting cotangent metrics h_t give simultaneous quotient metrics on the two complete flags. The common kernel of the A_j has uniformly positive Y_0 and the sum of the A_j is positive on its complement. The Schur complement gives $h_t \geq cI$ for small t. In a normalized frame $A_j \leq Ch_t/y_{j,t}$, since Y_0 is bounded. The Chern curvature formula then bounds the ambient curvature endomorphism in absolute value by a multiple of

$$\Omega_t = C_0\omega_0 + \sum_j dd^c(-\log y_{j,t})$$

for a fixed Kähler form ω_0 and sufficiently large C_0. Indeed its derivative terms use $|\partial y_{j,t}|^2/y_{j,t}^2$ and $|dd^c y_{j,t}|/y_{j,t}$. Differentiating $\log(\|s_j\|^2 + t^2)$ yields a bounded smooth term and a nonnegative rank-one term. The latter divided by $y_{j,t}$ and the squared gradient both occur with positive signs in Ω_t. This is a positive closed majorant in a fixed class. Quotient curvature has a nonnegative second fundamental term, so every prefix curvature is bounded below by $-C\Omega_t$ on pullback; the full determinant has the two-sided bound there. Powers of the signed curvatures can consequently be expanded into positive mixed products by adding multiples of the majorant and a fixed smooth positive form on the flag.

Let $R = 2 - \log\|s_D\|^2$ on the flag, with D the total boundary. With empty boundary take $R = 2$. The resulting positive factors can all be dominated by $\omega + dd^c w$ for a sufficiently large fixed Kähler form ω, with $-C(1 + \log R) \leq w \leq C$. For the compatible metrics the linear algebra bound $cI \leq h_t \leq C(1 + \sum_j y_{j,t})I$ bounds the absolute values of their relative weights by $C(1 + \sum_j \log y_{j,t})$. Adding a sufficiently large multiple of the majorant adds the potential $-C' \sum_j \log y_{j,t}$, which gives the asserted upper bound and preserves the lower bound. For the released scalar prefix metrics these inequalities follow directly from $\phi \leq \phi_t \leq C$ and the log-log lower bound on ϕ; the curvature loss is $O(\log\log(1/t)/\log(1/t))$. On a fixed joint resolution, the origin identifies both interiors with the abelian locus, so the positive multiplicities of their common boundary components give equivalent logarithmic size functions. These estimates also apply to a fixed auxiliary line with smooth metric. The actual scalar regularizations and their errors.

Here is a uniform integrability estimate. Set $\rho_T = \max(\log\|s_D\|^2, -T)$. For $0 \leq z \leq 1$ with $\omega + dd^c z \geq 0$, telescope its d-th power and integrate by parts. Since $dd^c\rho_T \geq -C_D\omega$, one obtains

$$\int (-\rho_T)(\omega + dd^c z)^d \leq \int (-\rho_T)\omega^d + C_D d \int \omega^d.$$

The right side is bounded independently of T by ordinary integrability of the boundary logarithm. Thus the mass of $R > T + 2$ is $O(1/T)$. On $2^j T < R \leq 2^{j+1}T$ truncate each w below $-C'(1 + \log(2^{j+1}T))$ so that it agrees there with its truncation. Divide by its oscillation bound and translate to $[0, 1]$. Positivity bounds the mixed measure by a fixed multiple of the d-th power of the oscillation bound times the Monge–Ampère measure of the mean of these normalized potentials. Summation over j bounds the integral of $1 + \log R$ on $R > T$ against every original mixed product in absolute value by $C(1 + \log T)^{d+1}/T$. The constants may depend on this fixed smooth flag. This proof applies to products whose factors have different smoothing parameters.

Keep all finite root lines fixed. The change of a smooth mixed arithmetic top intersection telescopes into integrals of one weight difference against d curvature factors. On compact subsets of the open both smooth families converge to the same Hodge weights; the absolute masses of the products are bounded. The tail estimate sends the remaining integrals to zero uniformly. The compatible block caps and the actual scalar roots thus define the same mixed arithmetic functional. For $L = \log(1/t)$ the released bounds on $R \leq cL$ give scalar weight error $O(t^{1/3} + (1 + \log L)/L)$ for a fixed sufficiently small c. Applying the tail estimate to its complement, first comparing two scalar parameters and then letting the smaller tend to zero, gives error $O((1 + \log L)^{d+1}/L)$. At $t = \exp(-1/s^2)$ this is $O(s)$ and the curvature loss is $o(s)$. Each metric-change integral on a marked cutoff model is the pullback of this fixed complex integral. The generic degree

and the number-field embedding factors cancel against the actual denominator $[c.F : \mathbf{Q}]\,\mathrm{Fl.deg}$. This proves cutoff uniformity at the wall and allows the use of compatible metrics in the rank-two reflection and the small Hodge test. Normalization by the component field and root cover and The convergence comparison at the wall scale.

The limiting intersection on a balanced class. Fix a nonterminal PEL cover in the preceding height statement. Put $S_T = \{l \le T : l \text{ prime}\} \cup \{l : l \mid 2\,\mathrm{disc}(E)M'\}$, $d = f\binom{N}{2}$, and $P_T = \sum_c I_{T,c} / \sum_c v_c$. Each $I_{T,c}$ divides out the coefficient-field degree and the generic degree of the root cover. Write η for the arithmetic functional of the product of Euclidean complete flags, D_{fl} for its geometric functional in degree d, and define

$$S(G) = D_{\mathrm{fl}}\left(\sum_{\tau,i\le p} \partial_{\tau i} G\right), \qquad U(G) = D_{\mathrm{fl}}\left(\sum_{\tau,i} \partial_{\tau i} G\right).$$

For every homogeneous root polynomial G of degree $d+1$ with $S(G) = U(G) = 0$ one has

$$\lim_{T\to\infty} P_T(G) = \eta(G). \tag{102}$$

The limit is taken separately on each fixed stage and weight. The two original arithmetic calibrations then give an absolute bound as the stage varies.

First consider the discrepancy on powers, $A_T(v) = (P_T - \eta)(\ell_v^{d+1})$. In each of the two signature blocks A_T is alternating. Forgetting an interior step there gives the projective bundle of the rank-two subquotient V_2 on the partial flag. For compatible Hermitian metrics its two roots are $x, c_1(V_2) - x$ in arithmetic degree one; the signs and reverse order of the conjugate block give the dual rank-two description. The sum of a power and its swapped power is symmetric in these roots. In the ordinary Segre pushforward it vanishes: a symmetric polynomial descends and a pullback pushes to zero in relative dimension one. The remaining arithmetic term is the fibre integral of the secondary form, an ordinary characteristic number of V_2 and the unchanged roots. Pull it back and multiply by the geometric quotient root of fibre degree one. It is now a top-degree root polynomial, so geometric proportionality makes it v_c times the same compact-dual number. This proves alternation after subtraction. The calculation applies on the genuine flag bundles of the two flattened image lattices on the marked base. Their signed roots coincide on a common normal graph with the prescribed saturated roots, by the height-one Weil differential comparison and then reflexivity. One may make a regular alteration of the base before taking these flag bundles. All contributions scale by the same generic degree. For the actual Hodge metrics the compatible smooth calculation passes to the mixed limit by the comparison of toroidal metric approximations. Mourougane, Theorems 3–4 and Section 8 and The adjacent reflection.

There is also a transitive label symmetry of A_T. The labels form a torsor for $\mathrm{Gal}(E/k_0)$ over the chosen embedding of $k_0 = \mathbf{Q}(\sqrt{-d_0})$. Twisting the action by such an automorphism preserves the determinant condition, the polarization and the split level ideal, which is extended from k_0. For completeness, the completed trace-perfect lattice is isometric to its twist: trace duality gives a skew-Hermitian perfect form with values in the inverse different. A split pair is a lattice and its dual. At an unramified quadratic factor take a trace-zero unit, also available in residue characteristic two, and a generator from the real field of the inverse different. Both forms become unimodular Hermitian forms. A nondegenerate residue Hermitian form has a vector of nonzero norm by polarization and nondegeneracy of the residue trace. Scale a lift to norm one using the surjective norm on units and split off its line. Induction gives the desired local isometry. The host has only these two kinds of finite factor. Write φ_n for the reduction of the resulting semilinear isometry modulo n. Precompose the full Weil markings by φ_n, leaving their cyclotomic generators fixed. At level M the marked map factors through the lattice quotient by $\mathfrak{m}$ and must agree with this Weil marking on $\bar{\mathfrak{m}}\Lambda$. Both ideals are preserved by the isometry; also $M\Lambda \subset \mathfrak{m}\Lambda$. Hence $\eta_{\mathrm{new}}(a) = \eta_{\mathrm{old}}(\varphi_M a)$ has exactly the prescribed kernel condition and restriction. The auxiliary level map uses φ_A in the same way (here $A = 1$). The opposite twist with inverse isometry supplies an inverse on the original moduli

scheme over its coefficient field. The resulting permutation of the components carries τ to $\tau \circ g$. On the joint normal marked graph the isomorphism of polarized semiabelian groups and its inverse extend from the generic fibre. The finite Weil characters generating the images at every $l \in S_T$ are permuted, so their projected differentials, the trace different, the two saturated flags and their signed duals are transported. On the ordinary fibres and at infinity this is the isometry of the same Lie spaces and polarization norms, at every embedding. The common mixed metric limit consequently transports too. No compatibility of choices of auxiliary isometry as g varies is needed; each label permutation preserves the component sum. The geometric normalization and the product of Euclidean flags have the same symmetry. The cutoff component sum under relabeling.

To determine the limiting cross-block walls, take a rational weight v, strictly dominant except at $v_{\tau p} = v_{\tau,p+1}$ at a single label, with $0 < v_{\sigma,N} \le \cdots \le v_{\sigma,1} < 1/2$ at every label σ. Put $\theta = 5f/T$, $s = \theta^2$, and add 2θ to the first p entries of this label to form v'. One may use any fixed ample model line with a smooth positive metric as perturbation $s\bar{A}$; the auxiliary exponent will return to zero. For $\mu = v' - sv$ with a fixed rational v (including $v = 0$), all cross-block gaps are at least θ for large T. The cutoff primes carry the framed complete flag in the trivial Tate bundle on every reduced special component. The dual prefix determinants are generated there. On the other special fibres the Hodge and conjugate filtrations give nonnegative weight degree whenever $(l-2)\min_\sigma(\mu_{\sigma p} - \mu_{\sigma,p+1}) \ge 4f\|\mu\|_\infty$. This holds for $l > T$ at the chosen probe. Thus the regularized line of class $\ell_\mu + sA$ is relatively nef. Its archimedean curvature is semipositive. In the expansion in the dual prefixes the coefficients are the nonnegative adjacent gaps and the positive final slot $\mu_{\sigma,N}$. These remain nonnegative at the chosen shifted probes for large T. The one-sided regularization loss at $t = \exp(-1/s^2)$ is therefore $o(s)$ on the fixed complex flag. The generic weight is nef by the same arbitrarily small curvature losses on curves. On the fixed reference the sum with sA is ample. On a dominating model its pullback is nef and big. The normalized geometric top degree is $O_v(\theta)$: the Weyl product at v' has a factor 2θ, and changing the generic class by $O(s)$ changes the fixed intersection polynomial by $O_v(s)$.

For all algebraic points the smooth reference lower bound and finite comparison are $-C(1 + \log(2 + H(z)) + \log T)$. The smooth regularization costs at most $C(1 + \log(1/s))$ from below. Adding sH and minimizing $sH - C\log(2+H)$ gives a lower point bound $-C'\log(T/s)$. A finite linear projection from the fixed reference has a dense inverse image of projective points with roots of unity as affine coordinates; on it H is bounded. The upper reference, finite and archimedean comparisons give cutoff height at most $C'\log T$ there and on its lifts. Shifting the metric by the lower bound produces an arithmetically nef line, whose top intersection is nonnegative by the arithmetic Hilbert–Samuel formula. The upper essential-minimum inequality bounds the original normalized intersection above by $(d+1)$ times its generic degree times the dense point bound. Thus the regularized intersection is $O_v(\theta \log T)$. This argument treats the full proper flag, including algebraic points on the boundary. Moriwaki, Corollary 5.5 and Ballaÿ, Theorem 1.5.

The mixed metric-change estimate at the prescribed scale costs $O(s)$ in this calculation after division by the covering and coefficient degrees. To remove the perturbation, write the limiting intersection as a homogeneous polynomial in (μ, s). Fix finitely many rational evaluations determining all its coefficients, in the bounded open set of strictly dominant weights with $0 < \mu_{\sigma,N}$ and $\mu_{\sigma,1} < 1/2$ at all labels, and sufficiently small positive s. Each has, for large T, the same relative nefness, positive curvature and the two point bounds with s fixed. Its generic degree is fixed and its normalized arithmetic value is $O(\log T)$. Inverting the fixed evaluation matrix bounds every coefficient by $O(\log T)$. Changing (μ, s) back to $(v, 0)$ costs $O_v(\theta \log T)$. We conclude that $P_T(\ell_v^{d+1}) \to 0$. The corresponding Euclidean line descends isometrically to the flag omitting the repeated-weight step, of relative dimension $d-1$, so $\eta(\ell_v^{d+1}) = 0$ as well.

The space of degree-$d+1$ polynomials in v with the indicated within-block signs and label symmetry is finite dimensional. If a polynomial in it vanishes on all rational probes in the indicated relatively open positive portion of each wall, it is divisible by the cross-wall equations. This tests the full polynomial and

hence also applies to balanced classes associated to signed weights. Within-block alternation supplies the within-block factors and conjugates the cross-wall equations to all $v_{\tau i} - v_{\tau j}$ with $i \le p < j$. Their product with the within-block factors is the full Vandermonde, of degree d. Its linear quotient is constant on each of the two slot orbits by within-block symmetry and label transitivity. Thus the kernel of all wall evaluations is spanned by the polynomials associated to S and U under polarization: on a power these are $-(d+1)D_{\mathrm{fl}}(\ell_v^d)$ times the large-slot sum and the all-slot sum. A finite set of rational evaluations has the same kernel. Evaluating a functional on a fixed balanced G annihilates this kernel and is therefore a fixed linear combination of those evaluations. Their convergence to zero proves (102). All constants from the wall argument disappear on taking this limit.

Return to the selected weight. For $F = \ell_u^{d+1}$ put $e = (d+1)D_{\mathrm{fl}}(\ell_u^d)$, $x = \sum_{\tau, i \le p} u_{\tau i}$, $y = \sum_{\tau, i} u_{\tau i}$, and $t = Nx/p - y$. Let H be the $(fpq+1)$st power of the small Hodge determinant times the top within-block root product, with $S(H) = 0$ and $U(H) = fqh > 0$. Let $J = p_{\tau_0,1} Q$ with $D_{\mathrm{fl}}(Q) = 1$, so $S(J) = p$, $U(J) = N$. Then

$$a = \frac{et}{fqh} \ge 0, \qquad b = -\frac{ex}{p}, \qquad F^\circ = F - aH - bJ$$

is balanced. For H the determinant of the small block with its quotient metric is exactly the tensor product of its signed graded lines. The within-block top product pushes to the ordinary flag fibre degree in arithmetic codimension zero. The mixed metric comparison therefore identifies its normalized value with the small determinant base height. In the J test the determinant monodromy is finite and the split CM point calibrates its ordinary height with the factor $(p-q)\log\Omega/2$. Use the two arithmetic tests $P_B(H) \ge -C_H fqh\Lambda_B$ and $|P_B(J)| \le C_J N\Lambda_B$, and the compact-dual facts $\eta(F) \ge 0$, $\eta(H) \le hfq\log N/2$, $\eta(J) = 0$. The established old-minus-new formula sums on the balanced class to $P_B(F^\circ) = \eta(F^\circ) + R_{>B}(F^\circ)$. Only its two nonsummable directions were removed. In particular the higher tail of J vanishes even though its individual degree-one jumps at type-interchanging primes generally do not vanish. At the actual cutoff, every small-slot distance has the stated lower bound g and $B + 1 > 4fN^2(1 + \|u\|_\infty/g)$, so the prime tail estimates are $|R_{>B}(F)| \le C_R eq\,\mathrm{spread}(u)$ and $|R_{>B}(H)| \le C_R h$. Since $0 \le t \le \|u\|_1$ and $N|x|/p \le 2\|u\|_1$, substitution yields

$$\frac{P_B(F)}{e} \ge -C\big(\Lambda_B\|u\|_1 + q\,\mathrm{spread}(u)\big)$$

with an absolute C. Here $B \ge \max(128, 4fN^2)$, so the $t\log N$ term is absorbed. This is the asserted component mean on the same root models. Equality of the component volumes gives the unweighted mean; the finite trace allowance changes it to the mean of the comparison slopes s_Z at the prescribed weight.

For the finite-degree application the finite trace corrections are absorbed in the same X as the mean lower bound. The weight budget gives $\ell|u'|_1 \le 834fmb$, $B_* \le 132b$, $q_j\,\mathrm{spread}(u') \le 6|u'|_1$, and $\log \mathrm{rd}(E) + |d\log L(1)| \le 9L$. Substitution in (101) gives $\ell X = O(fmbL)$. This is the bound on the mean of the comparison slopes s_Z. The next comparison relates these slopes to the Chow heights and the carrying section on the same cover. The weight budget.

To compare with h_Z, take the normalization of the schematic closure of the moving image in the projective space of the full moving section lattice. Give its tautological line $\mathscr{H}$ the quotient metric at infinity. After the metric shift e^{-t} for $t \ge c_s = \log\Lambda$, the moving generators are small and generate the line. The quotient metric has semipositive curvature. On any horizontal point the degree is nonnegative by evaluating a nonzero generator with all archimedean norms at most one, so the shifted Hermitian line is nef. Its arithmetic top intersection is $[K_{\mathrm{covf}} : \mathbf{Q}](D+1)d_Z(h_Z + t)$, where K_{covf} is the cover coefficient field. This is the Chow form height formula for the quotient metric, using the Gauss norms at finite places and the sphere integrals with their harmonic correction at infinity. The projection formula preserves this intersection on the normalization. Thus $h_Z + c_s \ge 0$. Bost–Gillet–Soulé, Theorem 4.3.2 and Lemma 4.3.4.

For a finite valuation test v of the de Rham lattice let $d_v(u_1)$ be the minimum valuation of the integral moving generators. Let $\mathscr{B}_Z(y)$ be the squared norm of evaluation at y in the Petersson space. Take sections f at nu_1 satisfying $v(f) \geq nd_v(u_1)$ and $\|f(y)\| \leq e^{nt}\mathscr{B}_Z(y)^{n/2}$. Their rational sections on the image are determined by dividing by the n-th power of any nonzero moving generator in its chart.

Detection on the normalized moving image. Let s_i generate the actual integral moving lattice on the selected component over F. The generic moving map is birational, so a rational section on the image can be read in the component function field. Suppose a section f at degree n has a pole in the tautological quotient line on the normalized image. Choose a divisorial valuation detecting it and a generator s_0 of least value in that frame. In the component function field the ratios s_i/s_0 are integral at the valuation and the nonzero ratio $c = s_0^n/f$ has positive value.

These inequalities specialize simultaneously to an algebraic point with a finite-place valuation. Take the finite-type $\mathcal{O}_F$-algebra generated by the ratios s_i/s_0 and c. This is a domain in the component function field. The ratios of a common moving frame generate that field. The valuation therefore has a centre whose closure lies in $c = 0$. Choose a closed point of the closure above a finite prime of F. For a horizontal centre such points exist since the image of its closure contains an open subset of $\operatorname{Spec}\mathcal{O}_F$. For a vertical centre take a closed point of its fibre. Write R for the local ring at the chosen closed point, and π for a uniformizer of the coefficient prime.

Shrink the generic open under consideration so that it lies in the framed moduli chart with s_0 and f nonzero. This is done by avoiding a nonzero regular element. Include its numerator and π in a product $s \in R$. Since R is a local domain, extend s to a system of parameters $(s, t_1, \dots, t_d)$, where $\dim R = d + 1$. A minimal prime of $(t_1, \dots, t_d)$ avoids s by the height bound and has one-dimensional local quotient, since adjoining s gives the maximal ideal up to radical. The curve through the chosen closed point is horizontal because it avoids π. Its function field is finite over F by the dimension formula over the Dedekind base. Normalize the curve and complete at a place above the closed point. This gives a discretely valued trait where all s_i/s_0 remain integral and c has positive value. Its generic point belongs to the chosen open.

Extend this trait to put the prescribed PEL fibre in the integral de Rham chart. The semiabelian extension and the splitting field of the eigenlabels are available after finite extension. For each label the intersections of the flag spaces with the integral module give a saturated filtration, so an integral basis in that label can be chosen adapted to the same flag. The rigidified theta line is trivialized over the trait. Changing to this frame multiplies all moving values by one scalar and the degree-n value by its n-th power. The ratios and the positive order of c are preserved, with ramification normalization. Thus this integral valuation test contradicts the depth inequality for f. Normality then supplies regularity on the image. The archimedean condition $\|f(y)\| \leq e^{nt}\mathscr{B}_Z(y)^{n/2}$ is exactly the quotient-metric bound on the open set, and extends by continuity. The finite specialization with a common frame.

Thus these sections inject into the small sections of the shifted $\mathscr{H}^{\otimes n}$. The normalized upper count in the limit is at most $d_Z(h_Z + t)$ for $t \geq c_s$. The nef upper count.

The carrying section for $8(u_1 - u')$ is the product $T_Z^{24}K_Z^{\mathrm{key}}S_Z^{\mathrm{flat}}$, using the nonzero level-five theta constant, its key correction and the flat determinant line. In the two signature blocks the squared theta factor has ordinary weights $(r, -r)$ and theta exponent r. Taking twelve factors gives $(12r, -12r; 12r)$. The $3r$-th power of the fourth-power key contributes $(0, 0; -12r)$, and the parallel flat section contributes $(-12r, -12r; 0)$. Their sum is $(0, -24r; 0) = 8(u_1 - u')$, with zero theta exponent, in the weight convention of the carrying section. Write θ_Z^{tr} for the resulting product. The three contributions to the weight.

The archimedean cost of the carrying section. For the auxiliary block polynomial in the carrying construction let $D = d_{\mathrm{pol}}$ and $n = n_{\mathrm{col}}$. It is evaluated at the pre-cover level $Q = M_0 = g_0 g_1$, between b^{16} and b^{70}. Use the torus notation of (134) and put $\varphi(y) = 2f/(\mathrm{rd}(E)^2\lambda(y)^2)$ at the first period minimum

$\lambda(y)$. Its raw marked forms have the full half-level indexing used in (137). The size restrictions $b \geq 1000$, $1 \leq D \leq b^2$, $frn \leq b$, and $\mathcal{R} \leq b$ give

$$\mathcal{R}^4 R_0^2 Dr \leq 180 b^9 \leq b^{16} \leq Q.$$

Thus the same numerical estimates for the short decomposition and the jets apply at this weight. If $M_{\mathrm{cov}}\lambda(y)^2 < b$ and $M_{\mathrm{cov}} \geq bQ$, then $Q\lambda(y)^2 < 1 \leq Q\tau^2$. Here $\Phi = \varphi(y) = 2f/(v^2\lambda(y)^2)$, and $\varphi(y) \geq b^{14}$ since $Q \geq b^{16}$ and $v \leq b$. The cover and jet parameters.

Apply the raw reading formula preceding (138), with the Lie factor $(2\pi\xi_R)^{-D/2}$ and the unscaled rules $f_{\mathfrak{b}}$. It gives $\sum_{\mathfrak{b}} \|f_{\mathfrak{b}}(y)\|^2 \geq (\pi Q)^{-D} J_{\mathrm{bnd}}$. There are Q^{frn} terms. At least one $f_{\mathfrak{b}}$ has

$$\log \|f_{\mathfrak{b}}(y)\| \geq -p_{14} D - \frac{\pi n C_\chi}{Q}\varphi(y) - \tfrac{1}{2}(D + frn)\log Q.$$

This implies the advertised carrying estimate with the allowance $(D+frn)\log Q$. Indeed $p_{14} \leq 3b^3$ follows from $r, n \leq b$, $d_* \leq 2b$ in its formula. Also $\pi n C_\chi/Q \leq 8100\pi{\cdot}16\, b^{-6} \leq 1/2500$. The normalized supremum bound for this same marked rule is $\log \|f_{\mathfrak{b}}\|_{\sup} \leq Db + 9000(3frL)^2 b^2$. Together with $\log Q \leq 70L$, the terms other than $\pi n C_\chi \varphi(y)/Q$ cost at most $90000b^6$, using $frL \leq bL$, $L = \log b \leq b$, and $D \leq b^2$. This is at most $\varphi(y)/2500$. The same norm ratio holds on the cover: the forgetting map preserves the first minimum and point norms, and the parallel flat correction scales point norms on the component by one positive factor. The component probability Petersson norm is bounded by the normalized supremum. Squaring the point-to-supremum ratio in the extremal characterization of the kernel proves $\mathscr{B}_Z(y) \geq e^{-\varphi(y)/625}$ in the thin range. The compatible supremum estimate and The probability kernel comparison.

In the thick region $M_{\mathrm{cov}}\lambda(y)^2 \geq b$, the congruence separation estimate makes the chart ball of radius $1/(100r^2)$ inject in the quotient. The weight gaps are at least ℓ, which dominates the fixed multiple of $(fr)^8$ needed in the weighted $\bar\partial$ estimate. Start with a cutoff equal to one near the chart center. The correction solving $\bar\partial$ with logarithmic poles at the translates vanishes at those points by weighted square-integrability. The resulting holomorphic peak has value one at the center and zero at its other translates. The local squared norm bound is $e^{O(|u_1|_1 + (fr)^8)}$ times the ball volume. Unfolding its cuspidal Poincaré sum costs at most $e^{O((fr)^8)}$ by the weighted Cauchy–Schwarz estimate. Since the ball injects, its volume is at most the quotient volume. The probability norm therefore gives $\mathscr{B}_Z(y) \geq e^{-B_{\mathrm{thick}}}$ with $B_{\mathrm{thick}} = O(|u_1|_1 + (fr)^8)$. In this range $\ell B_{\mathrm{thick}} = O(fmb)$ by the quotient weight budget. Chart injectivity, Weighted peak estimate, and Thick-region normalization.

For the fixed carrying product T_Z, the odd cubical evaluation gives the product of the normalized level-one theta values on its single primitive tuple, as in (133). Its square has the pure theta weight, so its point norm is the square of that product of absolute values. Put

$$A = \tfrac{1}{2}(625) r d_*, \qquad P_\eta = 2\cdot 625(p_3 + 2fr^2 \log 5) + A\log(AK).$$

On the thin torus, multiplying the logarithm of the bound on $\|T_Z^2(y)\|$ by 625 gives at most $2\cdot 625 p_3 + A\log(K\Phi) - 2\pi\Phi$. Since $A\log(K\Phi) \leq \Phi + A\log(AK)$, this is at most $P_\eta - \Phi$. The thick torus bound is at most P_η in the same units. Here $P_\eta = p_{16}(E, r, Q, m, 5, 1) = O(fr^2 L)$ by $d_* \leq 2fr$ and the logarithmic parameter bounds. In the thin part of the cover the function Φ is $\varphi(y)$. Thus the carrying comparison uses the same period minimum in both its upper theta bound and its lower Bergman bound. The pure theta square on the actual cover.

Taking the cube and using the two lower Bergman estimates gives $\|T_Z^2(y)\|^3 \leq e^{\beta}\mathscr{B}_Z(y)$, where $\beta = P_\eta + B_{\mathrm{thick}}$ and $\ell\beta = O(E_{\mathrm{cov}})$. These estimates apply to the same section after the flat metric change. Comparison of cubic norm with Bergman.

The finite cost of the carrying section. In a semistable test frame with period lattice Λ and monodromy quadratic form Q, the theta product is a product of null values at arguments y_b generating modulo 5^4. The nonzero values are chosen by translating in the level-five theta representation. Its translates

span the space of sections of the fifth power by irreducibility, and that power is generated. Thus the shifts can be chosen simultaneously in the prescribed classes, preserving generation modulo five. In the normalized Fourier expansion of the translated term at $y_b/25$, the coefficient at χ has valuation at least $Q(25\chi + y_b)/625$. The cross and constant terms here are those of the translated quadratic form in the rigidified theta norm. The ultrametric inequality gives the lower bound on its sum by the minimum over χ. Choose a minimizing representative z_b in each coset. Their reductions generate modulo 25, so the representatives span the period space over E. For a trace-one host element w_0, the functional $\beta \mapsto \sum_b \beta(w_0 z_b, z_b)$ on Hermitian forms is integral and strictly positive on every nonzero semipositive form. Since it takes the value $\sum_b Q(z_b)$ on the polarized monodromy form, this sum bounds from above the infimum Φ over such positive integral functionals. Thus $v(T_Z) \geq \Phi/625$. The normalized Fourier calculation is made over an algebraic closure of the trait field, and both orders scale with ramification. Nonvanishing translates, Valuations of translated terms, and The monodromy bound.

The reference calculation uses the base level $M_0 = g_0 g_1$. It is made in the same semistable trait and the same rigidified theta norm as the carrying product. At a level N prime to the residue characteristic, Fourier descent changes the coefficient valuation under $\chi \mapsto \chi + Nz$ by $(Q(\chi + Nz) - Q(\chi))/N$. In the balanced basis let $Q(b_i) \leq \Phi_0$, $|B_Q(b_i, z)| \leq M_* Q(z)$, with $B_Q(x, z) = Q(x + z) - Q(x) - Q(z)$, and let d_{pol} be the jet degree. For a mode $\chi = \sum_i a_i b_i$, $a_i \geq 0$, $\sum a_i \leq d_{\mathrm{pol}}$, one has $Q(\chi) \leq d_{\mathrm{pol}}^2 \Phi_0$. If $N > d_{\mathrm{pol}} M_*$, every other mode in the same coset has strictly greater order. On the abelian part, the very ample cube gives local parameters by taking ratios with a section nonzero at the origin. Products of these parameters through degree d_{pol} can be homogenized at level $N \geq 4d_{\mathrm{pol}} + 4$ using a section of the remaining ample power twisted by the given Poincaré line, nonzero at the origin. Cohomology and base change lift the resulting jets. Multiplying by the small toric modes spans all jets, since the monomials in $t_i - 1$ are integral combinations of the $t_i^{a_i}$. Nakayama and the bound on $Q(\chi)/N$ give a scalar cost at most $d_{\mathrm{pol}}^2 \Phi_0/N$.

Normalized finite-order translations preserve the lattice with minimum order $\min_\chi(v(F_\chi) - Q(\chi)/N)$. The linear valuation change in a translation of χ cancels against the quadratic weights, and any constant change is zero by finite order. When $N = \alpha\bar{\alpha}$ with coprime halves, $[\alpha]^*\theta_1$ is primitive in this lattice by $Q(\bar{\alpha}\chi) = NQ(\chi)$ and faithful flatness on the abelian part. The translates by $\ker \bar{\alpha}$ have distinct characters of $\ker \alpha$ by the perfect Weil pairing. Integral Fourier idempotents split these rank-one lines, since N is a unit. The primitive translates are therefore an integral basis. Taking one summand of an interpolated harmonic jet gives the same bound for its least order on these translates. Foster–Rabinoff–Shokrieh–Soto, Proposition 3.19, Spanning by the small modes, and The integral translate basis.

The finite cost at a type prime. Let g be the residue characteristic and $k = g'$ the other type prime, so $M_0 = gk$. Write $a\bar{a} = g$ for the prescribed split half and $\alpha_0 = aa'$ for the corresponding generator of the whole half level, with $a'\bar{a}' = k$. Use the same framed Raynaud chart as in the reference exponent. The integral $\mathcal{O}_E$-action splits into eigenframes there, since both type primes are totally split. In each conjugate pair, a acts by a multiple of g on one block and by a unit on the other. The leading-minor polynomial uses the first r columns for the first kind of block and the next r for the second. The conjugate blocks use disjoint columns, so every conjugate contraction kills the polynomial. The product of leading minors with exponents given by the successive weight gaps has the required harmonic weight. Let d_{v} be its total degree in the first set of columns, and d its total degree. The selected polynomial.

The Raynaud extension G has a g-divisible group $G[g^\infty]$, since multiplication on both its toric and abelian parts is finite flat. Its invariant differentials are those of G. For its Cartier dual, the Hodge–Tate map over the completed algebraic closure has image containing $c\omega_G$ after taking the valuation-ring span, where $v(c) = v(g)/(g - 1)$. Indeed the cokernel bound holds with $1/(g - 1) + \epsilon$ for every $\epsilon > 0$, and the finite matrix of the map has invariant factors over the valuation ring. Their valuations therefore have the stated bound. Modulo g, this map reads the differentials of the finite Cartier characters on $G[g]$. In the nonunit directions these restrict to $G[a]$, since the tangent vectors over R/g are killed by a. The projection to these prescribed coordinates is surjective on the integral cotangent module because the coordinates are

part of an integral eigenframe. All the finite characters and their differentials can be read after a finite extension of the trait. Choosing representatives for their differentials modulo g, their span contains c times the whole coordinate module. In fact the finite matrix expressing c times the basis vectors in this span modulo g has error gB for an integral matrix B, and $1 + (g/c)B$ is invertible. Scholze–Weinstein, Proposition 4.3.6, and Finite character columns.

Choose the columns in order at one label. Appending c times the next coordinate vector to a leading minor multiplies its determinant by c. Express this vector as an integral linear combination of character columns and expand the determinant in the last column. At least one such determinant has valuation at most $v(c)$ more than the preceding one. Iterate for the leading minors of that block. To combine the choices for all labels, precompose the finite characters with the split idempotents modulo g and multiply the characters. Logarithmic differentiation adds their coordinates and preserves the chosen columns in each block modulo g. The resulting product of minors in the first set of columns has valuation at most $d_{\mathrm{v}}v(c)$, which is strictly less than $v(g)$ by the degree bound. Each character of $G[a]$ extends to the generic $A[a]$: the injection of $G[a]$ into $A[a]$ follows from torsion-freeness of the period lattice, and both groups on the generic fibre are vector spaces over $\mathbf{F}_g$. By the perfect Weil pairing between the complementary halves it is the commutator with a point $x_b \in A[\bar{a}]$. The level-kg commutator on the g-part is the k-th power of the level-g commutator, so one uses the inverse of k modulo g in this choice.

For each selected x_b, consider sections obtained by pulling back a level-k section by a and translating by x_b. Their Fourier indices lie in a coset of $\bar{a}X$, since pullback acts by $\bar{a}$ on characters and the translation shifts them by one fixed index. Let z_b be a least vector for Q in this coset. In the basis $\bar{a}b_i$ round the coefficients of a representative independently to the two adjacent integers, with probabilities making each rounding error have mean zero. Each variance is at most $1/4$. Hence the expected value of the resulting square is at most $gt_{\mathrm{tor}}\Phi_0/4$ and $Q(z_b)$ is at most that value. Minimality also gives $|B_Q(z_b, u)| \le Q(u)$ for $u \in \bar{a}X$, by comparison with both $z_b + u$ and $z_b - u$. Consequently

$$Q(z_b + ku) - Q(z_b) \ge k(k-1)Q(u) > 0 \quad (u \neq 0).$$

Use a source Fourier family supported on the coset whose transformed indices are $z_b + k\bar{a}X$. The coefficient line on the abelian part is base-point free at level $k \ge 4$. Choose its coefficient at the selected source index with unit value at the specified abelian point and give it the scalar of valuation $Q(\chi)/k$ in the source frame. This puts the source family in $\nu_k \ge 0$. Pullback and the finite-order translation preserve the normalized coefficient valuations. The displayed strict inequality makes the selected index the single least term in the translated identity germ. Its unit leading value is preserved under the integral change of cubical frame. Thus the germ is $s_b\, W_a^k f_b$, where

$$0 \le v(s_b) = Q(z_b)/(kg) \le t_{\mathrm{tor}}\Phi_0/(4k), \qquad f_b \in R[[x]], \qquad v(f_b(0)) = 0.$$

The least Fourier term in the translated coset.

In the pulled-back frame the kernel of a fixes the source section. Its action on the translated section is by the specified commutator character. This identity extends from the generic fibre to the finite flat kernel in the integral formal chart, after taking out the scalar s_b. Evaluating at the identity plus a square-zero tangent vector modulo g gives

$$D_u f_b \equiv d\log \chi_b(u) f_b \pmod{gR[[x]]}$$

in every nonunit direction u. This congruence holds for successive invariant derivatives since their coefficients are integral. The gauge factor in the harmonic value becomes $\exp(kH \circ [a])$ after using W_a^k. All nonconstant invariant derivatives of its exponent in the nonunit directions have order at least $v(g)$: the invariant derivatives of H are integral on these directions, including order two by the unit difference $\tau(a) - \sigma(\bar{a})$; pulling back multiplies them by the corresponding eigenvalues of a. The repeated Leibniz rule for the exponential then proves the same congruence for the connection values at zero. The first

product of minors therefore has valuation at most $d_v v(g)/(g-1)$ before restoring the scalars s_b. The kernel congruence and The invariant-derivative comparison.

In the unit directions, conjugate adjointness again makes the restricted second derivatives of H integral, now because the eigenvalue of a is a unit and that of $\bar{a}$ is a multiple of g. At every order through d, changing between ordinary jets and lists of invariant derivatives is triangular with unit diagonal up to products of factorials at most $d!$; these factorials are units since $d < g$. Pullback in the unit directions is invertible on these lists and multiplication by the gauge exponential is triangular with diagonal one. Thus prescribe integral source jets giving identity columns for the unit minors. The small-mode interpolation at level k realizes the jets with total scalar order at most $\sum_b d_b^2 \Phi_0/k \le d^2\Phi_0/k$, where d_b is the order in column b. Now every chosen source section is in $\nu_k \ge 0$. Expand each in the integral translate basis at level k. The prime-to-g translations lift along a on that same torsion, so after pullback the pure products are translates of $[\alpha_0]^*\theta_1$ at the designated half. The ultrametric inequality applied to the expanded product of column evaluations selects one pure product with no larger valuation. With the marked frame equality, the bound in the original test is

$$\inf_f v(f) + v(C_{\alpha_0}) \le \frac{d\, v(g)}{g-1} + \frac{(r t_{\mathrm{tor}}/4 + 2d^2)\Phi_0}{g'}.$$

Here the extra allowance in the second term covers the interpolation scalars. By the same frame equality, normalized translates have integral germs in every integral Raynaud chart over that coefficient prime, and the conductor bound there is zero. Hence $v(f) \ge -v(C_{\alpha_0})$ on all such traits. The coefficient-field scalar C_{α_0} depends only on the fixed weight and the marked half generator. It cancels the two occurrences of its valuation. Multiplying by an additional scalar of order zero at the tested coefficient prime clears the finitely many denominators elsewhere. The selected level size absorbs the Φ_0 term into the boundary term, and $d/(g-1) < 1$ proves the reference-exponent bound one. The marked split-place bound and Scalar compensation in the reference cusp lattice.

4.3.3 Formal depths and the reference prime budget

Depth of a formal determinant. Let R be the valuation ring in a complete algebraically closed extension C of the fraction field of the test trait. Normalize the additive valuation by $v(p) = 1$. Let G be the commutative formal group of the test and let $u_1, \dots, u_i$ be the first i vectors of a typed integral flag. Their matrix has a unit minor. For a nonnegative integer c, write $d_G(u; c)$ for the infimum, over integral power series on G^i, of the valuations of the values at the identity of

$$\left[\det\left(D_{u_a}^{(j)}\right)_{a,j=1}^{i}\right]^c.$$

Here $D_{u_a}^{(j)}$ is the invariant derivative in direction u_a on copy j. This infimum is a minimum over monomial jets, since the value is linear in the germ and depends only on its jet of degree at most ic. Extension of the valuation ring preserves this minimum within a given chart. The release takes the infimum over integral charts and valued extensions for the factor at the point. We give the lower bounds in every chart and the upper bounds in an integral chart of the test. The invariant derivatives commute. Changing the frame basis by $\mathrm{GL}_i(R)$ multiplies the operator by a unit. Formal depth of a factor at the point.

Put

$$\mathcal{D}_{i,c}(Z) = \det(Z_{aj})^c, \qquad \kappa_{p,i,c} = \min_{\alpha \in \operatorname{supp} \mathcal{D}_{i,c}} v\big([Z^\alpha]\mathcal{D}_{i,c}\, \alpha!\big), \qquad \alpha! = \prod_{a,j} \alpha_{aj}!.$$

The integer $\kappa_{p,i,c}$ depends only on p, i, c. Write L_G for the formal logarithm with identity linear part. On a single copy, substitute $\exp_G(\sum_{a=1}^{i} u_a X_a)$. This puts the invariant derivatives in the indicated frame into the coordinate derivatives in X_a. Use i copies of the substitution for $\mathcal{D}_{i,c}(\partial_X)$. The unit minor gives integral linear forms on the original chart which restrict to the coordinates X_a to first order. Take the product of their powers prescribed by a minimizing α. Its homogeneous term of degree ic after substitution is X^α,

so the invariant operator takes the value $[Z^\alpha]\mathcal{D}_{i,c}\,\alpha!$. Every invariant derivative in an integral direction preserves integral series, by differentiation of the integral group law. Finally the diagonal monomial of $\mathcal{D}_{i,c}$ has coefficient one. Consequently

$$0 \le d_G(u;c) \le \kappa_{p,i,c} \le i v_p(c!) \le \frac{ic}{p-1}, \qquad d_G(u;c) = 0 \quad (c < p). \tag{103}$$

For $c \ge 1$, every monomial in the determinant power has row sums c, so each $\alpha_{aj} \le c$. For $n \ge 0$, Legendre's formula gives

$$v_p(n!) = \frac{n - s_p(n)}{p-1} \ge \frac{n}{p-1} - \log_p(n+1).$$

Indeed, writing $s_p(n) = a(p-1) + b$, with $0 \le b < p-1$, the smallest integer with that digit sum has a initial digits equal to $p-1$ and next digit b, so $n+1 \ge (b+1)p^a$. Concavity gives $\log_p(b+1) \ge b/(p-1)$. Therefore

$$\kappa_{p,i,c} \ge \frac{ic}{p-1} - i^2 \log_p(c+1). \tag{104}$$

The determinant bound on mixed types. This AI-proposed candidate concerns a commutative formal group with integral group law in characteristic zero and with $v(p) = 1$, $p \ge 3$. Fix a positive even integer h. The roots of unity of order $p^h - 1$ act by formal group automorphisms, and on the tangent space the type $s \in \mathbf{Z}/h\mathbf{Z}$ means the eigenspace for $\zeta \mapsto \zeta^{p^s}$. Assume that each odd type has rank at most Q_*. For any saturated integral frame u of rank i contained in a single even type and every integer $c \ge 1$, there exists $b_{\mathrm{bad}} \ge 0$, an integer, with

$$b_{\mathrm{bad}} \le i, \qquad b_{\mathrm{bad}} \le Q_*(3/2 + \tfrac{1}{2}\log_p c),$$

such that

$$d_G(u;c) \ge (i - b_{\mathrm{bad}})c\,\frac{p}{p^2-1} - i^2(1 + \log_p c).$$

For $i = 0$ the assertion is the empty determinant. The group need not have finite height. The scalar height-two estimate has the classical period scale $p/(p^2-1)$. The mixed-type assertion bounds the loss by the odd ranks uniformly over saturated frames. No independent prior theorem for that assertion has been identified. [116, Section V] [117, Appendix, Theorem (c)]

Curves of prescribed Frobenius type. First work in the application at p unramified in E. Extend the $\mathcal{O}_E$-action to $\mathcal{O}_E \otimes \mathbf{Z}_p$ in the integral formal group. This is possible because $[p](X) = pX + O(X^2)$ and $[p^j](X)$ tends to zero in the (p, X)-adic topology. For fixed degree its coefficients converge p-adically. The idempotents decompose the formal group into its local factors. In fact one can take integral curves whose tangents give bases of the Lie summands, apply the corresponding idempotents to those curves, and add the outputs in G. The resulting map is a formal coordinate isomorphism with integral inverse. The logarithm of each idempotent is its linear projector, so the logarithmic line of a typed curve stays in its local factor. For a factor of residue degree h, number the types by $\mathbf{Z}/h\mathbf{Z}$. A Teichmüller unit ζ acts on the Lie summand of type s by ζ^{p^s}. The shift $s \mapsto s+1$ acts on the reductions of the labels by Frobenius.

The construction below only uses the indicated Teichmüller action and the type labels. It therefore applies to the general assertion on that even cycle.

There is an integral tangent basis consisting of curves $\gamma_a(T)$ whose logarithms are of the form

$$L_G(\gamma_a(T)) = \sum_{n \ge 0} \lambda_{a,n} T^{p^n}, \qquad \lambda_{a,n} \in \mathrm{Lie}(G)_{s(a)+n} \otimes_R C.$$

To construct a curve through a chosen integral tangent vector, start with a linear curve in the original chart. For every prime $l \ne p$, subtract, in the formal group, its average over the l-th roots of unity acting

on the parameter. Multiplication by l^{-1} here is the integral inverse of $[l]$. In the logarithm this suppresses exactly the degrees divisible by l. Applying the operations in increasing order stabilizes every coefficient and leaves precisely the p-powers. For a desired type s use the additional average of $[\zeta]\gamma(\zeta^{-p^s}T)$ over $\zeta \in \mu_{p^h-1}$. In logarithmic degree p^n it keeps precisely the target type $s+n$. Both averages preserve the chosen tangent. The map $\Phi(T_1,\dots,T_d) = \sum_a^G \gamma_a(T_a)$ is an integral change of formal coordinates because its tangent matrix is a unit. Typed p-typical curves and their tangent bases.

In coordinates in any such tangent basis, every entry of $\lambda_{a,n}$ has valuation at least $-n$. The logarithmic differential equals the invariant differential with identity value at zero, so its pullback along γ_a is integral. This pullback has coefficient $p^n\lambda_{a,n}$ at T^{p^n-1}. If $\lambda_{a,1} = 0$, the bound improves to $1-n$ for $n \geq 1$. For this improvement the formal sum $\sum_{\zeta\in\mu_p}^G \gamma_a(\zeta T^{1/p})$ has integral coefficients in T by invariance under the roots of unity. Its logarithm starts at order p and has coefficient $p\lambda_{a,n}$ at $T^{p^{n-1}}$. It is a sum of curves $\gamma_b(x_{jb}T^{p^j})$, $j \geq 1$, with $x_{jb} \in R$. To see this, successively cancel the first logarithmic coefficient of the remaining curve in each degree p^j. That coefficient equals the first coefficient of the curve in the group coordinates and is therefore integral. Expand it in the integral tangent basis to obtain the x_{jb}. The sum is degreewise finite. At degree p^{n-1} all contributions use some $\lambda_{b,m}$ with $m \leq n-2$. This gives $v(p\lambda_{a,n}) \geq 2-n$. Valuation bounds on the typical logarithm coefficients.

For use in choosing the basis, a block of same-type curves may be rebased so that the coefficients $\lambda_{a,n}$ vanish modulo a prescribed C-linear subspace, except in at most the rank of their images. Indeed, form the linear map sending the coordinate vectors of the block to these images. Over a valuation ring the kernel in the C-space admits an integral basis extendible to a basis of the ambient free module. For a proof, choose a nonzero kernel vector, divide by one of its coordinates of least valuation to make a coordinate a unit, and split off its span. Iterate in the complementary coordinate hyperplane. Choose a matrix Y of this adapted integral basis. Take entrywise p^n-th roots x_{ab} of Y_{ab}. They are integral and the matrix $X = (x_{ab})$ has unit determinant, because modulo the maximal ideal $\det(X)^{p^n} = \det(Y)$. Replace the curves by $\sum_b^G \gamma_b(x_{ab}T)$. The tangent basis changes by X, while its n-th logarithmic coefficient changes by Y. This works over R of arbitrary rank-one value group.

A bound from the inverse logarithm. Use the basis of tangents of the γ_a and write the logarithm in the Φ-coordinates as

$$\mathcal{L}_b(T) = T_b + \sum_{a,\ n\geq 1} l_{anb}\,T_a^{p^n}.$$

Suppose nonnegative real numbers R_a satisfy

$$v(l_{anb}) + p^n R_a \geq R_b \qquad (l_{anb} \neq 0). \tag{105}$$

Write the coordinates of the frame vectors as $x_{\ell b}$ and suppose $\rho_\ell \geq 0$ with $v(x_{\ell b}) + \rho_\ell \geq R_b$ when $x_{\ell b} \neq 0$. For a series f in the X_ℓ, saying that it has Gauss valuation at least s means $v([X^\alpha]f) + \sum_\ell \rho_\ell\alpha_\ell \geq s$ for every α. The solution of $\mathcal{L}_b(T) = \sum_\ell x_{\ell b}X_\ell$ with $T(0) = 0$ then has each T_b of Gauss valuation at least R_b. Solve for T_b by moving the nonlinear terms to the right and iterating from zero. The iterates stabilize in every degree, because all nonlinear terms have degree at least two. The inequalities (105) and the valuation inequality for products prove the Gauss bound at each iteration. The map $\Phi \circ T$ is the exponential substitution along the frame. Pulling back an integral germ on the i copies of G therefore gives a series of Gauss valuation at least zero in the i sets of variables. The ℓ-th row has weight ρ_ℓ on every copy and total exponent c for all terms of the determinant power. Evaluation on the germ gives

$$d_G(u;c) \geq \kappa_{p,i,c} - c\sum_{\ell=1}^{i}\rho_\ell. \tag{106}$$

The lower bound on an alternating orbit. In the general setting the number of vectors of any odd type is at most Q_*. Frobenius on the labels changes parity. Work in the orbit of the chosen even type.

Fix an integer $J \geq 0$. We assign the curves stage indices $d_a \in \{-1, 0, \dots, J\}$. All odd types have stage -1. Each even type has at most Q_* vectors at each stage below J. At every nonnegative stage d_a the coefficient $\lambda_{a,1}$ vanishes and $\lambda_{a,2}$ is in the span of tangents in its target type with stage at least $d_a - 1$. To start, rebase the even types at logarithmic coefficient $n = 1$. The target is the next odd type, of dimension at most Q_*. Give the exceptional vectors index -1 and the rest index zero. The condition on $n = 2$ then holds with threshold -1. Given stages through j, work in the top stage j of each even type. Its second coefficients, modulo the span of tangents of stage j in the target type, can only occupy its stage $j - 1$, of dimension at most Q_*. Use the rebasing construction with $n = 2$ to kill those coefficients modulo the span except in at most Q_* vectors. Keep the exceptions at stage j and promote the other vectors to $j + 1$. All changes of basis are inside the old top stages, so the spans defined by thresholds at most j contain the corresponding old spans. This preserves the conditions on all stages. Induction on the stages of the typical curves.

Take the least $J \geq 0$ with $c \leq p^{2J+1}$, where $c \geq 1$. For $s \geq -1$ set

$$R(s) = \frac{1 + p^{-(2s+1)}}{p^2 - 1}, \qquad R(-1) = \frac{1}{p - 1}, \qquad p^2 R(s) - 1 = R(s - 1) \quad (s \geq 0).$$

The choices $R_a = R(d_a)$ satisfy (105). At stage -1, use $v(l_{anb}) \geq -n$ and $(p^n - 1)/(p - 1) \geq n$. At nonnegative stage the $n = 1$ term vanishes, and for $n = 2$ the target lies in stages at least $d_a - 1$ and the coefficients have valuation at least -1. For $n \geq 3$ the improved bound $1 - n$ suffices, since

$$p^n \geq (n - 1)(p^2 - 1) + (p + 1) \qquad (p \geq 3,\ n \geq 3).$$

This holds first at $n = 3$ and then follows by taking successive differences.

Perform integral row operations in the frame so that only its first b_{bad} vectors can have a nonzero projection to the span of tangents of stages below J. The saturated kernel argument used for Y applies to these operations, with b_{bad} equal to the rank of the projection over C. Thus

$$b_{\mathrm{bad}} \leq i, \qquad b_{\mathrm{bad}} \leq (J + 1)Q_* \leq Q_*\big(3/2 + \tfrac{1}{2}\log_p c\big).$$

The last inequality follows from the minimality of J, and is immediate when $J = 0$. Put $\rho_\ell = R(-1)$ on the first b_{bad} vectors and $R(J)$ on the others. Then (106) and (104) give the lower bound

$$(i - b_{\mathrm{bad}})c\,\frac{p}{p^2 - 1} - i^2 \log_p(c + 1) - (i - b_{\mathrm{bad}})c\,\frac{p^{-(2J+1)}}{p^2 - 1}.$$

The last error is at most $i/8$. Since $\log_p(1 + 1/c) \leq \log_3 2 < 7/8$, it can be absorbed using $i \geq 1$. This proves the asserted bound with b_{bad}. On the actual nonterminal cover all odd type ranks equal Q, by the inert place of the quadratic subfield and the signature. Taking $Q_* = Q$, the estimate in every chart gives

$$d_G(u; c) \geq \big(i - Q\big(3/2 + \tfrac{1}{2}\log_p c\big)\big)\, c\,\frac{p}{p^2 - 1} - i^2(1 + \log_p c). \tag{107}$$

The saturated even-type lower bound.

Keeping the individual radii. The same stage construction gives a further bound. For any $K \geq 0$, construct stages through K. In the coefficient matrix of the frame in the integral tangent basis, process the columns in increasing stage order. For a nonzero column of the remaining rows, choose a nonzero pivot of least valuation. Subtract its integral multiples to clear that column in every other remaining row, and leave the pivot row fixed thereafter. There is no division of the pivot row. Thus this is an integral invertible row operation. A row first selected in stage s only uses tangent directions of stages $\geq s$. At most Q_* rows can be selected at each $-1 \leq s < K$, since there are at most Q_* tangent basis vectors at that stage in the chosen type. The remaining rows have stage K.

Use $\rho_\ell = R(s)$ for a row at stage s. Write $E_s = p^{-(2s+1)}/(p^2-1)$, so $R(s) = 1/(p^2-1) + E_s$. The radius inequality gives

$$d_G(u;c) \geq \kappa_{p,i,c} - c\left(\frac{i}{p^2-1} + Q_* \sum_{s=-1}^{K-1} E_s + iE_K\right).$$

The left-hand side does not depend on the stage construction. Let $K \to \infty$ in this numerical bound and use $\sum_{s=-1}^{\infty} E_s = p^3/(p^2-1)^2$. With (104), this yields

$$d_G(u;c) \geq \frac{icp}{p^2-1} - \frac{Q_* c p^3}{(p^2-1)^2} - i^2 \log_p(c+1).$$

The reference prime calculation below uses the stated mixed-type bound (107). The stage dimensions.

An interval from the first stage alone. For use when $k \leq mp$, it is enough to compare with $\kappa_{p,i,c}$ itself. Rebase the curves of the starting type at $n = 1$ as before, so that at most Q of them have $\lambda_{a,1} \neq 0$. Give the curves with vanishing first coefficient radius $1/(p(p-1))$ and the others radius $1/(p-1)$. For the first set, $n \geq 2$ and $(p^{n-1}-1)/(p-1) \geq n-1$ imply (105). For the second set use the bound $-n$. The projection of the frame to the exceptional directions has rank at most Q, so row operations again give input radii with sum at most $i/(p(p-1)) + Q/(p-1)$. Together with (103) this yields

$$\kappa_{p,i,c} - \frac{c(i/p+Q)}{p-1} \leq d_G(u;c) \leq \kappa_{p,i,c}. \tag{108}$$

This bound uses the same integer $\kappa_{p,i,c}$ at every test and includes $c = 0$. The comparison with divided-power content on the first stage.

Character differentials at an algebraic test. At an algebraic test, whose trait has finite residue field, extend the trait so that the Raynaud chart is the formal completion of a semiabelian extension

$$0 \longrightarrow T \longrightarrow \widetilde{G} \longrightarrow B \longrightarrow 0$$

with T a split torus and B an abelian scheme. The $\mathcal{O}_E$-action and the Lie flag in this chart are the ones of the test. Use the additive valuation $v(p) = 1$. We bound the valuations of the colengths of the lattice of character differentials in the dual Lie lattice. The colength of a full-rank inclusion of finite free modules over a valuation ring is the sum of its Smith valuations. The Raynaud chart at an algebraic test.

The Tate module of the Cartier dual of $\widetilde{G}[p^\infty]$ gives integral formal characters χ of $\widehat{\widetilde{G}}$. A compatible system of points in the dual torsion defines compatible characters of each $\widetilde{G}[p^j]$. At any fixed precision modulo p^e, every nilpotent point of the formal group is killed by some p^j, since $[p](X) = pX + O(X^2)$. Evaluation consequently defines a formal character at this precision, and taking the inverse limit gives a character over R. Its derivative at the identity is the Hodge–Tate differential on the abelian part. The dual torsion sequence is exact, with abelian dual torsion as kernel and the dual torus torsion as quotient. Surjectivity passes to C-points and then to Tate modules, since their transition maps on abelian dual torsion are surjective over C and allow each difference between successive lifts to be corrected. Every torsion point so obtained is integral, since the torsion schemes are finite and R is integrally closed. Thus the chosen characters surject onto the character lattice of T after p-completion. Their differentials span the dual torus Lie lattice. In the dual Lie sequence the cokernel on $\widetilde{G}$ is therefore a quotient of the cokernel of the abelian Hodge–Tate map

$$\alpha_B : R \otimes_{\mathbf{Z}_p} T_p B^\vee \longrightarrow \omega_B.$$

This quotient comparison respects the completed coefficient action. Here the action on the dual group is by the dual endomorphism. Via the polarization the Lie multiplicity of type τ on $B^\vee$ is the Lie multiplicity of type $\bar{\tau}$ on B.

The Frobenius recurrence for the colength. Use the covariant Breuil–Kisin–Fargues module M of $G_B = B^\vee[p^\infty]$ over $A_{\mathrm{inf}} = W(R^\flat)$, and its map $\beta : M \to \varphi^* M$. Use the conventions

$$T_p G_B = \{x : \beta(x) = 1 \otimes x\}, \qquad \operatorname{coker} \beta = \operatorname{Lie} G_B, \qquad M/\xi M \twoheadrightarrow \omega_B,$$

where the last map restricts to α_B on the Tate lattice. Let $t = \varphi^{-1}([\epsilon] - 1)$ for the chosen system of roots of unity and $\xi = ([\epsilon]-1)/t$, so $\ker \theta = (\xi)$, $\theta(t) = \zeta_p - 1$, and $v^\flat(t \bmod p) = v(\theta(t)) = 1/(p-1)$. Howard, Theorem 2.2.2. The integral pairing with the dual module gives

$$tM \subset A_{\mathrm{inf}} \otimes_{\mathbf{Z}_p} T_p G_B \subset M.$$

Indeed, in the framed Tate pairing the matrix on the two Tate frames is t times a unit matrix for a perfect pairing. Pairing integral vectors with the dual Tate frame shows that multiplying an integral vector by t puts it in the Tate lattice. These are the same period and Hodge maps as in the modified lattice of the flag calculation.

At a prime unramified in E, index embeddings in a factor by $s \in \mathbf{Z}/h\mathbf{Z}$, with $\varphi \circ \tau_s = \tau_{s+1}$ on the Teichmüller lifts. Write X_s for the matrix of the Tate inclusion in the s-summand of M. The Tate module is free over the local integer ring, since it is finite and torsion free over that discrete valuation ring. Use its basis in every embedding of that factor. Then the invariant-vector relation reads

$$\beta_s X_s = \varphi(X_{s-1}).$$

In fact $1 \otimes x$ has type s in $\varphi^* M$ when x has type $s-1$ in M. If $j_s = \dim \operatorname{Lie}(G_B)_s$, the image of β_s contains ξ times its target and the cokernel is free over $R = A_{\mathrm{inf}}/\xi$ of rank j_s. In a basis lifting a decomposition of the target modulo ξ into this quotient and its kernel, the image is generated by the kernel basis vectors and by ξ times the quotient basis vectors. Thus $\det \beta_s$ is ξ^{j_s} times a unit. Reduction modulo p gives, for $D_s = v^\flat(\det(X_s \bmod p))$,

$$pD_{s-1} - D_s = j_s, \qquad D_s = \sum_{n \geq 1} p^{-n} j_{s+n}.$$

The first equality uses $v^\flat(\xi \bmod p) = 1$. For the second, iterate $D_s = p^{-1}(D_{s+1} + j_{s+1})$ and use boundedness on the finite orbit.

The determinant valuation D_s is also the valuation of the colength of the Tate inclusion in $(M/\xi M)_s$. Each Smith valuation of the inclusion after either reduction, modulo p over $R^\flat$ or modulo ξ over R, lies in $[0, 1/(p-1)]$, by the two inclusions with t. The induced matrices over $A_{\mathrm{inf}}/(p, \xi) = R/p$ coincide. The truncated valuations are the same under this isomorphism because $\theta([x^\flat]) = (x^\flat)^\sharp$ has valuation $v^\flat(x^\flat)$. Smith elimination in the common quotient, choosing an entry of least valuation at each pivot, determines all Smith valuations strictly below one in either lift. When the residual matrix in the quotient is zero, the remaining valuations in either lift are at least one and hence equal one. At $p = 2$ we have $\theta(t) = -2$ and $t \bmod p = \xi \bmod p$, so the valuation-one factors are included in the comparison. The Hodge quotient shows that the valuation ℓ_s of the character colength in the dual Lie lattice of $\widetilde{G}$ is at most D_s. The Lie ranks of the abelian quotient of type $\bar{\tau}$ are bounded by those of $\widetilde{G}$. For any starting label s in Ψ, we obtain

$$\ell_s \leq \begin{cases} \dfrac{Q}{p-1} & \text{if } p \text{ splits in the quadratic subfield,} \\ \dfrac{Q}{p-1} + \dfrac{(P-Q)p}{p^2-1} & \text{if } p \text{ is inert there.} \end{cases} \tag{109}$$

In the split case $j_{s+n} \leq Q$. In the inert case this upper bound is P for odd n and Q for even n, since the Lie ranks of G_B have conjugate labels. The unramified character-colength estimate.

Restriction to the flag. The saturated frame extends to a basis of the saturated typed Lie lattice. Restrict the finite list of character differentials to this frame. A full-size minor of minimum valuation on the typed Lie lattice has valuation ℓ_s. Expanding that minor along the frame rows gives an i-minor of the restricted matrix of valuation at most ℓ_s, since complementary minors are integral. In particular the columns giving this i-minor are differentials of actual integral characters $\chi_1, \dots, \chi_i$. On the test germ $\prod_j \chi_j(Y_j)$ the invariant determinant power takes the value

$$\det(d\chi_j(u_a))^c.$$

Indeed each character is an eigenfunction for every invariant derivative and has value one at the identity. Therefore the factor depth is at most $c\ell_s$ on algebraic tests. From character colength to factor depth.

The mean colength above a split prime. We use this estimate at the primes in the host's exceptional split set. Let w be a place of E over p whose embeddings all belong to Ψ, and assume the embeddings of its conjugate place belong to the other type. Write $n = [E_w : \mathbf{Q}_p]$ and define the normalized discriminant parameter

$$\delta_w = \frac{1}{n} v_p\big(\operatorname{disc}(\mathcal{O}_{E_w}/\mathbf{Z}_p)\big).$$

Use the whole w-idempotent of M, which is fixed under Frobenius on the coefficient ring. The determinant argument, without diagonalizing within w, gives $D_w = \dim \operatorname{Lie}(G_B)_w/(p-1) \le nQ/(p-1)$. Indeed the determinant identity now reads $\beta_w X_w = \varphi(X_w)$ with a $\mathbf{Z}_p$-basis on the w-factor of the Tate module. As before the specialized Tate colength is D_w. Let D be the dual Lie lattice of $\widetilde{G}$ in the w-block and $L \subset D$ the lattice of the chosen character differentials. This bounds $\operatorname{length}_v(D/L) \le nQ/(p-1)$.

The separate embedding projections can enlarge this colength. The dual Tate module of $\widetilde{G}[p^\infty]$ at w is generated by at most $P + Q$ vectors over $\mathcal{O}_{E_w}$. To see the bound, let ρ_w and $\rho_{\bar{w}}$ be the ranks of the completed torus character lattice over these local coefficient rings. The Lie multiplicities of B in these blocks are $P - \rho_w$ and $Q - \rho_{\bar{w}}$. The abelian Tate rank over $\mathcal{O}_{E_w}$ is the sum of these multiplicities, by the rational Hodge–Tate exact sequence for $B^\vee$ and the polarization. Add ρ_w for the torus quotient of the dual Tate module. The result is $P + Q - \rho_{\bar{w}}$. All three Tate modules at w are finite free over the local integer ring because they are torsion free there. It follows that L is generated over $A_w = R \otimes_{\mathbf{Z}_p} \mathcal{O}_{E_w}$ by $h \le P + Q$ vectors.

Choose an integral generator of the monogenic local order, with images a_σ in the n embeddings. Under these embeddings A_w is a sublattice of R^n of colength $d = n\delta_w/2$, by the Vandermonde determinant. For any lattice V write $\Delta V = \bigoplus_\sigma V_\sigma$ for the sum of its individual projected lattices. One has $\operatorname{length}_v(\Delta D/D) \le Pd$. Indeed there are P vectors in D projecting to bases in every D_σ at once. Take linear combinations of generators of D. Modulo the maximal ideal the nonvanishing conditions on their P-minors are finitely many nonzero polynomial conditions over the infinite residue field, and can be satisfied simultaneously. The A_w-span of the chosen vectors is a copy of A_w^P in D, with projected lattice ΔD.

For the h chosen generators of L, the linear map $(R^n)^h \to \Delta L$ they induce on the projected lattices is surjective. Its kernel is a direct sum of saturated free modules K_σ of ranks $h - P$. Set $b_\sigma = \prod_{\tau \ne \sigma}(a_\sigma - a_\tau)$. The element $b_\sigma e_\sigma$ of R^n belongs to A_w, by the Lagrange polynomial with numerator $\prod_{\tau \ne \sigma}(X - a_\tau)$. Thus $\bigoplus_\sigma b_\sigma K_\sigma$ lies in A_w^h. The image of the kernel in $(R^n)^h/A_w^h$ has valuation length at most $(h-P)\sum_\sigma v(b_\sigma) = 2(h-P)d$. The quotient of this last lattice quotient by the kernel's image is $\Delta L/L$. Hence $\operatorname{length}_v(\Delta L/L) \ge hd - 2(h-P)d$. Combining the two colength differences gives

$$\begin{aligned}
\sum_{\tau \mid w} \ell_\tau &= \operatorname{length}_v(\Delta D/\Delta L) \\
&= \operatorname{length}_v(D/L) + \operatorname{length}_v(\Delta D/D) - \operatorname{length}_v(\Delta L/L) \\
&\le \frac{nQ}{p-1} + (h-P)d \le nQ\left(\frac{1}{p-1} + \frac{\delta_w}{2}\right).
\end{aligned}$$

Here D_τ is the dual of the saturated Lie eigenspace, since restriction to that direct summand of the valuation-ring Lie lattice is surjective. The flag-minor argument applies in it. Summing over the distinguished places w gives the bound $FQ(1+\delta_p/2)$ on $\sum_{\tau\in\Psi}\ell_\tau$, where $F=|\Psi|$ and $F\delta_p=\sum_w[E_w:\mathbf{Q}_p]\delta_w$. We used $1/(p-1)\le 1$. Discriminant and mean character colength at split places.

The dyadic bound. For $p=2$ the distinguished factor of the host is split from its conjugate factor. The whole-place calculation applies even when that factor is ramified over $\mathbf{Q}_2$. The integral Hodge–Tate map on the w-part of $T_pB^\vee$ lands in $\omega_{B,w}$. The period determinant satisfies

$$D_w=\dim \mathrm{Lie}(B^\vee)_w=\sum_{\tau|w}\dim \mathrm{Lie}(B)_{\bar\tau}\le [E_w:\mathbf{Q}_2]Q.$$

The dual endomorphisms have conjugate Lie labels via the polarization, while the Hodge quotient uses the original w-label on ω_B. The Smith comparison at valuation one gives the same upper bound after the Hodge quotient. Including the torus characters and the embedding projections gives a finite sum of colengths at most $FQ(1+\delta_2/2)$. For each flag, a minor in the finite list of character differentials gives an integral test germ with depth at most $c\ell_\tau$. The dyadic upper bound on flag factors.

The gap exponents of the selected weight. At the actual nonterminal height stage j use the signature $P=k+j-1,\ Q=m-j$. Its flag rank is $N=P+Q=r$. In this calculation $F=[E:\mathbf{Q}]/2=f$ is the number of distinguished labels. At the nonterminal stage in question,

$$N+1=k+m,\qquad Q\le m\le k,\qquad m\ge 1,\qquad 2\le k\le P+1.$$

Use B_{gap} for the prescribed bound on the absolute weight coordinates, and S_{gap} for the bound on $\sum_{j<k-2}u_{\tau,j}$ at each distinguished embedding. The weights are decreasing, with $u_{\tau,k-2}=0$, and the column count is $n=n_{\mathrm{col}}$. For $\tau\in\Psi$, the two exponent rows after the prescribed shift, in zero-based order, are

$$\begin{aligned}a_{\tau,j}&=u_{\tau,j}-u_{\tau,P-1}+1 &&(0\le j<P),\\ a_{\bar\tau,j}&=\big(u_{\tau,P-1}-u_{\tau,N-1-j}-1-2\lfloor n/2\rfloor\big)_+ &&(0\le j<Q).\end{aligned}$$

Append zero to each row. The row at $\bar\tau$ uses the reversed small-block order and the opposite sign. Both rows are decreasing and nonnegative. The gap factor $f=(\tau,j)$ has rank $i_f=j+1$ and exponent $c_f=a_{\tau,j}-a_{\tau,j+1}$. Call f large if it is at a distinguished embedding and $i_f+1\ge k$. At each distinguished embedding there are at most $m+1$ large factors, since $P-k+2\le m+1$. Their exponents telescope to $a_{\tau,k-2}\le B_{\mathrm{gap}}+1$. For the remaining factors in that row,

$$\sum_{j<k-2}(j+1)(a_{\tau,j}-a_{\tau,j+1})=\sum_{j<k-2}(a_{\tau,j}-a_{\tau,k-2})=\sum_{j<k-2}u_{\tau,j}\le S_{\mathrm{gap}}.$$

In the conjugate row the entire weighted gap sum is $\sum_{j<Q}a_{\bar\tau,j}\le 2QB_{\mathrm{gap}}$. Thus the parameters for the oscillation bound at this stage are

$$M_*=2B_{\mathrm{gap}}+1,\qquad M_b=B_{\mathrm{gap}}+1,\qquad S_m=F(S_{\mathrm{gap}}+2QB_{\mathrm{gap}}),\tag{110}$$

with $c_f\le M_*$, per-distinguished-embedding large exponent sum at most M_b, and $\sum_{f\text{ remaining}}i_fc_f\le S_m$. The gap bounds for the shifted weight.

An integer baseline at each coefficient prime. Fix a rational prime p and a coefficient prime $\mathfrak{q}\mid p$. All depths are in units with $v(p)=1$. When using individual factor intervals, the chart inequalities give

uniform real endpoints A_f, B_f, with lower bounds on every integral framed test and upper bounds on the algebraic tests. For $v_y = \sum_f d_f(y)$, the common nonnegative integer

$$A_{\mathfrak{q}} = \left\lfloor \max\Big(0, \sum_f A_f\Big) \right\rfloor$$

is a lower bound on all the integral tests. Nonnegativity of the individual depths permits the maximum with zero. The upper bound on algebraic tests is $A_{\mathfrak{q}} + \sum_f (B_f - A_f) + 1$. Only one rounding loss is incurred. If $A_f = 0$ for all factors, use the bound $\sum_f B_f$ without a rounding loss. Rounding the common lower endpoint.

Suppose first that p is odd, unramified in E, inert in the quadratic subfield, and $mp < k$. A large factor with $c = c_f > 0$, $i = i_f$, lies in the starting even type. Use its lower bound (107) and its algebraic upper bound $c\ell_s$ from (109). Their difference is at most

$$c\left[\Big(P - Q - i + Q\big(3/2 + \tfrac{1}{2}\log_p c\big)\Big)\frac{p}{p^2-1} + \frac{Q}{p-1}\right] + i^2(1 + \log_p c).$$

Since $N - i \le m$, $Q \le m$, and $p/(p^2-1) \le 1/(p-1)$, this is bounded by

$$\frac{cm(4 + \frac{1}{2}\log_p M_*)}{p-1} + N^2(1 + \log_p M_*).$$

For $c = 0$ the interval is zero. For each factor outside the large range use (103). Using the exponent and count bounds gives the full width

$$e_{\mathrm{small}}(p) + 1, \qquad e_{\mathrm{small}}(p) = \frac{FmM_b(4 + \frac{1}{2}\log_p M_*) + S_m}{p-1} + F(m+1)N^2(1 + \log_p M_*).$$

The small-odd-inert factor sum.

At odd unramified inert p with $k \le mp$, use (108) for the large factors. In this range $i_f/p + Q \le (k + m - 1)/p + m \le (5/2)m$, so the width is at most $(6FmM_b + S_m)/(p-1) + 1$, including the single rounding loss. At odd unramified split p, use the Ψ-bound $\ell_s \le Q/(p-1)$ on the large factors, giving $(FmM_b + S_m)/(p-1)$ from baseline zero. In the exceptional split set $\mathcal{S}_0$, including 2, the mean colength bound $\sum_{\tau\in\Psi} \ell_\tau \le FQ(1 + \delta_p/2)$ and the per-embedding M_b give $Fm(1 + \delta_p/2)M_b + S_m$, also from baseline zero. For the two residual host primes $\mathcal{R}$ the crude degree bound costs at most $(FNM_b + S_m)/(p-1)$. For $p > M_*$ every factor has $c_f < p$, hence depth zero.

More precisely, $\mathcal{S}_0$ is the host's split set with $2 \in \mathcal{S}_0$, and $\mathcal{R} = \{q_{\mathrm{host}}, d_{\mathrm{host}}\}$. The host specifies the mean discriminant depth δ_p on $\mathcal{S}_0$, nonnegative, and every prime outside $\mathcal{S}_0 \cup \mathcal{R}$ is unramified in E. Putting the cases in their order of application gives

$$e(p) = \begin{cases} 0 & p > M_*, \\ Fm(1 + \delta_p/2)M_b + S_m & p \in \mathcal{S}_0,\ p \le M_*, \\ (FNM_b + S_m)/(p-1) & p \in \mathcal{R} \setminus \mathcal{S}_0,\ p \le M_*, \\ e_{\mathrm{small}}(p) + 1 & p \notin \mathcal{S}_0 \cup \mathcal{R},\ p \le M_*,\ mp < k, \\ (6FmM_b + S_m)/(p-1) + 1 & p \notin \mathcal{S}_0 \cup \mathcal{R},\ p \le M_*,\ k \le mp. \end{cases} \tag{111}$$

The widths are nonnegative, and the odd unramified split estimate is at most the assigned width in either of the last two cases. For each $\mathfrak{q} \mid p$ this produces the interval $A_{\mathfrak{q}} \le v_y \le A_{\mathfrak{q}} + e(p)$ with the upper and lower test quantifiers described above. Its application to the reference cusp lattice is at primes away from the selected type level $M_0 = g_0 g_1$. The all-prime interval.

Summing the widths with the cutoff. Write $\Lambda, \Lambda_x \ge 1$ for bounds satisfying

$$\log M_* \le \Lambda, \qquad \log N \le \Lambda, \qquad \log\log k \le \Lambda_x,$$
$$\sum_{p\in S_0} (1+\delta_p/2)\log p \le \tfrac{3}{2}\Lambda, \qquad \frac{N\log p}{p-1} \le m\Lambda \quad (p \in \mathcal{R}).$$

All logarithms without a subscript are natural. Put $A = FmM_b + S_m$. The S_0 terms of $\sum_p e(p)\log p$ are bounded by $3A\Lambda/2$. For each residual host prime the summand is bounded by $A\Lambda$, using $\log p/(p-1) \le 1$ for the S_m term. This contributes at most $2A\Lambda$. For the two general cases use the common allowance $(6FmM_b + S_m)/(p-1) + 1$ at all primes up to M_*. When $mp < k$, the only additional terms needed are

$$\frac{FmM_b \log_p M_*}{2(p-1)} + F(m+1)N^2(1+\log_p M_*).$$

Multiplication by $\log p$ turns $\log_p M_*$ into $\log M_*$. For the first correction use the sum of $1/(p-1)$ over primes at most k. For the second, use $\log M_* + \log p \le 2\Lambda$ and

$$(m+1)\#\{p \le M_* : mp < k\} \le (m+1)\left\lfloor \frac{k-1}{m} \right\rfloor \le 2N.$$

In particular the second correction contributes at most $4FN^3\Lambda$.

Here are bounds for the prime sums with absolute constants. For each integer $Y \ge 1$, $\vartheta(Y) = \sum_{p\le Y} \log p \le Y\log 4$. In the induction for $Y = 2a$ use $\binom{2a}{a} \le 4^a$ to account for the primes above a. For $Y = 2a+1$, the primes above $a+1$ divide $\binom{2a+1}{a} \le 4^a$, since the two central coefficients are equal. Applying the inequality $\lfloor Y/p \rfloor \le v_p(Y!)$ gives $\sum_{p\le Y} \log p/p \le \log Y + \log 4$. The remaining sum is bounded by

$$\sum_{n\ge 2} \frac{\log n}{n(n-1)} \le \frac{\log 2}{2} + \frac{\log 3}{6} + \frac{\log 3 + 4/3}{3}.$$

For $n \ge 3$, the term $\log(n+1)/(n(n+1))$ is at most the difference of $(\log n + 4/3)/n$ and $(\log(n+1) + 4/3)/(n+1)$, by $(n+1)\log(1+1/n) \le 4/3$. Together with $\log 2 < 3/4$ and $\log 3 < 6/5$, this proves

$$\sum_{p\le Y} \frac{\log p}{p-1} \le \log Y + 3.$$

The same bound for real $Y \ge 1$ follows by taking its floor. Abel summation at $k \ge 2$ gives

$$\sum_{p\le k} \frac{1}{p-1} \le 1 + \log\log k - \log\log 2 + \frac{3}{\log 2} \le \log\log k + 6 \le 7\Lambda_x.$$

Here $\log 2 > 2/3$, by convexity of $1/t$ on $[1,2]$, and $-\log\log 2 < 1/2$. The cutoff rounding sum costs at most $\vartheta(M_*) \le 2M_*$. Thus the common allowance is at most $24A\Lambda + 2M_*$, and the first small-prime correction is at most $(7/2)A\Lambda\Lambda_x$. Combining the bounds proves

$$\sum_p e(p)\log p \le 40\big((FmM_b + S_m)\Lambda\Lambda_x + FN^3\Lambda + M_*\big). \tag{112}$$

The weighted prime sum.

The allowances at the selected nonterminal stage. Let $\ell = \mathrm{lenT}(PP) \ge 1$ be the length factor of the selected row. At this stage $F = f$, $N = r$, and the exponent inputs are exactly (110). Write C_1, C_2 for the constants in the weight and size bounds, respectively. The weight bounds give

$$\ell(B_{\mathrm{gap}} + 1),\ \ell S_{\mathrm{gap}} \le C_1 b.$$

The size bounds give $L, b, \log\log x \geq 1$ and

$$\log b,\ \log r,\ \log(2Q_{\mathrm{host,S}}) \leq C_2 L, \qquad \log\log k \leq C_2 \log\log x,$$
$$\ell r^3 \leq C_2 b, \qquad FmbL\log\log x \leq C_2 E_{\mathrm{cov}}, \qquad \frac{r\log p}{p-1} \leq C_2 mL \quad (p \in \mathcal{R}).$$

Here $Q_{\mathrm{host,S}}$ is the host integer whose prime divisors form $\mathcal{S}_0$. At those primes $0 \leq \delta_p \leq v_p(Q_{\mathrm{host,S}})$. The discriminant sum is bounded by $(3/2)\log Q_{\mathrm{host,S}} \leq (3/2)\log(2Q_{\mathrm{host,S}})$ because $v_p(Q_{\mathrm{host,S}}) \geq 1$ there. With $K_0 = 2C_1 + C_2 + 1$, take $\Lambda = K_0 L$, $\Lambda_x = K_0 \log\log x$. In fact $2B_{\mathrm{gap}} + 1 \leq (2C_1+1)b$, so $\log M_* \leq \log b + 2C_1 \leq K_0 L$. Also $Q \leq m$ gives

$$\ell(FmM_b + S_m) \leq 4C_1 Fmb, \qquad \ell M_* \leq 2C_1 b.$$

After multiplying (112) by ℓ, the three terms are therefore bounded in total by

$$40(4C_1K_0^2 + C_2K_0 + 2C_1)C_2E_{\mathrm{cov}}.$$

The support is the primes at most $2B_{\mathrm{gap}} + 1$. The very same $e(p)$ supplies the pointwise bound at all primes outside the type level, so this is the allowance for the common reference-section estimate. The size bounds at the actual inputs.

The reference polynomial away from the type level. Fix a rational prime $p \nmid M_0$ and a coefficient prime $\mathfrak{q} \mid p$. Use $v(p) = 1$. At a framed integral test y, let ν_y be the sum of the least formal depths of its fundamental gap factors. For a factor of rank r_f and exponent c_f, its depth is the least value of the c_f-th power of the determinant of invariant derivatives in the first r_f directions of the prescribed eigenflag, applied to an arbitrary integral germ on r_f copies. The preceding calculation gives the interval

$$A_{\mathfrak{q}} \leq \nu_y \leq A_{\mathfrak{q}} + e(p), \qquad A_{\mathfrak{q}} \in \mathbf{Z}_{\geq 0},$$

for the selected gap exponents at all the required tests, with its upper bound on algebraic tests. Here the interval is the one for the formal gap factors. We must carry it to a single harmonic polynomial and to values of actual reference sections. The interval for formal depths and The sum of the factor depths.

Allocate $2r_f$ distinct auxiliary columns to every nonzero factor, as allowed by the prescribed column count. In $\mathcal{O}_E \otimes \mathbf{Z}_p$ choose u with $u\bar{u} = -1$. It exists by surjectivity of the norm on units in an unramified quadratic extension, and by using the two factors at split places. Use the formal map ι sending each input column z to the pair (z, uz) in its allocated columns and zero in unused columns. The completed endomorphism acts on the formal group. To see this, the group-law series of $[p]$ has linear part p and terms of degree at least two, so $[p^j]x$ tends to zero in the (p,x)-adic topology. The limit of integral CM endomorphisms is consequently defined coefficientwise on the formal group. The product map ι has an integral retraction by projection to its first entries and satisfies $\iota^*\iota = 0$ for the standard Hermitian form on columns.

For the product gauge H, the series $\exp(H \circ \iota)$ is an integral unit with constant coefficient one. Indeed, for any integral rectangular matrix B, addition and adjointness in the Poincaré frame express $H \circ B$, modulo logarithms of integral units, as

$$\sum_i h\big(x_i,\, w_0(B^*Bx)_i\big).$$

Approximate the entries of ι by integral CM elements. On every finite jet the expression on the right tends to zero by $\iota^*\iota = 0$. The exponential of the difference is an integral unit at every approximation. Its limit is integral and has constant coefficient one. This uses the same cubical multiplication as in (100).

Let P_{loc} be the product of the indicated determinant forms in the column directions of $d\iota$. The conjugate blocks are disjoint, so all conjugate contractions vanish. For an integral germ F on the auxiliary columns, differentiating after restriction by ι computes the same harmonic value as applying the product

of fundamental determinant operators to $(F \circ \iota) \exp(kH \circ \iota)$ at level k. The retraction and the integral unit show that the germs in this last expression range over all integral germs in the factor columns. The least order equals the sum of the factor orders: on each coordinate monomial the derivative values factor, and the minimum is reached by a product of minimizing monomials. Only the truncation of total degree at most d is needed. Since an integral frame change induces an isomorphism on the jet modules, this value is independent of the formal coordinates. Extension of the trait preserves it after renormalization, as can also be seen by the minimum over monomial jets. Thus the least order for P_{loc} is exactly v_y. The formal factor-depth comparison.

Choose a global integral CM approximation to u modulo p^N, using density of $\mathcal{O}_E$ in its completion at p. This gives the polynomial P over the existing coefficient field. Its column supports remain disjoint, so harmonicity is exact. The coefficients of the difference are in p^N times the integral coefficient lattice in every test over $\mathfrak{q}$. By the uniform conductor estimate, the difference of the two harmonic values on any integral germ has valuation at least $N - d\delta_p/2$ with $v(p) = 1$. Taking $N > A_{\mathfrak{q}} + e(p) + d\delta_p/2$ proves that P has the same depths wherever the upper bound on v_y holds, and the same lower bound throughout. Both comparisons use the same P at all tests over $\mathfrak{q}$. Globalizing the interval.

In the chart of the chosen test take a minimal jet for P. The small Fourier modes and integral translate basis at level M_0, a unit here, realize it up to scalar loss at most $d^2\Phi_0/M_0$ in unnormalized valuation units. The marked transport factor C_{α_0} is a unit at p. Linearity and the ultrametric inequality therefore give one marked rule with normalized valuation at most

$$A_{\mathfrak{q}} + e(p) + \frac{d^2\Phi_0}{M_0, v_{\mathrm{frame}}(p)}.$$

Every normalized marked translate has integral Raynaud germ in the same frame, so every such rule has order at least $A_{\mathfrak{q}}$ at all tests over that prime. The marked theta cusp sections then admit a single scalar of order $-A_{\mathfrak{q}}$ in the stated normalization. The order is available over the coefficient field since $A_{\mathfrak{q}}$ is integral and $v(p) = 1$. By weak approximation the scalar can also clear all denominators at the other finitely many coefficient primes. The finiteness follows from the fixed level and the conductor bound for the normalized translates. After this compensation and the selected boundary allowance, the reference exponent at p is at most $e(p)$ on the actual integral cusp lattice. The coefficient-prime compensation.

By the weighted bound (112) on the chosen weight, the ℓ-fold reference cost is $O(E_{\mathrm{cov}})$. Pullback preserves the bound on the same polarized fibre. The key and parallel weight change cost $O(f^2 r^3 L)$ in the weighted sum by the different bound and the cover's bad-prime support. This fits the budget since $\ell P_{\mathrm{cov}} \le \sqrt{b}$. The change of weight.

Combining the two kinds of estimate gives the metric and finite bounds

$$\|\theta_Z^{\mathrm{tr}}(y)\| \le e^{8t_\sigma} \mathscr{B}_Z(y)^4, \qquad v(\theta_Z^{\mathrm{tr}}) \ge 8\big(d_v(u_1) - e_v^{\mathrm{tr}}\big). \tag{113}$$

The sum of the finite costs e_v^{tr} and complex costs t_σ, with normalized weights of the cover field, is at most $c_{\mathrm{mult}} = O(E_{\mathrm{cov}}/\ell)$. In fact, the valuation of T_Z^{24} contributes at least $24\Phi/(625e)$ in a frame normalized by its ramification e, and the key has nonnegative finite valuation. Thus the flat valuation ρ_v contributes $-\rho_v/8$ to the finite loss after allowing the reference boundary term $\Phi/(2500e)$. At infinity the theta estimate and key bound give $t_\sigma = \beta/2 + 6fr^2 \log(2\pi)/8 + \log \|S_Z^{\mathrm{flat}}\|_\sigma/8$. Together with the finite flat valuations the last term contributes $-\deg(S_Z^{\mathrm{flat}})/8$ in normalized degree, so the key and flat allowances together cost $O(P_{\mathrm{cov}})$. The remaining finite costs are the reference-place budget above. The transfer from mixed integrality to de Rham tests costs at most $e_2(p)$ per unit of absolute weight on the prime support above. It is measured by the diagonal entries of the triangular change of eigenframes, whose denominators are bounded by a $(p-1)$-st root of p and the different. Set $c_c = c_{\mathrm{mult}} + |u'|_1 \sum_p e_2(p) \log p$. Then $\ell c_c = O(E_{\mathrm{cov}})$ by the size and prime-sum bounds. Carrying section, Its signed finite cost, Pointwise and finite comparisons, and Triangular frame cost.

For each $8 \mid n$, multiply a mixed-integral section f at nu' by $(\theta_Z^{\mathrm{tr}})^{n/8}$ and a nonzero scalar. Choose the scalar in the fractional ideal prescribed by the summed finite costs, rounding up. Minkowski's theorem chooses it with archimedean log sizes at most $n(c_c - t_\sigma) + O(1)$, because the product of these box bounds uses the same norm as the ideal, up to a fixed factor for the field and rounding. The product section satisfies the required depth and increases log norm by at most $nc_c + O(1)$ in the quotient metric. The multiplication is injective. Using the lower count at $t = c_s$ from the root model and then the upper count at $c_s + c_c + \varepsilon'$ for arbitrary $\varepsilon' > 0$ gives

$$\delta'_Z \max(0, s_Z + c_s) \le d_Z(h_Z + c_s + c_c).$$

Here $\delta'_Z \ge (1-\rho)d_Z$ by the comparison of geometric degrees. This is the mixed-to-moving height estimate. The finite test specialization and the quantitative estimates on the carrying section and the Bergman kernel concern these same component metrics. Scalar multiplication and count comparison.

Summing the target degrees. Set $Q = \binom{\ell-1}{D}$ and $R = \sum_Z r_Z > 0$. Multiply the count inequality by Q/R, sum over components, and replace the positive part by its argument. Using $(1-z)(1-\rho) \le Q\sum_Z \delta'_Z/R \le 1$ and $Q\sum_Z d_Z/R \le 1$ gives

$$R^{-1}\sum_Z q_Z t_Z \ge (1-z)(1-\rho)c_s - X - c_c.$$

Thus the flag bound is used with its geometric weights even though the target degree is divided by the total section rank. For $a \ge 0$, $A\log(2+a) \le a + 2 + A\log\max(A,1)$. Put $d_1 = 3D + 4 < \ell$. Since $|h_Z| \le t_Z + c_s$, the bracket in (98) is bounded below by $(\ell - d_1)t_Z - (\ell + d_1)c_s - O(b)$. Here $A_{\mathrm{HS}} \le 100\ell$ and $A_{\mathrm{HS}}\log\max(A_{\mathrm{HS}},1) + R_{\mathrm{HS}} = O(b)$, by the frame bounds. The coefficient of c_s after summing is at most

$$(\ell + d_1) - (\ell - d_1)(1-z)(1-\rho) \le 6D + 8 + \ell(z+\rho).$$

With $c_s \le P_{\mathrm{cov}}(\kappa/\ell + 1)$ this term is $O(zP_{\mathrm{cov}}\kappa + z\ell P_{\mathrm{cov}}) = O(b)$. The total lower bound per target rank is $-O(b + \ell X + \ell c_c) = -O(E_{\mathrm{cov}})$.

The lattice of sections at v splits by component. Its Petersson norm is the orthogonal sum of the squared component norms weighted by their probability masses at each embedding. Each mass belongs to $(0,1]$, so passage to the whole cover can only increase the determinant degree relative to the orthogonal sum of the component norms. Its rank is R. Consequently its squared mean slope is at least $2R^{-1}\sum_Z \widehat{\deg}_{\mathrm{n}} J_Z$, which is bounded below by $-O(E_{\mathrm{cov}})$. The component masses and determinant degree.

4.3.4 Descent and terminal truncation

The flat metric transfers. The determinant character for a constant weight at one CM type has ordinary degree

$$c_1 = -\tfrac{p_j - q_j}{2}\log\Omega + q_j(\log \mathrm{rd}(E) - \log d)$$

with the dual determinant convention. In the integral comparison, the finite de Rham correction at that type has normalized sum $q_j(\log \mathrm{rd}(E) - \log d)$, by the norm of the different and the correction at d. Thus multiplying by the flat generators to powers v_τ has normalized logarithmic comparison degree at least $-\tfrac{p_j - q_j}{2}\sum_\tau v_\tau \log\Omega$ before rounding. The residue formula and rational trivialization are used on the actual cover. The flat generators and their inverses have constant point norms and preserve the cusp condition. Signed flat degree and Finite residue correction.

Let $m_{\mathfrak{p}}$ bound the finite valuation from below in all tests over $\mathfrak{p}$. Twist the source by $\prod_{\mathfrak{p}} \mathfrak{p}^{-\lfloor m_{\mathfrak{p}} \rfloor}$ to make multiplication integral. The lost degree from rounding is at most the sum of the normalized log norms of the bad primes. The unramified comparison restricts them to primes over $2d\,\mathrm{disc}(E)M_{\mathrm{cov}}$, of total rational log norm at most $40L$. Since the map has full rank, enlarging its lattice image to the target can only increase degree. For the first transfer the powers are $v_\tau = -\kappa_{\mathrm{can},\tau,k-2}$, obtained by comparing the

anchors of a_ℓ and the lowered weight. They satisfy $|v_\tau| \le 3r$. Using $|c_1| \le 9rL$ and the residue allowance bounds the ordinary degree loss per rank by $120(frL)^2$. This gives the cost already used for the component comparison. Integral multiplier and Support of the rounding cost.

For the second transfer the constant weights are the actual stage anchors α_τ. There is also the inverse fourth root of the theta key to the power $n_\theta/2$, where n_θ is its even theta exponent. If the fourth power is $a_Z K_Z^{\mathrm{key}}$, the local changes from a_Z cancel by the product formula. The key comparison gives a finite loss bounded by the different on the Lie eigenspaces and an archimedean loss of $fr\log(2\pi)/2$. Their sum per inverse root is at most $fr(\log \mathrm{rd}(E) + \log(2\pi)/2)$. Round only after multiplying the flat and key factors. In squared slope units the flat contribution is $-(p_j - q_j)\sum_\tau \alpha_\tau \log \Omega$. Replacing the anchors by $w_{j\tau}$ costs $O(E_{\mathrm{cov}})$, and the key contribution does too since $n_\theta frL \ll b$. The lower-weight cover lattice therefore has squared mean at least $-M_j - O(E_{\mathrm{cov}})$ once the target weight has mean at least $-O(E_{\mathrm{cov}})$. The inverse key root and The anchor comparison.

Descending at a nonterminal stage. Let A be the lower-weight cover lattice with squared mean at least $-M_j - (C_U + C_F)E_{\mathrm{cov}}$ from the two comparisons. For the actual type lattice B on the cover, the descent construction uses the Hodge shift line $W = c_0 H(h)$ in additive notation and the action of the descent group S of order n_0. Let T_g be its action on sections and χ the prescribed character.

Evaluation at a level orbit. For the Hodge shift $W = c_0 h$, work on the toroidal base of dimension $d_X = fp_j q_j$. Consider the component stabilizer in the chosen finite level subgroup. The selected split prime has residue characteristic $g > r + 1$ and the subgroup acts unipotently on that half-level. The finite-order stabilizer calculation gives freeness on the open and descent of the linearized Hodge line on the equivariant projective model. Use the finite quotient and resolve on the quotient and on the source, keeping the same open finite étale locus. The descended Hodge line H_0 is nef and big on the quotient resolution. Indeed its pullback is the nef and big Hodge line on the corresponding source modification. Choose an orbit over a point in the étale locus outside the support of an effective divisor in a Kodaira decomposition of H_0. Every positive-dimensional subvariety Y through that point then has $H_0^{\dim Y} \cdot Y \ge 1$. The Hodge linearization at the prescribed neat prime and The equivariant modification and descent of adjoint sections.

Put $a = \binom{d_X+1}{2}$ and $m_0 = a + 1$. Kollár's isolated non-klt point construction applies to $(a + \frac{1}{2})H_0$, since its e-th intersection root exceeds a on all subvarieties of dimension e through the given point, and $\sum_{e=1}^{d_X} e/a = 1$. It gives an effective rational divisor $\Delta \equiv (a + \frac{1}{2})H_0$ whose non-klt locus near the point is the point itself. [118, Theorem 6.4] If f is the resolved quotient map, the cosupport of the multiplier ideal of $f^*\Delta$ near the orbit consists of the orbit points, by étaleness. The adjoint twist on the source at weight W is $K + m_0 f^* H_0 + N'$ before multiplying by the boundary equation. The remaining twist N' is nef by the weight-gap criterion, since $c_0 - m_0 \ge r$. Subtracting $f^*\Delta$ from $m_0 f^* H_0 + N'$ leaves a nef and big rational divisor. Nadel vanishing makes restriction to the multiplier-ideal quotient surjective. Its isolated local summands at the orbit points map onto their residue fields, so one can prescribe their values simultaneously. The canonical bundle comparison for the resolution pushes those sections to the toroidal base, and multiplying by the boundary equation yields cusp sections without changing surjectivity at the orbit. If the orbit meets several components apply this on the component stabilizer of each and use the direct sum of the component cusp spaces. [119, Theorem 2.4] The numerical criterion for the isolated point, The pullback and Nadel argument and Application to the level orbit.

The evaluation bound at the determinant shift. The short generators in the determinant construction belong componentwise to the cusp lattice on the actual cover at $W = c_0 h$, with theta exponent zero. The orbit-separating construction uses the toroidal boundary equation to put its sections in that cusp space. The lattice generators are chosen in the same space. In the analytic description at a prescribed CM embedding, this space maps into the span of the bounded holomorphic flag sections whose central

Fourier–Jacobi coefficient is zero at every integral isotropic line. This is the cusp condition in the definition of that lattice. Boundedness and the vanishing are linear conditions. Every rational isotropic line has an integral generator after clearing its denominator. Changing the generator multiplies the transvection coordinate by the squared scale and preserves its normalized central average. Hence the cusp condition applies to every transvection tested in the estimate at a point of more than half the supremum. The component cusp space and The boundedness and central cusp condition.

Apply the local argument in the principal cusp norm comparison on this cover at W. For the arithmetic stabilizer of a cover component use the principal transvection period $M_{\mathrm{cov}}\,\mathrm{N}\mathfrak{c}_K\,d$. That congruence acts trivially on the marked period module and the determinant-one transvections have trivial weight character. The period lattice and its Hermitian self-dual twist are the same after forgetting the extra marking. The flag spread $w = 1 + \sum_\tau \mathrm{spread}(W_\tau)$ has $\log w = O(L)$ since $|W|_1 \le (fr)^5$. The cutoff for the minimum at a point of more than half the supremum and the holomorphic chart radius therefore have inverse logarithms $O(L)$, using $\log M_{\mathrm{cov}} \le 30L$. The return count on the period lattice and the submean inequality in complex dimension $f\binom{r}{2}$ give a squared evaluation bound with logarithm $O((frL)^3)$, up to the mass used to normalize the quotient.

Use the component probability mass. The unnormalized covering mass V_Z of the complex points of the component Z is at most $V_j[K(M_0) : K(M_{\mathrm{cov}})]$, where V_j is the mass at the base principal level with the same real measure. The index is at most $M_{\mathrm{cov}}^{(2fr)^2}$ by taking the action on the underlying marked rank-$2fr$ period lattice modulo M_{cov}. This also bounds the mass if Z contains several geometric components. Thus $\log^+ V_Z = O((frL)^3)$. In the point norm and in the component probability L^2 norm the Lie factor is the same positive scalar for this weight. The ordinary ratio therefore satisfies

$$\|F_\sigma\|_{\sup,Z} \le \beta_W \|F_\sigma\|_{\mathrm{prob},Z}, \qquad \log \beta_W = O((frL)^3),$$

where $\beta_W \ge 1$ is the square root of the squared evaluation bound. This calculation uses the scalar line norm on the flag. Conjugation transports both measurements to the other embeddings of the cover coefficient field. The upper bound is uniform in σ and gives the point norm bound $\beta_W \Lambda$ for each generator of probability L^2 length at most Λ. The component probability norm and The bounded evaluation ratio at the actual weight.

Prescribe the value one at the identity and zero at every other point of such an orbit. The matrix $D_u = (T_g T_h u)_{g,h\in S}$ then evaluates to a permutation matrix. Its determinant is therefore a nonzero polynomial in u at one point of each component. Express a rational section satisfying these finitely many conditions in a short integral spanning set, with each generator supported on one component. Replace the coefficients one at a time by integers from 0 to n_0^2. Each coefficient affects at most n_0 determinant polynomials of degree at most n_0. Some integer keeps all their values nonzero. For the resulting integral u, if B_u is its common supremum norm bound, the short frame and point-norm estimate at W give $n_0^3 B_u^2 \le e^{2b}$. Indeed its pointwise bound is n_0^2 times the component generator count times $\beta_W \Lambda$, using the evaluation bound just proved. The short generator estimate gives $\log^+ \Lambda = O((frL)^9)$, including the logarithm of their number on a component. Integral coefficient selection and The size at the Hodge shift.

For $a \in A$, use the components $\sum_{g\in S} \chi(g)^{-1} T_g(a\, T_h u)$ for $h \in S$. The unnormalized projector is integral and lands in B. If all components vanish, the adjugate of D_u^{T} shows that $\det D_u$ kills the vector with entries $\chi(g)^{-1} T_g a$, so $a = 0$. The squared norm is at most $n_0^3 B_u^2$ by the triangle inequality and invariance of the Petersson norm. Conversely, $(b_h)_h \mapsto \sum_h b_h T_h u$ injects B^{n_0} into the cusp space at the upper shifted weight. Apply every T_g, divide by $\chi(g)$, and use D_u to recover injectivity. On each component the upper shifted weight is bounded in the nef direction by $(\ell + D + 4)u_1$, up to its constant character. The difference has gaps $3b_i - c_i$ minus the canonical gap, with a further subtraction of $2c_0$ at the cross-block gap. They are nonnegative by the bounds on b_i, c_i above. The jet-plane argument thus gives upper rank $d_Z\binom{a_+ + D - 1}{D}$, with $a_+ = 1 + \lfloor(\ell + D + 4)/(1 - \rho)\rfloor$. The ratio to the lower rank $d_Z\binom{\ell-1}{D}$ is at most

$$\left(\frac{1 + (3D + 4)/(\ell - D)}{1 - \rho}\right)^D \le \exp\left(\tfrac{9}{2}(D^2/\ell + D\rho)\right) \le 1 + z.$$

Here $D \geq 3$, $\ell \geq 1000D^2$, and $z \leq 1/2$. Thus the two maps give

$$\operatorname{rk}(B^{n_0}) \leq (1 + b/K_{\mathrm{cov}}) \operatorname{rk} A, \qquad K_{\mathrm{cov}} = P_{\mathrm{cov}}\kappa.$$

The unnormalized character projector gives integral spanning vectors in B from the short frame at its weight, adding at most $\log n_0$ to their log lengths. The prescribed bound is K_{cov}, with $K_{\mathrm{cov}} \geq \max(b, |M_j|)$. Character maps, Shifted rank, and Projected frame.

To see the slope estimate, saturate the image of A in B^{n_0}. Taking the top exterior power bounds its squared degree below by $\operatorname{rk}(A)(\mu(A) - 2b)$. Project an integral short spanning set of B^{n_0} onto the quotient and pick a rational basis from it. Its determinant bounds the quotient's squared degree below by $-2K_{\mathrm{cov}}$ times the quotient rank. If $t = -M_j - (C_U + C_F)E_{\mathrm{cov}} - 2b$, division by the rank of B^{n_0} now gives

$$\mu(B) \geq t - \frac{b}{K_{\mathrm{cov}}}\big(|t| + 2K_{\mathrm{cov}}\big).$$

The same-space comparison to the core downstairs costs at most $2b$ in squared slope. Using $b \leq E_{\mathrm{cov}}$, $|M_j| \leq K_{\mathrm{cov}}$, and $b/K_{\mathrm{cov}} \leq 1$ gives $\mu_{\mathrm{core}} \geq -M_j - (2(C_U + C_F) + 9)E_{\mathrm{cov}}$. Integral descent and metric comparison and Slope transfer.

Terminal stage and truncation. At $j = m$ the second signature is zero. The core slope comes from the finite definite space. The coefficient field splits its components. Evaluation at a framed point identifies the rational space supported there with a power n_θ of the theta line tensored with the dual Weyl representations of the Lie blocks. The integral fibre model maps into the integral cusp lattice in the prescribed point norm. Split terminal components and integral fibre comparison. Let $p = p_m$, write λ_τ for the highest weight on the rank-p Lie block of type τ, and d_τ for its dual Weyl dimension. The flat determinant line has ordinary degree $c_1 = -p \log \Omega/2$ for each type in the chosen CM type Ψ. For the ordinary degree of the theta line, the fourth-power theta key compares its norm with the product of these determinants. At infinity its constant is $(2\pi)^{2fp}$ in the determinant formula. At the finite places the loss is bounded by the colength from the different in the integral Lie decomposition. The product formula cancels the scalar of the key and gives

$$d_\theta \geq \tfrac{1}{2} \sum_{\tau \in \Psi} c_1 + fp\big(\tfrac{1}{4}\log(2\pi) - \tfrac{1}{2}\log \operatorname{rd}(E)\big).$$

The theta key at the two kinds of place and the product formula. The determinant of the action on a dual Weyl space is $(\det B)^{|\lambda_\tau| d_\tau/p}$, as is seen by evaluating the determinant character on scalars. Each unitary minor has absolute value at most one. The integral bideterminants therefore have squared norm at most d_τ in the d_τ-scaled Haar metric, and Hadamard bounds the Gram determinant by $d_\tau^{d_\tau}$. Tensoring the fibre factors gives the lower bound

$$n_\theta d_\theta + \sum_{\tau \in \Psi} \big((|\lambda_\tau|/p)c_1 - \tfrac{1}{2}\log d_\tau\big)$$

for the ordinary mean slope of each lattice in the fibre model. The Gram determinant.

Choose one fibre in each free orbit of the terminal level group G. Send a vector from the fibre at o to the sum of its translates with coefficients $\chi(g)^{-1}$, where χ is the prescribed character and the input vector is extended by zero on the other fibres. There is one term at each point of the orbit. Restriction to the representatives inverts this map on the character space. The map is integral and its probability Petersson norm is bounded by the orthogonal sum norm on the representative fibres. Taking determinants doubles the ordinary fibre slope in squared slope units. The leading term is $-\sum_{\tau,t}(\lambda_{\tau,t} + n_\theta/2)\log \Omega$. Replacing these entries by their anchors and bounding the logarithms of the dimensions by the Weyl product gives $-M_m - O(E_{\mathrm{cov}})$ using the terminal rank and gap bounds. Definite eigenspace and terminal main term.

The rank and index of the exceptional subspace. At a nonterminal stage, let V_a be the complex cusp space on the cover at Hodge shift ac_0. The type space is $V_0[P,\lambda]$, the λ-character space for the row subgroup P. The two split half-level subgroups are H_0, H_1, with the actual actions on V_0 twisted by characters. For either $H = H_k$ use the same shift line and the translate determinant for H to choose u_H. The map

$$V_0 \longrightarrow \mathrm{Fun}(H, V_1^H), \qquad f \longmapsto \left(\sum_{g\in H} T_g(f\, T_h u_H) \right)_{h\in H}$$

is injective by the adjugate argument for the character maps. Reindexing g shows that $T_s f$ corresponds to shifting the argument h to $s^{-1}h$. Thus the type action embeds in $n_H = \dim V_1^H$ copies of the regular representation, since twisting a regular representation by a character is removed by diagonal multiplication of its coordinates. The inverse rank comparison by products with the translates of u_H gives $n_H|H| \le \dim V_2$. On the P side the joint injection of V_{-1} into $V_0[P,\lambda]^{|P|}$ gives the other inequality

$$n_H|H| \le \dim V_2 \le \tfrac{3}{2} \dim V_{-1} \le \tfrac{3}{2}|P| \dim V_0[P,\lambda].$$

The middle bound is the shifted-rank estimate with upper shift $2c_0$. Its jet calculation uses $(\ell + D + 5)u_1$ as the nef upper weight, with $4b_i - c_i$ minus the canonical gap and at the crossing a further subtraction of $3c_0$. These gaps are nonnegative by the same choices of b_i, c_i. The ratio to the lower binomial bound is at most $((1 + (3D + 5)/(\ell - D))/(1 - \rho))^D$. This is at most $\exp(5(D^2/\ell + D\rho)) < 3/2$ in the cover range. The regular embedding and Rank growth at $2c_0$.

At stage j, the auxiliary signature is $(k + j - 1, m - j)$, and the finite type rank is its sum $r = k + m - 1$, including at the terminal stage. Write $G_k = \prod_{v|g_k} \mathrm{GL}_r(\mathbf{F}_{g_k})$ and $U_k = \prod_{v|g_k} U_r(\mathbf{F}_{g_k})$, with one chosen split component per place of F over g_k. The split level map of the cover identifies its finite group with $N \times G_0 \times G_1$, where N is the kernel of the map to the type group. The row character subgroup P is the preimage of $U_0 \times U_1$, the base level P_0. The half-level subgroup H_k is the preimage of the product with U_k replaced by G_k. After absorbing the scalar level twist into the action, the row character is inflated from the product of the generic characters of the U_r factors. This gives the base-level Whittaker projectors on both halves. The symmetric sharp level P_s uses the projectors e_J in the distinguished row count. The two half levels and the base character, The actual split action and nonterminal rank identifications, and The terminal rank identifications.

Let B_k be the central projector on the bad tuples at g_k, and $\mathcal{S}_k$ their share in the product of the base Whittaker laws at those f places. The identity coefficient of $B_k e_{P,\lambda}$ in $\mathbf{C}[H_k]$ is $\mathcal{S}_k/|P|$. Indeed the kernel average and the other half's Whittaker average contribute $1/|N|$ and $1/|U_{1-k}|$. The factors at g_k contribute the Whittaker weights divided by $|U_k|$. The weights are the same for every nontrivial additive character, since diagonal conjugation changes the additive character and fixes the central projectors. Regular trace and the preceding rank estimate give a relative bad rank at most $\frac{3}{2}(\mathcal{S}_0 + \mathcal{S}_1)$. At the terminal stage each free H_k-orbit of fibres gives copies of the regular representation, and its λ-dimension is the orbit's total dimension divided by $|P|$. The relative bound there is $\mathcal{S}_0 + \mathcal{S}_1$. The product trace on the character space.

The actual good selector is the product of the two good-tuple idempotents. Its tuple condition is the one in the distinguished row count. The labels are squarefree at each place, satisfy the root separation and two divisor-degree conditions there, and the prime-to-ℓ root test uses each prescribed twist value and its inverse. Both runs use the same host primes q, d, g_0, g_1, with $qd \mid g_k - 1$, and the count is uniform in the prescribed twist values at those places. The squarefree law from Section 3 therefore identifies the good-set probability for the product of Whittaker laws with its probability in the permutation model. The laws both have mass one, so the probabilities of the complementary sets also agree. The permutation estimates used in (66) give

$$\mathcal{S}_k \ll \frac{f r^3 \log g_k}{\sqrt{g_k}} + \frac{f r^2 m \log g_k}{\min(q,d)} + r(7/8)^f + \exp(-cf/\log r).$$

Here the squarefree failure itself is included in the first term. The size conditions used there, with $e_z = \varepsilon_{\mathrm{row}}$, give the uniform bound

$$2(S_0 + S_1) \le \exp(-x^{e_z/7}) \tag{114}$$

at a common threshold. This bounds the exceptional rank fraction at the actual base level at every stage. The actual central selector and Both root-twist signs.

The core in the primitive row realization. The row coefficient field K contains F_0. The untwisted semilinear realization ι_j takes the actual rational core cusp forms to the corresponding arithmetic rows. With the scalar normalization of the level action, the character equations are defined over F_0. Consequently the base type lattice in the row realization is $\mathcal{O}_K\, \iota_j(J_{\mathrm{core}})$. In the rational coordinates of the core this follows by tensoring the kernel of the character equations on the core lattice with $\mathcal{O}_K$. This extension is flat. Since the equations have root-of-unity coefficients and preserve the rational form, the resulting kernel is the intersection with the K-linear type space. The integral type subspace and The character equation on the selected rows.

The row measure and the probability norm. Fix j, and let U_j be the unitary group of the rank-r auxiliary row of signature (p_j, q_j). The principal level has modulus $M = g_0 g_1$. Write $K_j(M)$ for its finite subgroup and $K_j^{\max}$ for the full stabilizer of its self-dual lattice. Use the row measure dv_j, with $K_j(M)$ of volume one and the prescribed invariant densities at infinity. Put $V_j = v_j([U_j])$, where $[U_j] = U_j(F)\backslash U_j(\mathbf{A}_F)$. The probability quotient measure is dv_j/V_j.

For a core cusp section F and an embedding σ, let F_σ be the scalar holomorphic section in the framed dictionary of the core cusp space, and $H_\sigma(g)$ for its normalized adelic reading at the chosen unit lowest-type tensor. At an orthonormal frame the modulus of $H_\sigma(g)$ is the point norm of F_σ multiplied by the positive Lie factor $L_{j,\sigma}$. The component and determinant factors in the reading are units of modulus one. At an embedding over the opposite inclusion of the host quadratic field take the conjugate of the reading at the conjugate embedding. Thus in both cases

$$\|F_\sigma\|_{\mathrm{prob}}^2 = \int_{[U_j]} |H_\sigma(g)|^2 \, \frac{dv_j(g)}{V_j}.$$

Both this norm and the supremum of the point norm used below include $L_{j,\sigma}$. Probability norm and the Lie factor · The two embeddings in the same row model.

The Schur form on the multiplicity space is the integral of $|H_\sigma|^2$ in the row measure, since the evaluating tensor has norm one. The row metric multiplies this Schur form by the product of the lowest compact-type dimensions. At τ its factor is

$$d_{j,\tau} = \prod_{\substack{a<b \\ a,b \text{ in the same signature block}}} \frac{\kappa_{\tau,a} - \kappa_{\tau,b} + b - a}{b - a}$$

in the actual ordered lowest weights κ_τ. The two blocks have sizes p_j, q_j, with an empty product for a zero block. Relabelling the embeddings and taking the contragredient within a block preserve the product. It follows that the same ratio at every used embedding is

$$\|\iota_j F\|_{\mathrm{row},\sigma}^2 = R_j^{\mathrm{met}} \, \|F_\sigma\|_{\mathrm{prob}}^2, \qquad R_j^{\mathrm{met}} = V_j \prod_\tau d_{j,\tau}. \tag{115}$$

A finer marked level pulls back the probability norm isometrically, since the unnormalized volume and the integral of the pulled-back function acquire the same index. The character-level factor in the subsequent coordinate lattice comes from the separately normalized finite level inclusion. The row Schur form and the dimension product · Evaluation of the Schur form · The ratio at each embedding.

Here is a mass computation in the row measure. For one real factor, the invariant density in the independent imaginary diagonal and complex strict lower entries of the Lie algebra is one. Indeed $\mathfrak{u}(p_j, q_j)$ and $i\mathfrak{u}(p_j, q_j)$ are orthogonal for the real Frobenius pairing. Each strict lower entry has squared Frobenius length twice its squared absolute value, canceled in the volume by the factor $2^{-\binom{r}{2}}$ of the row density. The compact dual has the same density. The successive first-column fibrations of its unitary groups send these coordinates to the ordinary sphere coordinates and to the lower-rank coordinates in the fibre. Therefore its volume in this density is $C_r = \prod_{s=1}^{r} 2\pi^s/(s-1)!$. The real density and finite normalization · Compact-dual scale.

Let $\mathfrak{M}_r$ be the mass of the standard positive definite self-dual genus of rank r. The auxiliary row and this definite space have isometric integral Hermitian lattices at every finite place. Indeed E/F is split or unramified there. In the split case take a basis and its dual on the two halves. In the inert case a non-degenerate Hermitian form over the residue field has a vector of nonzero norm, by surjectivity of the residue trace. Lift the vector, scale by an unramified unit norm, and split off its unimodular orthogonal complement. Iteration gives the standard lattice. Hermitian self-duality also implies maximality for integral norms here. If a larger $\mathcal{O}_E$-lattice has integral norms, polarization with all scalar multiples of one argument makes each cross pairing lie in the trace dual of $\mathcal{O}_E$ over $\mathcal{O}_F$. This is $\mathcal{O}_E$ at the unramified places, including residue characteristic two. The larger lattice would therefore lie in the Hermitian dual of the given one. For the trace-self-dual polarization model first divide its lattice by the fixed fractional ideal from the polarization. The resulting Hermitian lattice is self-dual, as in the chart calculation below.

For these two inner forms the Tamagawa number is 2. Their invariant top forms correspond under inner conjugation, which has Lie determinant one. The finite measures on the isometric stabilizers agree and the real measures correspond by the compact-dual density just computed. The convergence factors agree since the two forms have the same character motive. The definite double-coset formula gives mass $\mathfrak{M}_r$ when the compact duals and full stabilizers have volume one. Rescaling to the row density at infinity and to principal volume one at the finite places gives

$$V_j = C_r^f \mathfrak{M}_r [K_j^{\max} : K_j(M)].$$

The index is at most M^{2fr^2}, by reduction in the actual split principal charts. This argument also applies at the terminal stage, where the real group is the compact dual itself. [63, (2.2)–(2.6) and §4] Index at the actual principal level.

The orthonormal genus class gives $\mathfrak{M}_r \geq (|\mu(E)|^r r!)^{-1}$. Indeed an integral unitary matrix for the standard positive form has one nonzero entry in each column, since each nonzero algebraic integer has total squared absolute values at least $2f$ over the $2f$ embeddings. Their absolute values at all embeddings are one by the equality in the arithmetic–geometric mean inequality. They are roots of unity by Kronecker's theorem, and $\varphi(u)^2 \geq u/2$ gives $|\mu(E)| \leq 8f^2$. For the upper bound, the self-dual mass formula is $2\,2^{-rf} \prod_{s=1}^{r} |L_F(1-s, \eta_{E/F}^s)|$. The functional equation bounds its logarithm above by

$$O\big(fr^2 \log(r+1) + r^2 \log|\Delta_F|\big) + \sum_{s=1}^{r} \log L_F(s, \eta_{E/F}^s).$$

For $s \geq 2$ use $L_F(s, \eta_{E/F}^s) \leq \zeta(2)^f$. At $s = 1$, the analytic class number formula for ζ_E/ζ_F gives $L_F(1, \eta_{E/F}) \leq \exp(O(f))h_E$. Here $h_F \geq 1$, and $R_E/R_F \leq 2^{f-1}$ for the regulators, by using the units of F as a finite-index subgroup of the free units of E and doubling their archimedean logarithms. Minkowski's ideal-class bound and $\sum_{\mathrm{N}\mathfrak{a} \leq X} 1 \leq X^2 \zeta_E(2)$ give $\log h_E = O(f + \log|\Delta_E|)$. [63, (4.2)–(4.3) and Propositions 4.4–4.5]

In this family $\log|\Delta_E| = O(fL)$, $\log M = O(L)$, and $\log(fr) = O(L)$. The actual shifted weights have total absolute size $O(frLb)$, and the theta exponent is at most a fixed power of x. Thus the logarithm of the largest weight spread plus one is $O(L)$. Each lowest compact dimension is at least 1, and the displayed Weyl product gives $\log \prod_\tau d_{j,\tau} = O(fr^2 L)$. The two mass bounds and the formula for C_r now give

$$|\log V_j| + |\log R_j^{\mathrm{met}}| = O(fr^2 L).$$

Since f, r, L are bounded by fixed powers of x and $b = \exp(x^{\varepsilon})$, this is $O(E_{\mathrm{cov}})$. The comparison of squared normalized degrees uses precisely this logarithm. The shifted weights in the dimension estimate · Bounds on the actual weights and level.

Supremum and rank on the principal cusp space. Write $W = 1 + \sum_{\tau}(\max_a \lambda_{\tau,a} - \min_a \lambda_{\tau,a})$ for the spread bound of the row flag weight, including its theta shift. The same size bounds give $\log W = O(L)$ at each embedding. The component of a finite class representative h is $\Gamma_h \backslash U_j(F_\infty)/T_\infty$. Here T_∞ is the diagonal compact torus and Γ_h is the image of $U_j(F) \cap hK_j(M)h^{-1}$ under the dictionary identification. The Hermitian matrix in the flag model is the negative transpose of the row matrix. Its orthonormal frames have Gram $J = J_{q_j,p_j}$, with the positive entries first. On the full flag the coefficient is the scalar section in the highest-weight line. The diagonal and finite level characters have unit modulus, so evaluation in its unit fibre gives the absolute lowest-type value.

The geometric principal cusp space at this weight consists of the algebraic sections whose readings at the good embeddings lie in the span of tuples of framed cusp sections. In each component such a section is holomorphic with bounded point norm and zero central Fourier–Jacobi coefficient along every nonzero integral isotropic line and at every frame. Both conditions are preserved by finite linear combinations. In particular the supremum over the finitely many components is finite, and the zero coefficient condition holds on the entire scalar space used below. The negated transpose in the flag model · The actual weight and principal cusp space · The framed cusp sections · The central coefficient and cusp condition.

Choose a point of more than half the supremum, and take an orthonormal frame x_0 there. In its component the period lattice Λ_h is self-dual for $\mathrm{Tr}_{E/\mathbb{Q}}(-\xi_0 H)$. With $(\xi_0)\mathfrak{a}\bar{\mathfrak{a}}\mathfrak{D}_E = (1)$ from the host polarization, $L'_h = \mathfrak{a}^{-1}\Lambda_h$ is Hermitian self-dual. In fact trace self-duality gives $\Lambda_h^{\sharp} = (\xi_0)\mathfrak{D}_E\Lambda_h$, and $(L'_h)^{\sharp} = \bar{\mathfrak{a}}\Lambda_h^{\sharp} = L'_h$. Both lattices are preserved by Γ_h. Let $u_\tau(v)$ be the squared Euclidean length of v_τ in the frame x_0, the Rosati majorant. Then every nonzero $v \in L'_h$ satisfies

$$\sum_{\tau} u_\tau(v) \geq e_*, \qquad e_* = \min\left(1, \frac{\pi}{2M\sqrt{d}\,(fW+2)\log 2}\right).$$

To verify this, suppose all $u_\tau(v) < e_*$. The integral value $H(v,v)$ has absolute value less than one at every place of F_∞, so it is zero. In particular this case cannot occur on the definite terminal group.

For such an isotropic line use the rank-one holomorphic transvection in each frame. Its nilpotent part is $-ik_0\, v_\tau v_\tau^* J$, with $k_0 = 2|\tau(\xi_0)|$ and the sign selected by the holomorphic polarization. It has square zero. The periods in its real variables include $M'\mathcal{O}_F$, where $M' = M\,\mathrm{N}\mathfrak{c}_K\, d$ and $M'k_0 = 2M\sqrt{d}$. Indeed $2\xi_0\,\mathrm{N}\mathfrak{c}_K\, d$ is integral and purely imaginary, and $H(v, L'_h)$ is integral. The resulting rational transvection for a real period is congruent to one modulo M on L'_h and Λ_h. Its determinant is one, so the character on the principal subgroup is trivial.

For complex parameter z_τ, the s-th absolute leading Gram minor is multiplied by

$$1 + 2k_0 \operatorname{Im}(z_\tau) \sum_{a \leq s} J_{aa}|v_{\tau,a}|^2.$$

This follows from the rank-one determinant identity. The partial sums increase in the positive block and then decrease to zero, with maximum $u_\tau(v)/2$. For diagonal $z_\tau = z$ with $\operatorname{Im}(z) = -t$, $t = (2k_0 e_*)^{-1}$, the frame remains in the domain and all these factors lie in $[1/2, 1]$. Subtract the smallest weight in each factor when computing the line metric, since the full determinant is unchanged. The sum of the positive logarithmic increments of the leading factors is at most $\log 2$, and the sum of the negative ones has its opposite value. Thus the ratio of the metric factors is at most $2^{W/2}$.

The holomorphic scalar on the transvection tube has a Fourier expansion for the full real period lattice. The bounded point norm gives at most polynomial growth of the scalar as the imaginary parameters increase separately, since all the Gram factors are linear and positive on that tube. The corresponding Fourier frequencies therefore have nonnegative components. The coefficient of frequency zero is the

central average. To apply the defining cusp condition choose an integer $a > 0$ with $av \in \Lambda_h$. The transvection with generator av is that with generator v at a^2 times its argument. The scaled period lattice and $M'\mathcal{O}_F$ are commensurable. Their normalized averages agree with the average over a common finite-index sublattice by partitioning a fundamental domain into translates. The coefficient is thus zero by the defining condition on av. Smoothness on the real period torus gives absolute summability there, and hence convergence to zero as all imaginary parameters rise together. On the diagonal the function is M'-periodic and vanishes at the origin of the punctured-disk coordinate $\exp(2\pi i z/M')$. Its modulus on the horizontal line at $-t$ is at most $2^{1+W/2}$ times its value in modulus at 0, by the half-supremum choice and the metric ratio. The disk Schwarz bound would give $1 \le 2^{1+W/2}\exp(-2\pi t/M') < 1$, since $2\pi t/M' \ge (fW+2)\log 2$. This proves the required lower bound for the majorant length. The actual self-dual lattice and transvection period · The point metric and cuspidal coefficient · The bound at a near-maximum frame.

Use the holomorphic chart $x_0(1+Z)$, with Z strictly lower triangular at each place. It has complex dimension $D_{\mathrm{fl}} = f\binom{r}{2}$. Take the Euclidean ball of radius $\rho = (100r^2(W+f+1))^{-1}$. In one factor, the s-th leading Gram block is

$$(1+Z_{\le s,\le s})^* J_{\le s}(1+Z_{\le s,\le s}) + Z^*_{>s,\le s} J_{>s} Z_{>s,\le s}.$$

The first term has absolute determinant one, and its inverse has operator norm less than 2 on this ball. The second term has operator norm at most ρ^2. Thus the signs of the leading minors are fixed and $|\log|\Delta_{\tau,s}|| \le 4s\rho^2$. The full determinant has absolute value one. Telescoping the point metric in the leading minors bounds its squared ratio to the metric at x_0 between $1/2$ and 2, using $\sum_\tau \mathrm{spread}(\lambda_\tau) \le W$.

The invariant density in these lower-triangular coordinates is $\prod_\tau \prod_{s<r} |\Delta_{\tau,s}|^{-2}$ times its density at zero. Indeed the upper triangular factor in a Gauss decomposition acts on the lower-triangular tangent quotient with complex determinant $\prod_{a>b} b_{aa}/b_{bb}$. The leading Gram minor identity for an isometry gives $\prod_{a\le s} |b_{aa}|^{-2} = |\Delta_{\tau,s}|$. Taking the inverse real Jacobian gives the asserted density. The bound on the leading minors puts this density ratio between $1/2$ and 2. The independent Lie coordinates of the row measure have density one, so its quotient by the compact torus has density $(2\pi)^{rf}$ at zero when the torus is pushed forward with its Haar mass. The finite level has volume one. These density statements also hold for the compact flag at the terminal stage.

The multiplicity with which the chart covers the component quotient is bounded by the number of $\gamma \in \Gamma_h$ returning the ball to itself. If γ sends a flag in this ball to another one in the ball, it identifies their positive planes and hence their Rosati majorants. Those majorants are between $\frac{1}{2}$ and 2 times that at x_0. For example, the positive plane is the graph of $Z_{>q_j,\le q_j}(1+Z_{\le q_j,\le q_j})^{-1}$, of norm at most 2ρ. The Hermitian orthogonal projector onto that graph is obtained by inverting $1-A^*A$, for this graph map A. Substitution in $2H(\mathrm{pr}_+ v, \mathrm{pr}_+ v) - H(v,v)$ gives the claimed comparability. Thus each such γ has operator norm at most 2 in the majorant at x_0.

In this majorant L'_h is a Euclidean lattice of real dimension $n_L = 2fr$ and covolume $2^{-fr}|\Delta_E|^{r/2}$. In a pseudo-basis the factor multiplying this covolume is the product of the ideal norms times the absolute norm of the Hermitian determinant, which is one by self-duality. Minkowski's second theorem, with first minimum at least $\sqrt{e_*}$, bounds its last minimum β by

$$\beta \le \frac{2^{n_L}\,\mathrm{covol}(L'_h)}{\mathrm{vol}(B_1^{n_L})\, e_*^{(n_L-1)/2}}.$$

Choose n_L independent lattice vectors of length at most β, which determine γ by their images. All images lie in the lattice and have length at most 2β. Disjoint balls of radius $\sqrt{e_*}/2$ give at most $(1+4\beta/\sqrt{e_*})^{n_L}$ possible images of each one. The number of returns is bounded by $N_{\mathrm{ret}} = (1+4\beta/\sqrt{e_*})^{n_L^2}$. It follows that $\log N_{\mathrm{ret}} = O((fr)^3 L)$.

The holomorphic submean inequality gives the value squared at the chart centre bounded by $D_{\mathrm{fl}}!/(\pi^{D_{\mathrm{fl}}}\rho^{2D_{\mathrm{fl}}})$ times the Euclidean integral of the scalar square. The two metric and density comparisons and the return bound convert this to the probability quotient integral. The return bound accounts also

for finite stabilizers, by summing the translates of the chart over a fundamental set. Incorporating the half-supremum factor and the common Lie factor gives

$$\|F_\sigma\|_\infty^2 \le C_{\mathrm{sup}}\,\|F_\sigma\|_{\mathrm{prob}}^2, \qquad \log C_{\mathrm{sup}} = O((frL)^3). \tag{116}$$

For instance one can take the maximum of 1 and $16V_jN_{\mathrm{ret}}D_{\mathrm{fl}}!/((2\pi)^{rf}\pi^{D_{\mathrm{fl}}}\rho^{2D_{\mathrm{fl}}})$, using the upper bound for β above. Here $\log(1/\rho)$ and $\log(1/e_*)$ are $O(L)$. The estimate is for the full principal cusp space at the actual stage weight, before either finite level cut. The proven inequality at a half-maximal point of each form gives a uniform inequality at every point. It passes to the conjugate embeddings by the same norm identification as in (115). Submean in the flag chart · The return count · Direction of the squared estimate.

For a probability-orthonormal family $f_1, \dots, f_a$ at a fixed embedding, evaluation in a unit fibre frame has squared operator norm $\sum_{l=1}^{a} |f_l(z)|^2$. By (116) this is at most C_{sup} everywhere. Its probability integral is a, so $a \le C_{\mathrm{sup}}$. The algebraic cusp space over F_0 injects after scalar extension into this space by the framed realization. For any of its integral sublattices of positive rank, the $\mathbb{Z}$-rank is therefore at most $[F_0 : \mathbb{Q}]C_{\mathrm{sup}}$. The selected core field satisfies $\log[F_0 : \mathbb{Q}] + \log \mathrm{rd}(F_0) = O((frL)^2)$. This bounds the logarithmic rank by $O((frL)^3)$. In particular the lower estimate for the mean log supremum gives a lower estimate for the mean log Petersson norm with cost at most $\frac{1}{2}\log C_{\mathrm{sup}}$. This is the direction needed before applying the lattice argument to the exceptional cut. Size of the selected core field · Rank at the actual scalar weight.

Taking determinants in (115) gives $\mu(J_{\mathrm{type}}) \ge \mu_{\mathrm{core}} - O(E_{\mathrm{cov}})$. The same comparison applies at the terminal stage, where the core lower bound comes from the definite fibres.

An integral CM frame detecting a cusp section. Use the principal-level cover X_o and the coefficient field K_0 of the given stage. Put $g = fr$, and for a weight w write n_θ for its theta exponent, $s = \sum_{\tau,i} a_{\tau,i}$ for the sum of its nonnegative tangent weights, and

$$W_0 = 1 + s + gn_\theta, \qquad Y = fr + L + \log|\operatorname{disc} E| + |\log \Omega|.$$

For all weights of this stage, with the two nonnegative tangent tableaux $a^+_{\tau,h}$, $a^-_{\tau,h}$, use the weighted degree

$$\mathrm{S}(w) = 1 + \sum_\tau \left(\sum_h a^+_{\tau,h} + \sum_h a^-_{\tau,h} \right) + fr\max(1, q_j)n_\theta. \tag{117}$$

Here $q_j = m - j$ counts the negative directions in the auxiliary row signature, and n_θ is the theta weight. Thus $\mathrm{S}(w) \ge W_0$. The weighted degree of the section.

Let S_σ be the supremum of the point norm at the embedding $\sigma : K_0 \hookrightarrow \mathbf{C}$, including the Lie factor in the probability norm. For a nonzero section P of the geometric integral cusp lattice, we prove

$$\overline{\log S_\sigma(P)} \ge -C_* Y^2\,\mathrm{S}(w)(1 + \log \mathrm{S}(w)). \tag{118}$$

Here and below a bar in a norm estimate is the mean over embeddings of its coefficient field. The constant C_* is absolute on the selected host range.

At the distinguished embedding choose an orthonormal adapted frame x_0 at more than half the supremum. Use the actual self-dual Hermitian lattice L' of that component and its majorant Q at x_0, as in the principal-level minimum calculation. In that calculation the sum of the flag weight spreads is at most $2W_0$, because the flag weights $\lambda_{\tau,i}$ are the signed tangent weights with shifts $\pm n_\theta/2$. It gives a lower bound e_* for $Q(v)$ for $0 \ne v \in L'$, with

$$\log e_*^{-1} = O(1 + \log M + \log d + \log f + \log W_0), \qquad M = g_0 g_1.$$

The same lower bound applies to a definite component by integrality of $H(v,v)$. The real dimension is $d_0 = 2fr$ and the Euclidean covolume is $2^{-fr}|\operatorname{disc} E|^{r/2}$. Indeed the trace pairing has Gram determinant

of absolute value $|\operatorname{disc} E|^r$ by Hermitian self-duality. Passing to the positive majorant preserves the determinant in absolute value and using one embedding from each conjugate pair divides the trace metric by two. Minkowski's second theorem therefore bounds the last successive minimum $\lambda_{d_0}(L')$ by

$$\log(1+\lambda_{d_0}(L')) \ll r\log|\operatorname{disc} E| + fr\big(\log(2fr)+\log e_*^{-1}\big).$$

Choose d_0 independent lattice vectors of lengths at most twice this last minimum. Rounding the coordinates in their real basis approximates any vector to distance at most $d_0\lambda_{d_0}(L')$.

Use $\rho = (1000rW_0)^{-1}$. A common integer $A \le 2 + 2rd_0(1+\lambda_{d_0}(L'))/\rho$ then gives columns $u_i \in L'$ with $\|u_i/A - (x_0)_i\|_Q \le \rho/r$. At each archimedean factor the resulting full frame is $x_0(1+E_\tau)$ with $\|E_\tau\|_{\mathrm{op}} \le \rho$. On the holomorphic line of frames $x(t)_\tau = (x_0)_\tau(1+tE_\tau)$, $|t| \le 5$, every leading Gram matrix changes by at most $2|t|\rho + |t|^2\rho^2$ in operator norm. Each retains its signature. Its log absolute determinant changes by at most $2r(2|t|\rho + |t|^2\rho^2)$. The flag metric is the product of consecutive leading determinant ratios to powers $-\lambda_{\tau,i}/2$. Telescoping its logarithm, using $\sum_{\tau,i}|\lambda_{\tau,i}| \le W_0$, bounds its inverse ratio by $\exp(2rW_0(2|t|\rho + |t|^2\rho^2)) \le e^{1/8}$. Thus the scalar reading has modulus at most $2e^{1/8}|P(x_0)|$ on the disk $|t-1| \le 4$. If the rational frame at $t=1$ were a zero, Schwarz's lemma would bound its value at $t=0$ by $2e^{1/8}|P(x_0)|/4$, a contradiction.

Obtain integral orthogonal columns b_i by multiplying each Gram–Schmidt step by the leading Gram determinant. Explicitly, if G_{i-1} is the Gram of $u_1,\dots,u_{i-1}$ and $c_i = (H(u_\ell,u_i))_{\ell<i}$, with H linear in the second slot, take

$$b_i = \det(G_{i-1})u_i - [u_1,\dots,u_{i-1}]\operatorname{adj}(G_{i-1})c_i, \qquad b_1 = u_1.$$

Integrality of H gives $b_i \in L'$. This triangular change preserves the full flag and its nonzero reading. At every factor $\|u_i\| \le 2A$, so Hadamard on the determinants and adjugates gives, with $B \ge 1$,

$$|\tau H(b_i,b_i)| \le B^2, \qquad \log B \ll r\log(2rA) \ll Y^2(1+\log W_0).$$

For the last bound use $\log M + \log d + \log(2fr) + \log \operatorname{rd}(E) \ll f$ in the selected host range and $\log|\operatorname{disc} E| = 2f\log\operatorname{rd}(E)$.

The period lattice of the detected point is $\mathfrak{A}L'$, where $\mathfrak{A} = I$ in the classical host. In each column the map $\alpha \mapsto \alpha b_i$ from $\mathfrak{A}$ gives the corresponding period direction. The period reading uses coordinates in the negative block and Hermitian pairing in the positive block, so these images are the tangent slots for the two CM types extending the two embeddings of K_d. Give each copy of $\mathfrak{A}$ the corresponding trace form with $\pm\xi_0$. It is principally polarized by the host ideal identity. The direct sum of the period maps gives an isogeny to the detected abelian variety, with finite kernel since the b_i form an E-basis. The rank-one CM model is $\mathfrak{A}\otimes_{\mathcal{O}_{K_d}} A_0$, with A_0 the elliptic CM curve of the indicated type. This tensor construction is the summand of a power of A_0 associated with a projective module presentation of $\mathfrak{A}$. Over a common finite extension K/K_0 the elliptic curves have everywhere good reduction and the isogeny kernel and the actual level are algebraic. Quotient the product good model by the finite flat closure of the kernel in its torsion, obtaining the good model $\mathcal{A}$ of the same marked point. The action and the principal polarization extend. Its rigidified symmetric principal line extends uniquely: extend a divisor on the regular model, rigidify along the identity and extend the generic cubical isomorphism, whose vertical divisor is zero by restricting to the identity on the products. Its square has the finite flat polarization kernel $\mathcal{A}[2]$. The two period characters in the adapted frame.

Write T for its integral theta line with the Gaussian L^2 norm for probability Haar, and give Lie the π-scaled polarization metric. Let ℓ_j, $1 \le j \le g$, be all the saturated tangent axes in the rank-one reference varieties. The discriminant-normalized elliptic differential generates the good cotangent line at every finite place, since its discriminant is a unit and scales with the inverse twelfth power. The Lie module of the tensor model is $\mathfrak{A}\otimes_{\mathcal{O}_{K_d}} \operatorname{Lie} A_0$. Its axis ideals and squared norms are therefore exactly those in the CM period computation. In ordinary normalized degree $\deg_n \ell_j = \frac{1}{2}\log\Omega$, also for the conjugate type. Pulling back the polarization along the i-th isogeny multiplies the reference Hermitian form by the totally

positive element $\pm H(b_i, b_i)$. The lattice identity implies this identity of morphisms and thus the bound B on the tangent maps at every embedding of K. The integral maps on Lie give, by Hadamard,

$$\deg_n \det \operatorname{Lie} \mathcal{A} \ \geq \sum_j \deg_n \ell_j - g \log B.$$

In Bost's key formula the degree of the dual theta line is $\frac{1}{2} h_F + \frac{g}{2} \log(\pi \sqrt{2})$ in the Faltings convention of [130, (12.1)–(12.2) and Theorem 13.1]. For the principally polarized form the cotangent integral metric there has squared norm π^g times the squared dual of the π-scaled Lie metric. Hence

$$\deg_n T^\vee = -\tfrac{1}{2} \deg_n \det \operatorname{Lie} \mathcal{A} + \tfrac{g}{4} \log(2\pi).$$

The good model and cubical line above give the hypotheses of the formula for this theta line. [123, Theorem 4.10(v)] The Gaussian and Lie metrics and integral traits.

Take rational generators t of T and v_j of ℓ_j. Evaluate the section in matrices of Lie frame changes about the images of the v_j, keeping the generator t. It is a polynomial in the tangent matrices of total degree s, by the nonnegative highest weights of the polynomial GL representations. Its dependence on the theta generator is homogeneous of degree n_θ. It is nonzero because the detected frame lies above the same abelian point.

At a finite place start with an integral theta generator and bases in the saturated eigenmodules of Lie $\mathcal{A}$. These are the integral framed tests of the geometric cusp lattice on this generic marked point. A coefficient of negative valuation would, after scaling, give a nonzero polynomial modulo the place. It is nonzero on some tuple of invertible matrices over a finite residue extension, since its product with the determinants is nonzero. Lift that tuple to an algebraic framed trait above the place. The evaluation then has negative valuation, a contradiction. This uses the generic level on the trait, as in the integral evaluation rule. Substituting the integral images of the local generators of the ℓ_j preserves coefficient integrality. Consequently each coefficient c_μ of the original polynomial represents an integral section of

$$L_\mu = (T^\vee)^{\otimes n_\theta} \otimes \bigotimes_j (\ell_j^\vee)^{\otimes \nu_j}, \qquad \sum_j \nu_j = s,$$

where ν_j is the degree in the entries of the j-th reference axis.

At an embedding σ of K, the value on an orthonormal Lie frame is at most $\|t\|_\sigma^{n_\theta} S_{\sigma|K_0}(P)$. The norm of the uniformizing theta generator is one and its orthogonal Lie frame has squared lengths $2\pi|\xi_0|$, so this is the point norm with the same Lie factor as above. By singular value decomposition and the maximum principle in each singular value, the same bound holds for contractions in the Lie blocks. In each entry of axis j take the polydisc of radius $(r^2 B \|v_j\|_\sigma)^{-1}$, whose matrices give contractions. Cauchy's estimate gives

$$\|c_\mu\|_{L_\mu,\sigma} \leq (r^2 B)^s S_{\sigma|K_0}(P).$$

Apply the product formula to any nonzero integral coefficient. Its mean log norm is at least $-\deg_n L_\mu$. The preceding degree and B bounds give (118), since $W_0 \leq \mathrm{S}(w)$. The mean over K agrees with the mean over K_0 for the supremum values, since every embedding of K_0 has the same number of extensions. Use of the geometric evaluation lattice · The scalar point-norm comparison.

For the exceptional cut, apply (118) to the actual principal-level lattice and (116) to its probability norms. The chosen stages have weights $\mathrm{S}(w) \leq x^{A_1} b$ and $Y + \log \mathrm{S}(w) = x^{O(1)}$. Thus every nonzero vector of the principal-level cusp lattice has mean log Petersson norm at least $-x^A b$ for a fixed A.

Applying Minkowski to any sublattice in the Euclidean sum of the archimedean spaces bounds its squared slope above by $2x^A b + O(\log n + \log \mathrm{rd}(F_0))$, where n is its integral rank. Indeed the geometric mean bound gives a lower bound for the length of every nonzero vector, and the volume of the Euclidean ball of that radius is a lower bound for covolume up to 2^n. The rank bound from (116) gives $\log n = O((frL)^3)$. Normalized Hermitian base change preserves maximal slope. Over a Galois extension the

maximal destabilizing subspace of largest rank is unique and invariant, so it descends. Flatness identifies its saturated lattice with the extension of the saturated lattice downstairs, and extension preserves its normalized degree. For a general finite extension take a Galois closure and use the extension of each candidate subspace. The mass ratio gives the same polynomial bound in the ambient row norm, which also bounds $\mu_{\max}(J_{\mathrm{type}})$. The principal-level detection estimate and The polynomial bound at the stage weight.

The good projector is a central spectral idempotent in the finite level algebra. In the finite character formula, multiplication by $g = |G_{\mathrm{lev}}|$ clears the denominator. The character values are integral, and inversion of a group element conjugates each coefficient after every complex embedding. Thus the unitary action makes the projector orthogonal, and the integral lattice is preserved after multiplication by g. Here $2\log g \le b$. Finite character coefficients. Let J_{good} and J_{exc} be its two integral intersections. The complementary projection injects

$$J_{\mathrm{type}}/(J_{\mathrm{good}} \oplus J_{\mathrm{exc}}) \longrightarrow g^{-1}J_{\mathrm{exc}}/J_{\mathrm{exc}}.$$

Its kernel consists of elements whose two projections are already integral. The finite index therefore costs at most $2\,\mathrm{rk}(J_{\mathrm{exc}})\log g$ in squared normalized degree. By orthogonality, subtracting the squared degree of J_{exc} as well gives

$$\mu(J_{\mathrm{good}}) \ge \mu(J_{\mathrm{type}}) - \frac{\mathrm{rk}(J_{\mathrm{exc}})}{\mathrm{rk}(J_{\mathrm{good}})}\big(\mu_{\max}(J_{\mathrm{type}}) + 2\log g - \mu(J_{\mathrm{type}})\big).$$

For rank zero the exceptional contribution is zero. The exceptional rank fraction is $O(e^{-x^c})$ for a fixed $c > 0$, whereas $\mu_{\max}(J_{\mathrm{type}})$, $|M_j|$ and E_{cov} are bounded by fixed powers of x times b. Inserting the lower bound $-M_j - O(E_{\mathrm{cov}})$ for $\mu(J_{\mathrm{type}})$ costs $O(E_{\mathrm{cov}})$. The integral orthogonal projector.

The absolute stage frame in the division estimate. Let $b_{L,j}$ be the rational basis of the primitive good space on the base face of the linked row model. Its size is the positive common row rank d_{row}, by the dimensions of the primitive summands. The coordinate lattice used for division is exactly the pullback of J_{good} under $c \mapsto \sum_k c_k b_{L,j,k}$. For each infinite place, the complex frame of division is the Gram frame of the complexified $b_{L,j}$, with the factor $1/\sqrt{n_P}$ in its vector norm. Here $n_P \ge 1$ is the number of elements in the level group P of this base face. The Gram form is positive definite on the actual good space. Thus

$$\frac{\mathrm{lcv}(J_j^{\mathrm{row}})}{d_{\mathrm{row}}} = -\mu(J_{\mathrm{good}}) - \log n_P \le -\mu(J_{\mathrm{good}}).$$

The finite components in this equality are the same coordinate lattice. Indeed, with normalized Arakelov degree and squared slope, the normalized log covolume of a rank-d lattice is $-2\widehat{\deg}_{\mathrm{n}}$, and multiplying its vector norms by $n_P^{-1/2}$ changes its log covolume per rank by $-\log n_P$. The Gram frame normalization and The actual coordinate lattice and rank.

The ordinary and terminal core estimates, the mass comparison and the exceptional rank and index bounds now give $L_{3,j} = M_j + C_3 E_{\mathrm{cov}}$ in the base face of division. This is the very same stage lattice and complex frame used for the determinant identity at that face. The transfer to the linked bundle takes place at its fixed bundle point and uses its transported primitive row realization, with the usual comparison at the places where the local tuple differs. In particular, with the uniform threshold of the chosen family,

$$\sum_{j=1}^{m} \frac{\mathrm{lcv}(J_j^{\mathrm{row}})}{d_{\mathrm{row}}} \le \sum_{j=1}^{m} M_j + O(mB_0 \log\log x) = \sum_{j=1}^{m} M_j + O\big(m(B_0+N_c)x^{e_D\varepsilon}\big).$$

The last bound uses the chosen $e_D = 2$ and fixed $\varepsilon > 0$. In the row-dual determinant inequality this is the contribution $\sum_j L_{3,j}$. The row metric normalization and the weighted finite Gram and transport terms at

the actual tuples enter through their own parts of the determinant identity. The same linked bundle point and Insertion in the base row-dual inequality. This proves (99) in the chosen frame.

Let $\mathcal{F}$ be the weighted sum of the reference exponents and Q_w the weighted sum for w, using $2e_v\lambda_v$ as above.

The complex rank of the common kernel of the f residue-unit averages was computed in the ramified column count. For the following argument require the integral and metric donor identification at every w. It identifies H_w with the integral sum-zero class lattice in M_w and identifies the actual N_w isometrically with a full sublattice of this model in the column chain's zero-stage metric.

In the identified model, put $q = T_w - 1$. On a residue-unit factor use the basis $e_i - e_q$, $1 \le i < q$, of the integral kernel of summation in the q unit coordinates. Its Gram is $I + \mathbf{1}\mathbf{1}^{\mathsf{T}}$, of determinant q. Compared with its unrestricted residue factor, the mean squared slope of this kernel therefore changes by $-\log q/(q-1)$. On each split class summand the identical coefficient and class metric factors cancel from this difference. The f independent unit directions use the tensor product of these kernels. Tensor determinants add the corrections per rank, and the orthogonal summands have the same retained rank ratio. In the metric of its integral inclusion this gives

$$d_M = \operatorname{rk} M_w, \quad d_r = (1-\theta_w)d_M, \quad \theta_w = 1-(1-1/(T_w-1))^f,$$
$$\mu(H_w) = \mu(M_w) - f\frac{\log(T_w-1)}{T_w-2}, \qquad \mu_{\max}(H_w) \le \mu_{\max}(M_w).$$

The last inequality follows by saturation in M_w.

Let $e_v\ell_{wv}$ be the local length of H_w/N_w. The same donor identification must supply

$$(N_w)_v \subseteq \mathfrak{p}_v^{e_v\tau_v}(H_w)_v, \qquad \ell_{wv} \le d_M(\tau_v + \iota_v) \qquad (v \nmid T_w),$$

where τ_v is a nonnegative integer. At $v \mid T_w$ the conditions are $\tau_v = 0$ and

$$\frac{2}{d_r}\sum_{v\mid T_w} e_v\lambda_v\ell_{wv} \ll N_T m(a_E/k + b')\log T_w.$$

These conditions concern the actual undivided lattice in the zero-stage norms. For a full sublattice the determinant formula subtracts $2e_v\lambda_v\ell_{wv}/d_r$ at v from the mean. The ramified rank, metrics and local lengths.

Set $g'_{wv} = e_v^{-1}\lfloor e_v g_{wv}\rfloor$. The determinant computation after division gives

$$\mu(N'_w) = \mu(H_w) + 2\sum_v e_v\lambda_v(g'_{wv} - \ell_{wv}/d_r)$$

in these common metrics. For a comparison through an isomorphism the difference must be bounded in the prescribed norms, as in the release's mean-slope comparison at the local cutoff. For $v \nmid T_w$,

$$\frac{\ell_{wv}}{d_r} - \tau_v \le \frac{\iota_v}{1-\theta_w} + \frac{\theta_w}{1-\theta_w}\tau_v.$$

Bernoulli's inequality gives $\theta_w \le f/(T_w-1) \le 2f/T_w$. In the level range used here $2f/T_w \le 1/2$, the index allowance therefore contributes at most $4\mathcal{I}(u)$. The remaining rank correction is bounded by $(4f/T_w)\,2\sum_v e_v\lambda_v\tau_v$. The sum $2\sum_v e_v\lambda_v\tau_v$ is at most the sum of the stage gains plus the error in (97), by nonnegativity of the ramps. The gain bound below makes this $O(x^c b)$ for a fixed c, so it is at most b^2 for large x. Since the selected levels satisfy $T_w \ge b^3$ and f grows at most polynomially in x, the rank correction plus $f\log(T_w-1)/(T_w-2)$ is at most 1 beyond a common threshold. Lower level range for this comparison. The rounding cost up to the local comparison cutoff X'' is $O(X'')$. The weight above one rational prime satisfies $\sum_{v\mid p}\lambda_v \le \log p$, and the sum of $\log p$ over $p \le X''$ is $O(X'')$. The remaining support consists of the level primes and the exceptional conductor primes. Each such prime costs at most

$2\log p$. In the selected size range their total logarithmic norm is $O(X'')$ as well. Level-prime rounding budget. Since

$$2\sum_{v} e_v\lambda_v(g_{wv}-\tau_v)=\mathcal{G}(u)+\mathcal{F}-Q_w,$$

these bounds give the lower slope estimate. For the upper slope estimate, intersect the local inclusions with a nonzero rational subspace of N_w. Its lattice is contained in the restriction of H_w multiplied locally by $\mathfrak{p}_v^{e_v\tau_v-\lfloor e_v g_{wv}\rfloor}$. Taking the degree on each such subspace gives an upper bound by $\mu_{\max}(H_w)+2\sum_v e_v\lambda_v(g'_{wv}-\tau_v)$. As $g'_{wv}\le g_{wv}$, this yields the maximal-slope inequality. Application in the prescribed lattices. Local slope inputs and weighted slope deductions. The resulting inequalities in the prescribed metrics are

$$\begin{aligned}\mu(N'_w)&\ge \mu(M_w)+\mathcal{G}(u)+\mathcal{F}-Q_w-R_{\mathrm{loc}}(w),\\ \mu_{\max}(N'_w)&\le \mu_{\max}(M_w)+\mathcal{G}(u)+\mathcal{F}-Q_w+O(mB_0),\end{aligned}\tag{119}$$

with

$$R_{\mathrm{loc}}(w)\ll \mathcal{I}(u)+N_T m(a_E/k+b')\log T_w+X''+1.$$

Here b' is the smaller scale for the weights. The fractional-ideal inclusion from the frozen gain statement supplies the map into C.

4.3.5 Height of the divided map

Use the zero-stage isometry in the ramified donor identification. An upper bound h_w for the divided map's squared height uses the operator norm in those source metrics and in the common target coordinates. If $V_{w,\sigma}>0$ bounds the squared norm at each embedding, then integrality gives

$$h(B_{K,w})\le [K:\mathbf{Q}]^{-1}\sum_{\sigma:K\hookrightarrow\mathbf{C}}\log V_{w,\sigma}.$$

At the base twist the tested period factors through a sequence of column maps $J_{w,l}$, followed by a bilinear evaluation $P_{w,\sigma}$ on the resulting column and a top row before quotienting. Give the top row quotient its quotient norm. The evaluation vanishes on the kernel of the row projection. Taking the infimum of the norms of lifts of a quotient vector bounds the operator norm of the descended evaluation by the original operator norm. The actual row identification satisfies an inner-product identity on the whole quotient. The pairing of the image of a vector in block α with the image of any quotient vector is its quotient pairing multiplied by $\kappa_{R,\sigma}\prod_j\lambda_{j,\alpha,\sigma}>0$. Hermitian symmetry makes the groups with different scalars orthogonal. On each group with the same scalar the whole form is rescaled. Thus the inverse identification has squared norm equal to the reciprocal of the least factor. Submultiplicativity gives

$$\|B_{K,w,\sigma}\|^2\le\frac{\|P_{w,\sigma}\|^2\prod_l\|J_{w,l,\sigma}\|^2}{\kappa_{R,\sigma}\min_\alpha\prod_j\lambda_{j,\alpha,\sigma}}.$$

If $H_{P,w}$ bounds the log squared norm of the evaluation at every embedding and $H_{w,l,\sigma}$ bounds that of the l-th column map, the height is at most

$$H_{P,w}+\sum_l\overline{H_{w,l,\sigma}}-\overline{\min_\alpha\sum_j\log\lambda_{j,\alpha,\sigma}}-J_\kappa,\qquad J_\kappa=\overline{\log\kappa_{R,\sigma}},$$

where the bar averages over all coefficient embeddings. Chain and quotient norm bound. In the specialized height the word term $\mathrm{SZbP}(w)$ and the kernel place means bound the log squared norm before the inverse row identification. If $L_{6,j}$ bounds each stage log Petersson eigenvalue below, the height expression is

$$h_w=\mathrm{SZbP}(w)+\sum_j\mathrm{kernelPlaceMeanBound}_j(w)-\sum_j L_{6,j}-J_\kappa.$$

The rational comparison must carry the actual lattices and metrics, so this bounds the height of the map in the rank argument. Fractional-ideal division keeps the archimedean donor metric. Metric comparison on the types and Height in the linked family. We now bound both norms on the same class operator, before applying the inverse row identification. Denote the column ranks by $N(j)$ and their largest value by n_*. Set $S_a = \sum_{t\in T} a_t$.

The terminal pairing in the fixed norms. Write $N = n_*$ for the top column rank, so the row coefficient representation has rank parameter $N-1$. At a CM place ρ write $\lambda = \lambda_\rho$ and $\alpha = \sigma_\rho$ for the terminal column and row weights. Number the entries from zero. Interlacing gives $\lambda_i \ge \alpha_i \ge \lambda_{i+1}$. The determinant shifts agree. Put

$$A_i = \alpha_i - \lambda_{i+1}, \qquad B_i = \lambda_i - \alpha_i, \qquad d_0 = \sum_{i<N-1} A_i, \qquad d_1 = \sum_{i<N-1} \min(\lambda_i - \lambda_{i+1}, d_0).$$

Thus $d = |\lambda| - |\alpha| = \lambda_0 - d_0$, where bars on a weight denote the sum of its entries. Factoring out the common determinant shift, use the usual polynomial models on the first $N-1$ columns. Write Δ_i for the minor on rows and columns $0,\dots,i$, and M_i for the same minor with row i replaced by the last row. The branching highest vector is

$$\xi = \prod_{i<N-1} \Delta_i^{A_i} M_i^{B_i}.$$

It has the column weight of the upper highest vector h_λ and lower row weight α. Both assertions follow by summing the row and column degrees of the minors. The upper triangular operators for the lower group fix it, since adding an earlier row to a later selected row leaves each minor unchanged. Extraction of the monomial $\prod_i x_{N-1,i}^{B_i}$ in the last row maps it to h_α. Indeed, first take all required factors in the last eligible column, and continue in decreasing column order. The resulting powers are $\Delta_i^{A_i+B_{i+1}}$, with zero exponents beyond the range.

In the coefficient norms each Fischer form is divided by the squared Fischer norm of its highest vector. Let $k_\rho = \|\xi\|_{\mathrm{F}}^2 / \|h_\lambda\|_{\mathrm{F}}^2$ and

$$c_\rho = \prod_{i<N-1} \prod_{\substack{\lambda_{i+1}\le u\le \min(\lambda_i,\lambda_{i+1}+d_0)\\ u\ne\alpha_i}} \prod_{l<N-1} \big((u-i)-(\alpha_l-l)\big).$$

This is the coefficient of the prescribed branching operator, which is $c_\rho k_\rho$ times last-row coefficient extraction followed by Fischer projection to the lower summand. The projection is orthogonal in the chosen Fischer form. Extraction commutes with the lower row action, and the irreducible lower type occurs only once. This multiplicity assertion follows by restricting the Schur character to $\mathrm{diag}(x_1,\dots,x_{N-1},z)$: the boxes labelled N in a semistandard tableau form the horizontal strip λ/α. Consequently the adjoint of the branching operator sends h_α to $c_\rho\xi$. The operator times its adjoint is $c_\rho^2 k_\rho$ on the whole lower space. These identities persist under an orthogonal change of frame in the lower space. Scaling the length squared g of the added orthogonal coordinate scales the upper branching metric, relative to the lower metric, by g^d. Including the scalar t^{-d} in the chosen branching operator gives the squared norm $(|t|^2 g)^{-d} c_\rho^2 k_\rho$.

At every embedding the sum of frames used in the terminal class pairing has $|t|^2 g = T_w$. In the unitary coefficient coordinates this reads $g = 1$, with $t = \rho(\overline{\pi_{T_w}})$ and $\pi_{T_w}\overline{\pi_{T_w}} = T_w$. This last identity is algebraic and holds for the conjugate coefficients as well. The finite pairing is the sum of the coefficient pairings over the actual terminal classes with the inverse stabilizer weights of the row model. These are the same weights in the two input squared norms. Thus weighted Cauchy–Schwarz introduces no class-count factor in the terminal pairing. The trivial-tuple row map has the fixed row norm by the class Gram identity. If V is the matrix of its good row basis in the principal-level class model and H the principal class Gram matrix, the row Gram is $|P|^{-1}V^*HV$ after restriction to level P. Both the class mass and the row frame use this same factor. The identity of the evaluated class vectors follows from the good-part projector fixing

the image of V. With the source terminal class norm and the row norm before the quotient comparison, this gives

$$\log \|P_{w,\varsigma}\|^2 \le -\log T_w \sum_{\rho} (\lambda_{\rho,0} - d_{0,\rho}) + \sum_{\rho} \log(c_\rho^2 k_\rho).$$

The product in c_ρ has precisely $(N-1)d_{1,\rho}$ factors, since α_i belongs to the indicated interval. For $l \neq i$ the shifted intervals for the two rows are strictly ordered, and for $l = i$ the erased value is the sole possible zero. The magnitude of each factor is bounded by $\lambda_0 - \lambda_{N-1} + N \le 2a_E$. For the Fischer ratio, use a standard complex Gaussian N-by-$(N-1)$ matrix and its QR decomposition UR, where the unitary frame is Haar and independent of R. Every minor on the first $i+1$ columns has the same factor $\prod_{l \le i} R_{ll}$. The radial factors in the two Fischer norms therefore agree, and $|\xi(U)| \le 1$ by the minor bound for an orthonormal frame. The function $h_\lambda(U)$ is the normalized highest matrix coefficient in the tensor product of the indicated exterior powers. Its squared Haar integral is the inverse dimension of its highest-weight summand by Schur orthogonality. This proves

$$0 < k_\rho \le \dim V_{\lambda_\rho} \le (2a_E)^{N(N-1)/2}.$$

The last inequality follows by bounding each of the $N(N-1)/2$ factors in the Weyl dimension formula by the spread bound above.

For the actual shifted weights,

$$\begin{aligned} &\lambda_{\rho,0} = a_\rho - (N-k)/2 \qquad (\rho \in T), \qquad \lambda_{\rho,0} \ge c_* - (N-1)/2 \quad (\rho \notin T), \\ &d_{0,\rho} \le c_* + m\eta_* + \mathbf{1}_{\rho \in T} b_\rho, \quad d_{1,\rho} \le 3d_{0,\rho} + c_* + \eta_*, \qquad b_\rho \le 2b'. \end{aligned}$$

Here $b_\rho = -u_{k-1} - (k-1)/2$ uses the lower endpoint on the sampled embeddings. For the d_0 bound, the first $k-1$ interlacing differences telescope to at most $\alpha_0 - \lambda_{k-1}$. Substituting the half-sum shifts gives $v_0 - u_{k-1} - k + 3/2$. In each of the m appended blocks, the sum of the differences is precisely $e_B - b_B - (t_{\mathrm{cut}} + 1)/2$, with all but the first difference of that block zero. This contributes at most $m\eta_*$. Outside T, use $u_{k-1} = v_{k-2} - 1/2$ and $0 \le v_{k-2} \le v_0 \le c_*$. For d_1 bound the gaps at $i = 0, k-1$ by d_0 each, and telescope the remaining gaps to $(\lambda_1 - \lambda_{k-1}) + (\lambda_k - \lambda_{N-1})$. In the first difference insert α_{k-2} and bound $\alpha_{k-2} - \lambda_{k-1} \le d_0$. The shifted cloud difference $\lambda_1 - \alpha_{k-2}$ is at most c_*. The appended blocks are constant on each block, with first-to-last difference $e_{m-1} - e_0 - (m-1)(t_{\mathrm{cut}} + 1) \le \eta_*$ by the bounds on the level nodes. These estimates and $\log T_w \ll mL$, $m \le N$, $\log(2a_E) \le 2L_a$ give the common bound

$$\log \|P_{w,\varsigma}\|^2 \le -S_a \log T_w + O\big(\big(N_T b' N + f(c_* + m\eta_*)N + fN^2\big)L_a\big) = \mathrm{SZbP}(w),$$

with a uniformly large constant in the definition on the right. The scaled coisometry in the terminal metrics, The terminal weighted class pairing, Restriction in the fixed row norm, Equality of the two normalized Gram forms, and The appended-block shift.

Norm of the chosen column kernels. Work at the trivial auxiliary tuple, in the class-sum realization of the rational column chain. Let $p = N(j-1)$, $q = N(j)$, and $N = n_*$. Write $g = g_0 g_1$ and $c = T_w g$. At each place the archimedean coefficient norm is the Fischer norm divided by the squared Fischer length of the extremal polynomial. Every class has measure the inverse of its stabilizer order. The required zero-stage isometry identifies the actual metric of N_w with this reference column metric.

The finite datum for an ordinary matrix entry is the product of the conditions on the decorated lattices at T_w and the unfolded functions at the selected halves of g_i. At T_w the factor is the indicator that the map preserves the integral lattice and its residue decoration. It is invariant under addition of T_w times an integral map. At a place of norm g_i the moved symmetric-level function takes values of modulus at most one under every embedding. In its scaled lattice frames it requires integrality of each entry after multiplying it by the ratio of the basis scalings. On that support its character factors are roots of unity of the

residue group, and are unchanged by adding g_i times an integral map. The same support and periodicity survive imposing a fixed coupling column and evaluating the last coupling kernel. The representatives of the coupling cosets are integral in the chosen frames, so both of these last finite functions still have absolute value at most one. The absolute values of the lattice determinant character factors are zero or one.

For the ordinary unfolded function, the extra coordinates are summed modulo the prescribed period. With $h = m - j$ added frame coordinates at the source, the basis scalings have exponents $0, 1, \dots, h$. The support bounds the denominator of every summation coordinate by g_i^h. In a sum over an interval with g_i^B integral lifts, at most g_i^{B+h} values per coordinate remain. The normalization is g_i^{-B} per coordinate, and the summands have modulus at most one, since the unfolded additive character has order dividing g_i on the support. There are at most N^2 coordinates. The inverse volume of the reduction stabilizer is bounded by the number of its residue matrices, at most $g_i^{N^2}$, because the stabilizer contains the kernel of reduction in the scaled frame. The unfolded bound is therefore at most $g_i^{(h+1)N^2}$ per selected half. These are bounds on the actual entries, and apply to roots of unity after arbitrary coefficient embeddings. Multiplying the places uses $\prod_{P\ mselected} NP = (g_0 g_1)^f$. Thus the matrix-entry function Ξ on the semilocal module of maps $\mathcal{H} = \mathrm{Hom}_S(L, L')$ satisfies

$$\operatorname{supp}\Xi \subseteq \mathcal{H}, qquad\Xi(x + cy) = \Xi(x) \quad (y \in \mathcal{H}), \qquad |\varsigma(\Xi(x))| \le \beta = g^{O(fm^2N^2)}.$$

The last two kernels allow $\beta = 1$. The finite symmetric-level function, Support and size after unfolding, The bundled matrix entry, The first coupled entry, and The last coupled entry.

For the matrix entry the initial congruence sum uses modulus cc', where $c' = g^{fB-1}$, and an outside allowance c_A. Support and invariance let one regroup it over $\mathcal{H}/c\mathcal{H}$. Indeed, within the support the factor Ξ is constant on each fibre of $\mathcal{H}/(cc')\mathcal{H} \longrightarrow \mathcal{H}/c\mathcal{H}$, and the coset indicators partition the coarser indicator. In the coefficient field the class normalizing character factors have modulus at most one. On the self-dual torus of $\mathrm{Hom}_E(L, L')$, the indicator of one such coarse residue is a sum of at most $c^{[E:\mathbb{Q}]pq}$ torsion translates of the same normalized vacuum, with coefficients of modulus at most one. The refinement uses only the selected halves, which are disjoint from their conjugate halves. The level modulus is odd, the pairing ideal is a unit at these places, and the frame ranks are even. In the finite Heisenberg representation one applies the residue characteristic idempotent by Fourier inversion on the annihilator of the residue subgroup. For the chosen refinement, a coefficient is the finite Fourier average of a lattice indicator against a character, times the two quadratic phases for the translation and their commutator phase. The phases are roots of unity by rationality of the trace pairing on the torsion points. The volume factors $\mathrm{vol}^{-1/2}\,\mathrm{vol}^{-1/2}\,\mathrm{vol}$ cancel in the normalized indicator, and the class vacuum normalizer cancels its inverse. The index of a rank-pq lattice of integral maps modulo c is at most $c^{[E:\mathbb{Q}]pq}$, both for the translates and the original coarse summands. Consequently the sum of the absolute coefficients in the normalized translated-vacuum expansion is at most $\beta c^{2[E:\mathbb{Q}]pq}$. Coarsening the support of the entry, The torsion translates of the coarser indicators, and The expansion of the chosen kernel.

Use an embedding ι of the common CM theta coefficient field extending the embedding of the column coefficient field. Let $V_{j,\iota}$ be the normalized squared torus norm of the conjugate standard vacuum. In the class lattice a moved vacuum has the same norm by the unitary theta-frame change and the primitive translation formula. The conjugate of each torsion translate is the corresponding conjugate theta section with the same metric, including a possible unit-modulus theta-group factor. Denote by a_ι the area of the conjugate elliptic period lattice of the induced CM torus. The normalized complex Euclidean variable is $Z_\rho = (\pi/a_\iota)^{1/2} z_\rho$. In these variables the entire function for the matrix entry has majorant

$$|G((a_\iota/\pi)^{1/2}Z)| \le S \exp\left(\tfrac{1}{2}\sum_\rho \|Z_\rho\|^2\right), \qquad S^2 \le \beta^2 c^{4[E:\mathbb{Q}]pq} V_{j,\iota} \exp(O(fpq\log(2fpq\,\mathrm{rd}\,E))).$$

For the last factor, after translating to the point, apply the holomorphic submean bound on a flat coordinate ball in the quotient by the period lattice. In the Euclidean metric normalized by elliptic area the

Hom period lattice is self-dual for the imaginary Hermitian pairing, so its Euclidean covolume is one. Its positive squared lengths are positive integer multiples of $1/\sqrt{d}$, by the trace calculation on the period ideal of the host of discriminant $-d$. Hence its minimum squared length is at least $1/\operatorname{rd}(E)$. Use a ball of radius a fixed fraction of this minimum length. Its reciprocal volume in real dimension $2fpq$ costs $\exp(O(fpq\log(2fpq\operatorname{rd}E)))$. The Hermitian Gaussian is bounded above and below by absolute constants on that ball. The resulting squared estimate is also obtained directly with the constant $3(1+2fpq\operatorname{rd}(E))^{fpq}$. The translated-section bound in the actual torus metric, Transport at both types of embeddings, and The majorant on the actual conjugate entry.

Here is the operator bound from this majorant. Write $v_i = v_{w,j,\rho,i} \geq 0$ for the shifted source partition at ρ, and $d_\rho = \sum_i v_i$. The jet of the substitution operator uses the determinant shift to make its source weight v and then pads that weight with zeroes in the target rank. These operations send extremal vector to extremal vector. Substitution at a point Z_ρ has squared norm bounded by $\prod_i s_i(Z_\rho)^{2v_i}$, with singular values in decreasing order. For a diagonal rectangular matrix the padding is isometric in the Fischer norms and the diagonal operation acts by the Schur weights in the source. Their maximum on a decreasing nonnegative diagonal is the highest weight. Singular value decomposition gives the same bound for all Z_ρ, since both unitary actions preserve the metric. It follows that the scalar matrix coefficient $P(Z)$ of this substitution, between two unit vectors, is multihomogeneous of degrees d_ρ and satisfies $|P(Z)| \leq \prod_{\rho,i} s_i(Z_\rho)^{v_i}$.

Let $H(Z) = G((a_\iota/\pi)^{1/2}Z)$, and put $D = pq$, the number of complex variables per block. The coefficient pairing with P is the Taylor-number pairing of the jet, with factorial $\gamma!$ on a monomial Z^γ. For each block with $d = d_\rho > 0$, put $\epsilon = (2d+1)^{-1}$ and use the Gaussian weight $\exp(-(1+\epsilon)\|Z_\rho\|^2)$ for the Taylor coefficients of H. Monomial orthogonality and Cauchy–Schwarz bound the square of the coefficient pairing by

$$S^2 \prod_\rho (1+\epsilon)^{D+d}\epsilon^{-D} \int |P(Z)|^2 e^{-\sum_\rho \|Z_\rho\|^2} \prod_\rho \frac{dZ_\rho}{\pi^D}.$$

A block of degree zero is evaluated at zero and omitted from this estimate. Maximize

$$\prod_i s_i^{2v_i} \exp\left(-(1-\epsilon)\sum_i s_i^2\right)$$

over nonnegative s_i, using the Hilbert–Schmidt identity for $\sum_i s_i^2$. The maximum is $(1-\epsilon)^{-d}\exp(\sum_i \phi(v_i))$, where $\phi(u) = u\log u - u$ and $\phi(0) = 0$. The remaining Gaussian integral in that block is ϵ^{-D}. The identities $(1+\epsilon)/\epsilon = 2d+2$ and $((1+\epsilon)/(1-\epsilon))^d = (1+1/d)^d \leq e$ give the bound $e(2d+2)^{2D}$ for the residual multiplier. Finally the normalized derivative uses the algebraic differential generator in that block. Its multiplier has squared absolute value $\Omega_{\iota,\rho}^{-1}$ in the variable Z_ρ, where

$$\Omega_{\iota,\rho}^{-1} = \pi|\iota(c_{D,\rho})|^2/a_\iota.$$

Thus the logarithm of the squared matrix-entry norm is bounded by

$$K_{j,\iota} = \sum_\rho \left(\sum_{i<p} \phi(v_{w,j,\rho,i}) - d_\rho \log \Omega_{\iota,\rho}\right) + \log V_{j,\iota} + O\big(fN^2L_a + fm^2N^2\log(gT_w)\big).$$

Here $\log(2d_\rho+2) \ll L_a$ by the actual weights and $\log(2fpq\operatorname{rd}E) \ll L_a$. The padded source and target weights, The same normalized jet operator bound, and The normalized differential coordinate.

In the passage from entries to the class operator $J_{w,j}$, a source class contributes its inverse stabilizer weight times its matrix entry. All stabilizer weights are at most one. Cauchy–Schwarz then bounds the squared norm by the entry bound times the two numbers of classes, including their label sets. In the last operator, the two actual entries are multiplied before summing over the intermediate class and the coupling coset. Use Cauchy–Schwarz on that sum as well. This uses $K_{m-1,\iota} + K_{m,\iota}$ and the square of the

count of intermediate classes and coupling cosets. It charges two stage bounds for the one coupled map. The label sets inject into products of scaled-frame residue matrices, with at most g^{fN^2} possibilities, and the coupling cosets inject into the lattice modulo g, with at most g^{2fN} possibilities. The class sets are covered by the self-dual lattice genera with their local level labels. The definite mass bound and the stabilizer bound give the unlabelled class bound $\log h_j \ll fN^2 \log(2N \operatorname{rd} E)$. The counts of additional frames at the selected finite places are at most the sizes of their matrix reduction spaces. Their logarithms are $O(fN^2 \log(gT_w))$. In particular all these counts fit with the displayed $O(fm^2N^2 \log(gT_w))$ term. Class counts with level labels, Reduction of the index set, The coupling cosets, and Cauchy–Schwarz on the actual class sums.

The column-chain maps are defined over the tied coefficient field and its conjugate embeddings have the same squared operator norm when regarded as one infinite place. For each such place take the arithmetic mean of the $K_{j,\iota}$ over the extending embeddings, counting both complex-conjugate restrictions. Each is a bound on the same operator in its prescribed coefficient metrics. These averages have overall mean precisely the normalized mean over the larger field, by constancy of the number of extensions. In the CM normalization the mean of $\log \Omega_{\iota,\rho}$ equals $\log \Omega$ for each ρ. The mean of $\log V_{j,\iota}$ is $O(fpqL)$ by the normalized elliptic vacuum norm and its tensor power. Taking the mean of the bound thus gives

$$\sum_{\rho} \left(\sum_{i<p} \phi(\nu_{w,j,\rho,i}) - d_\rho \log \Omega \right) + O\big(fN^2 L_a + fm^3 N^2 L\big) \le \operatorname{kernelPlaceMeanBound}_j(w)$$

with the constant in the last expression chosen sufficiently large. This proves the archimedean column contribution to the height on the fixed chain, with the terminal pairing already estimated in its own row and column norms. Embedding average and the CM normalization.

Petersson eigenvalues. Section 2.6.5 applies the placed inner-product formula with the class-sum, lowest-type and principal-level metrics. The auxiliary theta pair there has target rank r and definite source rank $n = r + t$, as distinct from the growing grid ranks. The quantitative bound (44) gives the stage allowance

$$L_{6,j} = \sum_{\rho} \log H^{+}_{j,\rho} - C_{\mathrm{sz}} f n^2 \log(b \operatorname{rd}(E) g_0 g_1 n), \qquad C_{\mathrm{sz}} \ge 8,$$

where $H^{+}_{j,\rho}$ is the positive Fock hook product at the base word. This product has the same sum of logarithms after the host-field embeddings permute the weights. The negative hook logs are nonnegative. The bound applies on the orthogonal complement of the kernel of the actual principal projection, using the source and target Gram identifications at each conjugate embedding. Consequently taking the least stage product on the nonzero summands and averaging over the embeddings retains $\sum_j L_{6,j}$. Dropping the nonnegative hook contribution gives the coarse bound used below for the compensated fitted evaluation.

The quotient forms and their determinants. The comparison of height and mean target slope uses the base-level row dual $C^{(0)}$ with the same row normalization. Let $R(t)$ be the row lattice at twist t in the common coordinates of that level, and let $t = 0$ denote the base twist. The row lattice dual to $C^{(0)}$ is locally $R(t(v))$, with $t(v) = t(\bar{v})$ at conjugate places. For lattices in one rational space write $i_v(A, B) = \operatorname{ord}_v(\det B / \det A)$, using the ratio of determinant lattices. At stage j write $A_j(t)$ for the twist lattice transported to the coordinates of J_j^{row}, and put $i_{jv}(t) = i_v(J_j^{\mathrm{row}}, A_j(t))$. Write $J_{jv}(t)$ for minus the valuation of the rational stage Gram determinant in a local basis before transport, and $G_v(t)$ for the quotient Gram determinant in a row basis before transport.

Let Q_t be the row quotient and $Z_{j,t}$ the selected orthogonal complement in the principal-level auxiliary stage. The prescribed maps to the common row space and the j-th definite auxiliary space are ι_t and $\vartheta_{j,t}$. The complex reads at the conjugating embedding are the assembled map on the row quotient and the one-step auxiliary theta map on $Z_{j,t}$. The latter is restricted to the image of its adjoint by the complement

coordinates, with the insertion of the lowest-type coefficient in the same finite level. The placement (40) identifies its matrix on the original level lattice, for all twists. These are the maps used to transport the lattices $R(t)$ and $A_j(t)$. The maps on the complement coordinates and The placed stage metric.

Here is the form calculation after these identifications. A common Hecke type is an eigenspace of a commuting family of Hecke operators closed under adjoints. The corresponding projections are orthogonal, including on an invariant complement or quotient. Indeed T^* acts by the conjugate of the eigenvalue of T on each eigenspace, and distinct eigenspaces are orthogonal by the adjoint identity.

The norm of the assembled row map. Work first at the distinguished complex realization, in the principal-level measures of Section 2.6.2. Use t for the rank increment and $\boldsymbol{\mu}$ for the tuple. The growing row arrow at index i goes down from Y_i to Y_{i-1}, with $Y_m = X'$ in the specified top frame. Let $D_{i,\boldsymbol{\mu}}$ denote its downward multiplicity map. The auxiliary operator $T_{i,\boldsymbol{\mu}}$ at the same index is the good theta map of (39), whose source V_{i-1} has rank $r+t$ and target W_{i-1} has rank r. Write S_0 for the fixed exceptional finite places and $S = S_0 \cup Q(\boldsymbol{\mu})$ for the exceptional set at the tuple. At a common spherical label p, write B_p for the standard rank-r parameter of the conjugate of the analytic Y_0 vector, in the lattice frames used in the doubling integrals. Define

$$z_{i,p,\boldsymbol{\mu}} = L_E^S\left(\frac{t+1}{2}, B_p \otimes \mu_{\boldsymbol{\mu},i}^{-1}\right).$$

In this notation an omitted F-place removes every factor above it in E. The bottom uses the dual double-coset label and W_{i-1} uses the ordinary label. If $H_d^* = H_{d^\vee}$, a joint unitary eigenvalue on the conjugate of a bottom evaluation is the H_d-eigenvalue on the matched W_{i-1} evaluation. The match in Section 6 identifies these eigenvalues at every place away from S_0. The spherical character determines its Satake class, so the two standard parameters agree away from S. In particular the E_S of (43) for this auxiliary pair is $z_{i,p,\boldsymbol{\mu}}$, since its V_{i-1}-splitting character is $\mu_{\boldsymbol{\mu},i}$. Thus $z_{i,p,\boldsymbol{\mu}}$ is positive on an occurring matched summand by that formula.

Set $a = r + (i-1)t$, $n_i = r + it$, and $d = (i-1)t$. The transferred parameter in the growing-chain doubling integral has rank a. The parity cancellation in Section 2.6.2 makes its critical numerator $z_{i,p,\boldsymbol{\mu}} H_i^S$, where

$$H_i^S = \prod_{h<i} \prod_{n=1}^{t} L_E^S(n, \mu_{\boldsymbol{\mu},h}/\mu_{\boldsymbol{\mu},i}).$$

The equality is first an Euler-product identity in a right half-plane, with the shifts $s + \frac{1}{2} + (t-1)/2 - \ell$, $0 \le \ell < t$. It extends to a neighborhood of $s = t/2$ by the entire continuation of the nontrivial quotient-character factors and absolute convergence of the rank-r factor there. The distinct real exponents make all the quotient characters nontrivial, so the factors at 1 and at $2, \dots, t$ are nonzero. In $H_i^{S_0}$ both the characters and the omitted places are fixed across all tuples, because the common twist cancels in each quotient. The shifted Euler product and its critical value.

For a source coefficient vector v_i and polynomial ϕ_i in the fixed real models, let $J_{i,\boldsymbol{\mu}}$ be the real doubling coefficient integral with the full level test retained and with finite level volume one on the lower group. It includes the oscillator squared vacuum norm in the rational-covolume-one oscillator measure. It is the coefficient of the lower multiplicity pairing after integration on the full level. Choose v_i, ϕ_i once on the base pair with $J_{i,0} \neq 0$. The real nonvanishing argument in Section 2.6.2 provides such a choice, with convergence in a neighborhood of the critical point. At infinite places the tuple twist is trivial. In the polynomial model all the tests use the same real oscillator action, the same joint-harmonic map and the same test vectors. The vacuum at infinity lies in the same normalized Gaussian line, so the real coefficient is fixed after dividing by the squared global vacuum norm. The real Haar measure is fixed because the level has volume one in its finite factor. The common real measure · The same tested real coefficient in every tuple.

We now keep the constants in the integral at the tuple primes. For $v \in Q(\boldsymbol{\mu})$, let q_v be its residue cardinality in the split chart, and put

$$C_v(u, w) = \prod_{u<n\le w} (1 - q_v^{-n}).$$

The X-support of (H) in Section 2.6.2 has first $r \times r$ block invertible modulo v, and $r(n_i - r)$ entries below it zero modulo v. All values on the support have modulus one. Thus the squared oscillator norm ratio to the full integral vacuum at v is

$$U_{i,v} = C_v(0, r) q_v^{-r(n_i-r)}.$$

The upper level has parahoric shape $(r, n_i - r)$. Counting the invertible diagonal blocks and the free $r(n_i - r)$ entries in the upper-right block gives its index

$$I_{i,v} = \frac{q_v^{r(n_i-r)} C_v(0, n_i)}{C_v(0, r) C_v(0, n_i - r)}.$$

This scales the upper quotient volume by $I_{i,v}$. In particular, writing $\mathcal{V}_{i,\boldsymbol{\mu}}$ for that volume, one has

$$J_{i,\boldsymbol{\mu}} = J_{i,0} \prod_{v\in Q} U_{i,v}, \qquad \mathcal{V}_{i,\boldsymbol{\mu}} = \mathcal{V}_{i,0} \prod_{v\in Q} I_{i,v}.$$

To evaluate the factor for extending the lower finite integration at v from its parahoric J to the full local group, normalize its Haar measure by $\mathrm{vol}(J) = 1$. Use the induced model in the calculation of (I) in Section 2.6.2. The J-character vector is supported on $P_{r,d}J$ and its inducing value at the identity has norm one. In the compact realization with maximal compact volume one its squared norm is $\mathrm{vol}_K(J)$. The normalized matrix coefficient averages over J with that same divisor. On replacing g by kg and X by Xk^{-1} under this average, the characters cancel and the average contributes one. The resulting group integral is on $P_{r,d}J$. There $g = \mathrm{diag}(A, D)uk$ has the product integration measure, with $\mathrm{GL}_r(\mathcal{O}_v)$, $\mathrm{GL}_d(\mathcal{O}_v)$, the integral upper unipotent and J each given volume one. Indeed $P_{r,d} \cap J = P_{r,d} \cap K$ and the product measure uses left Haar on $P_{r,d}$ in the order diagonal then unipotent. Its normalization on J fixes the integration constant. The two integral eliminations in (I) now leave the unrestricted integral rectangular lower block and require A and u to run over their integral compact subgroups. They have mass one. The initial integration on the X-support contributes its squared norm, already included in $J_{i,\boldsymbol{\mu}}$. The remaining integral is the spherical doubling integral for the suffix of rank d, with target rank $d + t$. For $d = 0$ it is the empty factor, since the coefficient is supported on the lower level itself. Consequently the exact multiplier at $s = t/2$ is

$$B_{i,v} C_v(t, n_i - r), \qquad B_{i,v} = \prod_{h<i} \prod_{n=1}^{t} L_{E,v}(n, \mu_{\boldsymbol{\mu},h}/\mu_{\boldsymbol{\mu},i}),$$

with both halves of the split E-factor included in $L_{E,v}$. This computation applies to the full occurrence space. The character line in each local induced representation is one-dimensional and the normalized lower-block coefficient is determined by the suffix Satake list. This list depends on the earlier splitting characters and the tuple. It is the same for all lower-rank representations that occur in the selected cut at this step. Local integration therefore multiplies the Petersson coefficient at the identity by the stated factor, also for a sum of occurrences. Absolute convergence near $t/2$ is the spherical suffix estimate for (I). The integral with the lower parahoric of volume one · Multiplication of the removed local factors.

Let d_i^S be the unramified doubling denominator for ranks a, n_i at $s = t/2$, with exactly S omitted. The split formula gives $d_{i,v}^{-1} = C_v(t, n_i)$. Restoring the omitted factors at Q gives

$$H_i^{S_0} = H_i^S \prod_{v\in Q} B_{i,v}, \qquad d_i^S = d_i^{S_0} \prod_{v\in Q} C_v(t, n_i).$$

The density identity $I_{i,v} U_{i,v} C_v(t, n_i - r) = C_v(t, n_i)$ follows by cancellation of the three consecutive products. Apply the convergent Siegel–Weil norm formula in the prescribed measures and with the transferred

label, as in Section 2.6.2. Its source is cuspidal by the stated first-occurrence calculation, the upper global Witt index is less than t, and the local integrals at the critical point are the continuations of the unfolded integrals in the right half-plane. Those on the level and the tuple-prime integrals converge on a neighborhood of this point, as above. At S_0 the coefficient support and the character cancellation in (F) give the full-level factor. Before passing from tested functions to multiplicities, the upward norm has multiplier

$$\mathcal{V}_{i,\boldsymbol{\mu}} \frac{z_{i,p,\boldsymbol{\mu}} H_i^S}{d_i^S} J_{i,\boldsymbol{\mu}} \prod_{v \in Q} B_{i,v} C_v(t, n_i - r) = z_{i,p,\boldsymbol{\mu}} \boldsymbol{\mathcal{K}}_i, \qquad \boldsymbol{\mathcal{K}}_i = \mathcal{V}_{i,0} \frac{H_i^{S_0}}{d_i^{S_0}} J_{i,0},$$

with simultaneous complex conjugation in the conjugate-linear theta convention. Each $\boldsymbol{\mathcal{K}}_i$ is nonzero. The placed norm identity · Exact cancellation of the volume, norm and density.

Let w_i be the adjoint of the real Howe coefficient contraction at ϕ_i, evaluated on v_i. The real pairing and contraction on the fixed modules determine w_i uniquely and independently of the tuple and the label. To identify the upward integral on $u(v_i)$, pair it with each upper evaluation $y(w)$. Fubini identifies this pairing with that of $(D_{i,\boldsymbol{\mu}} y)(C_{\phi_i} w)$ with $u(v_i)$, where C_{ϕ_i} is the coefficient contraction. Schur orthogonality gives the product of the multiplicity pairing of $D_{i,\boldsymbol{\mu}} y$ with u and the coefficient pairing of $C_{\phi_i} w$ with v_i. Taking both adjoints gives the pairing of $y(w)$ with $(D_{i,\boldsymbol{\mu}}^* u)(w_i)$. The upward integral lies in the prescribed upper type and level by first occurrence, the real transfer and the level character law. Its equality with this evaluation follows by testing against the span of the upper evaluations. Schur orthogonality expresses the pairing of two such evaluations as $\|w_i\|^2$ times their multiplicity pairing. For a nonzero reached lower vector, the norm identity just proved has a nonzero right-hand side. Hence $w_i \neq 0$ and the adjusted constant $\boldsymbol{\mathcal{K}}_i / \|w_i\|^2$ is positive real, by taking the same vector in both slots and using $z_{i,p,\boldsymbol{\mu}} > 0$. For the terminal step use the actual class metric h_X of X'. The fixed coefficient evaluation j_X into the top multiplicity model gives $h_X = \delta j_X^* h_m j_X$, for a positive δ independent of the tuple. In fact choose a pure vector v in the top real type whose evaluation is nonzero on the definite coefficient representation C_X, and write this evaluation as $\langle \cdot, a_X \rangle$ there. Compact Schur orthogonality of the evaluation gives $\|a_X\|^2 / \dim C_X$ times the class form. Schur orthogonality in the multiplicity model gives $\prod_\tau \|v_\tau\|^2$ times the multiplicity form for the same evaluation. Therefore

$$\delta = \frac{\dim C_X \prod_\tau \|v_\tau\|^2}{\|a_X\|^2}.$$

The adjoint of $D_{m,\boldsymbol{\mu}} j_X$ is $\delta^{-1} j_X^{-1} D_{m,\boldsymbol{\mu}}^*$. Its squared norm in the class metric acquires the factor δ^{-1}. Define

$$b_i = \frac{\boldsymbol{\mathcal{K}}_i}{\|w_i\|^2} \quad (i < m), \qquad b_m = \frac{\boldsymbol{\mathcal{K}}_m}{\delta \|w_m\|^2}.$$

The polarized adjoint pairing at every step is thus $b_i z_{i,p,\boldsymbol{\mu}}$ times the lower pairing for all reached vectors of the transferred label. The level at Y_0 and its Petersson measure are the base ones, since the twist modifies only the suffix there and it has rank zero. The good central projector is orthogonal and fixes the bottom input. The level and spherical transfer identities in Section 2.6.2 keep every successive upward image in its prescribed upper cut. This makes the composition of these adjoints the adjoint of the assembled downward map $\Theta_{\boldsymbol{\mu}}$ on the good bottom space. For the spectral calculation at a tuple one can restrict the spherical character to places outside S. The original bottom label embeds in this cut, and the norm calculation applies to its whole joint eigenspace there. The transfers inject on the spherical labels by cancellation of the appended strings. Hence the upper images of bottom vectors with different restricted labels are orthogonal at the last step. Multiplying the polarized identities gives the pairing of $\Theta_{\boldsymbol{\mu}}^* u$ with $\Theta_{\boldsymbol{\mu}}^* u'$ for u, u' in the same restricted label. Orthogonality supplies the cross terms against every other restricted label. It follows that

$$\Theta_{\boldsymbol{\mu}} \Theta_{\boldsymbol{\mu}}^*\big|_{E_p} = \kappa_0 \prod_{i=1}^m z_{i,p,\boldsymbol{\mu}}\, I, \qquad \kappa_0 = \prod_{i=1}^m b_i > 0.$$

The top coefficient normalization · Composition on the reached subspaces and the bottom cut.

The auxiliary formula gives $\lambda_{i,p,\boldsymbol{\mu},\sigma_0} = a_i z_{i,p,\boldsymbol{\mu}}$ with $a_i > 0$ equal to the mass of V_{i-1} times its coefficient dimension and the indicated real hook, lowest-type dimension and coefficient-integral factors, divided by the full auxiliary doubling denominator. That denominator is $d_{\mathrm{off}} d_S$ in (43), independent of the omitted set after the level and vacuum norm cancellation. Each a_i is therefore fixed. It follows at every tuple that the assembled scalar is

$$\kappa_R \prod_{i=1}^{m} \lambda_{i,p,\boldsymbol{\mu},\sigma_0}, \qquad \kappa_R = \frac{\kappa_0}{\prod_{i=1}^{m} a_i}.$$

The matching of the level-character cuts and the isotype dimensions in (173) identifies the dimensions of the bottom E_p and every auxiliary $W_{i,p}$. Their targets are independent of the tuple. Thus the two kernel complements have the same dimension on this label, since each is mapped isomorphically to its corresponding occurring summand. A missing summand has dimension zero and contributes one in the determinant product. If a stage scalar is recorded for such a summand, use the prescribed positive fallback for the absent part. The application has positive total bottom rank, so the formula for κ_R is determined by an occurring label. The common scalar on the matched blocks.

Return to t as the tuple index in the determinant comparison. The rational coordinate transition (41) uses the base-image row frame. Its square-constant ratios express the principal test values in a common coefficient field. The positive scalar just computed uses the actual analytic adjoints in that frame. A unit phase in the auxiliary lift has squared modulus one. In particular no bottom-triangle constant is multiplied into the adjoint scalar.

To pass from an adjoint scalar $a > 0$ on a target type to the full source pairing, take u in the prescribed complement with Tu in that type. Then $T(u - a^{-1}T^*Tu) = 0$. Both terms lie in the complement, by adjoint equivariance for the level cut. Thus $u = a^{-1}T^*Tu$, by injectivity on that complement, giving $\langle Tu, Tv\rangle = a\langle u, v\rangle$ for every v there. Apply this for the assembled map and for each auxiliary map. At the conjugating embedding the row form is the assembled form divided by the principal level order $|P|$. On the auxiliary side the inserted lowest-type frame multiplies squared norms by $|P|$, which cancels the same factor in the auxiliary source comparison. This leaves $\kappa_{R,\sigma_0} \prod_j \lambda_{j,\alpha,t,\sigma_0}$ on the row and $\lambda_{j,\alpha,t,\sigma_0}$ on the auxiliary summand. Here $\kappa_{R,\sigma_0} = \kappa_0 / \prod_j a_j$, the same at every twist. Adjoint scalars and the level complements.

At the remaining embeddings only the base twist is used. Its row form is defined by

$$H_{R,\sigma} = \kappa_R \sum_{\alpha} \Big(\prod_j \lambda_{j,\alpha,0,\sigma} \Big) \pi_\alpha^* H_{Q_0,\sigma} \pi_\alpha$$

in the base quotient coordinates, where the π_α are the orthogonal type projections. The same argument evaluates the entire form on a first vector of type α. On the auxiliary space the base theta comparison in Section 2.6.7 uses the corresponding conjugate source and target Gram forms, so its polarized comparison there gives the same stage scalar. Thus the form calculation uses the actual level insertion and Gram transport at (t, σ_0) and $(0, \sigma)$, as well as the equality of the occurring type dimensions on Q_t and each $Z_{j,t}$. That equality follows from the common type identification when the maps on complements are bijective on the prescribed types. The rationality and these image identifications for the actual maps are needed in this use of the calculation. Row form at both sets of embeddings, Auxiliary form on the same complement, and The base row frame at the other embeddings.

Choose bases adapted to the common type decomposition for computing determinants. If type α has dimension d_α, its row factor is $(\kappa_{R,\sigma} \prod_j \lambda_{j,\alpha,t,\sigma})^{d_\alpha}$, and each stage has factor $\lambda_{j,\alpha,t,\sigma}^{d_\alpha}$. Multiplication shows that the determinant of the transported row form divided by that of the source quotient is the corresponding product of stage determinant ratios times $\kappa_{R,\sigma}^{d_{\mathrm{row}}}$. The ratio is basis independent. Compare this identity at t and 0 at the conjugating embedding, where κ_{R,σ_0} cancels exactly. Use rational bases semilocal at

the paired places $v, \bar{v}$. If γ_t changes the two transported row bases and δ_{jt} changes the corresponding transported stage bases, the ratio identity is then algebraic:

$$\det(\gamma_t^*) \det \gamma_t \, \frac{G_v(0)}{G_v(t)} = \prod_j \left(\det(\delta_{jt}^*) \det \delta_{jt} \, \frac{G_{jv}(0)}{G_{jv}(t)} \right).$$

Here G_{jv} is the stage Gram determinant with $J_{jv} = -\operatorname{ord}_v G_{jv}$. The fixed complex frames cancel on each side, and the embedding on the rational quantities is injective. Valuating this equality, using $\operatorname{ord}_v(a^*) = \operatorname{ord}_{\bar{v}}(a)$, gives

$$\begin{aligned} i_v(R(0), R(t)) + i_{\bar{v}}(R(0), R(t)) &= \mathcal{D}_v(t) - \mathcal{D}_v(0), \\ \mathcal{D}_v(t) &= \operatorname{ord}_v G_v(t) + \sum_j \big(i_{jv}(t) + i_{j,\bar{v}}(t) + J_{jv}(t) \big). \end{aligned}$$

Exact determinant comparison for both sets of embeddings.

The base quotient comparison uses a rational projection P with full row rank and a positive Hermitian form H over the CM subfield. The archimedean quotient Gram matrix is $(P_\sigma H_\sigma^{-1} P_\sigma^*)^{-1}$. The rational quotient matrix $(PH^{-1}P^*)^{-1}$ uses the coefficient involution. Determinant Cauchy–Schwarz for the rows of P at σ and $\sigma^\# = \overline{\sigma \circ *}$ bounds the square of the absolute value of the rational Gram determinant of $PH^{-1}P^*$ by the product of the two positive archimedean determinants. Here $H_\sigma = H_{\sigma^\#}$ and the full rank ensures nonsingularity at the conjugating embedding. Inversion, averaging and the product formula bound the squared logarithmic covolume of the base quotient by $\sum_v \lambda_v \operatorname{ord}_v G_v(0)$. The stage Gram determinants at the base give their covolumes by the product formula since the base stage metrics are rational Hermitian. Under the base metric comparison the stage indices and Gram determinants therefore contribute $\mathcal{D}_v(0)$, and the normalization adds $d_{\mathrm{row}} J_\kappa$. Replacing the base lattice by the local twists changes covolume by $2\sum_v \lambda_v i_v(R(0), R(t(v)))$. Because $\lambda_v = \lambda_{\bar{v}}$, this replaces $\mathcal{D}_v(0)$ by $\mathcal{D}_v(t(v))$ in the sum. Thus

$$\mu(C^{(0)}) - J_\kappa \le \sum_j \frac{\operatorname{lcv}(J_j^{\mathrm{row}})}{d_{\mathrm{row}}} + \frac{1}{d_{\mathrm{row}}} \sum_v \lambda_v \mathcal{D}_v(t(v)).$$

Base quotient determinant bound and Weighted local identities.

Finite Gram terms at the two levels. We evaluate the finite Gram terms using the class-sum model. At a class c with stabilizer Γ_c, denote the coefficient lattices by N_c, D_c. Here N_c is the dual Weyl lattice ∇ given by integral column equations, and D_c is the Weyl lattice Δ generated by integral divided powers. With the rational coefficient form and its dual $^\#$ the integral comparison is

$$sN_c^\# \subseteq D_c \subseteq N_c, \qquad sD_c \subseteq N_c^\#,$$

for the prescribed positive integer s. The same notation for duality at v pairs the lattice at v with its conjugate at $\bar{v}$. Let $L = \bigoplus_c N_c^{\Gamma_c}$, with the class form of weight $1/|\Gamma_c|$. For a vector $y \in L^\#$ one can test against invariants supported at c. The dual of the invariant intersection for the coefficient form is the orthogonal average of $N_c^\#$. In fact restriction of an integral functional to the invariant intersection is surjective on the integral dual, since the intersection is saturated over each DVR. Its representing map on the rational spaces is the orthogonal average. Hence $y_c/|\Gamma_c|$ is in that average of $N_c^\#$, and $sy_c \in D_c^{\Gamma_c}$, since Γ_c preserves D_c. This proves $sL^\# \subseteq L$. The classwise duality.

For the row at principal level take the cut of L by the product of the central good-type idempotent e_X and the level-character idempotent e_P. This product is self-adjoint. The ordinary character expression for e_X clears its denominator with $|G|$, since the character values are algebraic integers. The averaging expression for e_P clears its denominator with $|P|$. In the row form, which is the class form divided by $|P|$, the dual of the cut is $|P| e_X e_P L^\#$, by the same restriction argument. Consequently $s|G|$ takes this dual into

the cut lattice. In an orthogonal quotient the dual of the projected lattice is the intersection with the dual of the source in the orthogonal section. Indeed evaluation on the projected lattice is evaluation on any lift, since the section is orthogonal to the kernel. The inclusion descends to $s|G|Q_t^{\#} \subseteq Q_t$ on the row quotient lattice. Taking determinants gives

$$\operatorname{ord}_v G_v(t) \le d_{\mathrm{row}} \operatorname{ord}_v(s|G|).$$

The sign follows from $i_v(Q_t, Q_t^{\#}) = -\operatorname{ord}_v G_v(t)$. The weighted contribution per row is at most $L_G = \lambda_0 + \log s$ by the product formula for $s|G|$. Cancellation of the principal level order and The quotient Gram bound.

For the auxiliary stage write T_{jt} for the projection of the integral level lattice to $Z_{j,t}$, and $L_{jt}^{\cap}$ and $D_{jt}^{\cap}$ for its intersections with the level lattice and with its Weyl comparison lattice. Let $\eta_{jv}^{\mathrm{cut}} = \operatorname{length}_v(T_{jt}/L_{jt}^{\cap})$. Put $g_{jt} = |P| \operatorname{lcm}_c |\Gamma_c|$, using the auxiliary classes of this stage and twist. The form on its level coordinates weights the coefficient forms by $(|P|\|\Gamma_c|)^{-1}$. To pair an element of $D_{jt}^{\cap}$ against T_{jt}, pair it against an integral lift using orthogonality. The inclusion $sD_c \subseteq N_c^{\#}$ then gives $sg_{jt}D_{jt}^{\cap} \subseteq T_{jt}^{\#}$. At a selected twisted place use the lifting exponent J_{jv}^* with $\mathfrak{p}_v^{J_{jv}^*} L_{jt}^{\cap} \subseteq D_{jt}^{\cap}$. The total determinant change of the two inclusions gives

$$J_{jv}(t) \le \eta_{jv}^{\mathrm{cut}} + d_{\mathrm{row}}\big(J_{jv}^* + \operatorname{ord}_v(sg_{jt})\big).$$

This formula applies also when the projection and intersection lattices coincide. The projected, cut and comparison lattices and Lifting exponent and Gram contribution.

The full good eigenspaces on the auxiliary stage. Fix $1 \le j \le m$ and an actual tuple $\boldsymbol{\mu}$ used at a twisted coefficient place. Write V_j for the definite member of rank $N = r + t$ and W_j for the member of rank r of the one-step auxiliary pair. Denote their selected source and good target multiplicity spaces by X and Y. The source X is the full class space at the given character level (P, ψ). Denote by $T : X \to Y$ the lowest-type downward theta map and use the Petersson measures in (40). In particular its complex adjoint T^* is defined using the same source class form. The finite part ϕ_{f} on the oscillator is the one in that column identity.

At a generic good point x of any residual factor at the chosen place, let $X(x) \subset X$ be the sum of the *full* spaces of the automorphic types with that point's normalized Hecke character. The compatible base change at that point has the splitting-character twist of a rank-r parameter Σ and the specified length-t rank-one string. This description applies to every occurrence at x. The allowed split operators give the characteristic polynomials of the compatible standard transfers in the normalized algebraic convention. The primes of E over split places of F have density one in E because the ordinary inert primes have square norm. After removing the finitely many excluded places, these Frobenius classes still suffice for Chebotarev and Brauer–Nesbitt. This gives equality for the semisimple transfers, hence the same unramified polynomials and isobaric base change by strong multiplicity one. The base change determines the type labels at the two marked primes on this fibre. The good projector therefore fixes every vector of $X(x)$. Under the usual unramified base-change convention the Satake multiset at every ordinary prime of E, including the inert ones, has the string

$$\left\{ \chi_V(\varpi_w) q_w^{(t-1)/2-a} \,:\, 0 \le a < t \right\},$$

where χ_V is the auxiliary splitting character of V_j. Here the complement has entries bounded in absolute value by $q_w^{(r-1)/2}$. Indeed it is the unitary rank-r cuspidal complement with the indicated unitary twist. The base change determines the good ordinary Satake polynomials also at places omitted from the test used to identify x. The Satake string on the actual realized eigenvectors · The whole good fibre in the good finite-type part.

Injectivity on the good fibre. Let y be any nonzero vector in $X(x)$. Use the adelic extension F_y of its class values in the fixed coefficient space C for the definite real group. For $\ell \in C^\vee$, write $u_\ell(h) = \ell(F_y(h))$.

Each nonzero u_ℓ is principal-congruence-fixed at the marked places and has the above ordinary Satake data. The extended functions on the definite group are cuspidal and the quotient is compact in the adelic unitary realization used here.

First consider all Schwartz tests in the downward lift from V_j to W_j. For some ℓ at least one of these tests gives a nonzero lift of u_ℓ. Here is the residue calculation with the local tests and its dependence on y. Enlarge W_j in its Witt tower to $B = W_j \oplus \mathbb{H}^t$, with $\mathbb{H}$ a hyperbolic plane. Then $\operatorname{rk} B = r + 2t$ and $(\operatorname{rk} B - N)/2 = t/2$. Keep the splitting character for this tower. Choose a nonzero joint harmonic vector in the real model of the actual pair, using the fixed finite vector ϕ_{f} of the column identity. On its real coefficient factor choose $w \in C$ giving a nonzero oscillator vector, and take $\ell = \langle w, -\rangle$, linear in the argument. For the chosen test on V_j, W_j, integration over the definite real group in the doubling integral of u_ℓ is a nonzero scalar times $\int |\langle w, \rho(b)w\rangle|^2 \, db$. The oscillator coefficient is proportional to $\langle w, \rho(b)w\rangle$. On the automorphic side, the integrated form on C obtained from $F_y(h)$ commutes with the definite real group by translation in h. Hence by Schur's lemma it is scalar on C, with positive scalar given by the trace divided by $\dim C$. The trace is the integral of $\|F_y(h)\|^2$ and is positive for this y. The two matrix coefficients in the product have opposite conjugations. Since $w \neq 0$, the integral of their absolute square is positive. This verifies the witness on an arbitrary vector of the entire fibre, allowing linear combinations of occurrences and of vectors within any multiplicity space. A real witness for every nonzero class vector · The scalar form and its actual spherical and principal levels.

Take the external product of this test with the hyperbolic vacuum. In the split charts at the marked primes the extra vacuum is a full-frame congruence coset in the half with the adjoint support convention. Two translates of that coset meet only when the source action is in the congruence kernel, as can be seen by multiplying the two frames by a common right inverse modulo the marking ideal. Its matrix coefficient is constant there and is constant on the definite real group since the hyperbolic real vacuum is fixed. The original test and u_ℓ are invariant under this congruence kernel. The height on the real group times this kernel is one. With the covolume-one oscillator measures the semilocal factor is a nonzero constant c_y, independent of s, from the real integral and the squared norm of the extra vacuum with the positive congruence volume. At the remaining places the enlarged test is spherical in the self-dual frames. The doubling Euler identity on a right half-plane for the V_j integral is

$$Z_y(s) = c_y \, \frac{L_E^S(s + \tfrac{1}{2}, \mathrm{BC}(u_\ell) \otimes \chi_V^{-1})}{d_{N,r+2t}^S(s)}.$$

Here S consists of the places of the prescribed exceptional levels, with all factors above S omitted on E, and $d_{N,r+2t}^S$ is the spherical denominator for the enlarged pair. At $s = t/2$ its arguments are $t+1, \dots, t+N$, so it is holomorphic and nonzero there. The rank-r numerator factor converges absolutely near $t/2$ by the bound on the complement and $t \geq r + 5$. The string factors have arguments $s + \frac{1}{2} - (t-1)/2 + a$ for $0 \leq a < t$. At $t/2$ exactly $a = 0$ gives the pole of $\zeta_E^S(1)$, and the remaining factors are nonzero. Thus the residue of Z_y at $t/2$ is nonzero.

Apply the complementary first-term identity to the doubled Weil section at $s = t/2$. [19, Theorem 7.3(i)] In its unitary notation $d(n) = N$, $m = r + 2t$ and $m' = r > 0$, so $m + m' = 2d(n)$. The spaces are in the same Witt tower and the splitting character has the same required parity on both. The coefficient of $(s - t/2)^{-1}$ of the B-Eisenstein series is a nonzero constant times the smaller regularized theta coefficient of the Ikeda image of the compact-averaged test. Its order is 0 for the anisotropic W_j and -1 for positive Witt index.

Suppose all the unrestricted downward tests against u_ℓ on W_j vanish. In the doubled Schrödinger model the restriction of a product test to the two copies of V_j pairs as the product of the two downward theta integrals, with conjugation on the second. This follows by separating the two rational sums in the theta kernel. Finite tensor sums and Schwartz density give zero for every doubled test paired against u_ℓ and its conjugate, at each $h \in U(W_j)(\mathbb{A}_F)$. Continuity here is on the compact definite source quotient. The compact projection and the Ikeda integral give an allowed Schwartz test. In the isotropic case the reg-

ularizing operator on the doubled side acts on that test by the Weil representation, leaving it a Schwartz test. The absolutely convergent regularized integral therefore has zero pairing, by performing the compact source integrals first. The regularizing scalar and the auxiliary Eisenstein weight are the same as in the regularized integral. [19, §§3.4–3.5] The paired meromorphic function is zero, so its Laurent coefficient vanishes. For anisotropic W_j, the unregularized compact integral gives the same conclusion. The paired coefficient on the B side is the residue of Z_y by the doubling and unfolding identity in the covolume-one oscillator models. This contradicts the nonzero residue. Thus some unrestricted test on W_j is nonzero. The complementary first-term calculation here is on input V_j. The positive scalar in (43) is a convergent-range doubling value on input W_j. The real and principal-congruence support factor · The complementary coefficient vanishing with the stated ranks · The residue calculation for the definite source · The denominator and its arguments.

To compare with the fixed lowest-type map, write

$$I_y(\phi, \ell; g) = \int_{[U(V_j)]} \theta_\phi(h, g)\, u_\ell(h)\, dh \qquad (g \in U(W_j)(\mathbb{A}_F)).$$

Use the chosen linear and conjugate conventions of the column identity throughout. Since y is in the good part, the good central idempotent on W_j already fixes its unprojected lowest-type theta map. In fact each term of the idempotent translates the test and pairs with the corresponding character-corrected translate of the definite vector. Sum those terms on V_j, where the idempotent fixes y. This is the same source- and target-level equivariance used to place the columns in (40). Thus if $Ty = 0$, the entire unprojected lowest-type contraction is zero.

For a lowest-type vector b on W_j, the chosen joint harmonic tensors $b \otimes w'$, $w' \in C$, give the tests $\phi_{b,w'} = \phi_{\mathrm{f}} \otimes \phi_\infty(b, w')$. For fixed b, g, define $A_{b,g} \in \mathrm{End}(C)$ by $\ell(A_{b,g} w') = I_y(\phi_{b,w'}, \ell; g)$. Changing variables on $U(V_j)(F_\infty)$ and using its action on the C-argument of the joint harmonic map proves that $A_{b,g}$ commutes with this compact group. The projection to the lowest type is the trace of $A_{b,g}$, obtained by expanding in a basis of C and its dual. A vanishing contraction therefore gives $A_{b,g} = 0$ by Schur's lemma. For a general polynomial the same change of variables makes its pairing with ℓ a source compact equivariant map into C. Averaging the polynomial against the character projector for the source C-isotype preserves this map. In the joint-harmonic Fock description of this pair, the C-isotype is generated by the displayed tensors under the compact actions and the Lie algebra of $U(W_j)(F_\infty)$, using the same lowest type as in (43). Under the integral the source action transfers to ℓ and the target action to g. Differentiating the target identity preserves vanishing for every g. Thus all polynomials with the fixed finite factor give zero. Real Hermite approximation in a Sobolev norm for which the theta integral is continuous gives zero also for Schwartz tests with that finite factor. A test satisfying the finite Heisenberg-translate relations of ϕ_{f} lies on this finite line in the Schrödinger tensor model. Indeed on each real slice the relations give an intertwiner from the cyclic finite oscillator module generated by ϕ_{f}, hence a scalar multiple of its finite vector by irreducibility. This formulation also applies after identifying the global oscillator with its tensor model. The good projection commutes with the fixed contraction · Trace and Schur's lemma · Joint-harmonic generation for the tested source coefficient type · Real polynomial generation · Sobolev passage with fixed finite part.

For finite generation use the source character and the two type-prime groups of the actual tuple. Let J_V and ε denote the source level and its finite transformation character on u_ℓ, so $u_\ell(hk) = \varepsilon(k) u_\ell(h)$. Changing h in I_y kills the differences $\omega(k, 1)\phi - \varepsilon(k)^{-1}\phi$, $k \in J_V$. Averaging the smooth finite test against this character writes it as a test of the inverse source level type plus a finite sum of these differences. For the first-block finite group at a type prime, use the two mirabolic subgroups whose elements act trivially on the quotient by a specified line or act pointwise trivially on a specified hyperplane. In each average the weight is $\chi_{j,w}(\det h)^\epsilon$ for $\epsilon = 1, -1$, where $\chi_{j,w}$ is the designated tame character in the type-level datum at the chosen half w. On u_ℓ these averages are taken using first-block lifts and dividing by the flag-level character of the lift. They annihilate u_ℓ because a nonzero such average would give a constituent

in the Harish–Chandra restriction to the block Levi of sizes $1, r-1$ or $r-1, 1$ with the prescribed scalar character on the size-one block. This would put the respective excluded scalar root in the first-block label. Both exclusions hold at each type prime for $X(x)$. Equivalently, each of the prescribed mirabolic averages composed with the good central idempotent is zero. The source right-translation formula carries this vanishing to the integral applied to the corresponding average on the oscillator. Its terms are the same flag-character-corrected linear combination of $\omega(k^{-1}, 1)\phi$ for source lifts k. Averaging in the finite group then decomposes a test of the inverse level type into the common kernel of the oscillator mirabolic averages plus a sum in their images. For example, on the finite orbit span the averages are orthogonal projections after scaling by subgroup orders. The orthogonal complement to their common kernel is the sum of their images. Both determinant averages are killed by the good type · The source translation under the integral · The finite average decomposition.

The inverse-level-type tests in the common kernel are spanned by $U(W_j)(\mathbb{A}_{F,f})$-translates of tests with finite part ϕ_{f}. For this generation use the actual adjusted half of the split Schrödinger chart. At a type prime, let U be the residue of the saturated column space of X in the standard integral frame. The matrix has a right translate in the frame coset $X_0 + \varpi M_{N,r}(\mathcal{O})$, $X_0 = \binom{I_r}{0}$, exactly when U is the full top coordinate space $k_v^r \oplus 0$. Otherwise its intersection with the top space is contained in a proper hyperplane L. If $a \in \mathrm{GL}_r(k_v)$ fixes L pointwise, the difference between $\mathrm{diag}(a, 1)$ and 1 on U factors through the bottom coordinate. It is therefore realized on U by an upper-block unipotent matrix fixing the top space. Lift this matrix on an integral complement of the saturated column space, and use $\mathrm{diag}(\widetilde{a}, 1)$ on that column space. The result is invertible by its reduction and is in the source level with trivial character there. All terms in the corresponding hyperplane average on the Schrödinger value of X equal that value. The average vanishes on the common kernel, so the value is zero. On a frame coset, completing its first r columns by the standard bottom columns shows that the principal congruence kernel acts transitively. The surviving pieces are thus constant on the right translates of the frame coset. Two such translates that meet have the same coset, since their relative change of frame is congruent to the identity. On a compact support, their locally constant nonzero pieces therefore form a finite union. At a tuple prime the conjugate-half Schrödinger function transforms under the upper-block parahoric by $\xi^{-1}(\det h_{11})$ in the adjusted chart, where ξ is the nontrivial tuple character at that prime. For a matrix whose saturated column residue is outside the frame cosets, use the same hyperplane L and take a top block a fixing L with $\xi(\det a) \neq 1$. An upper-right block can cancel $a - 1$ on the column residue because it vanishes on the intersection with the top space. The resulting block matrix lifts to one fixing the original columns, by choosing its action on their saturated integral module to be the identity. The character on this lift forces the value to be zero. On each frame coset the covariance determines the coefficient up to its value at one frame, with multiplier $\xi^{-1}(\det X_{\mathrm{top}})$ in that coset's coordinates. These pieces are exactly the target translates of the installed ω^{Q} test, up to their nonzero translation scalars. At an unmarked split place use the integral rectangular lattice chart. After fixing a bound on denominators and a constancy modulus, a left integral general linear orbit is specified by its row image in the resulting finite module. Choose rows giving a basis modulo the uniformizer of that image. They generate by Nakayama's lemma, so the additional rows can be eliminated by row operations. Two such minimal generating lists are related by an integral change of basis invertible modulo the uniformizer. The indicators that this row image is contained in a given submodule span all functions of the row image by finite poset inversion. In the local field these indicators are the lattice indicators for the inverse images of those submodules. They are right translates of the rectangular lattice vacuum. At unmarked inert places use the spherical generation of the unramified Weil representation by the other member acting on the vacuum. In the restricted tensor product only finitely many places differ from the fixed finite part. Applying the local generation at those places gives the assertion. This time the translations act on the *target* W_j variable g in I_y. Vanishing for all g with the fixed finite factor therefore

kills these terms. Together with the source differences and mirabolic averages this proves vanishing for every test, contradicting the nonzero residue. We conclude

$$\ker(T|_{X(x)}) = 0.$$

Hyperplane cancellation outside the frame cosets · Generation in the first-block frame · The tuple test on the conjugate half · Determinant covariance and the rectangular tuple-frame support · The four finite-place cases on the same tensor product · The change of variables for target translates · Using the fixed column on an entire good eigenspace.

Exhaustion on a fixed point. Let Y_Σ be the orthogonal sum of all row parts in Y with parameter Σ. For every occurring row part, the adjoint is injective by (43) and has the unramified transferred label, obtained by adjoining the specified string and applying the common twist. At the chosen coefficient place these eigenvalues determine the generic point in the full record. Consequently $T^* Y_\Sigma \subseteq X(x)$. For any other row label, cancel the common string and reverse the twist at each permitted split place. Equality there would give the same rank-r isobaric base change by the transfer and split determination above. Its adjoint image is thus orthogonal to $X(x)$ for the actual source Petersson form. By adjunction, $TX(x) \subseteq Y_\Sigma$. The two injections give both dimension inequalities, so

$$T^* Y_\Sigma = X(x), \qquad TX(x) = Y_\Sigma.$$

The dimensions here include all occurrences of all automorphic types with the point's normalized character. The same argument also shows that every occurring good row part contributes to a good point at the given coefficient place. Transfer gives its prescribed global block and preserves the good first-block type labels. The joint eigenspace occurs in the full auxiliary space by injectivity of the adjoint and the complete residual decomposition. Hence its point is among the actual records at (j, v). Taking all row parts proves that the adjoint image is the sum of all full good eigenspaces in those records. Adding the string and determining the split point · Orthogonality against the other row parts · The two dimension comparisons · A good point at this coefficient place for every occurring row part.

The integral class reading at a twisted place. Fix j, v and the actual near tuple $\boldsymbol{\mu} = \boldsymbol{\mu}(v)$ as above. Let K be the coefficient field, O_v its completed DVR, and $e_v = \mathrm{ord}_v(p)$ for the rational prime p below v. Write X_K for the full auxiliary level space and Λ for the level lattice obtained by taking integral class values in the prescribed coefficient lattices. Thus X_K is the space before taking the theta or central good cut.

Use the same adelic class representatives g_c, right level J_V, and character as in the placed one-step column identity. In the coefficient convention of that identity a class value a_c extends at a point $\gamma\, g_c\, u$ by

$$\rho(\gamma) a_c\, \varepsilon(u), \qquad \gamma \in U(V_j)(F), \quad u \in J_V.$$

The finite stabilizer conditions on a_c are precisely the equations making this extension independent of γ, u. For the conjugate reading the equation is conjugated in its entirety, including the coefficient action and the character. The extension and evaluation at the g_c are inverse maps. In the complex class realization of (40), the level-character coordinates are inserted in the same principal-level model. Its level transports bring the coefficient lattice of the representative to the coefficient lattice at the translated class and multiply by the prescribed roots of unity. The character factor is inverted on reading back the class value. Thus on each orbit the original integral coordinates are exactly the representative value in its coefficient lattice. The comparison of class realizations preserves both directions of this criterion. The factors $|P|^{1/2}$ and $|P|^{-1/2}$ of the Hilbert level embeddings cancel in the algebraic reading. In particular the inverse image of the span of the prescribed integral class generators of the record model is exactly Λ. The normalized finite-place level lattices are obtained by localization in this same frame. With finitely many c, denominators prime to v in the separate coefficient values can be cleared by one such denominator. Equivalently the integral class-value condition is a kernel in the sum of coefficient modules, so it commutes with the flat

extension to a completed coefficient DVR. The same statement applies to the comparison Weyl lattices, with their own coefficient lattices at the same representatives. The localized class-value criterion · The level coordinates, Gram, and integral reading together · The reading of both spans by the same coordinate map.

Use the normalized split Hecke operators of (174), with the nonspecified half determined by the unitary condition. Their allowed places are away from the rational prime p, the fixed exceptional levels and the tuple levels. A right double-coset formula evaluates the same integral class values at the finitely many translated arguments. In the fixed coefficient action the translating matrices and their inverses are integral after inverting the rational prime below the translating place. This preserves the localized coefficient lattice at the two arguments. The normalizing power of the residue cardinality is a unit at v. Thus the Hecke operators act integrally on $\Lambda_v = O_v\Lambda$.

Extend O_v to a common complete splitting DVR R for these operators and the coefficient fields of the records at (j, v). Use the splittings of the given normalized Hecke algebra, and write Q_R for the fraction field. Write O_α for the DVR of the record M_α, completed when taking this common extension. The image algebra on $\Lambda_R = R \otimes_{O_v} \Lambda_v$ is finite over R, since it is an R-submodule of a finite matrix algebra. Its reduction has the Artinian decomposition by the residual Hecke kernels. Lifting the orthogonal idempotents in this complete algebra gives the projectors e_α onto the actual residual summands. All generic characters reducing to the same kernel belong to the same summand, and its integer lattice is $e_\alpha\Lambda_R$. One can equivalently do this at the coefficient DVR of the normalized record construction and base change to R, keeping all the idempotents above any factor together if it splits.

Here is the integral identification with the given record M_α. On the rational form of the record construction write i_V for the full class reading in $\Lambda_R \otimes_R Q_R$ and pr_α for the reading in the generic local factor. The comparison ψ_α from that factor to the split class space takes values in the summand $e_\alpha X_{Q_R}$. For each rational integral class generator a the difference $\psi_\alpha \mathrm{pr}_\alpha(a) - i_V(a)$ is in the sum of the other residual summands. Applying e_α gives the identity

$$\psi_\alpha \mathrm{pr}_\alpha(a) = e_\alpha i_V(a).$$

The model of M_α uses the span of the images of exactly these generators. The change of record over the same (j, v) preserves their set because it identifies the same stage and tuple with the same level and class coordinates. By the integral class reading above, the $i_V(a)$ span Λ_R. Taking R-spans proves

$$\psi_\alpha(M_\alpha \otimes_{O_\alpha} R) = e_\alpha\Lambda_R.$$

This also identifies the integral intersection in the α-summand by the same map. The argument allows several identical local generic eigenvalues on different automorphic occurrences. Residual kernels and the normalized characters · The projected integral class generators · The full lattice identity for each residual factor.

Spectral projection in the rational Gram frame. Let H be the rational Hermitian class-sum matrix on X_K, with term c weighted by $(|P||\Gamma_c|)^{-1}$ in the class coefficient forms. In the complex realization this is the source Petersson Gram up to the common positive factor given by the volume of the level. Indeed partition the adelic quotient into its right-level classes and divide the level volume of each by the order of its rational stabilizer. The factor $|P|$ converts the full principal-level to the character-level pairing in the placed coordinates. The same form is used on the projected stage in the auxiliary quotient Gram calculation. Distinct full automorphic isotypic spaces are therefore orthogonal for this H, by unitarity. The class identification above reads the rational H and these equations in the original coefficient frame with their indicated class weights. The class-sum and character-level Gram · Type orthogonality for the class Gram.

Over $\mathbb{C}$ let C_g be the central good orthogonal projector and $Z = (\ker T)^\perp \cap \mathrm{im}\, C_\mathrm{g}$. The fitted column identity gives $TC_\mathrm{g} = T$. Taking adjoints gives $Z = \mathrm{im}\, T^*$. The full good-fibre calculation above identifies

this image with the sum of the whole good generic Hecke spaces over all the residual factors at v. It also identifies $Z^{\perp}$ with the sum of the other full automorphic type spaces. In a fixed residual factor a labeled automorphic type has one generic character. If it is labeled in two residual factors at (j, v), there is an allowed normalized split operator whose evaluations have a unit difference in the common record DVR. Its eigenvalue on the type gives the same scalar in both complex readings. The comparison of coefficient embeddings across the two records sends their two evaluations to these scalars in the same field, so injectivity would make the difference zero.

The spectral statement descends with the actual coefficient embeddings. For clarity, the generic eigenspaces are common kernels of finitely many equations $T_u - \lambda_x(T_u)$ after adjoining the eigenvalues. A finite set of allowed operators separates the finitely many characters. These equations and the rational class coordinates agree in the complex and local readings. To check this on the α-summand, pull an eigenspace coordinate functional for x back to the full space along e_α. It satisfies the same eigenvalue equations on the dual since e_α commutes with the operators. The eigenprojectors and a basis of such functionals can be taken over the field generated by the algebraic eigenvalues over K, by row reduction. Extend these functionals to the complex class space using the record embedding extending the given embedding of K. If one of them is nonzero on a type, equivariance forces that type to have the functional's full normalized split character. Thus a type labeled in another residual factor gives zero projection to this one. The same check for every functional of every point shows that the rational good space maps into the local good sum and its orthogonal complement into the local nongood sum. They are complementary over the coefficient field. After scalar extension both inclusions are equalities. The column identity implies the source cut · The type description of the cut and its complement · Reading the split eigenvalues across the two realizations.

In the rational coordinates write E for the fitted basis matrix of the cut and B for the quotient coordinate projection of the column identity. Then $BE = 1$ and $E^{*}H(1 - EB) = 0$ in the coefficient field. Indeed the identities hold under the complex embedding, which is injective. Here the starred matrix uses the coefficient involution. These equations identify the rational projection $P_Z = EB$ and its complementary kernel for the same H. When v and its conjugate differ, the Hermitian pairing uses the conjugate coefficient module in the first slot, with the lattice at the conjugate place transported to v. The local orthogonality identities come from the stated rational matrix identities. Read them over the fraction fields of O_v and R on the same source lattices. In particular pairing a vector in the rational cut against a projected integral vector gives the same value as pairing against its integral lift. This is the orthogonal projection used in the stage Gram bound. Since both Z and $Z^{\perp}$ have the stated spectral readings, P_Z after extension is the generic spectral projector with kernel the entire nongood sum. It commutes with all residual projections. On the α-summand it is the exact projector P_α onto the full good generic sum in that factor. The fitted projection is the orthogonal projection · Both generic summands in the projector comparison.

The twisted length and its sum. The projected lattice at v is $P_Z\Lambda_v$ and the sublattice used in its Gram comparison is $\Lambda_v \cap Z_{K_v}$, where Z_{K_v} is the local rational cut. The intersection commutes with flat base change, by expressing it as the kernel of the map to X_{K_v}/Z_{K_v}. The projection commutes with base change as the image of the rational projection. Over R the quotient therefore splits as

$$\bigoplus_{\alpha \in \mathcal{A}_{jv}} \frac{P_\alpha(M_\alpha \otimes_{O_\alpha} R)}{(M_\alpha \otimes_{O_\alpha} R) \cap X_\alpha^{+}},$$

where X_α^{+} denotes the full good generic sum in that factor after scalar extension. A finite-length DVR module has a decomposition as a sum of cyclic modules. Base change multiplies its length by the relative ramification index and also multiplies the valuation of p by that index. Consequently the length divided

by the valuation of p agrees in the class DVR and the local record DVRs. With the δ_α of the actual Hecke record, this gives

$$\frac{\eta_{jv}^{\mathrm{cut}}}{e_v} = \sum_{\alpha \in \mathcal{A}_{jv}} \delta_\alpha = \eta_{jv}. \tag{120}$$

This includes factors with empty good or nongood part, for which the corresponding quotient has length zero. Local additivity with the normalized record valuations.

The actual $E_{2,j}$ is the weighted sum over twisted coefficient places of

$$\frac{\eta_{jv}^{\mathrm{cut}}}{D_0} + \mathrm{ord}_v(g_{j,\boldsymbol{\mu}(v)}) + \mathrm{ord}_v(s_j).$$

Here $D_0 > 0$ is the row rank in the Gram comparison, common to the actual tuples by the row-column comparison with the base level. The factor $g_{j,\boldsymbol{\mu}(v)}$ is $|P| \operatorname{lcm}_c |\Gamma_c|$ in the same class-value frame, and $s_j > 0$ is the stage depth of that comparison. Indeed the weighted class Gram pairing on the cut uses $g_{j,\boldsymbol{\mu}(v)}$ to clear these denominators and s_j for the chosen coefficient duality inclusion. The Weyl lifting exponent contributes to $E_{1,j}$. With $\lambda_v = \log|\kappa(v)|/[K : \mathbb{Q}]$, the first sum is at most η_j/D_0. The actual selected twisted places lie in the specified support of η_j, whose other terms $\lambda_v e_v \eta_{jv}$ are nonnegative. If g_{gr} is the specified common multiple of the stage group factors, use $\mathrm{ord}_v(g_{j,\boldsymbol{\mu}(v)}) \le \mathrm{ord}_v(g_{\mathrm{gr}})$. For any positive rational integer a the product formula gives $\sum_{v \text{ finite}} \lambda_v \mathrm{ord}_v(a) = \log a$. All terms in the two group and depth sums are nonnegative, so extending them to all finite places proves

$$E_{2,j} \le \eta_j/D_0 + \log g_{\mathrm{gr}} + \log s_j. \tag{121}$$

The place weights and the support of the sum · The local numerical contribution · Its weighted global estimate.

4.3.6 Weyl lattice comparisons on the auxiliary cut

Use the rank $N = r + t$ of the one-step auxiliary group here. For a dominant weight ν, realize the two coefficient modules as the integral divided-power lattice $\Delta(\nu)$ inside the polynomial dual Weyl lattice $\nabla(\nu)$ in the same rational space. If necessary add a determinant character to make the weight polynomial. Multiplication by this character identifies both lattice pairs and keeps the comparison exponent. Let $A_p(\nu)$ be the least exponent for which $p^{A_p(\nu)}\nabla_{\mathbf{Z}_{(p)}}(\nu) \subseteq \Delta_{\mathbf{Z}_{(p)}}(\nu)$. Over the local DVR O_v the actual comparison uses the tensors of these lattices in the same class-value frame for each factor ρ. Tensoring, taking the prescribed invariants at the level, and intersecting the rational cut gives

$$J_{jv}^* \le e_v J_A(p), \qquad J_A(p) = \sum_\rho A_p(\nu_\rho).$$

Here J_{jv}^* is in uniformizer units, as is seen from the largest Smith exponent of $D_{jt}^{\cap}$ in $L_{jt}^{\cap}$. Finite extension of the DVR multiplies this exponent and e_v by the same ramification index. The trace-dual coordinate frames in Δ and ∇ are calculated using the same selected embedding $\tau_\rho : E \hookrightarrow K$. The tensor comparison in the same frames, and The uniformizer exponent.

The parabolic comparison. Put $b_p(z) = \lfloor \log_p z \rfloor - v_p(z)$ for $z \ge 1$. At a cut after a of the weight coordinates,

$$A_p(\nu) \le A_p(\nu_{\le a}) + A_p(\nu_{>a}) + \sum_{i \le a < j} b_p(\nu_i - \nu_j + j - i). \tag{122}$$

Work over $\mathbf{Z}_{(p)}$ on the full flag of GL_N. For the standard block Levi let w_L be its longest element, and let w_0 be the longest element of the full group. Relative Kempf vanishing and Serre duality on the fibre of the parabolic projection identify the integral cohomology of $w_L \cdot \nu$ in degree $\ell(w_L)$ with the induced

Levi Weyl module. They identify $H^0(\nu)$ with the induced Levi dual Weyl module. Induce the primitive Levi inclusion and its scaled inverse. Both composites remain multiplication by $p^{A_p(\nu_{\le a})+A_p(\nu_{>a})}$. Choose a reduced extension of w_L to w_0 by left simple reflections of increasing length. At a step $w \mapsto s_\alpha w$, Andersen's integral reflection maps in degrees $\ell(w)$, $\ell(w)+1$ have both composites multiplication by $p^{b_p(\langle \nu+\rho, w^{-1}\alpha^\vee\rangle)}$. The roots added in this reduced extension are precisely the positive roots crossing the cut. At w_0 Serre duality gives the full Weyl lattice. The composite map from it to $H^0(\nu)$ has image containing the latter multiplied by the power in (122). It maps into the primitive polynomial Weyl lattice: it is equivariant, sends the top line to an integral multiple of the normalized top polynomial, and the source is generated from that line by integral divided powers. This gives the bound in the actual realization. [110, Propositions 1.15, 3.2 and 4.1] In particular for shifted width $w \ge 1$ the crossing cost is at most $a(N-a)\lfloor \log_p w \rfloor$. Induction gives $A_p(\nu) \le \binom{N}{2}\lfloor \log_p w \rfloor$. In rank one, including a zero shifted width, the exponent is zero.

The prime ranges and the improved allowance. Write $k_0 = k-1$, use the auxiliary rank $N = k_0+m+t$, and write $S_{\mathrm{cl}} = S_c$, $\eta_{\mathrm{cl}} = \eta_{\mathrm{sz}}$ for the cloud width and weight displacement. Let $h = b+S_{\mathrm{cl}}+\eta_{\mathrm{cl}}+t+1$, $H = \lfloor h \rfloor$. At a thick label ρ, in the shifted coordinates $U_{\rho,a} = \nu_{\rho,a} + N - 1 - a$ for $0 \le a < N$, the first k_0 entries are in a window $[x_\rho^-, x_\rho^+]$ of width at most S_{cl}. There are j_ρ entries before the central consecutive block of length t and q_ρ after it, counted after the first k_0 entries, with $j_\rho + q_\rho = m$. Put $B_\rho = U_{\rho,k_0+j_\rho}$, the top shifted entry of that block, $D_\rho = x_\rho^- - B_\rho$, and let $d_\rho^0 \in [1,p]$ be the representative of D_ρ mod p. Let $I_{\mathrm{win}} = 1$ if $[x_\rho^-, x_\rho^+]$ meets $B_\rho + p\mathbf{Z}$, and 0 otherwise. For a depth I put

$$C_\rho(p;I) = (q_\rho + 2j_\rho)\left(\frac{H+N}{p-1}+6\right) + A_p(\nu_{\rho,\le k_0}) + k_0\left(\frac{x_\rho^+ - x_\rho^-}{p} + I_{\mathrm{win}}\right) + \mathbf{1}_{d_\rho^0 < D_\rho}\min(k_0, t, I/d_\rho^0) + \frac{I}{p}(1+\log^+(H/p)).$$

The notation $\nu_{\rho,\le k_0}$ in this expression denotes the first k_0 coordinates.

Use the size profiles $X' = \exp(x^{\epsilon/3})$, $b' = xb/X'$, $n = f/2$, and define

$$\ell_J = \frac{mbL}{k(m+t)\log(3h)}, \qquad \ell_* = \frac{mb}{nt^2(k+m(t+1)+t)L^2 + kn^2(S_{\mathrm{cl}}+t+L)}, \qquad K_* = \lfloor 2b/\ell_* \rfloor.$$

The small or excluded primes use the full $J_A(p)$. The improved range is $\ell_J < p \le h$, excluding the designated at most five special primes and the primes dividing an integer in the interval

$$[\, b - (t-1)/2 - K_*,\ b + S_{\mathrm{cl}} + \eta_{\mathrm{cl}} + 1 + (t-1)/2 + K_* \,].$$

In that range $\ell_J^2 > 3b$, $b' = o(\ell_J)$ and $K_p = \lfloor H/(p-1) \rfloor \le K_*$. The first two follow since $h \le 2b$ and the other size factors are polynomial in x. The third follows from $\ell_J/\ell_* \gg L^2$. Enlarging the consecutive block by K_p at each end introduces no congruences with outside entries or their Frobenius-relabelled partners. In fact those outside the cloud have positive distance from the enlarged block by their separation $2b/\ell_* + t/2 + 3$, and distance less than p by their displacement bound η_{cl}. For a cloud entry the signed difference from an entry of the enlarged block belongs to $[\, e-(t-1)/2-K_p,\ e+S_{\mathrm{cl}}+(t-1)/2+K_p \,]$, where e is the cloud-to-block center separation. At a thick label $e \in [b, b+\eta_{\mathrm{cl}}+1]$. A prime dividing a difference there belongs to the excluded interval. At the other labels $e \ge b/(nX')$ and the distance is $O(b')$, so is strictly between zero and p. Reflection of the shifted arrays on the other CM half does not change these distance tests. Thus in this range use $I_\rho = \lfloor q_\rho h \rfloor$ at every thick label. Let $J_p(j)$ be the sum of the full exponents off the thick labels and of $\min(A_p(\nu_\rho), \lfloor C_\rho(p;I_\rho)\rfloor_+)$ on them, using $J_p(j) = J_A(p)$ outside the improved range. The actual two-branch table, The tests at the Frobenius-labelled windows, and The graded allowance.

Weights at the maximal ideals of the cut. Fix a stage and a twisted place $v \mid p$ in the improved range. Let N be the auxiliary rank, let O be the coefficient DVR of the class-value model at v, and write ϖ for its uniformizer. Write ν_ρ for the dominant polynomial weight at the label ρ, with coordinates numbered from 0. At a thick label put $s_\rho = k_0 + j_\rho$ and

$$U_{\rho,a} = \nu_{\rho,a} + N - 1 - a, \qquad d_\rho(w) = \sum_{a<s_\rho} (\nu_{\rho,a} - w_a).$$

The consecutive block of length t occupies $[s_\rho, s_\rho + t)$. Let $H = \lfloor h \rfloor$ be the integral bound for the shifted weights. At a maximal ideal of the Hecke algebra meeting the cut, every simple coefficient subquotient S of $\bigotimes_\rho \nabla_O(\nu_\rho)/\varpi(\bigotimes_\rho \nabla_O(\nu_\rho))$ with nonzero localized forms has a highest weight w_ρ at each label. We will prove that

$$\mathrm{pr}_{\rho,w_\rho} S \neq 0, \qquad d_\rho(w_\rho) \le q_\rho H \quad (\rho \in T_2). \tag{123}$$

Here $\mathrm{pr}_{\rho,w}$ is the projection onto row multidegree w in factor ρ, acting on S through the coefficient algebra.

The local automorphic model. Use the frame at the chosen v for the rank-N definite group, at its prescribed level and character. Its coefficient algebra A is generated by the product of the integral general linear groups and the row-multidegree projections in each factor. The level acts through the integral matrices of this frame times the level character. At the good places the Hecke generators act by finite right-translation sums with coefficients in O. They preserve forms valued in an A-submodule and commute with maps on coefficient values that respect the level action, provided both spaces use these same Hecke generators and normalizations.

The class set is finite since the group is definite. Each stabilizer has order prime to p in this range. In fact p is unramified in the CM host E and $p - 1 > N$. A stabilizer with an element of order p would have such an element in its faithful rank-N representation over E. A primitive p-th root has degree $p - 1$ over E, since $\mathbf{Q}(\zeta_p)/\mathbf{Q}$ is totally ramified at p. This rules out the element. Taking stabilizer invariants on the finitely many classes is therefore exact, by averaging.

Galois representations for the two weights. At each place above p the frame uses the full local stabilizer of a self-dual lattice. The level ideal is prime to that place and the level character of the selected twist tuple is trivial on its stabilizer. Indeed the generators of the level below the twist are supported at the large primes, and the additional level is carried by the conductor of the tuple character. The large primes are above the improved range. At p the fixed factor of the tuple character is unramified outside the special primes, and its varying factor is unramified at the places above the tuple's own p. It follows that evaluation at the same level and character realizes every characteristic-zero eigensystem of dominant weight w in a definite unitary representation with a hyperspecial fixed vector at each place above p. At an inert place the local extension over the real subfield is unramified and the unitary stabilizer of the self-dual lattice is hyperspecial.

Write $F = E^+$. The host extension E/F is unramified at every finite place, so $[F : \mathbf{Q}] \ge 2$. Labesse's Corollary 5.3 applies to this definite group. At infinity the group is compact, so every finite-dimensional irreducible type is in the discrete series and the trace sign condition $(*)$ holds. The weak base change is assembled from conjugate-self-dual discrete representations of general linear groups. It matches infinitesimal characters and the local base change at the good places. At a split place this is the given general linear component, and at an inert hyperspecial place it is the spherical base change. The Mœglin–Waldspurger description of the discrete spectrum writes each discrete term as a Speh representation. Its cuspidal seed is conjugate self-dual by uniqueness of the Speh data. At places above p the seeds are spherical: the spherical general linear representation at such a place has cuspidal support consisting of unramified characters, as do each of its translated Speh factors and their seeds.

At a complex embedding of the host, let λ denote the definite weight including the frame's determinant shift. The infinitesimal entries are $\lambda_a + (N-1)/2 - a$, which are pairwise distinct and belong to $(N-1)/2 + \mathbf{Z}$.

For each translated rank-r seed, write its entries as $z+\beta$, where z is an entry of the unitary seed and β is its Speh shift. Set $f = (N-1)/2-\beta$. The numbers $f-z$ are integers. Twisting the seed by $|\det|^{(r-1)/2-f_0}$ makes its infinitesimal entries distinct in $(r-1)/2+\mathbf{Z}$, so it is regular algebraic. This uses only dominance of w. The norm twist of the conjugate-self-dual seed is polarized, with the required archimedean sign obtained by the quadratic character of E/F if needed. The polarized attachment [121, Theorem 2.1.1(2)–(4)] gives its p-adic representation. It is crystalline at every spherical place above p, including the inert places.

Take the direct sum over the seeds and their Speh shifts. The local comparison of the attachments and of base change gives the good Hecke polynomials in degree N, with the common N-dimensional Satake normalization. The resulting Hodge–Tate jumps at an embedding are the numbers $(N-1)/2-(z+\beta)$, hence $a-\lambda_a$. Here the cyclotomic character has Hodge–Tate weight -1. If $\lambda_a = w_{\rho,a}+s^0_\rho$ at the archimedean half specified by the frame, the weight at the conjugate embedding is $-\lambda_{N-1-a}$. Tensor both systems with the crystalline algebraic character with jumps $H-N+1+s^0_\rho$ and $-s^0_\rho$ on these two embeddings, respectively. These integers have constant conjugate pair sum, so they are the jumps of an algebraic Hecke character of conductor prime to p. For existence, choose an auxiliary odd rational prime away from p at which the reduction of each nontrivial root of unity of E differs from 1. A unit congruent to 1 modulo this prime has quotient by its conjugate equal to 1, since the quotient is a root of unity. Its norm from F is 1 by reduction modulo the auxiliary prime. The balanced type is therefore trivial on these congruence units and defines a character of the principal ray ideals. Extend it to the group of ideals prime to the auxiliary prime using the finite ray class quotient. Its p-adic avatar is crystalline above p with the chosen jumps by local class field theory for algebraic characters. After the common twist the jumps at the conjugate of the archimedean half are the entries of U'_ρ. At the opposite half they are the reflected array $(H-U'_{\rho,N-1-a})_a$. The common twist and the common choice of inverse Frobenius multiply the two good Hecke systems by the same normalization factors.

For an eigensystem in the selected cut the prescribed consecutive Satake string has length t, centered at unit absolute value in the unitary normalization. Every unramified Satake root of a unitary cuspidal seed has absolute value 1. In fact purity in Theorem 2.1.1(2) shows that they have a common absolute value, and unitarity of the seed's central character fixes their product. A string endpoint of absolute value $|\mathcal{O}_E/\mathfrak{q}|^{(t-1)/2}$ thus requires a Speh term of length at least t. As $N < 2t$, its seed has rank one. Its infinitesimal entries form a consecutive run at every embedding. The prescribed run of length t is maximal and $N < 2t$, so that Speh run meets it and is contained in it. The Speh length is therefore exactly t and the jumps after the common twist are the prescribed block of U_ρ and its reflected block. Consequently the good Hecke polynomials of the cut are those of a complementary representation and t crystalline characters, with precisely these labeled jumps.

Lifting a surviving coefficient simple. The row-multidegree projections are orthogonal idempotents with sum the identity. The coefficient of z^b in an elementary root subgroup $1+zE_{ij}$, acting on the weight w, is the operator obtained by applying $1+E_{ij}$ and projecting to $w+b(e_i-e_j)$. Thus the projections and the integral general linear matrices generate the root coefficient operators integrally, including after reduction. A subquotient of the reduction stable for these operations is a rational module for the split product of general linear groups. Indeed the weight projections give the diagonal torus action, the coefficients give the root subgroup actions, and these generate on the open triangular cell. The rational submodule condition descends to the residue field by the resulting polynomial identities. Conversely a rational submodule respects the same projections and root coefficients.

By highest-weight theory S is the simple module of a dominant tuple w, and it embeds, compatibly with the torus and all the root subgroup actions, in the reduction of $\bigotimes_\rho \nabla_O(w_\rho)$. Its highest-weight projection is nonzero. Since the weights of each $\nabla_O(v_\rho)$ have total degree $\sum_a v_{\rho,a}$ and every coordinate in $[v_{\rho,N-1}, v_{\rho,0}]$, the same holds for w_ρ. This follows from dominance and Weyl group invariance of the weights. Hence every shifted entry of U' belongs to $[0,H]$ and $\sum_a U'_{\rho,a} = \sum_a U_{\rho,a}$.

The constituent also has the same shifted residues as ν_ρ, with their multiplicities. To see this, the coefficients of the column Capelli determinant

$$\operatorname{cdet}\bigl(z\,1 + (E_{ij}) + \operatorname{diag}(N-1,\dots,0)\bigr)$$

are integral expressions in the first root coefficient operators and the Euler operators $E_{aa} = \sum_w w_a \operatorname{pr}_w$. They act on a highest-weight vector of weight λ by the coefficients of $\prod_a (z + \lambda_a + N - 1 - a)$. The off-diagonal determinant terms have a raising operator to the right of all other off-diagonal operators, so they kill the highest-weight vector. By centrality the coefficients have these scalar actions on the entire characteristic-zero modules, and on their integral lattices by torsion-freeness. On S one can calculate them using either its realization as a subquotient at ν or its embedding into the reduction at w. The expressions agree on these realizations since root coefficient operators and the torus projectors agree. Thus

$$\{U_{\rho,a} \bmod p\}_a = \{U'_{\rho,a} \bmod p\}_a \quad \text{as multisets.} \tag{124}$$

The embedding at weight w gives an injection on forms. Taking the class values and averaging over their stabilizers shows that the forms with values in the reduction at w are the reduction of its lattice of forms. Thus this weight-w lattice has nonzero localization at the maximal ideal of the surviving forms. Its Hecke image is a finite O-algebra, since the forms lattice is finite free. The image is torsion-free over O and each maximal ideal in its support contains a minimal prime lying over 0. For example, a minimal prime containing ϖ would make ϖ nilpotent in the nonzero localization at that prime, contradicting injectivity of multiplication by ϖ. The corresponding generic character occurs on the characteristic-zero forms, since the forms are faithful over the Hecke image and finite over O. It is integral and reduces to the required maximal ideal. After a finite coefficient extension, one can align this reduction with the reduction of an eigensystem in the cut. Both reductions embed the same finite Hecke residue field into $\overline{\mathbf{F}}_p$. The change between them is a power of the residue Frobenius, which lifts to a coefficient Galois automorphism over $\operatorname{Frac}(O)$.

Use the Galois comparison on the two congruent systems. The common normalization gives congruent characteristic polynomials at the good degree-one Frobenius elements. By Chebotarev and continuity the reductions have the same characteristic polynomials on the whole Galois group. Brauer–Nesbitt identifies their semisimplifications, so the residual representation of the weight-w system has the reductions of the t crystalline characters as composition factors, with multiplicities. The same statement holds when restricting to the places at p, by restricting a stable filtration or refining it there. Equation (124) also holds at the conjugate labels, since reflection preserves the equality of residue multisets.

The crystalline depth bound. Fix the completion K of the CM host at the place inducing the label ρ. This is a finite unramified extension of $\mathbf{Q}_p$, including when the place is inert in the host over its real subfield. Write $\sigma i = i \circ \mathrm{Fr}_K$ for arithmetic Frobenius on the embedding labels, and write the t line weights at i as $u_{i,b}$, $0 \le b < t$. They form the consecutive interval $[b_i, b_i + t - 1]$. We use the Hodge–Tate convention in which the cyclotomic character has weight -1.

Apply Bhatt–Gee–Kisin [113, Theorem 5.6.2] to the crystalline representation furnished by the weight-w eigensection at K. It gives decreasing integer arrays $\lambda_i \uparrow U'_i$ and Weyl permutations for the inertial characters of the semisimplified reduction. The order in that theorem is the affine order for the ordinary action of the Weyl group and translations by p times the coroot lattice. In particular, λ_i and U'_i have equal totals and equal residue multisets modulo p, and the top partial sums of λ_i are bounded above by those of U'_i. It follows that all entries of λ_i lie in $[0, H]$, and the sum of the smallest t entries of U'_i is at most the sum of any t distinct positions of λ_i.

Take a multiple d of $[K : \mathbf{Q}_p]$ for which the combined Frobenius and Weyl permutation of labels and positions has order dividing d. Put $i_n = \sigma^n \rho$. In the fundamental-character formula of Bhatt–Gee–Kisin, (5.3.10) and (5.6.1), each character at i_0 has exponent $-\sum_{n=0}^{d-1} p^n \alpha_n$ for a fundamental character of

niveau d. The digit α_n is the entry of λ_{i_n} at the n-th position of the corresponding permutation path. Indeed evaluating the n-th geometric Frobenius action on the induced cocharacter at the identity reads the embedding $\sigma^n\rho$, with the Weyl action permuting the positions. Local class field theory gives the same exponent formula for the reduction of the crystalline line b, with digits $u_{i_n,b}$. The t line characters in the residual representation can be matched to distinct positions at i_0, with multiplicities. The corresponding digits $\alpha_{n,b}$ occupy t distinct positions at every i_n, since each step is a permutation.

For $\delta_{n,b} = \alpha_{n,b} - u_{i_n,b}$, equality of these characters gives $\sum_{n=0}^{d-1} p^n \delta_{n,b} \equiv 0 \pmod{p^d - 1}$. Every cyclic shift of the sum is also divisible by $p^d - 1$, so the numbers

$$\kappa_{n,b} = \frac{\sum_{\ell=0}^{d-1} p^\ell \delta_{n+\ell,b}}{p^d - 1}$$

are integers, with indices modulo d. Since $|\delta_{n,b}| \le H$, they satisfy

$$|\kappa_{n,b}| \le K_p = \left\lfloor \frac{H}{p-1} \right\rfloor, \qquad \alpha_{n,b} = u_{i_n,b} + p\kappa_{n+1,b} - \kappa_{n,b}.$$

In particular $u_{i_n,b} - \kappa_{n,b}$ has the residue of $\alpha_{n,b}$, and it belongs to the enlarged interval $[b_{i_n} - K_p,\, b_{i_n} + t - 1 + K_p]$.

At $n = 0, 1$ this interval has fewer than p integers and no residue of a shifted weight outside the line block, by the window bound at ρ and at its Frobenius successor. The window at the successor uses the reflected array if the successor lies on the conjugate half. The residues of λ_{i_n} are those of U'_{i_n}, hence also those of U_{i_n} by (124). The t chosen positions at i_n therefore give the t block residues once each. Their representatives in the enlarged interval are unique, so $\sum_b \kappa_{0,b} = \sum_b \kappa_{1,b} = 0$. The partial-sum inequality above now gives

$$\sum_{a=N-t}^{N-1} U'_{\rho,a} \le \sum_{b=0}^{t-1} \alpha_{0,b} = \sum_{a=s_\rho}^{s_\rho+t-1} U_{\rho,a}.$$

The bottom $t + q_\rho$ entries of U'_ρ thus have sum bounded by the original line-block sum plus $q_\rho H$, and the other q_ρ original entries are nonnegative. Since the full arrays have the same total, the prefix depth is at most $q_\rho H$, which proves (123). In particular the depth is at most the table's $I_\rho = \lfloor q_\rho h \rfloor$.

The integral graded Weyl bound. For a dominant polynomial weight λ write $\Delta(\lambda) \subseteq \nabla(\lambda)$ for the Weyl and dual Weyl lattices in their common realization, now over $\mathbf{Z}_{(p)}$. Write $A_p(\lambda)$ for the least exponent of this inclusion and pr_w for the torus projection onto weight w. For the thick-label shape at ρ and $H < p^2$, we prove that if $d_\rho(w) \le I$, then

$$p^{\lfloor C_\rho(p;I) \rfloor_+} \mathrm{pr}_w \nabla_O(\nu_\rho) \subseteq \Delta_O(\nu_\rho). \tag{125}$$

Here $C_\rho(p; I)$ is the label cost in the exponent table with depth input I, so $C_\rho(p) = C_\rho(p; I_\rho)$. The graded estimate concerns every torus weight of the dual Weyl lattice that satisfies this bound on depth.

Insertion of a constant block. First let $\lambda = (a^t, \kappa)$ have a *leading* constant block, where a is a nonnegative integer and κ is a dominant polynomial weight of rank r with $\kappa_0 \le a$. Work over $\mathbf{Q}$ in the simple highest-weight module $V(\lambda)$, and write V_i for the part where the first block has total weight $ta - i$.

This AI-proposed candidate concerns the two integral lattices on each weight of V_i. For fixed i, form $D_i(T)$ as the monic least common multiple of

$$P_\tau(T) = \prod_{(b,c)\in\tau/\kappa} (T - (c - b)),$$

where row and column indices start at zero, $|\tau/\kappa| = i$, $\ell(\tau) \le r$, each skew column has at most t boxes, and $P_\tau(a) \neq 0$. An empty least common multiple is one. For a weakly decreasing tuple ξ on the full set of coordinates, write $\partial(w) = \xi \cdot (\lambda - w)$. Suppose p^{e_L} takes the dual Weyl weight lattice of κ into the Weyl lattice on all its weights w' with $\partial((a^t, w')) \le I$. For every weight w of V_i with $\partial(w) \le I$, the claimed integral inclusion is

$$p^{\,e_L + v_p(D_i(a))} \nabla(\lambda)_w \subseteq \Delta(\lambda)_w.$$

It holds at every prime, including at singular values of the scalar parameter a. The same statement for a trailing block follows by reversing the contragredient and applying a determinant twist. Rank zero on the remaining block is the determinant representation. Nakahama's norm calculation gives the rational scalars used below. The precise integral inclusion with the inherited Levi cutoff and without a projector factorial is the additional statement whose priority remains open.

Put $W = V_0$. Its first Levi factor acts by the determinant to power a, and its second factor is $V(\kappa)$. Write $E_{\alpha b}$ and $F_{b\alpha} = E_{b\alpha}$ for the mixed raising and lowering operators, where α is in the first block and b in the second. The mixed lowering operators commute, as do the mixed raising operators. The former generate V from W, and the latter kill W. These facts follow from the ordered basis of the enveloping algebra with the mixed lowering operators first, applied to a highest-weight vector.

Use the contravariant form normalized to give the standard highest-weight vector norm 1. In this pairing the dual of the divided-power Weyl lattice is the dual Weyl lattice in the square-matrix polynomial model. Here is the integral identification. For a rational polynomial vector f, pairing against the standard highest vector gives $f(1)$. Evaluation at the identity kills the other row weights by right Borel equivariance, and the standard highest vector takes value 1 there. Pairing against a word of lowering divided powers therefore gives the coefficient of the transposed raising word acting on f and evaluated at the identity. These are integral for an integral polynomial. In the other direction, the coefficients for ordered words are exactly the coefficients on the product of the lower unitriangular root groups in those coordinates. That product parametrizes the lower unitriangular group integrally, by solving for entries in increasing distance from the diagonal. If the restriction of f to that group is integral and $p^a f$ is a primitive integral polynomial for some $a > 0$, its reduction would vanish on the lower unitriangular group. The column homogeneity and right upper unipotent identities then give vanishing on the dense lower-unitriangular times upper-triangular cell over $\overline{\mathbf{F}}_p$, a contradiction. This proves the claimed duality over $\mathbf{Z}_{(p)}$. The restriction of the form to W is the corresponding normalized form on $V(\kappa)$. The integral ordered divided powers give

$$\Delta(\lambda) \cap V_i = \sum_{|\mathbf{m}|=i} F^{(\mathbf{m})} \Delta(\kappa), \qquad F^{(\mathbf{m})} = \prod_{b,\alpha} \frac{F_{b\alpha}^{m_{b\alpha}}}{m_{b\alpha}!}.$$

Indeed, the Weyl lattice is generated by integral lowering divided powers on the highest-weight vector. Put the mixed ones first and then those of the two Levi groups. The nontrivial lowering divided powers in the constant block kill the highest-weight vector. The ones in the other block generate $\Delta(\kappa)$.

Let X and Y be the standard spaces of ranks t and r, and let $Z = Y \otimes X^*$, with basis indexed by (b, α). For a degree-i monomial $\mathbf{m}$ in this basis put $\mathbf{m}! = \prod_{b,\alpha} m_{b\alpha}!$. There are rational maps

$$M_i : \ \operatorname{Sym}^i Z \otimes W \longrightarrow V_i, \qquad z^{\mathbf{m}} \otimes u \longmapsto \prod_{b,\alpha} F_{b\alpha}^{m_{b\alpha}} u,$$

$$B_i : \ V_i \longrightarrow \operatorname{Sym}^i Z \otimes W, \qquad v \longmapsto \sum_{|\mathbf{m}|=i} z^{\mathbf{m}} \otimes \frac{\prod_{b,\alpha} E_{\alpha b}^{m_{b\alpha}}}{\mathbf{m}!}\, v.$$

The first map is onto. In the second, every coefficient belongs to W by its depth. The map B_i is injective, by contravariance and the span of the vectors $F^{(\mathbf{m})}W$. If $v \in \nabla(\lambda) \cap V_i$, then every coefficient of $B_i(v)$ is in $\nabla(\kappa)$, by pairing against $\Delta(\kappa)$. The two maps preserve weights for the full diagonal torus, including the determinant factor on W.

The operator B_iM_i has an integral polynomial family. In the homogeneous polynomial ring on the upper mixed block write $z_{\alpha b}$ for the variable dual to $E_{\alpha b}$ and put

$$\mathcal{D}_{b\alpha}(T) = Tz_{\alpha b} - \sum_{c} z_{\alpha c}E^{W}_{bc} - \sum_{\gamma,c} z_{\gamma b}z_{\alpha c}\partial_{\gamma c}.$$

Here E^{W}_{bc} acts on the $V(\kappa)$ factor of W. For $B(v)(z) = \mathrm{pr}_W \exp(\sum_{\alpha,b} z_{\alpha b}E_{\alpha b})v$, one has $B(E_{\alpha b}v) = \partial_{\alpha b}B(v)$ and $B(F_{b\alpha}v) = \mathcal{D}_{b\alpha}(a)B(v)$. The second identity follows by conjugating $F_{b\alpha}$ by the exponential and applying the projection to W. The constant $F_{b\alpha}$ term projects to zero, the first commutator acts by the two diagonal blocks, and the second commutator divided by 2 is the displayed quadratic derivative term. Higher commutators vanish. On V_i, the Taylor coefficients of B give B_i. Define $\overline{\mathcal{A}}(T)(z^{\mathbf{m}}u)$ to be $\prod_{b,\alpha}\mathcal{D}_{b\alpha}(T)^{m_{b\alpha}}u$. The differential operators commute, as is also seen from the mixed lowering action on the open cell for arbitrary T. Their product defines a monic operator polynomial of degree i preserving the integral coefficient lattice and the diagonal weights. In particular

$$B_iM_i = \overline{\mathcal{A}}(a). \tag{126}$$

The scalar of $\overline{\mathcal{A}}(T)$ on a Y-type $V(\tau)$ is the following polynomial:

$$P_\tau(T) = \prod_{(b,c)\in\tau/\kappa} (T - (c - b)),$$

where row and column numbers start at 0. To obtain this from Nakahama [114, Theorem 1.2], use the holomorphic model of $\mathrm{SU}(t, r)$ with fiber $1 \boxtimes V(\kappa)$ and parameter ℓ. Its character on the block diagonal is $\chi(\mathrm{diag}(A, D)) = \det A$. For $g = 1 + \epsilon F_{b\alpha}$, the upper cell coordinate of $g^{-1}\exp(w)$ is $w + \epsilon wF_{b\alpha}w$ to first order, with diagonal blocks $1 + \epsilon wF_{b\alpha}$ and $1 - \epsilon F_{b\alpha}w$. Thus in the coordinate $z = -w$ the differential actions are $\partial_{\alpha b}$ and $\mathcal{D}_{b\alpha}(-\ell)$ for $E_{\alpha b}$ and $F_{b\alpha}$, respectively. For ℓ real and sufficiently large the invariant Bergman form h_ℓ pairs $F_{b\alpha}$ against $-E_{\alpha b}$, by real conjugation on this unitary group. Its restriction to the constant fiber is the normalized fiber form. The Fischer form h_{F} uses the ordinary Gaussian on matrix coordinates. Its squared monomial norm is $\mathbf{m}!$. For homogeneous g of degree i it follows that

$$h_\ell(\overline{\mathcal{A}}(-\ell)z^{\mathbf{m}}u, g) = (-1)^i h_{\mathrm{F}}(z^{\mathbf{m}}u, g).$$

Indeed repeated differentiation of g at zero gives its $z^{\mathbf{m}}$-coefficient times $\mathbf{m}!$. Nakahama computes h_ℓ/h_{F} on every $V(\tau)$-summand as $\prod_{(b,c)\in\tau/\kappa}(\ell + c - b)^{-1}$. The formula holds for either order of the ranks. This gives the stated scalar for $T = -\ell$, and then for all T by polynomial identity. In particular it applies at $T = a$ even when a norm ratio has a pole there. On the symmetric coinvariants every component also has the following column bound. Every τ/κ occurring there has at most t boxes in each column. In fact the characteristic-zero Cauchy decomposition of $\mathrm{Sym}^i(Y \otimes X^*)$ has only partitions with at most t parts on Y, and a Littlewood–Richardson tableau for the tensor product with $V(\kappa)$ then uses at most t entries in any strict column.

Let $D_i(Z_0)$ be the monic least common multiple of the polynomials $P_\tau(Z_0)$ specified in the statement. The empty least common multiple is 1. Its evaluation at a is the nonzero integer

$$D_i(a) = \prod_{c\neq a}(a - c)^{m_c}, \qquad m_c = \max_{\tau} \#\{\text{added boxes of content } c\}.$$

By (126), the image of B_i is supported on the summands whose polynomials are nonzero at a. Divide $D_i(Z_0)$ by the monic operator polynomial $\overline{\mathcal{A}}(Z_0)$, taking the polynomial quotient $\mathcal{Q}(Z_0)$. Its coefficients

are integral polynomials in the coefficients of $\overline{\mathcal{A}}$, so they preserve the indicated lattice and weights. On each surviving summand that division is the ordinary exact division by P_τ. Consequently

$$\overline{\mathcal{A}}(a)\mathcal{Q}(a)B_i(v) = D_i(a)B_i(v).$$

Equations (126) and the injectivity of B_i give $M_i\mathcal{Q}(a)B_i(v) = D_i(a)v$.

More generally, consider a weakly decreasing tuple ξ on all coordinates and the depth condition $\xi\cdot(\lambda - w) \le I$. Every W-coefficient of $\mathcal{Q}(a)B_i(v)$ for a homogeneous vector of weight w has a weight obtained from w by adding the raising weights for the corresponding monomial. These are nonnegative for ξ. If p^{e_L} sends the dual Weyl lattice to the Weyl lattice on all such depth-$\le I$ Levi weights, with the constant block fixed to a, it sends the coefficients into $\Delta(\kappa)$. Since M_i is ordinary lowering, it sends these integral coefficients into $\Delta(\lambda)$. This proves the bound $e_L + v_p(D_i(a))$ at the stated weight. Each factor of $D_i(a)$ prime to p is a unit in the coefficient ring.

Two integral tests in degree one. Let $t = r = 2$ and $p \ge 7$. For $\lambda = (a, a, p-1, 0)$, $a = 2p-1$, use the weight $(a, a-1, p-1, 1)$. In the basis indexed by $e_1 \otimes x^{p-2}y$, $e_2 \otimes x^{p-1}$ with one active $X^\vee$ coordinate,

$$B_1M_1 = \begin{pmatrix} a-p+2 & -(p-1) \\ -1 & a \end{pmatrix}.$$

The degree-one PBW generators are independent and span, so B_1 identifies the dual lattice with the dual of the Levi coefficient lattice in these coordinates. Here the binomial coefficients in the Levi pairing are units at p. The two eigenvalues are p and $2p$. The determinant is $2p^2$ and an entry is a unit, so the largest Smith exponent is two. Retaining only the maximum of the two specialized valuations would give one.

The Levi and content factors may both occur in the same invariant factor. For $\lambda = (a, a, p, 0)$, $a = 2p$, at $(a, a-1, p-1, 2)$, the basis indexed by $e_1 \otimes x^{p-2}y^2$, $e_2 \otimes x^{p-1}y$ gives

$$B_1M_1 = \begin{pmatrix} p+2 & -(p-1) \\ -2 & 2p-1 \end{pmatrix}.$$

Its eigenvalues are $p, 2p+1$, and its Smith exponents are zero and one. The two Levi coefficient factors are $\binom{p}{2}, \binom{p}{1}$, each of valuation one. In the dual coordinates B_1 the Weyl lattice therefore has the comparison matrix $B_1M_1\,\mathrm{diag}(\binom{p}{2}, \binom{p}{1})$. Its largest exponent is two, the sum of the Levi and content exponents. Both tests have shifted width less than p^2.

Trailing blocks and one-row bounds. The same statement for a constant block at the end follows by reversing and complementing weights with a common integer $M \ge \lambda_0$. The automorphism $g \mapsto w_0{}^t g^{-1} w_0^{-1}$, with w_0 the reversal permutation, followed by the determinant twist, sends $\lambda_b \mapsto M - \lambda_{n-1-b}$, where n is the rank. It transports the two lattices by the highest vector and integral divided powers, and preserves their normalized pairing. It changes ξ to the negative of its reversal, again weakly decreasing. For a trailing block of length t and value a, the first block in the reflected weight has value $M - a$. Its depth i is the excess of the sum of the weights on the original trailing block over ta.

In the leading-block calculation, if a shape has no added box of content a, all its added boxes have content less than a. Otherwise the row containing a larger content already contains content a, because it starts from a row of length at most a. The number of boxes of any one content is at most the maximum column length of the skew diagram. For any boxes of the same content, the leftmost of their columns meets all their rows in the skew diagram, since the inner partition is decreasing. For a trailing one-row insertion each content thus has multiplicity at most one in D_i. If the new weight is a and the preceding rank is r, the differences $(M-a) - c$ from the reflected content satisfy $1 \le (M-a) - c \le \lambda_0 + r - a - 1$. This uses $c \ge M - \lambda_0 - (r-1)$. Hence $D_i(M-a)$ divides $(\lambda_0 + r - a - 1)!$. Under a shifted width bound H, the Legendre valuation inequality pays at most $\lfloor (H-1)/(p-1) \rfloor$ per inserted row. This assertion also holds when the rank starts at zero, since a one-dimensional determinant factor has no cost.

The consecutive block in the given weight. At a fixed thick label omit the subscript ρ, and put $s = k_0 + j$. Start with the first k_0 coordinates, with exponent $e_{\mathrm{cl}} = A_p(\nu_{\le k_0})$. Adjoin the following j rows using the one-row bound without a depth cutoff. Insert the constant block of length t at the end of rank $s + t$, and then adjoin the final q rows, retaining the same cutoff with $\xi_a = \mathbf{1}_{a<s}$. On restriction to rank $s + t$, this depth is the sum of the weights on the last t entries minus ta, by equality of total weight. It is nonnegative on every weight of the module by dominance. Thus for the constant block it suffices to keep D_i for $0 \le i \le I$. The estimates so far give

$$e_{\mathrm{cl}} + (q + j)\left\lfloor \frac{H-1}{p-1} \right\rfloor + \max_{0\le i\le I} v_p(D_i(M-a)).$$

D_i here is computed with the reflected first s rows as its inner partition. All the ranks used in this induction have shifted width at most H.

Let $B = U_s$ be the top shifted entry of the constant block in the full weight, and let $[x^-, x^+]$ be the first k_0 window. For an added box in row b and column c of the reflected skew diagram, write $\kappa'_b = M - \nu_{s-1-b}$, set $a' = s - 1 - b$, and assign the label

$$x = B + (M-a) - (c-b) = U_{a'} - (c - \kappa'_b + 1).$$

This equality includes the common shift from the last q rows. The labels satisfy $B < x < x^+$ and hence $x - B \le H$. Their multiplicities in the least common multiple are the maxima on individual shapes. At each label $x < x^-$, at most j boxes are drawn from the j rows below the first window. At $x \ge x^-$ none of those rows contributes because its shifted entry is below x^-. The remaining multiplicity is at most $\min(k_0, t)$, and for $x < x^-$ it is also at most $i/(x^- - x)$. Indeed, each counted row has an added box at step $U_{a'} - x \ge x^- - x$, so uses at least that many boxes. These row and column arguments hold for each shape and give the same separate upper bounds for the multiplicities in the least common multiple. The j low rows cost at most

$$j\, v_p((x^- - B - 1)!) \le j\,\frac{H-1}{p-1}.$$

Since $0 < x - B < p^2$, the valuation of $x - B$ is 1 at multiples of p and zero elsewhere. In $[x^-, x^+]$ there are at most $(x^+ - x^-)/p + I_{\mathrm{win}}$ such labels, where I_{win} indicates that the interval meets $B + p\mathbf{Z}$. Below x^- the distances $x^- - x$ at such labels are $d_0 + np$, where $1 \le d_0 \le p$ represents $x^- - B \bmod p$ and $n \ge 0$. The contribution for $n = 0$ is at most $\min(k_0, t, I/d_0)$ when $d_0 < x^- - B$. For $n \ge 1$, sum $I/(np)$ for at most $\lfloor H/p \rfloor$ terms. The harmonic sum is bounded by $I(1 + \log^+(H/p))/p$. Thus the integer exponent is at most

$$e_{\mathrm{cl}} + (q + 2j)\frac{H-1}{p-1} + k_0\left(\frac{x^+ - x^-}{p} + I_{\mathrm{win}}\right) + \mathbf{1}_{d_0 < x^- - B}\min(k_0, t, I/d_0) + \frac{I}{p}(1 + \log^+(H/p)).$$

Replacing $(H-1)/(p-1)$ by $(H+N)/(p-1) + 6$ gives $C_\rho(p; I)$ from the table. The integer exponent is bounded by its floor. Finally both polynomial lattices and the torus multidegree projections commute with base change to O. This can be seen from the integral standard monomial basis of ∇ and the divided-power generators of Δ. This proves (125).

The inclusion of the localized form lattices. Keep the class-value frame at the chosen $v \mid p$. For a coefficient submodule C let $\mathscr{F}(C)$ be the forms valued in C at the prescribed level and character. All the Hecke operators below use the same finite translation formulas in this frame. Put

$$B_\rho = \Delta_O(\nu_\rho), \qquad V_\rho = \nabla_O(\nu_\rho), \qquad V = \bigotimes_\rho V_\rho.$$

For any set S of labels, let D_S be the mixed tensor sublattice of V using V_ρ for $\rho \in S$ and B_ρ otherwise. The coefficient algebra A preserves each D_S, since each of the two integral lattices at each label is stable under the integral group action and the row-multidegree projections. The latter are weight projections of the algebraic torus, with a different idempotent for each weight as an integer tuple. For $\rho \in T_2$ put

$$c_\rho = \min\bigl(A_p(v_\rho), \lfloor C_\rho(p; I_\rho)\rfloor_+\bigr).$$

For the other labels put $c_\rho = A_p(v_\rho)$.

We show for every maximal Hecke ideal $\mathfrak{m}$ meeting the selected cut that

$$p^{c_\rho}\mathscr{F}(D_{S\cup\{\rho\}})_\mathfrak{m} \subseteq \mathscr{F}(D_S)_\mathfrak{m} \qquad (\rho \notin S). \tag{127}$$

When $c_\rho = A_p(v_\rho)$ this follows from the coefficient inclusion on factor ρ. Otherwise ρ is thick. Set $D = D_S$ and $B = D + p^{c_\rho}D_{S\cup\{\rho\}}$. It is enough to show that every form with values in B takes values in D after localization at $\mathfrak{m}$. Both lattices are A-stable, and B/D has finite length because D contains a power of ϖ times V. For example the full tensor coefficient inclusion sends V into $D_\varnothing \subseteq D$ after multiplication by $p^{\sum_\rho A_p(v_\rho)}$.

Consider a simple A-quotient of any adjacent terms of a composition series of B/D. It is isomorphic to a simple subquotient M of $V/\varpi V$. To see this, intersect the terms of the filtration with the images of $\varpi^n V$ in V/D. The induced finite filtration of a simple quotient has a single nonzero step, and the step is a subquotient of $\varpi^n V/\varpi^{n+1}V \simeq V/\varpi V$.

Suppose $\mathscr{F}(M)_\mathfrak{m} \neq 0$, and let w_ρ be its highest weight at ρ. By (123) its depth is at most I_ρ and the projector $e(\rho, w_\rho) \in A$ acts nontrivially on M. On V this is exactly the actual coefficient operator pr_{w_ρ} on factor ρ, with identity on the other factors. By (125) it maps $p^{c_\rho}D_{S\cup\{\rho\}}$ into D. It preserves D, so it annihilates B/D, and every one of its subquotients. This is a contradiction. Thus the forms with values in every simple factor of B/D have zero $\mathfrak{m}$-localization.

For an adjacent pair of terms of the coefficient filtration, the quotient on values gives a Hecke-equivariant map of forms whose kernel is the forms taking values in the smaller term. Exactness of localization, applied at that kernel, gives equality of the two localized form modules when the quotient forms vanish there. Iteration gives $\mathscr{F}(B)_\mathfrak{m} = \mathscr{F}(D)_\mathfrak{m}$ inside the localized form module with coefficients in V. This argument uses kernels and Hecke equivariance for the quotient maps. It acts by the torus projector on coefficient values in the proof of vanishing, with no need for a projector on the Hecke module of forms. This proves (127). Iterating over the labels proves

$$p^{\sum_\rho c_\rho}\mathscr{F}(V)_\mathfrak{m} \subseteq \mathscr{F}(D_\varnothing)_\mathfrak{m}.$$

We read this inclusion back on the selected lattice. Let $F \in L^\cap$, with integral class values in V in the frame, and let L be the submodule of $\mathscr{F}(V)$ generated by F under Hecke. The image of the Hecke algebra on L is finite and torsion-free over O. Its generic characters occur in the prescribed cut by the Hecke equivariance of the class-value map and stability of the rational cut. Thus each maximal ideal in the support of L is the reduction of an occurring generic character, by the minimal-prime argument for finite torsion-free Hecke images. The same maximal ideals cover the support of the image of L modulo $\mathscr{F}(D_\varnothing)$. The displayed localized inclusion makes the class of $p^{\sum_\rho c_\rho}F$ vanish at each such ideal, so it vanishes globally. The class-value test for the comparison lattice, and descent through the finite coefficient extension of DVRs, now give

$$p^{\sum_\rho c_\rho}L^\cap \subseteq D^\cap, \qquad J^*_{jv} \le e_v \sum_\rho c_\rho = e_v J_p(j).$$

Here $L^\cap$ and $D^\cap$ are the intersections of the dual Weyl and Weyl level lattices with the same rational cut. This is the improved prime branch in the table. At the other primes the full coefficient inclusion gives its stated branch directly.

The weighted cost of the comparison. The exponent just obtained satisfies

$$\sum_{v \text{ twisted}} \lambda_v J_{jv}^{*} \le \sum_{p} J_p(j) \log p = O(fmbL^2).$$

The first inequality uses $\sum_{v|p} \lambda_v e_v = \log p$. For a nonthick label its shifted width is $O(b')$. The full exponent from (122) and $\sum_{p^s \le u} \log p = O(u)$ give $O(fN^2 b') = o(b)$ for these labels. The prime-power estimate follows from the central binomial coefficient in dyadic intervals, adding the higher powers. For a thick label the full cost at $p \le 2\eta_{\mathrm{cl}} + t$ is $O(N^2(\eta_{\mathrm{cl}} + t + 1)L)$. For $2\eta_{\mathrm{cl}} + t < p \le \ell_J$, cut after k_0 entries in (122). The lower block has shifted width at most $2\eta_{\mathrm{cl}} + t$, so its exponent is zero. The full cost from the upper block is $O(k^2 S_{\mathrm{cl}})$ and the crossing cost is at most $k(m+t)\ell_J \log h \le mbL$. Beyond h the full exponent is zero. The designated special primes cost $O(N^2 L)$ per label. Each integer in the excluded interval has at most one prime divisor $> \ell_J$, since $\ell_J^2 > 3b$. Here $K_* = O(x^C(S_{\mathrm{cl}} + 1))$ for a fixed C by the formula for ℓ_*. The log cost on this interval is $O(fN^2 L(S_{\mathrm{cl}} + \eta_{\mathrm{cl}} + t + K_* + 1)) = o(b)$, using $S_{\mathrm{cl}} + \eta_{\mathrm{cl}} \le b/X'$.

In the improved range the upper block's exponent is zero. The first term of $C_\rho(p; I_\rho)$ costs $O(mbL)$, since $\sum_{p \le h} (\log p)/(p-1) = O(1 + \log h)$ by partial summation. The cloud width term costs $O(kS_{\mathrm{cl}}L)$. For I_{win} sum over the integers in its window. The prime divisors of any $u - B_\rho$ there have total log weight at most $\log h$, giving $O(k(S_{\mathrm{cl}} + 1)L)$. For the d_ρ^0-term group the primes by $d = d_\rho^0 < D_\rho$. They divide $D_\rho - d$ and have total log weight at most $\log h$. Their coefficient is at most I_ρ/d, and $I_\rho \le 2mb$. The harmonic sum gives $O(mbL^2)$. The final term contributes at most $I_\rho \sum_{p \le h} (\log p)/p \, (1 + \log^+(H/p)) = O(mbL^2)$. Together these prove the asserted bound. The sum of the two prime branches.

A weighted bound from the rational spectrum. The division estimate only needs the weighted bound, which also follows from the content spectrum and integral multiples of the Schur–Weyl projections. On an improved thick label reflect with $M = \nu_{\rho,0}$. The inner partition κ' of the reflected block has k_0 rows of length at most S_{cl} and another j_ρ of length at most H. The surviving degrees are $i \le q_\rho H$, so put $R_\rho = |\kappa'| + q_\rho H = O(mb)$. In every outer diagram of at most R_ρ boxes the multiplicity of content c is at most $\lfloor R_\rho/(1 + |c|) \rfloor$, by counting their distinct rows or, for $c < 0$, distinct columns. [115, Proof of Lemma 7.4] If $a' = M - a$ is the reflected block value, the nonzero content differences $a' - c$ in surviving diagrams are positive and at most H. They therefore give a common divisor bound by the integer

$$Q_\rho = \prod_{\substack{|c| < R_\rho \\ 1 \le a' - c \le H}} (a' - c)^{\lfloor R_\rho/(1+|c|) \rfloor}.$$

On the degree-i Levi tensor lattice the usual monomial and symmetric quotient realizes it as an integral quotient of $Y^{\otimes n} \otimes (X^\vee)^{\otimes i}$, $n = |\kappa'| + i$. Schur–Weyl gives the rational projections by the central idempotents of $\mathbf{Q}[\mathfrak{S}_n]$, whose denominators divide $n!$ by the character formula. The quotient has no torsion, so these integral multiples preserve the integral kernel and descend. Thus $R_\rho! Q_\rho$ clears the inverse scalar times each surviving projector on $B_i(V_i)$. The weight cutoff and the one-row insertions from the preceding proof give the exponent bounded by $v_p(R_\rho! Q_\rho((H-1)!)^{j_\rho + q_\rho})$ on this label. The absolute log of this integer is $O(mbL^2)$ by the harmonic sum in the exponent of Q_ρ and the factorial bound. Summing valuations by the product formula gives the same weighted order as above. The stronger factorial-free statement concerns the finer integral weight comparison for general shapes, with the indicated specialization.

4.3.7 Transport of the auxiliary stage lattice

The reference kernel at the own-prime support. At an own-prime split place in the chosen half, write k for its residue field, $q_0 = |k|$, r for the target rank and $s = r + t_{\mathrm{cut}}$ for the source rank. The two integral Schrödinger coordinates in the paired frames are $X \in \mathrm{M}_{s \times r}(k)$ and $Y \in \mathrm{M}_{r \times s}(k)$. Use the trace pairing $\mathrm{tr}(YX)$ with the primitive additive character $\mathbf{e}$ obtained from the uniformizer section. In the reference-kernel comparison the residue function on the conjugate half is zero unless $X = \binom{A}{0}$ with $A \in \mathrm{GL}_r(k)$,

where its value is $\xi(\det A)^{-1}$. The multiplicative character ξ is the nontrivial residue character of the fixed own-prime determinant correction. The target flat kernel on the other half has value $\xi(\det B)$ when the first r columns B of Y are invertible, and zero otherwise. Thus the character in this comparison belongs to the fixed reference kernel. The two residue matrix functions · Nontriviality of the tame correction.

To calculate the finite Fourier transform, put

$$S_\xi(B) = \sum_{A \in \mathrm{GL}_r(k)} \xi(\det A)^{-1} \mathbf{e}\big(\mathrm{tr}(BA)\big).$$

When B is invertible, the substitution $A' = BA$ gives $S_\xi(B) = \xi(\det B) S_\xi(I)$. When B is singular, choose $D \in \mathrm{GL}_r(k)$ equal to the identity on $\operatorname{im} B$, with $\xi(\det D) \neq 1$. Replacing A by AD leaves the trace unchanged because $DB = B$, and multiplies the character by $\xi(\det D)^{-1}$. Therefore $S_\xi(B) = 0$ in this case. Applying the square-matrix Fourier transform twice gives

$$S_\xi(I) S_{\xi^{-1}}(I) = q_0^{r^2} \xi((-1)^r)^{-1}.$$

Here the double transform is $q_0^{r^2}$ times evaluation at the negative matrix, by additive orthogonality. Both factors on the left are algebraic integers. Hence $S_\xi(I)$ and its inverse are integral away from the residue characteristic, and the same holds for $q_0^{-sr} S_\xi(I)$. For the rectangular Fourier expansion, the coefficient at $-Y$ is exactly $q_0^{-sr} S_\xi(B)$: the bottom block of X vanishes and only B pairs with its top block. Products of these equalities over the own-prime places give a scalar G_0 and its inverse integral away from p_o. Rectangular finite Fourier inversion.

In the marked completed pair lattice, the two residue coordinates are those of the lattice and its adjoint frame. A lift of Y on the chosen half is an element of the uniformly dilated lattice with reduction Y, and its pairing against a lift of X gives $\mathbf{e}(\mathrm{tr}(YX))$. Such lifts exist by lattice reduction in the chosen uniformizer frame. The integral lattice inside the dilated one is the kernel of reduction. Consequently, when the Fourier expansion is interpreted as a sum of Heisenberg translates of the lattice vector, evaluation at a point t in the dilated lattice on the chosen half selects exactly the coefficient at minus the reduction of ϖt. Heisenberg equivariance and the normalization at the half-indicator give the resulting half-function

$$t \longmapsto G_0 \, \Phi_{\mathrm{flat}}(\varpi t)$$

in the same Schrödinger model. The terms at the two type primes remain the prescribed type factors, since those places are disjoint from the own-prime support. This calculation uses the paired self-dual frames for the actual kernel on the chosen run. The common type and flat frames · Heisenberg translates in the completed pair lattice.

Apply the central source element z with factor $\varpi^{-1} I_s$ on the chosen half, its dual inverse on the conjugate half, and identity outside this support. The split Schrödinger action pulls t back to $\varpi^{-1} t$. Its scalar is the prescribed splitting character at the determinant of z times the square root of the index of the affected lattice. The index is a power of p_o. The finite-order part of the splitting character is a root of unity. For each of the unramified algebraic characters in the splitting, take a power making the supporting prime ideal principal. On a generator the local value to that power is the product of the algebraic infinity-type factors and a finite-order level value. The generator of that ideal and the normalization by its absolute norm are units away from p_o, so the character value and its inverse are integral away from p_o too. Thus $\omega(z, 1)\phi$ is the prescribed reference flat kernel times a nonzero scalar unit away from p_o. Theta covariance moves z onto the finite source argument as $h \mapsto hz^{-1}$. Translations at the type level in the target commute with it, since they act on the other member of the dual pair. They use the same character in the two readings. The source translation can change the coefficient lattice only at places above p_o, where a power of p_o clears the denominator. Central dilation in the same model · Splitting-character unit calculation · Covariance of the finite-level trace.

The pointwise comparison on the marked cover. Fix a run, its own prime p_o, an adapted source lattice Λ at h, and a short-section rule $F = F_{\Lambda,N,x_0,P}$ of the selected weight κ. The ranks are r on the target and $s = r + t_{\mathrm{cut}}$ on the source, with s even. Use the marked cover of this very rule, of level Mp_o over the level M. At a target argument g, take the dictionary point through a rational representative γg, its transported tangent frame, and the dual lifts of its finite marking in the symplectic period basis. This is the same framed fibre for the oscillator and the algebraic trace evaluation. The isomorphism of coefficient rings for this cover reads the polynomial coefficients via the distinguished embedding of the base field. The same marked fibre and lifts.

The cover coefficient factors into its old-level component and its own-prime component by the Chinese remainder theorem. On integral columns use the split residue maps ρ_P from the adjoint marked frame. At an old type place the condition is that their reduction equals $y_{0,P} = \rho_P(p_o x_0)$. Multiplication by p_o is invertible there and the admissibility of x_0 gives $y_{0,P}$ full column rank. At an own-prime place only the M-multiple component contributes. The first r vectors of the level marking are a basis exactly when the corresponding top square matrix A_P in the adjoint coordinate is invertible at every own-prime place. This follows from their generation of the free marked quotient at each factor of the Chinese remainder decomposition. The weight of the cover section vanishes on the other markings. On a generating marking compare those first r vectors with the pinned marking. The unique change of basis of the marked quotient has adjoint acting on the dual half. Its character is the product of the residue determinant characters ξ_P. The normalized weight consequently has value a fixed root of unity times $\prod_P \xi_P(\det A_P)^{-1}$. The root is the product of the determinant-character values at the pinned marking. This proves the equality $\widehat{\phi}_{\mathrm{free}} = \zeta\, w_{x_0}$, where w_{x_0} is the actual cover weight and ζ is a root of unity in the prescribed field. In forming the Fourier transform here one uses the same dual marking as in the geometric evaluation. The lattice reduction described above determines the translate with weight w_{x_0}, and the common Fourier denominator is kept in both evaluations below. The determinant character of the pinned basis · The old and own-prime components of the marking.

Write $\mathcal{T}(g)$ for the common weighted theta derivative in these coordinates, with polynomial P renamed by the tangent-frame labels, and with finite input $\widehat{\phi}_{\mathrm{free}}$. It is a finite sum of translated theta derivatives over the isotropic kernel and the marked level classes. Under the common identification their phases are the Weil pairings of the same two marking halves. Indeed orthogonality of the characters in the dual marking leaves the cosets prescribed by the cover weight. The polynomial derivative is unchanged by the quadratic exponential change between compatible complements, since the contractions annihilate the harmonic P. If d is its homogeneous degree, the dilation in the algebraic rule multiplies the derivative by N^d, canceled by its N^{-d} normalization. Thus the two evaluations use $\mathcal{T}(g)$ with the same polynomial and finite terms. The algebraic value in the classical theta frame · The same derivative in the oscillator.

Put $I = [\Lambda : \Lambda \cap \Lambda_{\mathrm{std}}]$, let m_Λ be the inner index $mIdx(N, r)$, and let C be the cardinality of the level marking used in the Fourier formula. The outer index is $z_\Lambda = m_\Lambda I^r$, by multiplying indices in

$$N\Lambda_{\mathrm{std}} \subseteq \Lambda \cap \Lambda_{\mathrm{std}} \subseteq \Lambda.$$

The geometric evaluation of F is

$$c_0'\, \chi(x_g)\, \beta_H(\Lambda)^{-1} \frac{z_\Lambda}{m_\Lambda C}\, \kappa'(g)\, \zeta^{-1} \mathcal{T}(g),$$

where x_g is the determinant-root idele of γg, χ is the fixed own-prime correction, and $\kappa'(g)$ is the ratio of the two pure-tensor vacuum realizations at this point. The factor ζ^{-1} converts the Fourier coefficient to the weight w_{x_0}. The ratio is independent of Λ, N, P. The factor $c_0'\chi(x_g)$ is the value of the cover's zeroth-weight determinant section in this frame. At the base point the cover coefficient is c_0'. Transporting the section multiplies it by $\chi(x_g)$ because the determinant root gives its action on the marking. The base factor c_0 and this c_0' belong to the period field and have a common positive power equal to one, by the

finite order identities on the two determinant sections. The two finite indices and the cover value · The product of indices · The base determinant section · The same calculation on the own-prime cover.

In the adelic evaluation the fixed harmonic map contributes the common factor ka. Let c_g denote the product of the normalized component constant and the base splitting factor at g, and let $\ell_\infty = (2\pi\xi_{\mathbf{R}})^{-1/2}$. The oscillator formula at the fixed lowest-type vector reads

$$\mathrm{ka}\,(-\delta')^d\,\ell_\infty^d\,c_g\,\chi(x_g)\,\kappa'(g)\,\chi_{\mathrm{split}}(\det_{\mathrm{half}} h)^{-1}\,\frac{\sqrt{I^r}}{C}\,\mathcal{T}(g).$$

Here δ' is the sign of the prescribed frame orientation. The factors ℓ_∞^d and $(-\delta')^d$ come from differentiating the Gaussian in those frame coordinates. The finite scalar is the unit-vacuum normalization of the changed source lattice: by the self-dual lattice model the finite vacuum overlap has modulus $I^{-r/2}$, and the split Weil covariance adds the displayed inverse determinant character. Rational equivariance of the theta sum permits using γg in place of g. The fixed harmonic map and both splitting factors are in these same coordinates. The harmonic normalization in the pointwise formula · The finite scalar and the Gaussian derivative.

In the geometric comparison with the phase-one read the pointwise value above is multiplied by $\mathrm{ka}(-\delta)^d\,\ell_\infty^d c_g c_0^{-1}$, with δ the base orientation sign. Since $z_\Lambda/m_\Lambda = I^r$, this value equals the adelic value multiplied by

$$(\delta\delta')^d\,\frac{c_0'}{c_0}\,\zeta^{-1}\beta_H(\Lambda)^{-1}\chi_{\mathrm{split}}(\det_{\mathrm{half}} h)\,\sqrt{I^r}.$$

This is an equality of multiples of the same derivative, including when $\mathcal{T}(g) = 0$. The central level move in the fixed phase-one realization is handled on the section by its inverse on the same marked cover. It translates x_0 by the prescribed level action, preserves the other short-rule data, and gives an additional finite-order scalar. On the fixed base short rule the common marked phase is (46), with the retained phase witnesses in the same framed fibre. Its determinant ideal is read by (48). This identity on the geometric class generators extends to the rational cusp lattice by linearity. Thus the scalar for the actual rational cusp reading is the displayed scalar times the prescribed finite-order factors. These are elements of the coefficient field with zero valuation at every finite place. Equality of the normalized multiples · The central translation on the marked rule · The phase-one comparison on the generators.

For the remaining class scalar, the split matrices of h have elementary exponents $b_1, \dots, b_s$ at a chosen half and $-b_1, \dots, -b_s$ at the opposite half. Their contribution to I is the residue size to $\sum_i |b_i|$. Let $\mathfrak{c}$ be the numerator ideal of the norm-one determinant, with $\mathfrak{c}$ and $\bar{\mathfrak{c}}$ coprime. Its local norm exponent is $|\sum_i b_i|$ at this pair. Since r is odd, $\mathrm{N}\mathfrak{c}/I^r$ is a positive rational square. Write it as q_h^2. On the split determinant idele, the restriction of the splitting character to fixed-field ideles is trivial, and its components at the supporting places are unramified. Thus its value is the ideal character of $\mathfrak{c}$, including the extra factor $\delta_{p_o}(\mathfrak{c})$ of the present reference kernel. In the prescribed reading the base character and $\sqrt{\mathrm{N}\mathfrak{c}}$ are read as the single algebraic coefficient $B(\mathfrak{c})$ of the split class. Consequently the class scalar is the image of

$$\beta_H(\Lambda)^{-1}B(\mathfrak{c})\,\delta_{p_o}(\mathfrak{c})/q_h$$

in the period field. This reads the square root jointly with its character. At a place where Λ is adapted, its determinant lattices and those of $\Lambda\cap\Lambda_{\mathrm{std}}$ both become standard after extension to every labelled coefficient factor. The defining generator $\beta_H(\Lambda)$ of the quotient of their weighted determinant fractional ideals is therefore a unit there. Also $\mathfrak{c}$ and I are prime to the residue characteristic. The algebraic infinity-type character on $\mathfrak{c}$ and the norm factor are units at v, by raising the ideals to principal powers. The scalar for the trace is therefore v-integral. This proves the pointwise comparison used in the off-own-prime trace for the actual marked orbit. The numerator ideal of the determinant · The joint ideal and square-root coefficient · The adapted genus scalar.

The short-section trace off the own prime. Fix $v \mid p$ with $p \neq p_o$ and use the corresponding geometric run at weight κ. A generator of its trace lattice comes from a short-section orbit f_u indexed by $u \in U$. In the adapted realization all its sections use one harmonic polynomial P, one source lattice Λ, and the marks $x_u = \ell(u)x_1$. There is a common relation $s_0 f_u = a_0 f'_u$ with a_0 integral and s_0 a v-unit. Here Λ is the source lattice adapted at v, and its conjugate is the lattice at the finite source argument h. In particular it is standard above p. This is the simultaneous change of the short-section orbit: localization is performed before taking the U-sum. The common scalar is unaffected by a further change of this orbit by a root of unity. The adapted orbit on the same cover.

Read the harmonic coefficient as w on the coefficient space of h. There is a v-unit s_1 making $s_1 w$ integral there. Indeed the two source lattices in the adapted realization agree at all primes over p. Choose an integer a prime to p clearing the comparison in both directions. A power of a clears all denominators in transferring P to the standard harmonic lattice, by polynomial homogeneity. The integral standard harmonic-to-polynomial comparison puts it in the dual-Weyl coefficient lattice. Transferring back to the class of h costs only another power of a, by the same covariance for the polynomial lattice. All these operations preserve the polarized weight and the fixed harmonic map. The prime-to-p coefficient comparison.

At the fixed lowest-type vector, the short-section evaluation on the marked chart gives

$$\mathrm{read}(f'_1) = c\ \mathrm{Series}_{\phi_{\mathrm{free}}}(h, w)$$

in the prescribed phase-one rational coordinates, with $c \in K$ integral at v. The input to this formula is the normalized equality of the geometric and analytic class evaluations. The factor c contains the genus-lattice scalar, the determinant character and the square-root index for Λ. In the marked frames, cancelling the common cover cardinality leaves the class scalar

$$\beta_H(\Lambda)^{-1}\, \chi_{\mathrm{split}}(\det_{\mathrm{half}} h)\, \sqrt{[\Lambda : \Lambda \cap \Lambda_{\mathrm{std}}]^r},$$

multiplied by the signs and root-of-unity ratios of the normalized frames. The square-root factor is algebraically integral. The determinant idele has a representative for its fractional ideal prime to p since Λ is standard there. The unramified algebraic character on that ideal can have denominators only at its support, by the principal-power calculation, and the tame factors on units are roots of unity. The genus scalar has valuation zero at an adapted prime by the marked isogeny comparison for the trivialized polarization bundle. Recovering the class scalar in the period field uses the prescribed algebraic readings of the simultaneous character and square-root factor. A norm-one frame scalar alone gives no valuation assertion. The scalar in the short-series equality · Cancellation of the cover count · The character and square-root reading · The genus scalar at the adapted prime.

The level translation in the marked short rule is the action by u^{-1} on the target read. Its determinant character is trivial on U. This gives the same scalar c for every f'_u. It extends from the prescribed vector to every vector of the lowest type because the respective compact groups generate the lowest-type factors from that vector, and both series are equivariant there. Thus reindexing $u \mapsto u^{-1}$ yields

$$F' = \sum_{u \in U} \psi(u)\, \mathrm{read}(f'_u) = c\ \mathrm{tr}_{U,\psi}\, \mathrm{Series}_{\phi_{\mathrm{free}}}(h, w).$$

The trace on the right uses $\psi(u)^{-1}$ multiplying the finite action of u. This is also the orientation of the character eigenspace in the transport calculation. The marked level translation · Extension through the lowest type · The unnormalized trace.

The free and reference kernels differ by translation on the source and a scalar c' with c', c'^{-1} integral away from p_o. At the own-prime places the change of variables scales the lower $s - r$ rows in the adjoint coordinate and leaves its top block fixed. Integrality of the scaled lower rows removes the lower-zero condition after reduction. The determinant character remains on the same top block. The scalar from the

split linear action is its determinant splitting factor times a power of the local residue size. It is a unit away from p_o, by the same calculation as for central dilation. At a type place a full-rank $s \times r$ matrix is obtained from the first r columns of the identity by an element of $\mathrm{SL}_s(k)$: extend the columns to a basis, then change the last added column to make determinant one. This is possible since $s > r$. Lift the matrix integrally at the split half and take its dual inverse at the paired half. There is no determinant scalar from this operation. It preserves the coefficient lattice and the reference vacuum normalization at the other places. Combining these translations gives one finite k with $\phi_{\mathrm{free}} = c'\omega(k,1)\phi_{\mathrm{ref}}$. For some $m \geq 0$, $w' = p_o^m s_1 w$ belongs to the coefficient lattice at hk. The lower-block change at the own prime · The determinant-one move at the type primes.

If $c = 0$, faithfulness of the lowest-type reading makes the trace zero. Otherwise define $F_0 = p_o^m s_1 (cc')^{-1} F'$ in the rational cusp space. Covariance of the theta kernel identifies its reading with $\mathrm{tr}_{U,\psi}\,\mathrm{Series}_{\phi_{\mathrm{ref}}}(hk, w')$. Because w' is integral at hk, this is a generator of the reference class-series lattice Θ. The original trace generator is

$$\frac{a_0}{s_0}\,\frac{cc'}{p_o^m s_1}\,F_0.$$

All displayed denominators are units at v, and c, c' are v-integral. This proves the localized trace inclusion $\Theta_0 \subseteq \Theta$ using the pointwise marked short-section comparison and its scalar integrality. The trace is the unnormalized U-sum. The integral scalar for the trace.

Membership of the reference generators. Use the common coefficient reading at v, with the fixed reference δ_{p_o} and the same level character as the base tests. The preceding own-prime comparison identifies the reading of a reference class generator, before coefficient normalization, with a scalar times the test kernel at hz^{-1}. The scalar and its inverse are integral away from p_o. A power of p_o clears the change of the source coefficient lattice. It therefore suffices to verify membership for that translated harmonic coefficient.

The same polynomial coefficient coordinates are used for both readings. In those coordinates the source dual-Weyl lattice at hz^{-1} is the base change of the test coefficient lattice. One can see this by using an integral basis for the underlying source lattice on each labelled factor, taking the polynomial generators in that basis, and tensoring their spans. The normalized harmonic map identifies its coefficient polynomial with the original one, with factor c_0 from the fixed polynomial frame. The class values are read after multiplication by ν^{-1}, with c_0/ν a root of unity in the phase-one normalization. Thus the coefficient used in the test generator is in its integral O'-span, and the extra scalar c_0/ν is an algebraic unit. This is the coefficient-class identification on the selected one-step model. The harmonic coefficient in the test lattice · The common normalization factor.

Choose a finite representative h' of the double class of hz^{-1} supported at split places outside the finite excluded set. For this existence, the determinant lies in the norm-one torus. Lift its finite idele by $x \mapsto x/\bar{x}$. The preimage of the determinant of the level contains a ray congruence subgroup. A prime ideal in the ray class of x, away from the finite set and split over the real field, gives an idele supported on the corresponding split pair with the required determinant class. Such a prime exists by ray class reciprocity and Chebotarev. The prime ideals of the CM field of residue degree greater than one over the real field have density zero, and the ramified places can be excluded. The rational norm-one determinant is realized by a rational unitary element acting on a nondegenerate Hermitian line. The remaining determinant-one idele is moved to one additional allowed split place by strong approximation for the special unitary group, using its noncompact split factor at that place. The subgroup to approximate at the other finite places is the conjugate of the prescribed open level by the determinant representative. This gives the required finite double class, including the original level.

Writing $hz^{-1} = \gamma h' k$ in that class, replace the coefficient by its translate under γ^{-1}. The finite class covariance carries the integral polynomial lattice to the one at h'. The right level element k contributes its finite-order character value. Hence this move realizes an admissible reference test coefficient times an

algebraic unit. It commutes with applying the target trace and good-type projector, since those act in the target coordinates. Split representative with the prescribed level · Covariance of the test datum.

The original ν^{-1}-normalized analytic generator is therefore an algebraic scalar u, integral away from p_o, times the reading of an O'-linear combination of the admissible tests. If that combination is nonzero, division of a nonzero coordinate in the common faithful basis puts u in the common fraction field. Integrality away from p_o places it in O' because p_o is a unit at the selected place. If the combination is zero, faithfulness gives a zero generator. Taking the spans and applying the same projector on the two coefficient models gives $O'R_v(e\Theta) \subseteq B_0$. The comparison uses the fixed geometric coefficient frame and the exact normalized read established for this reference tuple. The one-sided integral base comparison.

The two geometric section comparisons

The numerical choices common to the two section runs. The two geometric comparisons below use the explicit row arrays and embedded geometric field in the following design. To apply the comparisons at the indices of the period and trace constructions, require that these exact rows and this field be retained by both marked runs on the same chosen records, including i_0 and every $i(u)$ used in the assembly. This identification is the separate condition described in Section 6.1. The following arguments establish the estimates on the identified objects.

Use the translated witnesses with the explicit Harish–Chandra rows and the initial embedded coefficient field of the common construction. The formulas below concern that realization. Let $L = \log b = x^{\varepsilon}$, $Z_b = frL$, and $r = k + m - 1$. In this calculation $D = f\binom{r}{2}$ is the dimension of the full flag space, independent of the stage. Write $\mathrm{DF}(w)$ for the tableau degree of a weight w, and set $d_L = \mathrm{DF}(w_L)$. The weight size $\mathrm{S}(w)$ is defined in (117). For either own prime $p_o = p_{\mathrm{de}}, p'_{\mathrm{de}}$ the base level is $Q = g_0 g_1$, and the short rules use its prescribed refinement at p_o. Put $q_* = \min(g_0, g_1)$. The rank formula in Section 1 specifies

$$T = \lceil x^{377} \rceil, \qquad n_W = A_x, \qquad n_P = 8fr + 4, \qquad n_\kappa = r + t = Tn_W + n_P.$$

Here the copy numbers of w_L, w_F are n_W, n_P. The evenness of k, m, n_κ makes r, t odd. Set $\mathcal{P} = C_P Z_b^{12}$, with $C_P \geq 1000$ chosen in the first group of constants of the common selection. Keep it fixed when selecting x. Then $C_P^{20} \leq Z_b$ and $C_P \leq x^{1/10}$ eventually. The factor 10 in A_x is fixed. Rank and order prescriptions, The fixed exponent twelve, and The initial choice of C_P.

The original choice of Harish–Chandra rows uses two independent polynomial spacings. Denote them by $\Delta_{\mathrm{cl}} = \lceil x^{u_{\mathrm{cl}}} \rceil$ and $H_{\mathrm{node}} = 2\lceil x^{u_{\mathrm{node}}} \rceil$. Take the first exponent at least 805 and then the second at least 1000, increasing them to any other fixed lower bounds of the common construction. The word columns at each CM label have the same base values

$$v_a = 4\Delta_{\mathrm{cl}}(k - 2 - a), \qquad 0 \leq a \leq k - 2,$$
$$b_{\tau,l} = B_\tau + lH_{\mathrm{node}}, \qquad e_{\tau,l} = b_{\tau,l} + H_{\mathrm{node}}/2, \qquad 0 \leq l < m.$$

The word bits control their interlacing entries. The three choices for B_τ are $b_t + \lfloor b'/x \rfloor$ on the top labels, $\lceil b \rceil$ on the thick labels, and $\lceil w'' \rceil$ on the remaining labels, with the root-of-unity correction at one thick label. It can be made with a shift between zero and $|\mu(E)| - 1$ since that label embeds $\mu(E)$ faithfully. For the classical host $\mu(E) = \{\pm 1\}$, by the conductor and inertia of its quadratic and real dyadic factors. The parallel increments preserve the congruence since f is even. Set $\eta_{\mathrm{wt}} = mH_{\mathrm{node}}$. Then the rounded starts and last nodes stay in the three prescribed intervals, and $\log \eta_{\mathrm{wt}} = O(\log x) \leq x^{\varepsilon/4}$. The positive scales $xb, b'/x, w''$ eventually exceed every fixed power of x. In particular the top first entry exceeds the first cloud entry, and $\lfloor b'/x \rfloor \geq k$ puts the top last entry above $-B_\tau$. At the other labels the extreme entries bracket the cloud and $-B_\tau$ by construction. The interior entry at position $a+1$ is $v_a - 1/2$ or $v_a - 1/2 - \Delta_{\mathrm{cl}}$ according to the bit, both between v_a and v_{a+1}. These formulas give interlacing simultaneously for every word, using the already installed primes and the prescribed a_t, b_t. The original selection makes this choice

before tying the rows to the period data. The original cloud and node choices and The simultaneous choice for all words.

To check the uniform scale, the classical rank brackets give eventually

$$f \le x^{2.17}, \qquad k, r \le x^{4.71}, \qquad x^{.49} \le m \le x, \qquad L \le x^{.02},$$
$$Z_b \le x^{6.90}, \qquad D \le x^{11.59}, \qquad n_W \le x^{23}, \qquad \mathcal{P} \le x^{83}.$$

Also $Z_b \ge x^{6.5}$ and $m^4 \le f$. These estimates give $T \ge 40(D+2)$, $T \ge 11D^2$, and $T \ge (D+2)n_W + 2r$. They also give $T \ge r + 2(1250r)(D+1)$. Directly from A_x, $n_W \ge fr(r+1)$, $n_W \ge f(k-1)+2r$, and $n_W \ge n_P$. With the thickness allowance $d_2 \ge x^{-\varepsilon}$ the lower bound $\mathcal{P}^4 D^3 kfx/(md_2) \le T$ has left side at most x^{375}. All fixed polynomial upper bounds in $T, \mathcal{P}, D, f, r$ are eventually below $b^{1/2}$, including $\mathcal{P}^3 T^2$ and $TD^2 L$. The selected Δ_{cl} dominates $(T+D+2)(T+D+2+2N_m)$ with $N_m = k + m(t+1) \le x^{402}$. The selected H_{node} dominates $4T^2 + 2t + 8$ and the higher products needed in the numerical allowance. For example, $\mathcal{P}Dfm^2 N_m T \le x^{882}$ and $CTN_m f^2 m^2 kD\mathcal{P} \le x^{891}$ for any fixed C and sufficiently large x. The third product $\mathcal{P}(fmk)^3 k/m$ is at most x^{117}. The lower bound for the order and the adjacent rank bounds · The cloud-spacing comparison and The polynomial comparison for the node spacing.

The weight floors and the division margins. In signature $(p, \rho) = (k+j-1, m-j)$ the two tableaux of κ are obtained by cutting the following row after p entries.

$$(v_0, \dots, v_{k-2}, -b_{\tau,0}, \dots, -b_{\tau,m-1}).$$

On the first block add $e_{\tau,j-1}$ and subtract the staircase $p-1-a$ and the scalar $(t+1)/2$. On the second block, in reverse order with $\nu = \rho - 1 - a$, the entry is $b_{\tau,j+\nu} - e_{\tau,j-1} - (t+1)/2 - \nu$. Thus the last entry in each nonempty block is $H_{\mathrm{node}}/2 - (t+1)/2$. The within-block consecutive gaps are $4\Delta_{\mathrm{cl}} - 1$, $H_{\mathrm{node}} - 1$, or $B_\tau - 1$. In particular these gaps, increased by two, are at least $(T+D+2)(T+D+2+2r)$. The two block orders and the staircase.

Put $B_T = T + D + 2$. A within-block gap γ of κ becomes the gap $\ell = \lfloor (\gamma+2)/B_T \rfloor$ of w_L. Each last entry a becomes $\ell = \lfloor (a - A)/B_T \rfloor$, where $A = (D+2)n_W + 2r$ on the first block and $A = 0$ on the second. Summing the gaps from the last entry gives the two tableaux of w_L. Subtract Tw_L from κ to get w_F. All floors are positive. In the selected chambers the last entries of w_L grow at least as x^{600}, since $H_{\mathrm{node}} \ge 2x^{1000}$ and $B_T \le 2x^{378}$. The within-block gaps grow at least as x^{400}. The same floor formulas prove $d_L \ge 3$ and $\mathrm{DF}(\kappa) \le 4kfb$, using $B_\tau \le b + o(b)$ and $k\Delta_{\mathrm{cl}}, mH_{\mathrm{node}} = o(b)$. At a top label the upper bound follows from $b_t = O(b')$ and $b'/b = x\exp(-x^{\varepsilon/3})$. They give $\mathrm{S}(\kappa) \le 10kfb$ and $\mathrm{S}(\kappa) \ge b$ at a thick label. Also

$$T\,\mathrm{S}(w_L) \le 2\,\mathrm{S}(\kappa), \qquad \mathrm{S}(w_F) \le (D+2)\mathrm{S}(w_L) + 5fr^2 T.$$

For the latter bound, each residual gap is at most its $(D+2)$-multiple of the w_L gap plus $3T$. Summing along each block gives $3fr^2 T$. The remaining scalar term is bounded by $fr^2 T$. The weights w_F, $w_F + Dw_L$, κ have even copy number at least $8fr + 4$. The companion weights have the theta degrees $2n_W - 2$ and $(D+4)n_W - n_P$, and their tableau entries and copy numbers are bounded by those of κ. The floor and companion identities.

In the angular flag metric the cross-block central gap of a weight is the sum of the two last entries and its scalar copy number. If this gap for w_L is C_L^{gap}, that for w_F is at least $(D+2)C_L^{\mathrm{gap}} + 2r + n_P$, by the first-block subtraction of A. The canonical angular row is $K_{Y,a} = 2a - (r-1)$, so its difference $K_{Y,a} - K_{Y,a+1}$ is -2. Therefore a same-block gap of $w_F - K_Y - Dw_L$ is $2\ell + (\gamma + 2 - B_T\ell) \ge 2\ell$. At a crossing the canonical difference subtracted is at most $2(r-1)$, and the cross-block gap is at least $2C_L^{\mathrm{gap}} + n_P$. In particular $C_L^{\mathrm{gap}} \ge 19T/10$. It also exceeds the polynomial allowance $B_{\mathrm{an}}^2 f \min(p, \rho) + B_{\mathrm{an}}$ with $B_{\mathrm{an}} = 5fr^2$, since $C_L^{\mathrm{gap}} \ge x^{600}$. This proves the strict mixed margins with the loss $2T$. Subtracting a sufficiently small positive multiple of w_L preserves all the strict inequalities. These are the margin comparisons for interpolation and division. At the definite endpoint the first block is the whole row and the same within-block reserve applies. Curvature weights for the chosen rows and The interpolation and division margins.

The installed primes and the cutoff. On the classical hosts $v = \mathrm{rd}(E)$ satisfies $f^2 \le v$, $v \ge 16$, and $\log v = o(L)$. Indeed $v \ge \sqrt{d}$ by the discriminant formula in the tower, and the lower bound on d exceeds every polynomial in x. With $\mathcal{R} = v^7$, put

$$c = \frac{\mathcal{R}\mathcal{P}d_L^2}{q_*}, \qquad g_* = 2 \cdot 625r, \qquad c_1 = \frac{1}{2}\mathcal{R}Tc\,(2f\rho)^2\rho.$$

The lower bound $q_* \ge (kb)^8$ absorbs $\mathcal{R}^4\mathcal{P}^3d_L^2$, because $d_L \le 4kfb$ and $\mathcal{R} = b^{o(1)}$. Also $\mathcal{P} \ge 90(fr+1)$, $\mathcal{P} \ge r$, and $\mathcal{P} \ge 16200\pi n_W m^4(2fm)^4$, since $n_W \le Z_b^4$ and the last expression grows more slowly than Z_b^{12} on the displayed rank brackets. For $R_0^2 = 90(fr + d_L)$ this gives $Q \ge \mathcal{R}^4R_0^2d_Lr$. It gives $Q \ge (D+4)d_L + 2$ and $g_i \ge 2\,\mathrm{DF}(\kappa) + 2$ on each half, so the available marked theta jets cover the same weights. For either i, $Td_L/(g_i - 1) < 1/2$. The upper bound on the installed primes gives $\log g_i \le 9L + 1$, and their lower bound precedes any local choice of half prime. The own primes have $\log p_o < L$. The analytic torus parameters at the actual marked level.

In particular $c \le b^{-5}$ eventually. It follows that $Tcg_* \le 1$, $cg_* \le 1/2$, and $fm^3\mathcal{R}^2c \le 1/2$. For the auxiliary decay constant $a_1 = (2\cdot 625v^2)^{-1}$ on the indefinite charts we have $c_1/a_1 \le 1$. Thus $r(1+c_1/a_1) \le 2T$, exactly the loss to absorb in the shifted weight above. The same cutoff c_1 is used in the Fourier–Jacobi vanishing condition and in that loss. When $\rho > 0$, $c_1 \ge 2fTc$, and therefore $Tc\Phi \le c_1/\lambda_1^2$ in the first-minimum units of the marked boundary potential. At the terminal stage $c_1 = \Phi = 0$. The cutoff in the half-trace convention and The normalized auxiliary decay constant.

All logarithmic upper bounds used for the section generators and finite norming sets now follow with their stated exponents. For example $\log(1 + \mathrm{DF}(\kappa)) \le 2L$, $\log n_\kappa \le L$, and the same two bounds hold for the companion and divided weights. The discriminant bound of the host and the unit class ideal give the remaining logarithmic bounds for the half-supremum frame. The degree and discriminant bounds of the geometric coefficient field come from its initial field construction. The full generating set and the finite norming set then give the specific logarithms in the Chow coefficient estimate. For example $n_W \le Z_b^4$ gives both $\mathrm{nbR} = O(Z_b^5)$ and $C_5 = O(Z_b^5)$. The size of the CM lower-norm coefficient is $O(Z_b^2)$. Together with $\log h \le 1041Z_b^3$ and $\log(1 + \mathrm{S}(w_L)) \le 2L$, this gives $P_0' = O(Z_b^7)$ at the exact Chow multiplier. The height allowance $P_H = O(Z_b^{10})$ and the degree-four or degree-five section costs are therefore dominated by $\mathcal{P}^2D$ in the aggregate comparison. The generator counts and large-multiple height constant.

The choices are ordered as follows. Fix the design constants and the two row exponents before choosing the scale. Take one lower threshold in x after those constants are fixed. Every bound above is then uniform over the labels, stages and words. The two runs change only the own-prime part of the same marked construction. Passing to a tied translate keeps the a_t, b_t arrays and uses the same formulas at its labels. The two type primes, both own primes and the designated auxiliary prime retain their installed values throughout.

The coefficient field in the section heights. Use the geometric K_0 carried by the original marked selection. Its initial construction is over the fixed two type primes and the two choices of own prime. The data needed for its size estimate are the embedded E, the Weber values at Qp_o for both p_o, the indicated roots of unity, and the ideal and type-norm generators. In the original construction the Weber values lie in the ray fields of $\mathbf{Q}(\sqrt{-d})$ of moduli Qp_o. The chosen field is installed before making the tied stages. The coefficient identification sends its entire ring of integers to the chosen embedded integral coefficient ring. Conversely, those coefficient values are precisely the images of its integers. Taking fractions identifies the embedded fields. This is the geometric K_0 used on both runs whenever the two tied records retain this original coefficient identification. The field chosen over already fixed levels and Identity with the geometric coefficient field.

Write $s(K) = \max(\log[K : \mathbf{Q}], \log \mathrm{rd}(K))$. The degree formula and the different under base change give $s(KK') \le s(K) + s(K')$. A cyclotomic field of order a has $s \le \log a$. For either Weber ray field, the ray exact sequence bounds its relative degree by the class number times the unit residue group order,

and the conductor-discriminant formula bounds its relative root discriminant. These two bounds give $s \le 7 + \frac{3}{2}\log d + 3\log(Qp_o)$, by the elementary quadratic class-number bound. Its order parity and the condition $d > 4$ hold on the selected hosts. The two Weber fields.

Let h_E be the ideal class number of E. Minkowski's bound represents every ideal class by an integral ideal of norm at most $\mathrm{rd}(E)^f$. Counting these ideals by the absolutely convergent sum at 2 gives $\log h_E \le 2f \log \mathrm{rd}(E) + 2f \log \zeta(2) \le fL$ at a common threshold. The orders of the roots used in the geometric field are

$$\prod_{i=0}^{1} |\mathrm{GL}_r(\mathbf{F}_{g_i})|^f, \qquad g_i(g_i - 1), \qquad 81, \qquad 625,$$
$$p_{\mathrm{de}}, \qquad p'_{\mathrm{de}}, \qquad h_E, \qquad 2h_E|(\mathcal{O}_E/(Qp_o))^\times|$$

with both i and both p_o in their respective factors. The general linear group contribution is at most $fr^2(18L + 2)$. The two last contributions together are at most $2\log h_E + 4f(19L + 2) + 2\log 2$. The roots of unity in the original field.

For the principal generators, enlarge by roots in class-group steps. If a class of order a in the current quotient has $\mathfrak{a}^a = (u)$ after extension, adjoin an a-th root of u. This kills that class and preserves all killed classes. The product of the successive a's is at most h_E. At each finite place divide the root by a generator of the localized $\mathfrak{a}$, making its a-th power a unit. Since the base already contains μ_a, the relative degree d' divides a and the resulting root has d'-th power a unit in the base DVR. The derivative of its monic polynomial bounds the relative discriminant contribution per degree by $\log a$. Thus principalization adds at most $\log h_E$ to either logarithmic size.

The same step works for the stage type norm $\prod_\tau \tau(\mathfrak{a})^{p_j}\bar{\tau}(\mathfrak{a})^{\rho_j}$ in each of the m signatures. The class subgroup now records generators with squared absolute value $N\mathfrak{a}^r$ at each embedding. Products and inversion preserve this condition. For an element generating a principal ideal the formula for the norm gives it, since the exponents of each conjugate pair sum to $p_j + \rho_j = r$. An a-th root preserves the condition when it kills a class in the current quotient. Thus these m passes add at most $m \log h_E$ to either logarithmic size. The class-group steps and The type-norm version with the metric normalization.

Using $m \le f$, $L \ge 4$, and the bounds on the installed primes gives the following generous total:

$$s(K_0) \le fr^2(18L + 2) + (m + 4)fL + (79f + 190)L \le 19fr^2L + 198f^2L = P_O.$$

This uses the geometric field before adjoining the values of the moving ray characters. Those values and subsequent fields for finite tests extend the given embedding of K_0. The degree and root discriminant in the geometric height allowance are $[K_0 : \mathbf{Q}]$ and $\mathrm{rd}(K_0)$. Normalized heights of these sections are unchanged under subsequent scalar extensions by the local degree formula. The size calculation in the original field construction.

The scaled products and their local index. Work on one of the two marked runs, at stage j, and write $\kappa = Tw_L + w_F$ and $w_0 = Dw_L + w_F$. The dimension of the full flag over the rank-r target is $D = f\binom{r}{2}$ for the signature $(k + j - 1, m - j)$. Here T is the positive integral Hodge-line multiple of the run. Let A_w and S_w be its full and short theta modules. Put $G = g_0 g_1$, $h_w = frn_w$ and $z_w = p_o^{\mathrm{DF}(w)+2h_w}$, where p_o is the own prime of the run. With $n(w)$ the integral CM norm product and $U_w = Q_S^{\mathrm{DF}(w)}(qd)^{\lceil \mathrm{DF}(w)/2 \rceil}$, the four scaled modules are

$$\begin{aligned}
A'_L &= n(w_L)U_{w_L}A_{w_L},\\
S'_F &= z_{w_F}n(w_F)U_{w_F}G^{h_{w_F}}S_{w_F},\\
S'_0 &= z_{w_0}n(w_L)^D n(w_F)U_{w_L}^D U_{w_F}G^{h_{w_0}}S_{w_0},\\
S'_K &= z_\kappa n(w_L)^T n(w_F)U_{w_L}^T U_{w_F}G^{h_\kappa}S_\kappa.
\end{aligned}$$

The sum $\mathcal{T} = (A'_L)^T S'_F + (A'_L)^{T-D}S'_0 + S'_K$ is a full-rank sublattice of the integral supremum lattice M_0. The product of two modules here is the span of their products. Pointwise valuation additivity and the

integral scaled section bounds give the inclusion, and S'_K spans the rational space. After dividing by $\sigma = n(\kappa)G^{h_{w_F}} U_{w_L}^T U_{w_F}$, the three respective scalar factors are z_{w_F}, $z_{w_0} G^{frDn_{w_L}}$, and $z_\kappa G^{frTn_{w_L}}$. The short-product relations on the marked cover place the divided lattice inside the short lattice at weight κ. The sum of three scaled section products · Short products on the marked cover · The scalar and the placed lattices.

The section families on the two marked covers. Fix a stage j, the geometric coefficient field K_0, and one of the two choices of own prime p_o. Write $G = g_0 g_1$ for the base principal level and C_w for the rational cusp space of weight w on that level. Its short rules are read on the specified refined level Gp_o, with the cover's determinant section and its finite Fourier weight. The two coefficient fields of the cover and the base are identified in this rule. Write $n = n_w$, $h_w = frn$ and $s_w = \mathrm{S}(w)$. Use $A(w) = \mathrm{DF}(w)$ for the total tableau degree and $d_w \le A(w)$ for the homogeneous degree in the tangent letters. In these notations $s_w \ge 1 + A(w) + h_w$.

For each $w \in \{w_F, w_0, \kappa\}$ there is a finite list of elements of S'_w spanning C_w over K_0, where S'_w denotes S'_F, S'_0, S'_K, respectively. Each element is integral for the supremum lattice of its own weight and has logarithmic supremum at most $G_0 s_w$ at every embedding. The full module A'_L has an integral generating list $(w_k)_{k=1}^N$ with a common bound m_σ for $\log \|w_k\|_\sigma$ and

$$\frac{1}{[K_0 : \mathbb{Q}]} \sum_\sigma m_\sigma \le G_0 L_0, \qquad \log N \le P_w(L_0 - 1). \tag{128}$$

The full list can retain the block polynomial and all its marked translates. The choices of these lists are made before forming the generic equations or passing to a local valuation ring.

The prescribed refinement in the spanning argument. The stable theta kernel proof applies in even copy rank $n \ge 8fr+4$. The relative bound on the sum of off-diagonal overlaps uses $\sum_{a\ge2} a^{4fr-n/2} \le \sum_{a\ge2} a^{-2} < 1$. The doubled integral is in the convergent range, with definite auxiliary space and doubled target rank $2r$. These conditions hold for w_F, where $n_F = 8fr + 4$, and for w_0, κ by their additive copy ranks. The chosen weights have dominant tableaux and positive last entry in the prescribed positive tableau block. For example, if a is that last entry of κ, its entry in w_L is the floor of $(a - (D+2)n_L - 2r)/(T + D + 2)$. It is positive, and subtraction of T times that entry from a leaves a positive last entry for w_F. The same inequalities apply when the other tableau block is empty. The floor and positivity checks and The stable rank on the three short weights.

Here is the level identification for applying that kernel proof. Use the half-ideal convention of the dual marking. Let $\mathfrak{m}$ be its old half, let $\mathfrak{j}$ be the added half over p_o, and put $\mathfrak{a} = \mathfrak{m}\mathfrak{j}$. Thus $\mathfrak{a}\bar{\mathfrak{a}} = (Gp_o)$ and $\mathfrak{a} + \bar{\mathfrak{a}} = \mathcal{O}_E$. At the finite split factors in $\mathfrak{a}$ the period lattices have free reductions of rank r. In one reference lattice choose a basis β there. Write J for the kernel of reduction from $\mathrm{GL}_r(\mathcal{O}_E/\mathfrak{a})$ to $\mathrm{GL}_r(\mathcal{O}_E/\mathfrak{m})$. The Fourier weight on the added half is the function of the first r rows which vanishes on nonbases and otherwise takes the inverse adjoint character of the change from the pinned basis. Under the Chinese remainder isomorphism it is a fixed unit multiple of $k \mapsto \lambda(k)$ on J, where $\lambda(k) = \psi(\mathrm{adj}\, k_\mathfrak{j})^{-1}$ in the coefficient embedding. The fixed multiplier is the value on the reference basis. This follows by writing the basis matrix as the product of its change of basis and the pinned basis matrix. The character factors through determinant: over each tame local factor the determinant-one group is generated by elementary matrices, each a commutator with a diagonal unit whose residue differs from 1. The lower $n - r$ rows are free on this added half. The old Fourier character still fixes the top reduction to β mod $\mathfrak{m}$ and the lower reductions to zero. Both statements concern the Fourier transform of the same weight w_{x_0} in the short rule. The two half ideals and the determinant character and The actual Fourier weight.

Choose the first r members of x_0 to give the old dual basis in the reference lattice, with the factor p_o included, and take the rest zero. The same x_0 works on all components after rational unitary moves of their lattices. Indeed, lift the given old dual basis, and a basis at the added half, simultaneously by the Chinese remainder theorem. A unitary move can carry these vectors arbitrarily close to β at the chosen

split places. For weak approximation one can use the Cayley transform of a skew-adjoint matrix. On the dense open where its inverse exists, the local matrix at a chosen half determines the opposite one by adjoint inversion. Approximate the independent skew-adjoint coordinates over the maximal real subfield and apply the transform. This moves the full local lattice to the reference lattice at the marked places, since the vectors form a local basis. Self-duality gives the corresponding equality at the conjugate places. The moved lattices are therefore in the family used by the overlap operator. Equality of two such lattices by an allowed rational move gives equality of the old dual marks modulo $\mathfrak{m}$. The perfect polarized pairing recovers the primal marks modulo G, so the two components are the same. The first r entries of this x_0 generate the required old quotient, also when admissibility permits all n entries. Approximation on the marked basis and The common admissible vector and disjoint component orbits.

For a rational stabilizer of a moved lattice the old half congruence and inverse adjoint give congruence to 1 mod G on the full lattice. Its determinant is an algebraic integer of norm-one absolute values at all embeddings. If the determinant differs from 1, the norm of its difference from 1 is at most $2^{[E:\mathbb{Q}]}$ and at least $G^{[E:\mathbb{Q}]}$. Thus the determinant is 1, and the cover character on this stabilizer is trivial. The shifted vacuum has the same trivial stabilizer multiplier. To see the theta phase, write $w_{\mathrm{tr}} = (1 + \sqrt{-d})/2$, with $d \equiv 7 \bmod 16$. Modulo 8, $e = w_{\mathrm{tr}}^4$ is an idempotent and its CM conjugate is $1 - e$, by the two distinct roots of its quadratic equation modulo 2. Its two summands of the shifted period lattice modulo 8 are dual Lagrangians for the alternating form. A symplectic basis adapted to them modulo 8 lifts to an integral symplectic basis by elementary symplectic elimination. The rational stabilizer has off-diagonal symplectic blocks zero modulo 8. Its first block has determinant 1 mod 8: the full linear map $eg + (1 - e)1$ has E-determinant congruent to 1 mod 8, by idempotence and $\det_E g = 1$, and taking the rational determinant reads that first block. The other block has the inverse determinant. A block-diagonal symplectic change from $\mathrm{SL}_{fr}(\mathbb{Z})$ can therefore reduce both blocks to the identity modulo 8, by lifting elementary matrices. The diagonal change is an exact relabeling of the theta series. The remaining change belongs to $\Gamma(8)$, where the squared theta transformation has only its determinant and Gaussian factors. They cancel with the change of period frame in the normalized vacuum. The characteristic is preserved, since it is the semicharacter of the same stable lattice. Taking the even n-fold tensor removes the square-root sign. Igusa, *Theta Functions*, Chapter V, §1, Corollary, and §2, Theorem 3, and The normalized multiplier.

Apply the stable graph reproduction with half-ideal $\mathfrak{a}$, and with the test whose first r rows have the reduction $k\beta$ and whose lower rows are zero modulo $\mathfrak{a}$, weighted by $\lambda(k)$. This is the test before freeing lower rows. All supported top marks are bases, which is the support used in the isotropic coinvariant argument. Its nonzero coefficients have modulus one, so the overlap bound and the equal diagonal mass in the reproducer hold with this finite family of cosets. Use auxiliary lattices that are standard at 2, at the excluded bad primes, and at Gp_o. The kernel proof obtains these conditions by imposing the single combined integer modulus on its auxiliary lattices. The disjoint component orbits and their stabilizers just checked allow extension of any tuple of bounded cusp sections to the lattice family, with the required scalar covariance and zero on the unused orbits. Invert the overlap operator and apply the Gaussian reproduction. This expresses the tuple as a finite linear combination of these theta kernels with the required highest-weight derivatives. The convergence uses the same Cartan bound in the prescribed real frames and holds also on the compact terminal quotient. The support of the compact test, The overlap bound, and Reproduction at the combined modulus.

The supported lower rows can now be freed precisely at $\mathfrak{j}$. Choose $\alpha \in \mathcal{O}_E$ with the valuations of $\mathfrak{j}$ at its support and unit at the other places of the combined modulus. Let H on the copies be diagonal, with entries 1 on the first r and $\bar{\alpha}/\alpha$ on the others. For a source lattice Λ just used, put $\Lambda'' = H\Lambda$. Patch it to the standard lattice over the places of $\mathfrak{j}\bar{\mathfrak{j}}$ and leave it equal to Λ'' elsewhere, obtaining Λ'. The patch is self-dual place by place. Let L be a moved period lattice and $L^\sharp = \mathfrak{a}\bar{\mathfrak{a}}^{-1}L$ its shifted lattice. The source vacuum for Λ'' is a normalized projection of that for Λ' on the selected local factors. The index of the two tensor lattices is $I_{\mathrm{loc}} = [\Lambda'' : \Lambda'' \cap \Lambda']^r$, independent of L. Choose representatives of the quotient of the tensor lattices with only lower residues at $\mathfrak{j}$. This is possible because Λ'' has lower factor $\mathfrak{j}^{-1}$ there,

lower factor $\bar{\mathfrak{j}}$ at the conjugate places, and no other difference from Λ'. The representatives can be taken in $2\Lambda'' \otimes (\mathfrak{a}\bar{\mathfrak{a}}L)$. The normalized projection is $I_{\mathrm{loc}}^{-1/2}$ times the sum of their Weyl translates. With these even representatives the commutator in each term of the shifted sum is trivial: the other translate lies in $2\Lambda' \otimes (\bar{\mathfrak{a}}L)$ after the move, and its pairing with such a representative is integral. The representatives cover exactly the free lower residues at $\mathfrak{j}$, once each. At the old half the lower residue still vanishes, and the top row and its coefficient are unchanged. Thus the moved packet is the free-lower-row packet of Λ' times a nonzero common scalar.

The change of copy variables by H preserves the Gaussian and the row weights and carries a harmonic derivative to a harmonic derivative. Its phase on the normalized vacua is common to the lattice family. Indeed, the partial Fourier transform of a doubled pair of normalized vacua on $\Lambda \otimes L_1^\sharp, \Lambda \otimes L_2^\sharp$ is supported on $\Lambda \otimes (L_1^\sharp + L_2^\sharp)$, with a quadratic trace character whose matrix is on the target side. Its scalar at zero is independent of Λ. One way to see this scalar is to use the sesquilinearity and isometry up to constant of the partial Fourier transform. Pair the doubled standard vacua with the doubled projected vacua. On the intersection of their supports the two quadratic characters coincide, so the positive volume of that intersection leaves the phase of the pairing equal to the phase of their scalar ratio. On the single-vacuum side the inner products with the standard vacua are positive by the positive normalization of the two orthogonal projections. The overlap moduli themselves agree because the relative indices between the two target lattices are independent of Λ. Thus the scalars agree. Reindexing the doubled sum by the rational unitary H preserves the trace character. The products of the resulting single theta values with the conjugates at arbitrary second frames and lattices are equal. Fixing one nonzero vacuum value gives the common unit scalar. This argument keeps the same scalar on every component, as needed in the span. The local patch and The free lower rows. The doubled comparison and The common scalar.

Read the freed packet using the actual short rule $F_{\Lambda',N,x_0,P}$ on the Gp_o cover. Choose N odd, divisible by Gp_o, with $N\mathcal{O}_E^n \subset \Lambda' \cap \mathcal{O}_E^n$. The permitted torsion points are the lifts of the prescribed dual marking. Fourier inversion with the Weil pairing changes the weighted sum over the full cover marking to the just-computed free-lower-row translates. In fact, if $W(y)$ denotes a Weyl translation and v the projected vacuum,

$$W(u)W(y)v = e^{-2\pi \mathrm{i}E(u,y)}W(y)v$$

for the even dual representatives u fixing v, where E is the alternating polarization form on the copies. Summing against each finite Fourier character gives exactly the translate labelled by its dual representative. The old factor of the cover Fourier coefficient fixes $p_o x_0$; the added factor has the basis indicator and inverse character derived above. The inner torsion average is the lattice vacuum projection, with index ratio $[\Lambda' : \Lambda' \cap \mathcal{O}_E^n]^r$; the factor N^{-d_w} cancels the derivative of dilation. The shifts at the cover level and the auxiliary lattice projection act at disjoint sets of primes. Consequently the comparison with the analytic kernel separates into $C_Z u_{\Lambda',N}$, where $C_Z \neq 0$ depends only on the marked component and $u_{\Lambda',N} \neq 0$ depends only on the auxiliary lattice and its rule. The latter contains the positive square root of the auxiliary index ratio and the prescribed $\beta_{\Lambda'}^{-1}$. The determinant section has finite order and is constant on each cover component. Its character cancels the change of the inverse Fourier character, so the product descends to the base. The principal theta comparison has constant modulus on unitary frames by the Gaussian torus mean. With the same holomorphic frame weight on both sides, its remaining scalar is constant on the connected component, by the constant-modulus identity as in the principal packet. The identity holds for all harmonics and sections, including at their zeros. The separated component comparison and The rational rule on the prescribed cover.

The harmonic equations and the row weight are linear equations over K_0 in the tangent-letter basis. Their complexification is the harmonic space in the analytic kernel. For fixed Λ' the integral symmetric algebra of its letters spans that coefficient space, since it contains the N-multiples of the standard letters. Clear denominators to obtain integral harmonics in the same weight. Each such polynomial gives an algebraic short section of the prescribed weight by the natural normalized torsion average. Its linearity

follows with the marking, N, the lattice and the determinant section held fixed. One may check this formula on common lifted framed points. The lifts extend faithfully flatly, so the equality there descends. This produces the complex spanning assertion using sections over K_0. The scalar extension of C_w injects into the space of framed readings. Indeed flat extension of the algebraic line-section space can be checked on a finite affine trivializing cover, using the kernel of the restriction differences. Uniformization at the good embedding reads all the complex geometric components of that space, on which pointwise evaluations separate sections. It follows by scalar extension of the quotient space that the actual short rules span C_w over K_0. At the other embeddings the section norms are given by their good partners in the same field, with conjugate coefficients and the identical Lie norm factors. Integral harmonics in the tangent frame, Linearity on the fixed rule and Descent of the span.

The scalar in the geometric field. The denominator β_Λ is taken in the already chosen field K_0. For the inclusion $\Lambda \cap \mathcal{O}_E^n \subset \Lambda$, let $\mathfrak{b}$ be the quotient of their determinant ideals, the smaller lattice's determinant divided by the larger one's. Local elementary divisors give $N\mathfrak{b} = [\Lambda : \Lambda \cap \mathcal{O}_E^n]$. The denominator generates the type norm ideal $\prod_\tau (\tau(\mathfrak{b})\mathcal{O}_{K_0})^{m_\tau}$ and has squared absolute value $(N\mathfrak{b})^r$ at every embedding. Here the tangent multiplicities satisfy $m_\tau + m_{\bar\tau} = r$.

The field construction before the runs principalizes these type norm ideals with precisely that absolute value. On a principal ideal (a) the generator is $\prod_\tau \tau(a)^{m_\tau}$, and the absolute value identity follows by pairing conjugate embeddings. The ideal classes with such a generator form a subgroup. If an additional class represented by $\mathfrak{b}$ has order k modulo the subgroup, adjoin a k-th root of the normalized generator for $\mathfrak{b}^k$. Ideal valuations and absolute values both divide by k. Locally one divides the root by the appropriate power of the uniformizer, so it is a root of $X^k - u$ with u a unit. The derivative of its monic minimal polynomial divides k in the resulting local ring, by differentiating the factorization of $X^k - u$. The different exponent is therefore at most the exponent of k. The degree and root-discriminant factors for this step are at most k. The successive indices multiply to at most the class number of E for each stage. This is the type norm part of the original geometric field construction. It makes the required normalizers available simultaneously, before any auxiliary lattice or polynomial is chosen here. The field chosen before the runs, The normalized generators, and The determinant quotient.

Integral harmonics with bounded coefficients. For each auxiliary lattice Λ in a short rule put the Euclidean norm on its Minkowski space by summing the squared Hermitian norms over all embeddings of E. The Euclidean dual is the trace-dual lattice, represented by $\mathfrak{d}_E^{-1}\Lambda$. For a nonzero vector v in this dual, $h(v, v)$ belongs to $\mathfrak{d}_E^{-1}\overline{\mathfrak{d}_E}^{-1}$ and has positive conjugates. Taking its norm to $\mathbb{Q}$ and using the arithmetic–geometric mean gives a first minimum at least $\sqrt{2f}/\operatorname{rd}(E)$. Thus transference gives $2fn$ rationally independent lattice vectors, each of length at most $n\sqrt{2f}\operatorname{rd}(E)$. [53, Theorem 2.2]

Use these vectors as a rational basis for the parameters of a minor-product family. At a pair of conjugate embeddings write $p + q = r$ for the two tableau sizes. With p upper vectors u_i, p trial vectors v_i and q lower vectors y_i, let $B = (h(u_i, v_l))_{i,l}$. Here h is linear in the second argument. Replace each y_i by

$$y_i^\perp = (\det B)y_i - \sum_l v_l\big(\operatorname{adj}(B)(h(u_k, y_i))_k\big)_l.$$

This lies in Λ and pairs to zero with all u_k, including at singular B. Pair the tangent-letter rows with the appropriate embedded upper and conjugate lower vectors. Take the product of the leading minors with powers the successive tableau differences. The row operations and row degrees give the prescribed highest weights. The cross contractions vanish because $h(u_k, y_i^\perp) = 0$. The coefficients are integral in the symmetric algebra of the embedded lattice letters.

These products span the harmonic space. Over $\mathbb{C}$ the orthogonal pairs of full row ranks form a dense orbit in the variety of pairs (X, Y) with $XY^{\mathsf{T}} = 0$. For fixed rank a of X, this variety has stratum dimension $a(p + n - a) + q(n - a)$. For $a < p$ it drops from the top dimension by $(p - a)(n - q - a) > 0$. The pq equations have their expected codimension. Their complete intersection is generically smooth on the full-rank open set, by solving for Y, and is reduced. In the Fischer pairing their ideal is orthogonal to

the joint kernel of the contraction operators. Thus evaluations at orthogonal pairs separate the harmonic polynomials. At the standard orthogonal pair the evaluation on the row highest space is represented by the indicated product of minors. One can see this by the row and column diagonal weights: their total degrees force the upper and lower polynomials to use only their separate standard column blocks. Evaluation at the standard pair transforms by the row character when a matching triangular matrix acts in the column block. Taking its adjoint in the Fischer pairing gives the opposite triangular covariance of the representing polynomial. On the Gauss cell the two covariances determine the polynomial by its value at the identity. This gives exactly the product of leading minors up to a nonzero scalar. The joint column action, with inverse transpose on the dual columns, preserves the equations and carries this evaluation to those on the full-rank orbit. Equivalently, the compact unitary orbit is Zariski dense there. This proves the span.

For the parametrization above the independent complexified coordinates of the rational lattice basis give all coordinates at all embedding pairs. On each full-rank cell choose the trial vectors with $B = 1$. The resulting parametrization contains the full-rank orthogonal pairs. Its polynomial degree in the rational parameters is at most $(2r+1)d_w$. Consequently evaluations with each parameter in $\{0, \dots, (2r+1)d_w\}$ span its values, by univariate interpolation in each parameter. The grid takes place in Λ. Because the equations of the harmonic space and the coefficients of the products are over K_0, the same grid spans over K_0. In degree zero use the constant polynomial when that space is nonzero.

For the coefficient bound, set

$$B_1 = n\sqrt{2f}\,\mathrm{rd}(E)\,(2fn)(2r+1)(1+d_w).$$

Every unprojected vector of the grid has length at most B_1 at each embedding. The page bounds give $\log B_1 \le 12L$, where $L = \log b$. The Gram adjugate bounds each projected vector by $(p+1)p!B_1^{2p+1}$. A linear form with coefficient vector v has coefficient sum at most $\sqrt{n}\|v\|$, and a minor of size i costs at most $i!$. Multiplying the minors gives

$$\log |P|_{1,\sigma} \le 140\, rL\, d_w \le 140\, rL\, A(w) \qquad \text{at every } \sigma.$$

The bound uses $\log(1+d_w) \le 2L$, $\log n \le L$, $\log f, \log \mathrm{rd}(E) \le 2L$, $r \le n$ and $L \ge 2$. It also holds for an empty tableau block. The adjugate and the interpolation grid and The lattice-letter norms.

Uniform norms and finite integrality of the short lists. For each allowed (Λ, N, x_0) retain the sections of the grid harmonics for these data. Linearity on the fixed rule shows that their union spans C_w. Its elements satisfy a common bound independent of Λ, N, x_0. To verify it, at a nonzero half-supremum choose the orthonormal frame of first minimum $\lambda \ge e^{-40L}$ from the transvection estimate used for the auxiliary sections. The rank-one bound depends on the actual geometric cusp condition and the page bound on the spread, so it applies to each individual section here at each good embedding. The definite case uses compactness and the same minimum bound for the self-dual norm. Let $I = [\Lambda : \Lambda \cap \mathcal{O}_E^n]$. The unnormalized auxiliary projection has norm $I^{-r/2}$ times the standard vacuum norm. The prescribed algebraic denominator satisfies $|\sigma(\beta_\Lambda)|^2 = I^r$ at every embedding. The inner index ratio is I^r, so these three factors cancel. In the outer average, each weight is bounded by $p_o^{h_w}$. Indeed the own-prime Fourier sum has at most $p_o^{h_w}$ summands, each zero or a root of unity. The old Fourier factor and the determinant section have modulus one. This calculation gives the archimedean bound for the same average as in the integral rule.

For a direct Gaussian estimate, normalize the torus metric so its alternating form is the imaginary part of the Hermitian form. The tensor lattice then has covolume one by self-duality. Its first minimum is at least $\delta = \lambda/(2f^{3/2}n^2\,\mathrm{rd}(E))$. Indeed, contract a nonzero tensor with the short rational basis of Λ just constructed. Each contraction belongs to the target lattice by Hermitian integrality, and one is nonzero. By Hermitian Cauchy–Schwarz its target length is at most the norm of the tensor times the maximum conjugate length of the chosen basis vector. This gives the stated lower bound in the tensor metric.

The probability mean of the squared norm of the product vacuum on the standard torus is one. The projected mean is I^{-r}. This follows also on the common period cover: the I^r translates in the projection are orthogonal because their differences give distinct characters under old period translations, and each has mean square one. Weyl translation of the projected section preserves its probability mean on the new torus. In Fock coordinates, translate the evaluation point to zero. Apply holomorphic submean on the ball of radius $R = \min(1, \delta/3)$, which injects into that torus. The weighted point norm is bounded by the probability squared norm's square root times $e^{\pi/2}\sqrt{h_w!}/(\pi^{h_w/2}R^{h_w})$. Here $\log(1/R) \le 49L$ and $\frac{1}{2}\log h_w \le 2L$, so the logarithm of this factor is at most $52h_wL$. The letter directions have length $\sqrt{t}$ and are orthogonal in these coordinates, where $t = 2\xi_{\mathbb{R}}$. For a derivative of multi-index γ, take Cauchy radii $(\gamma_i/(\pi t))^{1/2}$ in its active directions. The Gaussian weight and $\gamma_i! \le \gamma_i^{\gamma_i}$ give the bound $(\pi e\, t\, d_w)^{d_w/2}$ for the derivative of degree d_w, times the point norm bound. The factor $(\pi t)^{-d_w/2}$ in the section norm cancels the scale in that jet. This argument also applies to the degree-zero term without differentiation. Using the half-supremum factor 2, $\log p_o \le L$ and the page bounds gives

$$\log \|F_{\Lambda,N,x_0,P}\|_\sigma \le \log |P|_{1,\sigma} + 10LA(w) + 200h_wL \le 400(frL)s_w.$$

Zero sections may be dropped. The estimate for an unrefined full generator holds with $p_o^{h_w}$ replaced by 1. The finite Fourier sum bound, Cancellation of the three vacuum-index factors, The letter normalization, and The section norm at every embedding.

The finite-place inclusion uses the oriented integral chart of this run. For each valuation v and each such framed point, the short rule has order at least

$$-v(n(w)) - \mathrm{ramLoss}_v(w) - h_w v(Gp_o) - A(w)v(p_o).$$

Away from the marking, the adapted isogeny and trace in the integral primary average cancel the inner index and the determinant ideal of the auxiliary lattice. At a marked prime the auxiliary lattice is standard locally, so its determinant factor is a unit. The primary inner torsion multiplicity cancels the primary part of the inner averaging denominator. Removing the excess primary multiplication in N leaves one level-prime factor, since Gp_o is square-free. At g_i factor the remaining pullback into the conjugate CM halves, using the given orientation of the old marking. On the selected half the chain rule then costs $v\big(\prod_\tau \tau(\pi_{g_i})^{d_\tau}\big)$, with d_τ the harmonic degrees. At p_o use the rational multiplication by p_o and allow its full cost $A(w)v(p_o)$. The two old factors are bounded by $v(n(w))$. The canonical connection gives $\mathrm{ramLoss}_v(w)$, cleared by U_w, by the same integral harmonic evaluation. The remaining outer marking denominator has order $h_w v(Gp_o)$. The Fourier coefficients are sums of roots of unity and the determinant section is a unit. These are bounds on every framed algebraic valuation, so they pass to integral linear combinations. The full rule has the stronger bound $-v(n(w)) - \mathrm{ramLoss}_v(w)$ on the same charts. The pole bound for the actual short rule and The four scaled modules.

The integers $U_{w_L}^i U_{w_F}$ clear the ramified loss at $w = w_F + iw_L$. The degrees add, the half-degree ceilings are subadditive, and $n(w) = n(w_L)^i n(w_F)$. The factor G^{h_w} clears the old marking denominator, and $z_w = p_o^{A(w)+2h_w}$ clears the own-prime loss. This proves integrality after multiplication by exactly the factors of S_F', S_0', S_K'. For the archimedean bound put $Z = frL$ and use $\log|\sigma(n(w)U_w)| \le \lambda_\star s_w$, where $\lambda_\star = \log G + \log(f^2qd)$ is the page constant for the norm and conductor factors. There is no choice of a representative with a larger archimedean norm here: the CM norm product has $\log|\sigma(n(w))| = \frac{1}{2}A(w)\log G$. For $i = 0, D, T$,

$$i\, s_{w_L} + s_{w_F} = s_{w_F+iw_L} + i, \qquad i \le s_{w_F+iw_L}.$$

The logarithm of the total short scalar at the indicated weight is therefore at most $(2(\lambda_\star + 2\log p_o) + \lambda_\star)s_w$; the last allowance includes G^{h_w}. Adding $400Z s_w$ is within the page's $G_0 s_w$. The scalars are nonzero, so their multiples still span over K_0. Finite-dimensionality permits taking a basis from the union itself, with the same bounds on its individual elements. The exact scalar cost and The bounded spanning lists.

Integral generation at w_L. For the full module use the standard auxiliary lattice and its integral polynomial lattice. This lattice is the intersection of the rational harmonic space with the polynomials of integral coefficients. There are $h = h_{w_L}$ tangent letters, so the number of monomials of total degree at most d_{w_L} is bounded by $(1 + d_{w_L})^h$. Take the short grid above for the standard lattice. Multiply it by a set of $[K_0 : \mathbb{Q}]$ rationally independent algebraic integers whose conjugate norms are at most $[K_0 : \mathbb{Q}]^3 \operatorname{rd}(K_0)$. The trace-dual minimum and Euclidean transference give such a set. These products span the real coefficient space of the integral polynomial lattice and have Euclidean norm at most $\sqrt{[K_0 : \mathbb{Q}]}\,[K_0 : \mathbb{Q}]^3 \operatorname{rd}(K_0) \exp(140 r L d_{w_L})$, in the sum of the squared coefficient norms over all embeddings.

For this Euclidean lattice of rank a, choose an integral basis successively by projecting orthogonally off the span of the previous basis vectors. At each step choose a shortest nonzero projected vector. It is primitive and is no longer than the common bound for the spanning products. Choose a lift whose Gram–Schmidt coefficients along the previous vectors have absolute value at most $1/2$, by subtracting integral multiples in reverse order. Each lift has length at most $\sqrt{(a+3)/4}$ times the same bound. This produces a basis of the entire integral polynomial lattice. Returning to coefficient sums at one embedding costs at most the square root of the number of monomials. Since $a \le [K_0 : \mathbb{Q}](1 + d_{w_L})^h$, each resulting polynomial satisfies

$$\log |P|_{1,\sigma} \le 140 r L d_{w_L} + 2hL + 4 \log [K_0 : \mathbb{Q}] + \log \operatorname{rd}(K_0).$$

This step gives integral generation of the full harmonic module.

Use these polynomials with every point of the original half-level marking, for each copy. The full rule is linear in the polynomial, and its value depends on the marked point only modulo the chosen half ideal, by the normalized theta-group translation. These two facts show that the resulting sections generate A_{w_L} over $\mathcal{O}_{K_0}$. If the block polynomial is required, enlarge the list by that polynomial at every marked point. For a label of tableau size r_τ it uses disjoint sets of $1, \dots, r_\tau$ copies for the minors of those sizes. In total this uses $\sum_\tau r_\tau(r_\tau + 1)/2 \le f r(r+1)/2 \le n_{w_L}$. The paired blocks therefore share no copy, so their cross contractions vanish. The coefficient sum is at most $r^{d_{w_L}}$. The added sections belong to the same full module and keep its chosen block evaluations.

The number of marked points on n_{w_L} copies is G^h. Multiplication by the prescribed scalar $n(w_L) U_{w_L}$ gives the integral generators of A'_L. The count and the common logarithmic bound, for this whole list at once, satisfy

$$\begin{aligned} N &\le 2G^h [K_0 : \mathbb{Q}](1 + A(w_L))^h, \\ m_\sigma &\le \lambda_\star L_0 + 400 Z L_0 + 3hL + 5 \log [K_0 : \mathbb{Q}] + \log \operatorname{rd}(K_0) + 2. \end{aligned}$$

Here the second inequality uses the full Gaussian estimate and the integral polynomial basis bound. Here P_O bounds both field logarithms. Since $\log g_i \le 9L + 1$ and $\log(1 + A(w_L)) \le 2L$, the first inequality gives $\log N \le (38L + P_O)(L_0 - 1)$. The second is within $G_0 L_0$ for the common growth constant of the two runs. For example the prescribed G_0 dominates, for each own prime, the sum

$$1040 Z + 3L + 2 \max\{\log [K_0 : \mathbb{Q}] + \log \operatorname{rd}(K_0), \log [E : \mathbb{Q}] + \log \operatorname{rd}(E)\} + 2 f r(\lambda_\star + 2 \log p_o).$$

This sum also dominates the short-section allowance above after division by s_w. These are the constants in (128), at the geometric field and before extending to period coordinates. The full integral polynomial lattice, Integral generation by the marked rules, Adding the block evaluations, and One growth bound for both runs.

In the normalized fibre take independent images of the weight-w_F list over $K_0(z)$. After division, use the weight-w_0 list in the order-$T - D$ product span, and again select independent columns over $K_0(z)$. Each nonzero column $u^\alpha \theta_i^0$ has sum of coefficient norms at most $(N e^{m_\sigma})^{T-D} \|\theta_i^0\|_\sigma$. Thus the determinant-content estimate uses one simultaneous bound for all of the w_L coefficients, and one mean bound for each

weight-w_0 column. The projected weight-κ list gives independent columns in the low-coefficient quotient. Its bound is inherited from the same scaled list. At $j = m$ that quotient has rank zero. For the division inclusion take a complex specialization with algebraically independent generic parameters over K_0. The ranks of the product matrix and of the matrix with the deep columns adjoined are their ranks over $K_0(z)$, since this specialization is a field embedding. Thus the complex inclusion gives the same inclusion over $K_0(z)$ and over every Gauss completion. Each generator is integral at its own weight, so the products are in the local supremum lattice by additivity of framed point valuations.

Companion sections and the boundary product. Write $L_0 = \mathrm{S}(w_L)$, $n_W = n_{w_L}$ and $n_P = n_{w_F} = 8fr + 4$. The full copy rank is $n_\kappa = Tn_W + n_P = r + t$. The two companion weights are $w_1 = 2(w_L - w_{\mathrm{de}})$ and $w_2 = (D + 4)w_L - w_F$. For w_1 the nonnegative tableau entries on both blocks are twice those of w_L, and the theta degree is $2n_W - 2$. This uses the scalar weight w_{de} of theta degree one and zero tableau entries. The construction of w_1 is used at the nonterminal stages. The terminal calculation only uses w_2.

To check w_2, put $T' = T + D + 2$. For an adjacent gap γ of κ inside either block, the gap of w_L is $l = \lfloor(\gamma + 2)/T'\rfloor$. The corresponding gap of w_2 is $(T' + 2)l - \gamma = 2l + 2 - (\gamma + 2 - T'l) > 0$, by the chosen gap bound $l \geq T' + 2r$. For a last entry a of κ, the associated entry of w_L is $l = \lfloor(a - b)/T'\rfloor$, with $b = 0$ or $b = (D + 2)n_W + 2r \leq T$. Here $a \geq 2T^2$ and $T \geq 40(D + 2)$ give $l \geq 3T/2$. The last entry of w_2 is therefore $(T' + 2)l - a = 2l - b - (a - b - T'l) > 0$. Its theta degree $(D + 4)n_W - n_P$ is even and positive. Both weights lie below κ in the tableau entries and theta degree. The last positive-block entries are strictly positive, and the relative theta degrees in use are at least r. Their sizes satisfy $\mathrm{S}(w_1) \leq 2L_0$ and $\mathrm{S}(w_2) \leq (D + 4)L_0$. Floors and auxiliary weights and The weight and theta-degree bounds. The relative theta degree and jet inequalities are the ones in the companion-weight calculation. The two companion weights are in the theta range.

Counting the complex components. For the selected unitary lattice of each component, untwist the common perfect trace ideal. The fractional ideal I and the alternating trace axis ξ obey $(\xi)I\bar{I}\mathfrak{D}_E = \mathcal{O}_E$, so the lattice $I^{-1}\Lambda$ is self-dual for the integral Hermitian pairing. We bound the number of such unmarked lattices and then count the principal markings modulo the modulus M' of the run.

Each such lattice admits a compatible positive majorant of its Hermitian form for which the squared first minimum in the sum of the f complex metrics is at least f. Here is the reduction argument. The majorants all give the same real covolume. A maximizing sequence for the first minimum stays in a Mahler compact set of lattices. Choose integral bases with convergent positive Gram matrices. In these bases the integer trace matrix of the Hermitian pairing and the matrices for multiplication by an integral basis of $\mathcal{O}_E$ have bounded entries. Indeed the trace pairing is majorized by twice the Euclidean norm product, and multiplication has operator norm at most the maximum of its conjugate absolute values. A subsequence makes all these integral matrices constant. Thus the changes of basis along that subsequence commute with $\mathcal{O}_E$ and preserve the original Hermitian pairing by trace nondegeneracy. The limiting Gram matrix is a compatible majorant and attains the maximum. If its squared first minimum were smaller than f, the shortest vectors would be mutually Hermitian orthogonal. The arithmetic–geometric mean inequality and Hermitian Cauchy–Schwarz give absolute norm strictly less than one for each integral pairing. They would therefore span a totally isotropic E-space. In a Witt decomposition scale that space up and the dual isotropic space down, by reciprocal real factors at the complex labels. All other vectors had a strict positive gap above the shortest vectors by discreteness. For a sufficiently small scaling all the shortest vectors grow and every other vector stays longer. This contradicts maximality. Compatible majorants and The perfect trace ideal comparison.

Writing $d_E = |\operatorname{disc} E|$, the Euclidean dual of the lattice equals $2J(\mathfrak{D}_E^{-1}L)$ in this metric, where J is the norm-preserving Hermitian signature involution. This follows by expressing the real trace form as twice the pairing with J. Its first minimum is at least $2\sqrt{f}/d_E$, since $d_E\mathfrak{D}_E^{-1}L \subset L$. Transference and selection of

r linearly independent vectors over E now give vectors in L whose Hermitian Gram matrix has conjugate absolute values at most

$$B_{\mathrm{cl}} = \max\left(1, \frac{(2fr)^4 d_E^2}{4f}\right).$$

The weaker power four in this bound is enough, since Banaszczyk's transference bound in real rank $2fr$ bounds the longest of a real spanning set by $2fr/\lambda_1(L^*)$. Pack disjoint products of discs of radius $1/2$ to see that each entry of such a Gram matrix has at most $(1 + 2B_{\mathrm{cl}})^{2f}$ possibilities. Distinct algebraic integers are separated by at least one in the maximum of their conjugate absolute values, by the norm. Given a Gram matrix A, pairing with the r vectors realizes the lattice as a submodule between $A\mathcal{O}_E^r$ and $\mathcal{O}_E^r$. Each such submodule is determined by at most r generators of its image in the finite quotient. This follows by choosing up to r generators over each local principal ideal ring and combining them by the Chinese remainder theorem. Thus there are at most $|N_{E/\mathbb{Q}} \det A|^r \le (rB_{\mathrm{cl}})^{2fr^2}$ possibilities. These two data determine the unmarked Hermitian lattice, including its rational form.

A marking on one such lattice is determined by the images of r generators of the original lattice modulo M'. The rank-r projective module is free over the finite principal-level coefficient ring, which is a product of local rings. There are at most $(M')^{2fr^2}$ choices. The complex flag fibre is connected, so the marked period-lattice classification counts the components of the full flag as well. Using $f \ge 16, L \ge 2, \log f \le 2L, \log r \le L, \log d_E \le 4fL$ and $\log M' \le 30L$ gives

$$\log C_{\mathrm{comp}} \le 2fr^2(60 + 16f)L \le 20(frL)^2 \le 41(frL)^2.$$

In the terminal definite case the same proof applies. Counting forms and interposed modules, The principal-level markings, and The run parameters.

The two small companion sections. For either available companion weight w, choose at each component one nonzero reading of a full standard theta generator in the marked half-level family. The tableau conditions ensure nonzero harmonics in the allowed columns. The theta line on the sublattice $A = \mathfrak{m}L$ has degree $q = g_0 g_1$ relative to its principal polarization with pairing E_L/q, where $\mathfrak{m}\bar{\mathfrak{m}} = (q)$ and E_L is the perfect alternating pairing on L. Indeed $A^\vee = q^{-1}A$ for E_L. The odd-level semicharacter on A is the q-th power of the semicharacter for E_L/q. The line separates jets through degree $d = \mathrm{DF}(w)$, since $3d + 2 \le q$. The cube of the principal theta line is very ample. A section nonzero at the point and sections cutting independent cotangent directions generate the jets by products in degree d. The remaining degree at least two is globally generated. This applies also on the product of the marked columns. The translates of the old principal vacuum by $q^{-1}A/L$ give q^{fr} nonzero eigenvectors with distinct characters of L/A by the perfect pairing modulo A. This equals the dimension of the theta space on A. On the product, tensor these bases to obtain the finite half-level packet of the principal marking. One element has nonzero value under the selected nonzero harmonic functional. Replacing this harmonic by the short minor-product family gives a nonzero reading with an element of that family as harmonic. Its polynomial coefficient norm is at most $\exp(140rL\,\mathrm{absA}(w))$. The marked jet and its bound use the same harmonic and finite level on every component. The sublattice and its principal polarization, The finite translated basis, The marked half-level and the jet evaluation and The short harmonic coefficient bound.

For the half-supremum transvection bound, use the actual period $M_{\mathrm{tr}} = qN_AN_Cm_0$, the block-spread bound w_{sp}, and $k_{\mathrm{tr}} = 2\xi_R$. The half-supremum Fourier argument gives squared first minimum at least $k_{\mathrm{tr}}\epsilon/N_A^2$, where $\epsilon = \min(1, \pi/(M_{\mathrm{tr}}k_{\mathrm{tr}}(fw_{\mathrm{sp}}+2)\log 2))$. The chosen levels satisfy $\log M_{\mathrm{tr}} \le 30L$, $0 < k_{\mathrm{tr}} \le 2$, $\log(fw_{\mathrm{sp}} + 2) \le 5L$, $\log N_A \le 4L$, and $\log(1/k_{\mathrm{tr}}) \le 3L$ give $\epsilon^{-1} \le e^{36L}$ and first minimum at least e^{-40L}. In that frame the Gaussian reproducing kernel and Cauchy's formula give the unscaled bound

$$\log \|s_w\|_{\sigma,\sup} \le 140rL\,\mathrm{absA}(w) + 10L\,\mathrm{absA}(w) + 200fr\,w.nL \le 400frL\,\mathrm{S}(w).$$

For these generators the lattice of the tensor vacuum is the standard tensor lattice and the probability L^2 norm is one. Gaussian kernel unfolding and packing in complex dimension $s = fr\,w.n$ bound the point

norm by $\sqrt{2}(1+4\sqrt{2s}\,R_0/(\sqrt{\pi}\,\lambda))^s$, where $R_0 = 2f^{3/2}(w.n)^2\mathrm{rd}(E)$ and $\lambda \geq e^{-40L}$ is the first minimum. Indeed the tensor lattice has first minimum at least λ/R_0 by contraction into a period column. Integrate the disjoint-ball lattice count against the one-variable Gaussian and apply the L^{2s} triangle inequality. This gives the displayed bound with room in the factor four. Taking Cauchy radii $\sqrt{m_i/\pi}$ for a monomial of derivative orders m_i shows that the order-d jet costs $(ed)^{d/2}$ after the differential and metric scale factors cancel. The logarithm of the Gaussian parenthesis is at most $55L$, using $\log(w.n) \leq L$, $\log f \leq 2L$ and $\log \mathrm{rd}(E) \leq 2L$. With $d \leq \mathrm{absA}(w)$ and $\log(1+\mathrm{absA}(w)) \leq 2L$, this proves the claimed bound including the factor two from the half-supremum point. At a conjugate embedding the good-embedding formula gives the same bound by conjugating the reading. Multiplying by the prescribed standard theta scalar makes all these generators integral and adds at most $\lambda_\star \mathrm{S}(w)$ to their log sup. In the oriented chart the standard generator pole bound is the sum of the two split-prime content losses and the ramified-place depth. The split factors of $s^{(0)}(w)$ cancel the content losses, and its rational-integer factor cancels the ramified-place depth. The local bound passes to every integral combination by the valuation inequality. This scalar and the inclusion apply to all weights on the same oriented run, including these companion weights. Transvection period and block spread, The first minimum in the actual frames, The unit-frame Gaussian formula, Gaussian and coefficient factors, and The same integral scalar. The all-weight integral inclusion.

Keep the chosen vector for each component and form their linear combination with coefficients in $\{0,\dots,C_{\mathrm{comp}}\}^{C_{\mathrm{comp}}}$. One vanishing condition excludes at most $1/(C_{\mathrm{comp}}+1)$ of this box, by varying the coordinate belonging to its nonzero chosen reading. Thus one sum is nonzero on every component. Its additional log sup cost is at most $2\log C_{\mathrm{comp}}$. The prescribed P_I dominates $5(400frL+\lambda_\star)+2\log C_{\mathrm{comp}}/D$. The size bounds and $D+4 \leq 5D$ give $\log\|t_1\|_{\sigma,\sup} \leq 2P_IDL_0$ and $\log\|t_2\|_{\sigma,\sup} \leq P_IDL_0$. Both are integral on the marked run, so the orders used in the generic height calculation are nonnegative. Simultaneous nonvanishing with bounded coefficients and The two actual auxiliary bounds. The definition of the allowance.

The scalar boundary product. For $j<m$, let $R=5^4=625$. Use the full linear R-marking above the actual run and the group $G=\mathrm{Aut}_{\mathcal{O}_E}(\Lambda/R\Lambda)$. Choose r classes x_b^0 generating mod 5 and set $J=(\Lambda/5\Lambda)^r$. For $y\in J$ let $x_b(y)=x_b^0+5\widetilde{y}_b$ mod R with fixed lifts. Every such tuple, and every G-translate of it, still generates mod 5 and mod 25. Let $t(x)$ be the normalized null value of the $[25]$-pullback of the principal theta section at the marked point x, of scalar theta degree one. On the full marking define

$$\eta_0=\prod_{g\in G}\left(\sum_{y\in J}c_y\prod_{b=1}^{r}t(gx_b(y))\right). \tag{H12}$$

The normalized odd theta lift makes the null value a rational rule in the marking and compatible with field extension. The product is invariant under changing the marking since this permutes its factors. Faithfully flat descent gives the same rational rule over the coefficient field of the original run.

Fix one complex point on each component. For each pair of a component and g one of the products of null values in the sum is nonzero. In fact changing y_b varies $[25]gx_b(y)$ by all the 5-torsion points. Let $\mathcal{L}$ be the symmetric principal theta line on this abelian variety $\mathcal{A}$. The pushforward $[5]_*\mathcal{O}$ is the sum of distinct flat five-torsion lines by its deck character decomposition. The principal polarization identifies those lines with $t_z^*\mathcal{L}\otimes\mathcal{L}^{-1}$ for $z\in\mathcal{A}[5]$. By the projection formula the pulled-back translates form a basis of $H^0(\mathcal{A},[5]^*\mathcal{L})$, with one section per character. The line $[5]^*\mathcal{L}\simeq\mathcal{L}^{25}$ is globally generated. Therefore the translates have no common zero downstairs. Choose a nonvanishing translate in each column. Avoiding all $C_{\mathrm{comp}}|G|$ vanishing equations in the coefficient box gives $0\leq c_y\leq C_{\mathrm{comp}}|G|$ with every factor nonzero at the chosen points. The group G has order at most 5^{8fr^2}, by counting matrices, and $2R$ divides its order. For the latter the kernel of reduction modulo 5 has order 5^{6fr^2}, and -1 has order two modulo 5. Hence $i_\eta=|G|/(2R)$ is an integer. The scalar theta degree of (H12) is $r|G|=g_*i_\eta$ with $g_*=2Rr$. The exact full-marking product and Componentwise nonvanishing.

On an integral Raynaud trait, obtain the R-marking after finite extension and use the integral cubical normalization. If Q is the monodromy quadratic form and $\bar{x}_b$ is the projected period-lattice class of one argument, the translated theta expansion gives the order bound

$$v(t(x_b)) \geq \min_{\chi} Q(25\chi + \bar{x}_b)/R.$$

The compatible shifts and this inequality are the same ones used for the level-five carrying factors. They hold also in residue characteristic five after extending the trait. The projected classes still generate mod 25. Representatives realizing the minima therefore generate the rational period module. For such representatives z_b the trace-one form functional $\beta_{\mathrm{mon}} \mapsto \sum_b \beta_{\mathrm{mon}}(w_0 z_b, z_b)$ is integral on the integral forms and positive on all nonzero semipositive forms. Its value on the monodromy form is $\sum_b Q(z_b)$, which is at least Φ_{fin} in the marked boundary normalization. It follows that one spanning tuple's product has order at least Φ_{fin}/R. Since each c_y is an integer, the same holds for every sum in (H12). The product has $v_a(\eta_0) \geq |G|\Phi_{\mathrm{fin}}/R = 2i_\eta\Phi_{\mathrm{fin}}$ in base units. This conclusion descends, since both valuations acquire the same relative ramification factor. It gives integrality and the effective boundary order of the same section. In the finite depth estimate for the w_k replace $c\,\Phi_{\mathrm{fin}}$ by $c\,v_a(\eta_0)/(2i_\eta)$. The Raynaud order of the descended product and The base-valuation normalization.

For a horizontal divisor, take a toroidal chart of the marked model at its generic point, pulling back to the graph of the moving map. The canonical weight lines extend there using the invariant differentials and the rigidified theta frame. Shrink along the divisor so its local parameter z, all weight frames, and the unit parts of the finitely many sections are defined. At a smooth point of its reduction outside the finite sets of bad primes, nonunits of those unit parts, and the support of nus, choose a trait with z a uniformizer. This is possible by formal smoothness on this open and a finite residue-field extension. The generic point can be taken in the abelian locus. Orders of the sections in its marked frames are the same exponents of z as on the horizontal divisor. Thus a strict failure of the order inequality at the horizontal divisor would give the same failure on this trait. This proves $b_\zeta \leq c\,\mathrm{ord}_\zeta(\eta_0)/(2i_\eta)$, and transfers nonnegative auxiliary orders as well. At the terminal stage the same argument gives $b_\zeta = 0$, using the terminal finite bound with $\mathrm{nus}(a) = 0$. The horizontal inequality in the marked frame.

At infinity use the torus estimate for a spanning level-five tuple in the marked unit frame. Write $d_* = 2fm$,

$$p_1 = 1 + fr\log(1 + 6\sqrt{frq}), \quad p_2 = \log(2(3d_*)^{d_*}), \quad p_3 = r(p_1 + p_2 + d_*),$$
$$k_0 = 4\,\mathrm{rd}(E)^7 d_*^2, \qquad A = Rrd_*/2.$$

In the thin range the product of the normalized null values is bounded by $e^{p_3}(k_0\Phi)^{rd_*/4}e^{-\pi\Phi/R}$, where Φ is the run's archimedean boundary function. In the thick range its bound is e^{rp_1} and $\Phi = 0$. The inequality $A\log(k_0\Phi) \leq \Phi + A\log(Ak_0)$ changes the thin bound to $\exp(p_3 + (A\log(Ak_0) - \Phi)/(2R))$. The same upper bound dominates the thick bound at zero. Consequently (H12) satisfies

$$\log\|\eta_0\| \leq i_\eta(P_{16} - \Phi), \qquad P_{16} = 2R(\log|J| + \log(C_{\mathrm{comp}}|G|) + p_3) + A\log(Ak_0). \tag{H13}$$

The zero Lie tableau entries of η_0 make its Lie character factor one. Use the host bound $\log q \leq 70L$, the component count $\log C_{\mathrm{comp}} \leq 41(frL)^2$, and the preceding logarithmic bounds. Then $\log|J| + \log|G| \leq 20fr^2$, $p_3 \leq fr^2(10 + 43L)$, and $\log(Ak_0) \leq 10 + 24L$. Since $4fr^2 \leq (frL)^2$, $2fr^2L \leq (frL)^2$, and $A \leq Rfr^2$, these give $P_{16} \leq 200000(frL)^2$. Thus both P_η and C_2 in the height calculation may take that allowance. To see cuspidality, take the constant Fourier coefficient along a rational isotropic line. Holomorphy and the real transvection law make that coefficient constant along the complex flow. At positive imaginary parameter t, the scalar weight metric grows by $(1 + t/t_0)^{n_\eta/2}$ at the moving label, with $t_0 > 0$, by the rank-one determinant identity. The bounded point norm of the form bounds this coefficient in the growing metric, so it vanishes. The terminal run uses the determinant estimate with no boundary product. The normalized torus estimate, The product estimate, and The two degree-two allowances. The bounded scalar cusp rule.

Gaussian estimates in the marked section coordinates

In these paragraphs D is the sum of the tableau entries of w_L. The section index resumes with flag dimension $D = f\binom{r}{2}$.

The Gaussian torus and its short directions. Work at a complex embedding, with the positive Hermitian space V of dimension $g = fr$. The Hermitian form H is linear in its first variable. The form $E_H = \operatorname{Im} H$ is integral and perfect on the O_E-stable lattice L, where E is the CM field and $f = [E : \mathbf{Q}]/2$. CM multiplication has Hermitian adjoint given by complex conjugation. The normalized norm of a level-Q theta function s is

$$|s(z)|_Q = |s(z)|e^{-\pi Q|z|^2/2}, \qquad \|s\|_2^2 = \int_{V/L} |s(z)|_Q^2 \, dz.$$

The measure here is Haar probability, or Euclidean volume on a fundamental domain since $\operatorname{covol} L = 1$. The theta law is

$$s(z + \gamma) = \alpha_\gamma e^{\pi Q(H(z,\gamma)+|\gamma|^2/2)} s(z).$$

The unitary factors α_γ satisfy the theta cocycle. In the actual polarization data they are powers of the semicharacter $\epsilon_\gamma = (-1)^{E_H(\gamma, w_0\gamma)}$, where w_0 is the fixed CM multiplier defining the characteristic. They are 1 on an O_E-stable isotropic sublattice.

Write $\nu = \operatorname{rd}(E)$, $\mathcal{R} = \nu^7$, $d_* = 2fm$, and $\Omega = 2f^{3/2}\nu$. The integers m, Q here are the run parameter and the analytic level, respectively. With D the degree of the differential polynomial, put

$$R_0^2 = 90(g + D), \qquad \tau^2 = fR_0^2\mathcal{R}^3/Q, \qquad b_0 = Q/(R_0^2\mathcal{R}^4).$$

We will use $m^4 \le f$, $f^2 \le \nu$, $\nu \ge 16$, $b_0 \ge 1$, and $d_*^2\Omega\tau^2 < 1$. At least one CM eigenspace in V has dimension at most m. Set $\lambda = \lambda_1(L)$, and set $\Phi = 2f/(\nu^2\lambda^2)$ when $\lambda < \tau$ and $\Phi = 0$ otherwise.

Choose $2f$ rationally independent $\omega_i \in O_E$ with $\max_\rho |\rho(\omega_i)| \le \Omega$. To obtain them, use the real trace metric $\sum_\rho |\rho(x)|^2$ on E. Its trace-dual lattice is the inverse different, up to conjugation. Every nonzero element y of that lattice has

$$\sum_\rho |\rho(y)|^2 \ \ge \ 2f|N_{E/\mathbf{Q}}y|^{1/f} \ \ge \ 2f/\nu^2.$$

Banaszczyk's Euclidean transference inequality $\lambda_d(M)\lambda_1(M^*) \le d$ in rank $d = 2f$ bounds the last minimum of O_E by $\sqrt{2f}\nu$. Euclidean transference.

Starting with $W = 0$, grow a rational, O_E-stable real isotropic subspace of V. Write $I = L \cap W$ and define u_W by $\langle u_W(x), a\rangle = E_H(x, a)$ for $a \in W$. Maintain the property that nonzero vectors of $u_W(L)$ have length at least $1/(\Omega\tau)$. If $x \in L \setminus I$ has distance less than τ from W, write $x = p + y$ by perpendicular projection onto W. Then $u_W(x) = 0$ since $|u_W(x)| \le |y| < \tau < 1/(\Omega\tau)$. For each i, $E_H(x, \omega_i x) = E_H(y, \omega_i y)$, by the adjoint rule and the orthogonality to W for E_H. This integer has absolute value at most $\Omega|y|^2 < 1$. Adjoin the E-span of x and take its real span with W. This adds $2f$ real dimensions and preserves isotropy. A lattice vector with a nonzero old pairing still satisfies the invariant. If its old pairing vanishes but its new pairing is nonzero, some integer $E_H(v, \omega_i x) \ne 0$ has absolute value at most $|u_{W_{\mathrm{new}}}(v)|\Omega\tau$, which proves the same invariant.

Write ρ_0 for the E-rank of the resulting rational space, and $s_0 = 2f\rho_0 = \dim_{\mathbf{R}} W$. The complex span $W_{\mathbf{C}} = W + iW$ identifies with the complexification of the real trace space. Conjugation on $W + iW$ relative to W is antiunitary and interchanges opposite CM embeddings. Each σ-eigenspace in $W_{\mathbf{C}}$ has dimension ρ_0. Thus $s_0 \le d_*$. The construction gives $W = 0$ when $\lambda \ge \tau$. It also gives $W = 0$ when one of the CM eigenspaces of V vanishes, in which case $\lambda \ge \tau$.

Set $\Lambda = u_W(L)$ and $V_B = (W + iW)^\perp$. The primitive inclusion $I \subset L$ and perfectness of E_H imply $\Lambda = I^*$, the Euclidean dual of I. Indeed every integral functional on I extends to L and is induced by the symplectic pairing with a vector of L. Put $L^0 = \ker u_W \cap L$ and $L_B = \operatorname{pr}_{V_B}(L^0)$. This is a full lattice

in V_B because the kernel of projection on L^0 is I and W is rational. Fundamental domains in the three orthogonal coordinates W, iW, V_B give

$$1 = \operatorname{covol} L = \operatorname{covol} I \, \operatorname{covol} \Lambda \, \operatorname{covol} L_B.$$

The integral symplectic form on L_B is consequently perfect. Termination of the construction gives $\lambda_1(L_B) \geq \tau$. When $s_0 > 0$, the invariant gives $\lambda_1(\Lambda)^2 \geq b_0$ since $b_0 \Omega^2 \tau^2 = 4f^4/\nu^5 \leq 1$. The vectors $\omega_i x$ at each step project perpendicular to the preceding W with norms at most $\Omega\tau$. The covolume of the lattice they generate bounds that of I from above. Transference for I bounds the last minimum of I^* by s_0/λ. It yields the lower covolume bound

$$(\lambda/d_*^2)^{s_0} \leq \operatorname{covol} I \leq (\Omega\tau)^{s_0} \leq 1.$$

Use the convention that zero-dimensional covolumes are 1.

If $s_0 > 0$ and $x_1, \dots, x_r \in L$ have O_E-span mapping onto $L/5L$, put $u_j = u_W(x_j)$. The same assertion holds for u_j in $\Lambda/5\Lambda$, because u_W commutes with O_E-multiplication. In particular the u_j generate W over the real algebra $E \otimes_{\mathbf{Q}} \mathbf{R}$. For a shortest vector $\xi \in L$, we have $\xi \in I$ by termination. The functionals $\alpha \mapsto \langle u_j, \alpha\xi \rangle$ are integral on O_E, so their trace representatives y_j belong to the inverse different. At least one $y_j \neq 0$ by the spanning assertion. Comparing trace coefficients in the eigenspaces of the antiunitary conjugation on $W + iW$, and summing Cauchy–Schwarz over one embedding from each pair, gives

$$\sum_{\rho} |\rho(y_j)|^2 \leq \tfrac{1}{2} |u_j|^2 |\xi|^2.$$

Indeed, the coefficient on each embedding is the Hermitian pairing of the two projected vectors, and a real vector has half its squared norm on one embedding from each pair. The inverse different bound therefore implies

$$\tfrac{1}{2} \sum_{j=1}^{r} |u_j|^2 \;\geq\; \Phi. \tag{129}$$

Fourier expansion along the short directions. Choose real orthonormal coordinates on W, and write $z = z_A + z_B$ for the decomposition along $W + iW$ and V_B. Let

$$U_\gamma f(z) = e^{\pi Q(H(z,\gamma) - |\gamma|^2/2)} f(z - \gamma), \qquad T_\gamma = \alpha_\gamma U_\gamma.$$

The cocycle makes T an action of L. A theta function is fixed by this action. For $a \in \Lambda$, choose a level-Q theta function β on V_B/L_B with multiplier $\alpha_\gamma e^{2\pi i \langle a, \operatorname{Re} \gamma_A \rangle}$ for $\gamma \in L^0$. The multiplier descends modulo I and satisfies the theta cocycle on L_B. The function

$$F_a(z) = e^{\pi Q z_A \cdot z_A/2 - 2\pi i a \cdot z_A} \, \beta(z_B)$$

is fixed by T_{L^0}. Hence its sum $S_a = \sum_{\gamma \in L/L^0} T_\gamma F_a$ is a theta function on V/L. The summand with imaginary coordinate $\mu = u_W(\gamma)$ has norm

$$|\beta(z_B - \gamma_B)|_Q \; e^{\pi |a|^2/Q - \pi Q |u_W(z) - a/Q - \mu|^2}. \tag{130}$$

The sum converges normally by Gaussian decay. After dividing by the holomorphic factor $e^{\pi Q z_A \cdot z_A/2}$, its frequencies on W/I are $-(a + Q\mu)$. They are mutually orthogonal. Integrating over W/I, then over V_B/L_B and an imaginary fundamental region for Λ, unfolds the Gaussian in $u_W(z)$ over W. The V_B norm is invariant under translation on the quotient. With the same probability measures this gives

$$\|S_a\|_2^2 = \operatorname{covol} I \, (2Q)^{-s_0/2} e^{2\pi |a|^2/Q} \|\beta\|_2^2. \tag{131}$$

At level 1 every theta function is obtained by this construction with $a = 0$. Indeed, division by the indicated holomorphic Gaussian gives an I-periodic function whose Fourier coefficients satisfy the multiplier on V_B/L_B. Translation by L acts transitively on the frequencies in Λ, so all coefficients are determined by the zero coefficient.

The theta-null estimate at the five-primary marking. For a principally polarized torus of dimension h, lattice M and level k, the normalized evaluation functional satisfies

$$\|\mathrm{ev}_z\|^2 \le k^h \sum_{\gamma\in M} e^{-\pi k|\gamma|^2/2}. \tag{132}$$

To see this, reproduce a holomorphic function in the Gaussian metric on $\mathbf{C}^h$ using the kernel $k^h e^{\pi k H(z,w)}$. For a theta function the integral converges absolutely because its normalized norm is bounded on the torus. The reproducing identity follows by translating z to zero and taking the angular mean of the holomorphic function under the radial Gaussian. Folding the integral over M gives the kernel on the torus, with terms differing from the translated Gaussian kernels by unitary factors. The diagonal is the squared evaluation norm. Taking absolute values gives (132). For a lattice of spacing ℓ, the number of points in a ball of radius R is at most $(1+2R/\ell)^{2h}$ by disjoint balls. Integrating this bound in the layer representation of the Gaussian, and using Minkowski with the density $e^{-t}\,dt$, gives

$$\sum_{\gamma\in M} e^{-\pi k|\gamma|^2/2} \le \bigl(1+2\sqrt{2h}/(\sqrt{\pi k}\ell)\bigr)^{2h}.$$

For $k=1$, $h\le g$, and $\ell\ge\tau$, the square root is bounded by e^{p_1}, where

$$p_1 = 1+g\log(1+6\sqrt{gQ}), \qquad p_2 = \log\bigl(2(3d_*)^{d_*}\bigr), \quad p_3 = r(p_1+p_2+d_*), \quad K = 4\mathcal{R}d_*^2.$$

Here $\sqrt{Q}\,\tau\ge 1$. When $W=0$ apply this directly on V/L.

For $W\neq 0$, Λ has spacing at least 1. Disjoint balls of radius $1/2$ and the inequality $|y-u|^2\le 2|\mu-u|^2+1/2$ when $|y-\mu|\le 1/2$ give

$$\sum_{\mu\in\Lambda} e^{-\pi|u-\mu|^2/2} \le e^{\pi/8}(2\sqrt{s_0})^{s_0} \le e^{p_2}.$$

For the first inequality integrate $e^{-\pi|u-y|^2/4}$ over these balls and use the cube of side $1/\sqrt{s_0}$ in each ball. Extract $e^{-\pi\,\mathrm{dist}(u,\Lambda)^2/2}$ from the level-one version of (130). The covolume bounds and (131) now give

$$\frac{|s(z)|_1}{\|s\|_2} \le e^{p_1+p_2+d_*}(K\Phi)^{d_*/4} e^{-\pi\,\mathrm{dist}(u_W(z),\Lambda)^2/2}.$$

The numerical uses are $(d_*^2/\lambda)^2\le K\Phi$, $K\Phi\ge 1$, and $2^{s_0/4}\le e^{d_*}$. Take any $z_1,\dots,z_r$ representing marked points, with $25z_j\in L$ generating $L/5L$ over O_E. Choose $y_j\in L$ so that $u_W(z_j-y_j)$ is a shortest representative of $u_W(z_j)+\Lambda$. The $25(z_j-y_j)$ have the same spanning property modulo 5. Equation (129) applied to them gives $\frac12\sum_j \mathrm{dist}(u_W(z_j),\Lambda)^2\ge\Phi/625$. It follows that

$$\prod_{j=1}^{r}\frac{|s(z_j)|_1}{\|s\|_2} \le \begin{cases} e^{p_3}(K\Phi)^{rd_*/4}e^{-\pi\Phi/625}, & \lambda<\tau,\\ e^{rp_1}, & \lambda\ge\tau. \end{cases} \tag{133}$$

The level-625 marking gives the arguments z_j by multiplying its torsion points by 25. Their primitive spanning property is exactly the one used here.

Jets on the complementary torus. Keep the lattice and level notation of the preceding estimates. Write $h=\dim_{\mathbf{C}} V_B$, $t_0=Qb_0$, and put

$$M_1 = 90m^2d_*^2\nu, \qquad C_\chi = 8100m^4d_*^4\mathcal{R}, \qquad K_6 = 6(2d_*)^{d_*/2}.$$

The inequalities on the run parameters used below are

$$\begin{aligned} 1\le D\le Q, &\qquad M_1\le Q/2,\\ t_0 &\ge 4D(7+\log(QD))+d_*+1,\\ t_0 &\ge g\log 2+2d_*+3+(D+1)(r\log r+3r+3). \end{aligned} \tag{134}$$

For any unitary theta multiplier on L_B and $0 \le c \le D$, we can prescribe the Taylor polynomial through degree c at 0. Let P be a polynomial of degree at most c, with squared Gaussian norm $\int_{V_B} |P(z)|^2 e^{-\pi Q|z|^2}\, dz$. Its periodization is $\mathcal{P}P = \sum_{\gamma \in L_B} T_\gamma P$ using the given multiplier. The sum converges normally. Taylor truncation is Gaussian orthogonal projection onto polynomials of degree at most c. For two such polynomials P_1, P_2, the estimate

$$|\langle U_\gamma P_1, P_2\rangle| \le 2^{h+c} e^{-\pi Q|\gamma|^2/8} \|P_1\| \|P_2\|$$

follows by splitting off one quarter of the sum $|z|^2 + |z-\gamma|^2 \ge |\gamma|^2/2$ in the product of the two Gaussian factors and applying Cauchy–Schwarz with Gaussian parameter $Q/2$. Under that parameter the squared norm of a monomial of degree l increases by 2^{h+l}, and the monomials are orthogonal. For $\lambda_1(L_B) \ge \tau$, count the shell between $j\tau$ and $(j+1)\tau$ by $(2j+3)^{2h} \le 5^{2hj}$ for $j \ge 1$. Since $Q\tau^2 \ge 90(g+D)$, the sum of the overlap bounds for $\gamma \ne 0$ is at most $1/16$. Indeed it is at most

$$2^{h+c} \sum_{j \ge 1} \exp\bigl(2hj \log 5 - \pi Q \tau^2 j^2/8\bigr).$$

This is bounded by $2^{h+c} \sum_{j\ge 1} e^{-29(g+D)j}$. Hence the truncated periodization has the form $1 + A$, $\|A\| \le 1/16$, and is invertible. The squared torus norm of a periodized polynomial equals $\sum_\gamma \langle T_\gamma P, P\rangle$ by unfolding over L_B. The same overlap estimate applies to the integrals of absolute values and justifies the unfolding. The section with prescribed jet P_0 obtained from $P = (1+A)^{-1}P_0$ therefore has squared norm at most $(17/16)(15/16)^{-2}\|P_0\|^2 \le 2\|P_0\|^2$. For $P_0 = 1$ or $P_0 = H(z_B, \xi)^c$ with $|\xi| = 1$, the monomial norm gives $\|P_0\|^2 \le Q^{-h}$, using $c! \le Q^c$. The same shell sum in (132) is at most 2, so the normalized pointwise norm of the resulting β is at most 2 everywhere. In dimension zero use the constant section 1.

Characters for a CM row of the short part. Fix a CM embedding σ of E, and suppose $\rho_0 > 0$. Construct $v_j \in \Lambda$, $1 \le j \le \rho_0$, successively as shortest vectors whose σ-projection leaves the preceding complex span. Write $\mu_j = |v_j|$. By base change from E to $\mathbf{C}$, the same vectors generate a rank-j E-subspace, with real span W_j in W. In particular every vector of Λ shorter than μ_j is in W_{j-1}. Transference for $\Lambda = I^*$ gives

$$1 \le \mu_1 \le \cdots \le \mu_{\rho_0} \le d_*^2/\lambda.$$

Set $K_0 = d_*^2 \nu$. In W_j, take t_j supported at $\sigma, \bar\sigma$, with its σ component perpendicular to the preceding span and of length $(30\rho_0^2 + 1)K_0\mu_j$. The $\bar\sigma$ component is its real conjugate in W_j. The independent vectors $\omega_i v_k$ with $k \le j$ have length at most $\Omega\mu_j$. Successive rounding in their Gram–Schmidt coordinates gives a point of the lattice they generate within distance $\sqrt{d_*}\Omega\mu_j/2 \le K_0\mu_j$ of t_j. Let a_j be a nearest point of $\Lambda \cap W_j$.

For $y \in \Lambda$ shorter than μ_j, minimality against $a_j \pm y$ and $t_j \perp W_{j-1}$ imply $|2\langle a_j, y\rangle| \le |y|^2$. For longer y, use

$$|a_j| \le \bigl(\sqrt{2}(30\rho_0^2 + 1) + 1\bigr) K_0 \mu_j \le 45\rho_0^2 K_0 \mu_j.$$

Let R_j be the length of the σ-projection of a_j perpendicular to the preceding span. It is at least $30\rho_0^2 K_0 \mu_j$. Write $p_j = \mathrm{pr}_\sigma(a_j)/R_j$. The p_j differ from an orthonormal frame by vectors of length at most $1/(30\rho_0^2)$. The preceding spans for a_j and v_j coincide, since $R_j > 0$. These bounds imply

$$|2\langle a_j, y\rangle| \le M_1 |y|^2, \qquad |a_j| \le \tfrac{3}{2} R_j, \quad 2\pi R_j \ge 1, \quad |a_j|^2 \le C_\chi \Phi. \tag{135}$$

For the last inequality use $\rho_0 \le m$, $\mu_j \le d_*^2/\lambda$, and $d_*^4 \le 8 f \nu^3$, which follows from $m^4 \le f$ and $f^2 \le \nu$.

A column with the prescribed leading jet. For an order-c column in this row, with $1 \le c \le D$, use S_{a_j} from (131) and take the complement jet equal to 1 through order c. The zero summand, with $\mu = u_W(\gamma) = 0$, has c-linear jet along this row equal to

$$\zeta \prod_{l=1}^{c} H(x_l, p_j), \qquad \zeta = (-2\pi i R_j)^c.$$

Indeed $z_A \cdot z_A$ has no quadratic term on a single CM eigenspace. The associated bilinear pairing on $W + iW$ pairs that eigenspace with its conjugate. For a unit vector ξ in the complement in the same row, use S_0 and complementary jet $H(z_B, \xi)^c$. The corresponding jet is of the same form, with p_j replaced by ξ and $\zeta = c!$. An unused column takes $a = 0$ and the complementary value 1 at 0.

Here is a uniform error bound from all summands with $\mu \neq 0$. For $|z| \leq R$, (130) and (135) bound the ordinary absolute value of the summand by

$$2 \exp\bigl(\pi Q R^2/2 + 2\pi R(|a| + Q|\mu|) - (\pi/2)Q|\mu|^2\bigr).$$

Use $R = c/(2\pi|a| + 5\sqrt{Q(c+1)})$ for $c > 0$, and $R = 0$ for $c = 0$. The inequality $2R|\mu| \leq 20R^2 + |\mu|^2/20$ gives the exponent $20.5\pi Q R^2 + 2\pi R|a| - .45\pi Q|\mu|^2$. As $QR^2 \leq (c - 2\pi R|a|)/25$ and $Q|\mu|^2 \geq t_0 \geq 28c$, this is at most $1.07c - 1.33Q|\mu|^2$. For unit directions x_l apply Cauchy's inequality on $z = \sum_l t_l x_l$, $|t_l| \leq R/c$, to the coefficient of $t_1 \cdots t_c$. Its loss is $(c/R)^c$. Put $N = 2\pi R_j$ for a character column and $N = 1$ for the other columns. In each case $c/R \leq 6.5N\sqrt{Q(c+1)}$. For $c = 0$ use the value estimate without a Cauchy factor. From (134),

$$e^{1.07c}(6.5\sqrt{Q(c+1)})^c \leq e^{.14t_0}.$$

For example, the logarithm on the left is at most $c(3.3 + \frac{1}{2}\log(QD))$ since $c + 1 \leq 2D$ for $c > 0$. Counting shells in Λ of spacing at least $\sqrt{b_0}$ gives $\sum_{\mu \neq 0} e^{-1.33Q|\mu|^2} \leq 2{\cdot}5^{s_0} e^{-1.33t_0}$. Thus for all choices of unit directions the jet error is at most ϵN^c, where

$$\epsilon = 7 \cdot 2^g (4.3)^{d_*} e^{-t_0}.$$

This follows from $4 \cdot 5^{s_0} e^{-1.19t_0} \leq \epsilon$, using $t_0 \geq s_0$ and $5e^{-.19} < 4.3$. When $\Lambda = 0$ there are no such summands. The squared norm of each column is bounded by $2e^{2\pi|a|^2/Q}$ by (131). In particular it is at most $K_6 e^{2\pi C_\chi \Phi/Q}|\zeta|^2$ for a used column. For an unused column, the squared norm is at most K_6 and its absolute value at zero is at least $1/2$.

The determinant differential operator. Let the polynomial be a product of powers of determinant blocks, each block of size i in a CM row, with determinant power $c \geq 1$. Different blocks occupy disjoint columns among the n columns. The total degree is $D = \sum ic$. For a given row the p_j and orthonormal complement vectors form a basis whose column matrix M differs from a unitary matrix by rank at most ρ_0 and norm at most $1/(30\rho_0^{3/2})$ when $\rho_0 > 0$. For an orthonormal i-row projection A, the matrix AMM^*A^* differs from the identity by rank at most $2\rho_0$ and has eigenvalues at least $(1 - 1/(30\rho_0^{3/2}))^2$. Its determinant is thus at least $(1 - 1/(30\rho_0^{3/2}))^{4\rho_0} \geq (14/15)^2$. The bound also holds with an orthonormal basis in the absence of short directions. Cauchy–Binet and the fact that the row dimension is at most r give a choice of i columns whose projected minor has modulus at least $\delta_0 = 2^{-r/2}/2$. Make this choice in every block and take the tensor of the corresponding sections in their prescribed column slots.

Divide the evaluated jet by the product of the leading scalars ζ and the actual values of unused columns. The leading term is the product of the chosen minors to their block powers and has modulus at least δ_0^D. All leading column vectors have length at most 2, and the normalized column jet errors are at most ϵ since $|\zeta| \geq N^c$. Expanding determinants into permutations gives at most r^D terms. Telescoping within each term gives a total error at most $D\epsilon(4r)^D$. This is at most $\delta_0^D/2$ by the second lower bound on t_0 in (134). Indeed $\log(14D) \leq 3 + D$, $\log 4.3 < 2$, and $(3 + r/2)\log 2 + \log r + 1 \leq r \log r + 3r + 3$. Normalize the tensor in its L^2 norm. There is at least one used column. Thus the final factor $1/2$ and the factors $1/2$ for each unused column cost at most 2^{-n} in total. For the resulting level-Q section F of squared norm 1, the differential operator $\mathcal{D}_W$ satisfies

$$|\mathcal{D}_W F(0)|^2 \geq J_{\mathrm{bnd}} := (4K_6)^{-n} 2^{-(r+2)D} \exp(-2\pi n C_\chi \Phi/Q). \tag{136}$$

The parameters and norms on the two runs. At a stage j of either geometric run use its rank $r = k + m - 1$, signature $(k + j - 1, m - j)$, and the $n = n_{w_L}$ columns of its selected block polynomial of degree $D = \mathrm{DF}(w_L)$. The analytic theta level is $Q = g_0 g_1$. The two half ideals are generated by $\alpha, \bar{\alpha}$, where $\alpha\bar{\alpha} = Q$, and the chosen marking is modulo $(\bar{\alpha})$. This convention applies in the marked coordinates of each run.

In either run let $\mathcal{P}$ be the chosen polynomial bound for the run parameters and $q_* = \min(g_0, g_1)$. Its size and gaps satisfy

$$D \geq 3, \qquad \mathcal{P} \geq 90(fr + 1), \qquad \mathcal{P} \geq r, \qquad q_* \geq \mathcal{R}^4 \mathcal{P}^3 D^2,$$
$$\mathcal{P} \geq 16200\pi n m^4 d_*^4, \qquad c = \mathcal{R}\,\mathcal{P} D^2 / q_*.$$

In particular $Q \geq \mathcal{R}^4 R_0^2 D r$ since $R_0^2 \leq \mathcal{P} D$. The numerical implication from this lower bound on Q, with $f^2 \leq \nu$, $m^4 \leq f$, and $\nu \geq 16$, gives the conditions used in the short decomposition and (134). Also $\pi n C_\chi / Q \leq c$ by the last lower bound on $\mathcal{P}$ and $Q \geq q_*$. The selected weight parameters and The numerical implication for the analytic level.

For each complex embedding satisfying the prescribed CM choice, use the compatible marked uniformization and a unit frame adapted to its flag. Its Hermitian matrix is $2\xi_R I$, where $\xi_R > 0$ is the absolute size of the imaginary polarization scalar at every CM embedding. The change $\psi = \sqrt{2\xi_R}$ times the row identification makes the analytic torus Euclidean. It identifies its lattice with the period lattice and its symplectic pairing with the perfect integral Riemann form. Relabelling its CM embeddings preserves multiplicities and complex conjugation. In each conjugate pair the smaller multiplicity is at most $m - j \leq m$. The rank of the short isotropic space over E is bounded by that multiplicity. The resulting Φ is the boundary function of the marked domain at that point. The distinguished level-one theta function in this uniformization has squared L^2 norm 1, with Haar probability and the Gaussian norm. It comes from the normalized theta frame already used for the marked evaluations. The Euclidean change of frame and Relabelling the CM embeddings.

The marked nulls in the scalar product. For a lift $x \in L$ of a full 625-marking argument, the torsion point is represented by $x/625$. The odd cubical null at this point has absolute value $|s(x/25)|_1$. The factor 25 comes from the cubical isomorphism $[25]^*\mathscr{L} \simeq \mathscr{L}^{625}$. It is an isometry for the rigidified Gaussian metrics, since their curvatures agree and the ratio is 1 at the origin. The normalized torsion lift has finite order. It differs from the unitary Fock translation by a scalar, and finite order forces that scalar to have modulus 1. This gives the stated cubical null norm in the holomorphic frame. Every tuple of arguments in the product, including all full-marking translates, generates the marked period module modulo 625. Its lifts therefore generate $L/5L$. Formula (133) applies to every term of every factor. The scalar boundary weight has zero Lie tableau entries, so the theta frame already gives its pointwise weight norm. The marked product estimate applied to these absolute values gives $\log \|\eta_0\| \leq i_\eta(C_2 - \Phi)$ with the coefficient and factor count of the descended product. Cubical null evaluation and The arguments of every factor.

The entire marked basis at the selected derivative. On the Euclidean torus let π act by multiplication by α. It preserves the lattice, commutes with the CM action, and scales H by Q. The characteristic obeys $\epsilon_{\pi\gamma} = \epsilon_\gamma^Q$ because its exponent is the pairing with the fixed w_0. The pullback $s \circ \pi$ is therefore a level-Q section of squared norm 1, by Haar probability under the isogeny. For b ranging over $Q^{-1}L/\pi^{-1}L$, the translates

$$F_b = U_b(s \circ \pi)$$

remain level-Q sections of norm 1. They are pairwise orthogonal. Indeed U_a for $a \in \pi^{-1}L$ acts on $s \circ \pi$ by a unit scalar and is unitary at level Q. Its eigenvalues on $F_b, F_{b'}$ differ by the character $e^{-2\pi i Q E_H(a, b-b')}$. If this is 1 for all such a, then $Q(b - b') \in (\pi^{-1}L)^\# = \bar{\alpha}L$ for the symplectic dual. This is precisely $b - b' \in \pi^{-1}L$. The quotient has Q^g elements, which is the dimension of the level-Q theta space by Riemann–Roch and

ample vanishing. Tensoring these bases gives an orthonormal basis on V^n/L^n. The squared norm of a linear functional is its full squared basis sum. Applying this to $\mathcal{D}_W$ and the section of (136) gives

$$\sum_{\mathbf{b}} |\mathcal{D}_W(\otimes_{i=1}^n F_{b_i})(0)|^2 \;\geq\; J_{\mathrm{bnd}}. \tag{137}$$

The actual half-level classes index that entire basis. Let mk be the marking in the Euclidean period coordinates. It is O_E-linear and an isomorphism modulo Q. For a half-level class choose the defining even lift x_2 in $2\alpha^2$ times the marked module, congruent to the given class modulo $\bar{\alpha}$. Such a lift exists since Q is odd and $(\alpha), (\bar{\alpha})$ are coprime. Put $b = -\bar{\alpha}\, Q^{-2}\mathrm{mk}(x_2)$. For $x_2 = \alpha^2 y$, its class under multiplication by Q is $-\alpha\,\mathrm{mk}(y)$ in $L/\bar{\alpha}L$. This is a bijective reindexing of the given class, since α is invertible modulo $\bar{\alpha}$. The even lifts and the actual half-level representatives.

In the original Hermitian coordinates let a denote the period of $Q^{-2}x_2$, so $b = -\bar{\alpha}\psi(a)$. With $\theta = s \circ \psi$, substitution in the Fock translation gives the identity

$$F_b(\bar{\alpha}\psi(z)) = e^{-\pi Q^2(H_m(z,a)+|a|^2_{H_m}/2)}\theta(Q(z+a)), \qquad H_m = 2\xi_R I.$$

The algebraic marked jet uses this cubical germ and the factor Q^{-D} from the tangent map of multiplication by Q. This comparison uses chart independence for the fixed harmonic differential in characteristic zero. The change of exponential gauge pairs only conjugate CM directions, and these contractions vanish on the harmonic. Thus this is the same reading as in the prescribed algebraic chart. The harmonic chart comparison and The marked cubical germ. Every monomial in the chosen block polynomial has fixed degree in each CM row. The $\bar{\alpha}$ factors therefore form a common scalar of squared modulus Q^D. The squared coordinate factor from ψ is $(2\xi_R)^D$. If $f_{\mathbf{b}}$ is the raw marked theta rule, this proves

$$|f_{\mathbf{b}}|^2_{\mathrm{raw}} = (2\xi_R/Q)^D |\mathcal{D}_W(\otimes_i F_{b_i})(0)|^2.$$

Finally the actual family w_k contains the full half-level family $s_W f_{\mathbf{b}}$ at distinct indices. The stipulated scale satisfies $|\sigma(s_W)|^2 \geq Q^D$ at every complex embedding. In the unit flag frame the squared Lie normalization is $(2\pi\xi_R)^{-D}$. Writing $\beta^2 = \sum_k \|w_k\|^2$ in the marked weight norm, we obtain

$$\beta^2 \geq (2\pi\xi_R)^{-D}|\sigma(s_W)|^2(2\xi_R/Q)^D \sum_{\mathbf{b}} |\mathcal{D}_W(\otimes_i F_{b_i})(0)|^2 \geq \pi^{-D} J_{\mathrm{bnd}}. \tag{138}$$

The actual selected weight and scale and The pointwise weight normalization.

Taking logarithms and using $D \geq 1$ gives

$$-\log\beta \leq C_5 D + c\Phi, \qquad C_5 \geq p_{14} = \tfrac{1}{2}\big(\log\pi + n\log(4K_6) + (r+2)\log 2\big), \qquad c \geq \pi n C_\chi/Q.$$

Orthogonalizing an arbitrary adapted flag frame multiplies each reading by its weight character, which cancels the frame factor in the pointwise norm. Each remaining complex embedding is the conjugate of one with the prescribed CM choice. Conjugation carries the marked norm and period minimum to the companion frame and keeps the estimates pointwise. At the terminal stage the signature has a zero CM multiplicity, hence the short part is zero and $\Phi = 0$. The polarization scalar ξ_R is still positive, and (138) applies with that same normalization and J_{bnd} with $\Phi = 0$. At a nonterminal stage the null bound gives $\Phi \leq C_2 + i_\eta^{-1}(-\log\|\eta_0\|)$, supplying the pointwise comparison in the moving height calculation. Frame invariance and conjugation and The terminal value of the boundary function.

The normalized fibre of the scaled section system. Fix one run and one stage. Use $w_1, \ldots, w_N$ generating the scaled full module A'_L and the common weight-w_F cusp section g'_0 that is nonzero on every component. Let Ω_Z be the function field of a component Z of the full flag variety, and let D be its dimension. Read sections as functions in a generic frame. An admissible theta germ of weight w_F has a nonzero evaluation at a complex frame of each component, using the available w_F harmonic polynomial

and $3\,\mathrm{DF}(w_F) + 2$ within the marked level. Avoiding the finitely many proper vanishing subspaces gives the common g_0'. Scaling by a nonzero integral scalar at this step is harmless. The simultaneous nonzero cusp section.

Here is the rank argument for the actual w_L ratios. In the nonterminal case let $A = i_\eta g_*$ with $g_* = 2 \cdot 5^4 r$. Write $S = \sum_k \|w_k\|^2$ at the fixed good complex embedding. The lower bound (138) and the scalar product bound (H13) give $S \gg \exp(-2c\phi)$ and $\|\eta_0\|^2 \ll \exp(-2i_\eta\phi)$. The first-minimum function ϕ is bounded below by a positive multiple of the inverse squared minimum where the minimum is small. The inequality $2cg_* \le 1$ gives $i_\eta - Ac > 0$. Hence every positive superlevel of $\|\eta_0\|^2/S^A$ has bounded ϕ, and reduction theory makes it compact modulo the arithmetic group and change of frame. The generator bound gives $S > 0$, and η_0 is nonzero on every component, so the ratio has a positive interior maximum. The positive row gaps and scalar weight bound give strict dominance of $Aw_L - w_\eta$, where w_η is the weight of η_0. In holomorphic frames write the Hermitian metrics as $\|s\|^2 = |s|^2 e^{-\phi_w}$. At this maximum $\log|\eta_0|^2$ is pluriharmonic and

$$A\,\mathrm{i}\partial\overline{\partial}\log\bigl(\textstyle\sum_k |w_k|^2\bigr) \;\ge\; \mathrm{i}\partial\overline{\partial}\phi_{Aw_L - w_\eta} > 0.$$

The form on the left is A times the pullback of the Fubini–Study form. Its positive definiteness on the full flag proves rank D. In the definite terminal case use $\eta_0 = 1$ and $A = 1$ on the compact quotient, with w_L strictly dominant. In both cases a nonzero local Jacobian gives algebraic independence of D ratios over K_0. The first-minimum estimates on the selected family · Dominance from the selected row gaps · Componentwise rank in both signatures.

The full transcendence degree gives a finite extension of the ratio field in Ω_Z. Given $b \in \Omega_Z$, take an algebraic equation over this ratio field and clear its denominators homogeneously. For some $k > 0$ a nonzero homogeneous polynomial $P(w)$ of degree k then makes $P(w)b$ integral over the homogeneous section ring. Its monic equation has coefficients of weight $(d_{\mathrm{eq}} - i)kw_L$ in degree $i < d_{\mathrm{eq}}$, where d_{eq} is the degree of the equation. Normality in local frames extends the numerator over the flag charts. The bound for the roots of a monic polynomial gives boundedness in the metric of kw_L. The constant term in the cusp expansion of the monic equation is the d_{eq}-th power of that of the numerator. The other terms have positive-weight cuspidal coefficients and zero constant term. Thus $P(w)b$ extends to a cusp section of weight kw_L. This argument represents component constants as quotients of sections. The weight-zero component projectors give sections supported on individual components. Clearance in homogeneous weights · Integral homogeneous cusp sections · Component projection of sections.

Write z_{ik} for $(D+1)N$ independent variables, $u_i = \sum_k z_{ik} w_k$, and $y_i = u_i/u_0$ for $1 \le i \le D$. Use the generic common zero fibre of the y_i over $K_0(z)$. On a component choose b with $w_b \ne 0$. The incidence equations eliminate z_{ib} for $i \ge 1$ by

$$z_{ib} = -\sum_{k \ne b} z_{ik}\, w_k/w_b.$$

The resulting field is $\Omega_Z(z_{\mathrm{free}})$, where the free variables include the whole row $i = 0$. The image of the base variables is algebraically independent. Indeed, the wedge of their differentials over K_0 is, up to sign, the wedge of the free differentials followed by the D rows $\sum_{k \ne b} z_{ik}\, d(w_k/w_b)$. The latter wedge is nonzero because the ratios have full transcendence degree in characteristic zero and each row has its own free variables. The extension over $K_0(z)$ is therefore finite separable.

For the section ratio algebra generated by all s/u_0^k of weight zero, the open sets $w_b/u_0 \ne 0$ cover since $\sum_b z_{0b}\, w_b/u_0 = 1$. On such a chart and one component start with the subring of Ω_Z generated over K_0 by the s/w_b^k. Adjoin the parameters and invert u_0/w_b. Before passing to $K_0(z)$, eliminating the z_{ib} gives a polynomial ring in the free variables over this subring, with u_0/w_b inverted. Passing to $K_0(z)$ inverts nonzero elements by the algebraic independence just proved. The resulting subalgebra over $K_0(z)$ in the finite extension $\Omega_Z(z_{\mathrm{free}})$ is a field. Multiplication by a nonzero element is injective on a finite dimensional vector space and hence surjective. Its fraction field is $\Omega_Z(z_{\mathrm{free}})$ by the homogeneous presentations above. All nonempty charts on Z describe the same incidence field because their overlaps contain its generic

point. The component projectors separate these fields. Thus the reduction of the section ratio algebra is exactly the finite reduced product, also after extending coefficients.

The fibre projects generically to each component, so g_0' is nonzero throughout it. The auxiliary nonzero cusp section of weight w_2, with $w_F + w_2 = (D+4)w_L$, is nonzero there as well. Multiplying the numerator and denominator of $F/(g_0' u_0^k)$ by that section expresses the quotient in the ratio algebra localized along the fibre. In particular this identifies the localized section module with the free module on g_0'. A denominator invertible modulo $(y_1, \dots, y_D)$ has an inverse modulo every power by a finite geometric series. The same holds for its action on the truncated section module. Generic section with finite reduced fibre · Auxiliary section of complementary weight · Componentwise localization and the fibre generator.

At a coefficient place a, use $O_a' = \mathcal{O}_{K_{0,a}}[z]_{(\varpi_a)}$ and its fraction field $K_{0,a}(z)$. The Gauss field $K_0(z)$ is dense in $K_{0,a}(z)$ for this valuation and the two completions are the same. In a finite separable algebra the branches after completing are its local field completions. Consequently the branches over the Gauss place before and after this coefficient change correspond, with the same residue fields and normalized value groups. Write ζ for a branch on Z and ω_ζ for its restriction to Ω_Z.

Choose b at ζ so that all w_k/w_b are integral for ω_ζ. The residues of the full tuple z are algebraically independent over the residue field of a, and the eliminated residues are expressions over the residue field of ω_ζ in the free residues. The latter residue field has transcendence degree at most D over the residue field of a. This bound follows by lifting an algebraically independent set of residues to units of Ω_Z. Hence the free residues are algebraically independent over the residue field of ω_ζ. This determines ζ as the Gauss extension of ω_ζ in the free variables. In particular its value group already occurs on Ω_Z, its residue field is generated by the free residues over the residue field of ω_ζ, and distinct branches on a component restrict to inequivalent valuations. The free row $i = 0$ shows that u_0/w_b has order zero. Thus

$$x_k = w_k/u_0 \text{ is integral at every branch,} \qquad \sum_k z_{0k} x_k = 1.$$

Independence over the component residue field · Reading the generic point on its component.

The generic point on Z is read in the completion of the normalized DVR of ζ. At the chosen level the proper marked model extends the base moduli point over this DVR. The integral flags and frames are constructed type by type. For each τ, intersect the rational τ-eigenspace with the integral Lie lattice of the semiabelian model. Intersect the members of the rational τ-flag with that lattice. The intersections are saturated and their successive quotients are free over the DVR. Choose an adapted integral basis in each τ. Invertible pushforward of the marked theta bundle provides its local generator over this same DVR. The cubical/frame comparison attaches these bases and the theta generator to the framed valued point of the marked cover. Passing to the completion preserves a uniformizer. The change from the generic frame to the integral frame multiplies a section of weight w by its weight character. In the quotient of a weight-$(w_F + k w_L)$ section F by $u_0^k g_0'$, the characters cancel, because they multiply under weight addition. If c_ζ is the additive order of g_0' in the integral frame, the fibre order of F/u_0^k is

$$\operatorname{ord}_\zeta(F/(u_0^k g_0')) + c_\zeta = \operatorname{ord}_{\text{frame}}(F) - k \operatorname{ord}_{\text{frame}}(u_0).$$

Use $+\infty$ for a zero value. The integer $\operatorname{ord}_{\text{frame}}(u_0) = e_\zeta v_\zeta$ has the normalization of the scaled w_L module, since the x_k are integral and at least one is a unit in each branch. The pointwise depth bound applies in this frame even when the residue field has transcendence degree. Indeed the marked integral flag charts are finite type over the completed coefficient DVR. For a finite list of sections their frame values are given by rational functions on such a chart. A strict violation of integrality, or of an upper bound on the minimum generator valuation of the form $\nu + c' v(\eta_0)$ with $c' \geq 0$, is witnessed by finitely many rational valuation bounds. For the second assertion with $v(\eta_0)$ finite, choose upper and lower bounds $v(\eta_0) \leq t$ and $v(w_k) \geq u_k$ with $\min u_k > \nu + c' t$. Include bounds keeping all necessary denominators nonzero. After a finite ramified extension to clear the radii, these restrictions define a nonempty affinoid domain of the generic fibre of the completed chart. A closed point of this domain has field finite over the coefficient completion

by the affinoid Nullstellensatz. Its integral framed point has the same strict violation, contradicting the bounds on algebraic valued points. [120, Corollaries 12–13 of §2.2] This extension argument gives $v_\zeta \le \operatorname{nus}(a) + c\, v_\zeta(\eta_0)/(2i_\eta)$ for $j < m$, and $v_\zeta \le \operatorname{nus}(a)$ for $j = m$. Integral frames on the normalized branch · Integral tangent lattices by type · The integral theta frame · Frame comparison and additive orders · Normal finite-type charts and transfer from algebraic points · Depth of the distinguished generator.

Integral Hasse coefficients. Use Hasse derivatives of the coefficient parameters z_{ik}, holding the component functions and sections fixed. First consider the subring of ratio germs generated by integral Gauss coefficients and tuples of component constants integral at every branch. Adjoin inverses of elements that restrict to units at all branches. Every element of this ring has integral restricted Hasse coefficients. For the Gauss coefficients this follows by Taylor expansion of a polynomial fraction with primitive denominator. The components are constant under parameter translation. Multiplication and inversion of a series with unit constant coefficient preserve the assertion at every branch.

The image of this ring on the reduced fibre is the integral closure of O'_a there. To see surjectivity, use weak approximation on each Ω_Z at the finitely many distinct ω_ζ. It provides a component constant with prescribed order at one branch and order at least $e_{\zeta'}$ at every other branch. It also supplies constants with any prescribed residue in the component residue field at a branch, integral at every branch. Write an arbitrary residue at ζ as a quotient of polynomials in the free residues with coefficients in the component residue field. Lift the coefficients by this approximation. If P, Q are the lifted polynomials and h has residues 1 at ζ and 0 elsewhere, then $hQ + 1 - h$ is a unit at all branches. The germ $hP/(hQ + 1 - h)$ gives the requested residue at ζ and zero elsewhere. Combining this with the prescribed orders lifts each graded piece modulo ϖ_a of the integral closure, since its local component has ϖ_a-order e_ζ. The image is an O'_a-submodule of a finite O'_a-module. Surjectivity modulo ϖ_a gives surjectivity by Nakayama's lemma.

Use weak approximation again to multiply g'_0 by a tuple of component constants of order $-c_\zeta$ at every branch. This gives a generator of the normalized integral section fibre A_0. Its higher parameter derivatives vanish. Multiplying it by the ratio lifts gives representatives of every element of A_0 with all restricted Hasse coefficients integral. A representative in the localized section ring can be replaced modulo any $(y_i)^u$ by a representative before localization as above. A Hasse derivative of order s takes $(y_i)^u$ into $(y_i)^{u-s}$ for $s \le u$, by the product formula. Thus the finite-order properties used here are preserved in the original truncated jet module. Integral closure and approximation on all branches · Lifting the normalized section fibre.

Let $J = (y_i)$ and use jets of order less than $T - D$. In degree l the leading term of a representative in J^l is a sum $\sum_{|\beta|=l} y^\beta c_\beta$ modulo J^{l+1}, with c_β in the section fibre. Translate the variables z_{ik} for $i \ge 1$ in a single column k, keeping $z_{0k'}$ fixed. Then y_i changes by $t_i x_k$. The restricted coefficient of t^β in degree l is $x_k^l c_\beta$. The x_k generate the unit ideal in the integral closure. For $l > 0$ any maximal ideal containing all x_k^l would contain every x_k. For $l = 0$ the powers are all 1. Therefore there are integral coefficients a_k with $\sum_k a_k x_k^l = 1$. The resulting linear combination reads c_β integrally, with any allowed lower bound on each branch order. It also proves that a relation $\sum y^\beta c_\beta = 0 \bmod J^{l+1}$ has all leading coefficients zero. In the reverse direction, each integral c_β lifts as above, and multiplication by y^β preserves integrality of restricted Hasse coefficients. This uses the explicit parameter translation

$$y_i(z + t) = \big(y_i(z) + \textstyle\sum_k t_{ik} x_k\big)\big(1 + \sum_k t_{0k} x_k\big)^{-1}.$$

All statements about restricted coefficients descend to the quotient by J^{T-D} by the product formula.

For a scaled integral form $\varpi_a^{E_a} F$ of weight κ, the raw derivatives of $\varpi_a^{E_a} F / u_0^T$ have fibre order at least $e_\zeta(E_a - T v_\zeta)$. Indeed F in a fixed section basis has integral parameter coefficients, and the Taylor series of the inverse power of $1 + \sum_k t_{0k} x_k$ is integral. This gives the pole-allowed image with bound $m_\zeta = \max(0, e_\zeta(T v_\zeta - E_a))$. The restriction of N_1 contributes the classes $y^\beta \theta'_i$. Their coefficients at each degree span the restriction lattice generated by the short θ'_i in every monomial slot. Their rational linear independence on the fibre implies independence in the jet module by choosing the least degree of

a nonzero relation. The order-$(T-D)$ products vanish in that module. This verifies that the kernel on the rational span of the two sets of products is exactly the span of the latter set.

Filter by total degree. In degree $l < T-D$, the length of the quotient by the image of N_1 is at most $\binom{D+l-1}{l} I_a$. The length added by pole allowances is at most $\binom{D+l-1}{l} \sum_\zeta f_\zeta m_\zeta$, using the quotient of two rank-one balls in each branch. Summing gives the two bounds with multiplier $\binom{T-1}{D}$. The argument is unchanged on the full flag of the terminal definite stage. At that stage the bound on v_ζ gives $m_\zeta = 0$ for every ζ. Integral reading in coefficient columns · Leading Hasse coefficients · Kernel and pole-allowed image.

The rank of the two section systems. Fix a run and use its scaled sections on the marked cover. Let $V = C_\kappa$, let K_0 be its coefficient field and put $t_0 = T - D$. Denote the generic intersection by J. Its finite separable algebra over $K_0(z)$ has dimension h. Its points at a complex specialization with algebraically independent coefficients are the reduced intersection of $u_1, \dots, u_D$. They lie in the open flag varieties and u_0 is nonzero at each point. The projective boundary separation proved below ensures that this same generic zero set is in a compact part of the marked flag model and is transverse on every component. It also gives $A \geq b_J S$ away from small neighbourhoods of the zero set, for $A = \sum_i \|u_i\|^2$, $S = \sum_k \|w_k\|^2$ and a positive constant b_J.

At a nonterminal stage define $C_{\mathrm{deep}} \subset V$ by requiring the vanishing of all Fourier–Jacobi blocks in the cut. For a cusp and its positive index form, a block is in the cut if some nonzero vector of the dual cusp lattice has index trace at most

$$c_1 = \tfrac{1}{2}\,\mathrm{rd}(E)^7\; Tc\,(2f\rho)^2\rho, \qquad \rho = m - j.$$

Equivalently, its readings at rank-one isotropic directions have no modes of trace at most c_1, in the half-trace convention. The identification follows by grouping the Fourier expansion according to the restriction of its index form to the chosen line and applying uniqueness of the remaining Fourier expansion. At the terminal stage the signature is definite and there are no such directions. Then $C_{\mathrm{deep}} = V$. Write $r_Q = \dim_{K_0}(V/C_{\mathrm{deep}})$. The Fourier–Jacobi cut · Identification of the two expansions.

Rank-one Fourier estimates at the marked level. For a nonzero lattice vector ξ isotropic for the Hermitian form h, use the holomorphic rank-one flow $x \mapsto x_s$ at the archimedean factors. At τ it is the transvection

$$v \longmapsto v - ic_\tau s_\tau\, h_\tau(v, \xi_\tau)\xi_\tau, \qquad t_\tau(x) = \frac{1}{c_\tau\ \mathrm{maj}_{x,\tau}(\xi_\tau)}.$$

Here $\mathrm{maj}_{x,\tau}$ is the positive definite majorant from the positive and negative orthogonal decomposition at the flag x, and $c_\tau > 0$ is the fixed factor in the flow normalization. Write $y_\tau = \mathrm{Im}\, s_\tau$. The change of the Gram matrix on the first j directions is a rank-one update. Its determinant is multiplied by

$$1 + 2c_\tau y_\tau m_{\tau,j}, \qquad m_{\tau,j} = h_\tau(\mathrm{pr}_{V_{\tau,j}}\,\xi_\tau, \mathrm{pr}_{V_{\tau,j}}\,\xi_\tau).$$

For $j \leq p_0$ these partial projection lengths increase, and for $j \geq p_0$ they decrease to zero by isotropy. Thus $0 \leq m_{\tau,j} \leq m_{\tau,p_0}$ and $\mathrm{maj}_{x,\tau}(\xi_\tau) = 2m_{\tau,p_0}$. In particular all flag signs persist for $y_\tau > -t_\tau(x)$. The rank-one inverse formula on the positive plane gives $m'_{\tau,p_0} = m_{\tau,p_0}/(1 + 2c_\tau y_\tau m_{\tau,p_0})$, hence $t_\tau(x_s) = t_\tau(x) + y_\tau$. Normalized flow and heights · The domain of the flow.

The real periods contain the embeddings of a fixed integral multiple of $\mathcal{O}_{E^+}$, uniformly in the lattice vector ξ. Indeed one common multiple clears the Hermitian denominators and the principal level in the transvections, and multiplying by the order of the weight character kills the character. This gives the full real period lattice used in the rank-one readings. On its dual frequencies the expansion of a section P is

$$P(x_s) = \sum_\nu a_\nu(x)\, e^{4\pi i \sum_\tau \nu_\tau s_\tau}.$$

The factor 4π uses the half-trace convention for ν. Integrating $P(x_s)e^{-4\pi i\sum_\tau \nu_\tau s_\tau}$ on the real period torus and shifting it to imaginary part y gives

$$|a_\nu(x)|\, e^{-4\pi\sum_\tau \nu_\tau y_\tau} \le \sup_{u\in\Pi} |P(x_{u+iy})|, \tag{139}$$

where Π is a bounded period parallelepiped. This is Cauchy's theorem with cancellation on the opposite faces. The norm factor $n_w(x)$ is defined by $\|P(x)\|_w = |P(x)|n_w(x)$. It is a fixed product of real powers of the absolute Gram minors. As y_τ grows along any of the positive coordinate directions, the minor formula makes its reciprocal grow at most polynomially, uniformly on Π. For a bounded section (139) therefore forces $\nu_\tau \ge 0$ at every place whenever $a_\nu(x) \ne 0$. Products of bounded sections can be expanded using a common period sublattice. The Fourier frequencies add. This proves that multiplication by the weight-w_L sections preserves vanishing at every half-trace at most c_1. The common real periods · Fourier integration after imaginary translation.

For $w = w_F + kw_L$, $k \ge 0$, the indefinite rows are nonincreasing within each sign block and satisfy $w_{\tau,0} < w_{\tau,r-1}$ at a real place of mixed signature. Flow only at that place with $y \ge 0$. Put $L_j = \log(1 + 2c_\tau y m_{\tau,j})$. The logarithm of $n_w(x)/n_w(x_s)$ is $\frac{1}{2}\sum_j w_{\tau,j}(L_{j+1} - L_j)$. By summation by parts the sum before the factor $1/2$ is at most $(w_{\tau,0} - w_{\tau,r-1})L_{p_0}$: the differences at indices other than p_0 are nonnegative and $L_j \le L_{p_0}$, with $L_0 = L_r = 0$. Consequently the norm factor ratio is bounded by $(1 + y/t_\tau(x))^{-1/2}$. Using (139) with $\nu = 0$ and the bounded Petersson norm forces $a_0(x) = 0$. This is the cusp condition in the marked flag model. The definite components have no isotropic directions. The marked gaps, including the shifted weights · The constant coefficient bound.

Decay and the initial weighted remainder norm. Let P be bounded and have no coefficients of half-trace at most c_1. In (139) shift down by $y_\tau = -3t_\tau(x)/4$ at all places. All the Gram determinant factors are in $[1/4, 1]$, so the norm factor ratio is uniformly bounded, even when x or ξ varies. This gives

$$\|a_\nu(x)\|_w \le C\exp\left(-3\pi\sum_\tau \nu_\tau t_\tau(x)\right).$$

Write $\lambda_1(x)^2$ for the minimum nonzero value of the squared majorant $\sum_\tau c_\tau \operatorname{maj}_{x,\tau}$. For sufficiently small $\lambda_1(x)$ a nonzero lattice vector attaining this minimum is isotropic. Indeed $|h_\tau(\xi_\tau, \xi_\tau)| \le \operatorname{maj}_{x,\tau}(\xi_\tau)$, so after clearing the fixed denominator, all absolute conjugates of $h(\xi, \xi)$ are less than one. Its algebraic norm then forces it to be zero. At this vector $t_\tau(x) \ge \lambda_1(x)^{-2}$ for all τ. For $\lambda_1^2 \le 1$ and half-trace $s = \sum_\tau \nu_\tau > c_1$ one has $3\pi s/\lambda_1^2 \ge \pi c_1/\lambda_1^2 + (3\pi/2)s$. The nonnegative frequencies with bounded sum lie in a compact simplex in the fixed dual lattice. Their exponential sum with factor $e^{-3\pi s/2}$ is finite. Summing the Fourier expansion therefore gives

$$\|P(x)\|_w^2 \le C_P e^{-2\pi c_1/\lambda_1(x)^2}$$

when the minimum is small. Boundedness provides the same estimate with a larger constant elsewhere.

On the chosen marked run, $S = \sum_k \|w_k\|^2$ satisfies $e^{-2c\phi} \ll S$ and $Tc\phi \le c_1/\lambda_1^2$. For a section of weight $w_F + Tw_L$ satisfying the strengthened cusp condition, the decay estimate gives $\|P\|^2 \ll S^T$. Each subtracted leading term $u_0^{T-|\alpha|}u^\alpha\theta_i'$ obeys the same bound, since each $\|u_i\|^2 \ll S$ and θ_i' is bounded. Thus the remainder after subtracting the jets has $\int_Y \|g\|^2 S^{-T}\, d\mu < \infty$. The constants in the decay and integral are used for finiteness in the rank argument. At the compact endpoint this follows directly from boundedness and the positive minimum of S. The isotropy of a sufficiently short vector · Decay at the stated half-trace cutoff · Comparison with the scaled generator norm.

The potential from the marked theta values. Use the principally polarized torus attached to the Hermitian lattice at an indefinite flag x. Write its lattice as L_x and its complex dimension as $g = fr$. The isometry from the flag majorant uses $\sqrt{c_\tau}$ on each orthogonal archimedean summand and conjugates

the negative block. In particular the squared norm of v in this torus lattice is $\sum_\tau c_\tau \operatorname{maj}_{x,\tau}(v_\tau)$, so its first minimum is the same $\lambda_1(x)$ as in the rank-one flow. Write E_{pol} for its integral perfect alternating polarization form. The CM half generator ω_E satisfies $\omega_E + \bar{\omega}_E = 1$. Thus

$$q(v) = E_{\mathrm{pol}}(v, \omega_E v) \pmod 2$$

is a quadratic refinement of E_{pol} on the lattice. Its polar form is $E_{\mathrm{pol}}(u, (\omega_E - \bar{\omega}_E)v)$, congruent to $E_{\mathrm{pol}}(u, v)$ because $\omega_E - \bar{\omega}_E = 2\omega_E - 1$. It specifies the symmetric principal theta line. The marked arithmetic subgroup preserves the pairing and commutes with ω_E, so it preserves the refinement. The torus lattice and its characteristic · The parity of the CM generator.

Choose a nonzero principal theta section s_x on this torus. Let $\mathcal{T}_x$ be the finite set of ordered r-tuples of points $z_b \in A_x[25]$ whose representatives $25z_b$ generate the $O_E/5O_E$-module $L_x/5L_x$. Changing a representative changes $25z_b$ by an element of $25L_x$, so the condition is well defined. Each such point is obtained by multiplying a point of the full 625-marking by 25. For its fibrewise Gaussian norm and Haar probability integral put

$$\psi_0(x) = \log \sum_{\mathbf{z} \in \mathcal{T}_x} \prod_{b=1}^{r} \frac{|s_x(z_b)|_1^2}{\|s_x\|_2^2}.$$

The ratio is independent of the scalar of s_x. It depends only on the maximal positive plane and its orthogonal complement, so is invariant under change of triangular flag frame. An arithmetic translation induces a polarized torus isomorphism preserving the refinement and permuting the generating tuples. The function consequently descends to the whole marked quotient.

To see strict positivity inside the logarithm, choose an $O_E/5O_E$-basis of $L_x/5L_x$, which exists because L_x is projective of rank r over O_E. All level-25 lifts of one basis vector form a coset of $A_x[5]$ in $A_x[25]$. At least one lift has nonzero theta value. In fact let $\mathscr{L}$ be the principal theta line. The translates of its section by $A_x[5]$ give generators of the one-dimensional spaces $H^0(A_x, \mathscr{L} \otimes P)$ for all $P \in \widehat{A}_x[5]$, by principal polarization. The decomposition of $[5]_*\mathcal{O}_{A_x}$ into the characters of the deck group identifies their direct sum with $H^0(A_x, [5]^*\mathscr{L})$. If the section vanished at every point of the coset, all sections in this direct sum would vanish over its representative. This contradicts global generation of $[5]^*\mathscr{L} = \mathscr{L}^{25}$. Choosing such a lift separately for each basis vector gives a positive summand in the logarithm. Nonvanishing on each coset · Strict positivity of the tuple sum.

For the Levi form, choose a locally constant symplectic basis of the period lattice and use the holomorphic normalized Siegel matrix Z with $Y = \operatorname{Im} Z > 0$. Write a flat torsion point as $a + Zb$ with fixed rational a, b. The Gaussian factor in the squared theta norm is $e^{-2\pi b^{\mathrm{t}} Y b}$, which is the squared absolute value of $e^{\pi i b^{\mathrm{t}} Z b}$. A theta with the fixed characteristic has squared L^2 norm $\det(2Y)^{-1/2}$ with Haar probability. Indeed integrating in the real half of the period parallelepiped gives orthogonality of its frequencies, and integrating in the second half unfolds their Gaussian to $\mathbf{R}^g$. The normalized pointwise square is therefore

$$2^{g/2}(\det Y)^{1/2} \left| e^{\pi i b^{\mathrm{t}} Z b}\, \theta(a + Zb; Z) \right|^2.$$

The same Y-factor occurs for every flat torsion point. In normalized periods $\det Y$ is the Hodge determinant metric. In the flag coordinates it equals $e^{-\phi_{w_{\mathrm{de}}}}$ times the squared modulus of a nowhere zero holomorphic change-of-frame determinant and a positive constant. To compare these metrics, write the positive plane as the graph of a matrix U in an orthogonal splitting. Its Hodge basis and the conjugate of the negative orthogonal basis have Gram determinants $\det(I - U^*U)$ and $\det(I - UU^*)$, respectively. Their product over the embeddings is $e^{-2\rho_*}$. The flag weight $w_{\mathrm{de}} = (-1, \dots, -1, 1, \dots, 1)$ has $\phi_{w_{\mathrm{de}}} = 2\rho_*$ in these frames. The change to the symplectic period normalization has holomorphic determinant because both Hodge frames are holomorphic. Thus the log of the tuple sum is smooth. Adding $(r/2)\phi_{w_{\mathrm{de}}}$ leaves $\log \|F\|^2$ up to a pluriharmonic function, where F is the vector of holomorphic theta products in the chosen period chart. Its Levi form in a direction v is

$$\frac{\|F\|^2 \|dF(v)\|^2 - |\langle dF(v), F \rangle|^2}{\|F\|^4} \geq 0.$$

Consequently

$$\partial\bar\partial\psi_0 \geq -(r/2)\partial\bar\partial\phi_{w_{\mathrm{de}}}.$$

The common holomorphic norm factor · The Levi lower bound.

For the decay take the same Gaussian boundary scale $\Phi = 2f/(\mathrm{rd}(E)^2\lambda_1^2)$ below the short-lattice threshold. The Gaussian null-product estimate (133) applies to every summand because the tuple generates modulo 5. After squaring the product and summing at most 25^{2gr} tuples, it gives

$$\psi_0 \leq C + B\log(K\Phi) - 2\pi\Phi/625$$

in that region and a finite constant above the threshold. In particular ψ_0 is bounded above. For every nonzero isotropic lattice vector ξ one has

$$\lambda_1^2 \leq \sum_\tau c_\tau \,\mathrm{maj}_{x,\tau}(\xi_\tau) = \sum_\tau t_\tau^{-1}, \qquad \lambda_1^2 \min_\tau t_\tau \leq f.$$

With $a_1 = 1/(2\cdot 5^4\,\mathrm{rd}(E)^2)$, this gives $2\pi\Phi/625 \geq 8\pi a_1 \min_\tau t_\tau$. Above the threshold $\min_\tau t_\tau$ is bounded by the same first-minimum inequality. Absorbing that bound and splitting the logarithm yields on the whole component

$$\psi_0(x) \leq C' + B'\log(1+\lambda_1(x)^{-1}) - 8\pi a_1 \min_\tau t_\tau(x).$$

The square of the marked theta product and the inverse height $t_\tau^{-1} = c_\tau\,\mathrm{maj}_{x,\tau}(\xi_\tau)$ are both needed for this coefficient. The generator rate c uses Φ in the same units. In particular $Tc\Phi \leq c_1/\lambda_1^2$ at the prescribed cutoff, since $c_1/(Tc) = 2\,\mathrm{rd}(E)^7 f^2(m-j)^3 \geq 2f/\,\mathrm{rd}(E)^2$ at every nonterminal stage. These comparisons apply with the Gaussian parameters of each of the two runs. At the compact endpoint take $\psi_0 = 0$. The exact height and first-minimum comparison · The marked potential cutoff.

Curvature margins on the marked rows. At an orthogonal frame let d_j be the absolute leading Gram minor of order j, and let ϵ_a be the sign of the a-th basis direction of the form. The quadratic term of $\log d_j$ in lower unitriangular coordinates is $\sum_{a<j\leq b}\epsilon_a\epsilon_b|z_{ba}|^2$. Indeed for the first j columns their top square block is unitriangular, of determinant one. Divide it out before taking the Gram determinant. The lower rectangular block in the resulting matrix is unchanged to first order, and the Gram matrix is its signed quadratic contribution added to the signature matrix of the first j directions. Expand the logarithm of the determinant to second order. Since $\phi_w = \sum_a w_a\log(d_{a+1}/d_a)$, telescoping gives the Levi coefficient $\epsilon_a\epsilon_b(w_a - w_b)$ on z_{ba}. The mixed entries are the ones contributing to $\partial\bar\partial\rho_*$, with coefficient one at the orthogonal reference frame. This Levi form is invariant, since the unitary fractional-linear action on the matrix ball changes ρ_* by the logarithm of the squared modulus of a nonvanishing holomorphic determinant. The Petersson curvature is also invariant, so the comparison holds at every point. In the same coordinates $\log d_\mu = -2\sum_{j=1}^{r-1}\log d_j$ is the metric potential of the canonical row K_Y, whose entries are $2a-(r-1)$. The canonical row · The invariant mixed-direction curvature comparison.

Write $G = w_F - K_Y - Dw_L$ as a row. On the two marked indefinite runs the signed gap inequalities give within each block $G_a - G_b \geq 2(w_{L,a} - w_{L,b})$ for $a<b$. For $a<p_0\leq b$ they give $w_{L,b} - w_{L,a} \geq 19T/10$ and the endpoint comparison $G_{r-1} - G_0 \geq 2(w_{L,r-1} - w_{L,0})$. Adding the within-block comparisons from 0 to a and from b to $r-1$ gives $G_b - G_a \geq 2(w_{L,b} - w_{L,a})$. The curvature of G consequently dominates $(19T/5)\partial\bar\partial\rho_*$, and that of w_L dominates $(19T/10)\partial\bar\partial\rho_*$. The extra row w_{de} has entries -1 on the first block and 1 on the second, hence its curvature is $2\partial\bar\partial\rho_*$. For $\psi = \tau_0\psi_0$, $\tau_0 = 0$ or $1 + c_1/a_1$, the loss permitted by its Levi lower bound is $\tau_0\,\mathrm{rr}\,\partial\bar\partial\rho_* \leq 2T\,\partial\bar\partial\rho_*$. Thus the interpolation margin can be taken as $M_{\mathrm{int}} = 9T/5$. There is also a reserve of $\frac12\partial\bar\partial\phi_{w_L}$ in the same comparison. The mixed gaps of $G - (\tau_0\mathrm{rr}/2)w_{\mathrm{de}} - \frac12 w_L$ are at least $(3/2)(19T/10) - 2T = 17T/20$, and its within-block gaps are positive by the first gap inequality. Division uses $G + Dw_L$, so one may take $M_{\mathrm{div}} = (19/5 + 19D/10)T$. Adding kw_L for $k\geq 0$ preserves these positive curvature comparisons. The within-block and mixed gaps for the two marked runs · Propagation of the endpoint comparison.

For $b_{\mathrm{P}} = 5fr^2$, $D = fr(r-1)/2$, $r \geq 8$, $f \geq 1$ and $T \geq 11D^2$, one has $b_{\mathrm{P}} \leq 20D^2/49$. Also

$$b_{\mathrm{P}}^2 f \min(p_0, r-p_0) \leq \tfrac{25}{2} f^3 r^5 \leq 20D^3.$$

The second inequality follows from $5r^2 \leq (r-1)^3$ for $r \geq 8$. It follows that $4fr^2 < b_{\mathrm{P}} < M_{\mathrm{int}}$ and $b_{\mathrm{P}}^2 f \min(p_0, r-p_0) + b_{\mathrm{P}} < M_{\mathrm{div}}$. For the cusp condition, the within-block gaps of $w_F + kw_L = G + K_Y + (D+k)w_L$ are nonnegative, since the canonical adjacent difference is -2 and the within-block gap inequality reserves twice a strictly positive integer gap. Its last entry exceeds its first. This proves the row conditions used in the rank-one constant coefficient argument for every shifted weight. For the compact endpoint $p_0 = 0$, the within-block inequality holds on the whole row, G supplies the same reserve, and $\rho_* = \psi = 0$. There is no cusp test there. The strict numerical margins · The definite row.

Weighted Poincaré summation on the flag. Work on one of the marked complex components $Y = \Gamma\backslash\mathcal{F}$, where $\mathcal{F}$ is the open unitary flag orbit. Put $p = p_0$ and $q = r - p_0$. At each of the f archimedean factors choose an orthonormal frame for the form of signature (p, q). The orbit consists of complete flags with positive first p directions, and is homogeneous under $\prod_\tau U(p,q)$ with compact diagonal stabilizer. The marked level gives a free arithmetic action on these flags. The invariant measure has finite mass on Y by arithmetic covolume and compactness of the flag fibre. Use its density d_μ in each flag chart and the squared Petersson metric $e^{-\phi_w}$ in a frame of weight w. Both are homogeneous. In the lower unitriangular chart the invariant density is the product of the inverse squares of the leading absolute Gram minors of orders $0, \dots, r-1$. Indeed a change of chart by a unitary element followed by an upper triangular matrix with diagonal $b_1, \dots, b_r$ changes these minors by the squares of the partial diagonal products. The flag Jacobian acts on the lower triangular entries with diagonal weights b_i/b_j for $i > j$. Its squared absolute determinant cancels the change of the density since $\left|\prod_i b_i\right| = 1$. The invariant flag chart and radial function · The free action at the marked level · The invariant quotient volume.

In the big cell based at x_* write the lower unitriangular matrix in blocks as $\left(\begin{smallmatrix} A & 0 \\ C & D_1 \end{smallmatrix}\right)$. Positivity of the leading p directions is equivalent to $I - Z^*Z > 0$ with $Z = CA^{-1}$. The remaining directions in the quotient by this positive plane are negative definite, so every remaining Gram minor has the required sign. Thus the cell is biholomorphic to a product of the matrix balls with affine triangular coordinates. In these coordinates put

$$\rho_* = -\sum_\tau \log\det(I - Z_\tau^* Z_\tau), \qquad k_{b_{\mathrm{P}}} = e^{-b_{\mathrm{P}}\rho_*}.$$

The function ρ_* extends to all of $\mathcal{F}$ by taking the graphs of the maximal positive planes relative to the reference orthogonal splitting. If z denotes the full lower triangular coordinates, $\rho_* + \log(1 + |z|^2)$ is a strictly plurisubharmonic exhaustion of the cell. Bounded sublevels bound the coordinates and keep every matrix ball coordinate away from the boundary. In fact for a matrix tangent vector H the Levi form of the ball potential is $\|(I - ZZ^*)^{-1/2} H (I - Z^*Z)^{-1/2}\|_{\mathrm{HS}}^2$. The pairing of this matrix with Z is the derivative of the potential in direction H. Cauchy–Schwarz on the product gives

$$|\partial\rho_*(v)|^2 \leq f \min(p,q)\, (\partial\bar\partial\rho_*)(v, \bar v). \tag{140}$$

Here $\|Z\|_{\mathrm{HS}}^2 \leq \min(p,q)$ on the ball. If a nonzero holomorphic function u is excluded, adding $|u|^{-2}$ to the exhaustion gives a strictly plurisubharmonic exhaustion of its nonvanishing locus. The matrix ball expression · The bound on the derivative of the ball potential.

Let W be an invariant measurable squared metric for weight w, finite almost everywhere, with $W \geq c\, e^{-\phi_w}$ for some $c > 0$. For the chosen $b_{\mathrm{P}} > 4fr^2$, the upstairs norm $\|P\|_{W/k_{b_{\mathrm{P}}}}^2 = \int_{\mathcal{F}} |P|^2 W/k_{b_{\mathrm{P}}}\, d\mu$ controls Poincaré summation. The estimate follows from a uniform count. Use the maximum of the operator norms over the f factors, and write gx_* for a flag with g unitary. If $A_\tau(g)$ is the positive square block of g_τ, then

$$\rho_*(gx_*) = 2\sum_\tau \log|\det A_\tau(g)|.$$

All singular values of $A_\tau(g)$ are at least one and the unitary block identities give $\|g_\tau\| \le 2\|A_\tau(g)\|$ in the indefinite case. The definite case has $\|g_\tau\| = 1$. Consequently $k_s(gx_*) \le 2^{2s}\|g\|^{-2s}$ for $s > 0$. Discreteness in the matrix group gives a fixed radius $\epsilon > 0$ excluding every element of $\Gamma \setminus \{1\}$ from the identity ball. If $\|\gamma g\|, \|\gamma' g\| \le R$ and $\gamma \ne \gamma'$, then

$$\epsilon \le \|(\gamma' g)(\gamma g)^{-1} - 1\| \le R\|\gamma' g - \gamma g\|.$$

We used $\|h^{-1}\| = \|h\|$ for a unitary matrix of this signature. Packing in the real $2fr^2$-dimensional matrix space gives $O(R^{4fr^2})$ such points, uniformly in g. Summing over dyadic annuli gives

$$\sup_{x \in \mathcal{F}} \sum_{\gamma \in \Gamma} k_s(\gamma x) < \infty \qquad (s > 2fr^2). \tag{141}$$

On translates of a fixed small flag ball, $k_{b_{\mathrm{P}}}$ changes by a bounded factor from its value at the centre. To see this, choose bounded unitary representatives h for the small ball at x_*. In the indefinite case $A(gh) = A(g)(A(h) + A(g)^{-1}B(g)C(h))$ and $\|A(g)^{-1}B(g)\| \le 1$ by the unitary identities. Apply the determinant bound first with h and then with its inverse. This is uniform in the translating element g.

For holomorphic P of finite upstairs norm, apply the local holomorphic mean value bound on the orthonormal flag ball of radius ϵ_1 at x, using homogeneity of the Petersson norm and measure. It holds at the centre and uniformly on a smaller ball for any fixed sufficiently small radius. Converting its integral to the $W/k_{b_{\mathrm{P}}}$ norm and using the comparison of $k_{b_{\mathrm{P}}}$ on the ball gives at the centre

$$\|P(x)\|_w \le C_{\epsilon_1}\|P\|_{W/k_{b_{\mathrm{P}}}} k_{b_{\mathrm{P}}/2}(x) \left(\sup_{y \in B(x,\epsilon_1)} \frac{e^{-\phi_w(y)}}{W(y)} \right)^{1/2}. \tag{142}$$

The ratio here is intrinsic. The same bound on the smaller ball with the centre value of $k_{b_{\mathrm{P}}/2}$ gives a locally uniform summable majorant for the translates by (141). The periodization $\mathcal{P}_w P(x) = \sum_\gamma \chi_w(\gamma)^{-1} P(\gamma x)$ in frame notation is therefore holomorphic and automorphic. Its norm is globally bounded, because χ_w is unitary and W is bounded below by the Petersson metric up to a positive constant. The quotient norm also satisfies

$$\int_Y |\mathcal{P}_w P|^2 W \, d\mu \le \left(\sup_x \sum_\gamma k_{b_{\mathrm{P}}}(\gamma x) \right) \|P\|^2_{W/k_{b_{\mathrm{P}}}}. \tag{143}$$

This is Cauchy–Schwarz with the factors $k_{b_{\mathrm{P}}}(\gamma x)$, followed by integration over a fundamental set. For an automorphic h of finite quotient W-norm, unfolding the same sum gives

$$\langle \mathcal{P}_w P, h \rangle_{Y,W} = \int_{\mathcal{F}} P\bar{h} \, W \, d\mu, \qquad \int_{\mathcal{F}} k_{b_{\mathrm{P}}} |h|^2 W \, d\mu \le \left(\sup_x \sum_\gamma k_{b_{\mathrm{P}}}(\gamma x) \right) \|h\|^2_{Y,W}.$$

The integrals for the pairing converge absolutely by Cauchy–Schwarz and the second inequality. Multiplication by an automorphic section u commutes with summation in the corresponding weights, because $u(\gamma x) = \chi_{\mathrm{wt}(u)}(\gamma) u(x)$ in frames. In particular $\mathcal{P}_{w+\mathrm{wt}(u)}(uP) = u\mathcal{P}_w P$. The uniform orbit sum · The translated local mean bounds · The unfolded pairing.

Surjectivity in the singular metric. Suppose $W = e^{-\Phi}$ on the big cell away from the divisor of a nonzero holomorphic function u, with Φ smooth there, and

$$\partial\bar{\partial}(\Phi - \log d_\mu) \ \ge \ M \, \partial\bar{\partial}\rho_*, \qquad M > b_{\mathrm{P}}^2 f \min(p, q) + b_{\mathrm{P}}.$$

On the chosen weights every bounded automorphic holomorphic section belongs to the cusp space. The image of $\mathcal{P}_w$ therefore lies in the finite-dimensional cusp space of finite quotient W-norm, and is closed there. Take h in its orthogonal complement. Put $I = \int_{\mathcal{F}} k_{b_{\mathrm{P}}} |h|^2 W \, d\mu$, which is finite by unfolding. On the

big cell with $u \neq 0$, solve the equation $\bar\partial v = \bar\partial(k_{b_{\mathrm{P}}} h)$ with the Euclidean weight $\Phi - \log d_\mu - b_{\mathrm{P}}\rho_*$. Its Levi form dominates $(M - b_{\mathrm{P}})\partial\bar\partial\rho_*$. The right side is $-b_{\mathrm{P}} k_{b_{\mathrm{P}}} h\, \bar\partial\rho_*$, so (140) implies that its squared dual Levi norm is at most $\vartheta |k_{b_{\mathrm{P}}} h|^2$, where $\vartheta = b_{\mathrm{P}}^2 f \min(p, q)/(M - b_{\mathrm{P}}) < 1$. Use the weighted $\bar\partial$ estimate in top holomorphic degree [125, §5.A, Theorem 5.1]. Trivializing by $dz_1 \wedge \cdots \wedge dz_D$ cancels the Kähler density in the squared norm of the solution. In bidegree $(D, 1)$ the inverse curvature operator gives the squared dual Levi norm in the right side of the estimate. The strict plurisubharmonic exhaustion supplies a complete Kähler metric, which can be enlarged by the positive Levi form of the weight. This makes the right side square-integrable for the complete metric as well. For semidefinite Levi form first add $\varepsilon \log(1 + |z|^2)$ to the weight. The same dual bound holds for each regularization. A weak limit on compact exhaustions as ε decreases to zero gives $\|v\|^2_{W/k_{b_{\mathrm{P}}}} \leq \vartheta I$ and preserves the $\bar\partial$ equation.

The holomorphic function $F = k_{b_{\mathrm{P}}} h - v$ has finite upstairs norm. Both the deleted divisor and the complement of the flag cell are analytic sets. In local frames along either set the weight $W/k_{b_{\mathrm{P}}}$ dominates a positive smooth metric, so F extends by local square-integrable removability. Its norm stays finite since the deleted sets have measure zero. Pairing h with $\mathcal{P}_w F$ and unfolding gives $I = \int_{\mathcal{F}} v\bar h\, W\, d\mu$. The right side has absolute value at most $\sqrt{\vartheta} I$. It follows that $h = 0$. Thus every cusp section with finite quotient W-norm has a holomorphic lift of finite upstairs norm whose single periodization is the given section. At the compact endpoint $p = 0$ one has $k_{b_{\mathrm{P}}} = 1$, the arithmetic subgroup of the compact matrix group is finite, and the lifted section itself, divided by the order of that subgroup, is a periodization preimage. The weighted cusp-space orthogonal projection · The strict contraction in the adjoint equation.

Separation of the cut from the cusp. For the selected combinations put $A = \sum_i \|u_i\|^2$ and $S = \sum_k \|w_k\|^2$. To control A/S away from J, consider the projective coordinates $[w_1 : \cdots : w_N]$ on each component. The w_k have no common zero there. The projective coordinate closure has dimension at most D. In fact any $D + 2$ sections of the same weight have a homogeneous relation over the coefficient field of the marked model: after trivializing the line generically, $D + 1$ ratios are functions on the D-dimensional algebraic flag model. The homogeneous relation obtained by clearing the denominator evaluates on the complex component. Applying this to each $D + 2$ chosen coordinates bounds the transcendence degree of each irreducible component of the affine cone by $D + 1$. As $u_0, u_1, \ldots, u_D$ are generic linear combinations, they have no common zero on this closure. In particular u_0 is nonzero at every point of J. Homogeneous relations in the marked section algebra · Transfer of the relations to the flag readings.

The projective limits from $S \to 0$ lie in a closed algebraic subset of dimension at most $D - 1$. For this use the actual nonzero section $\Theta = \eta\, \theta_1^{i g_\theta/2}$ of weight $k w_L$, with $k = i g_\theta$, $i = i_\eta$ and $g_\theta = g_*$ the even theta genus parameter. Use $\|\eta\|^2 \ll e^{-2i\phi}$, boundedness of θ_1, $T c g_\theta \leq 1$, and $e^{-2c\phi} \ll S$. Since $T \geq 2$, $i \geq 1$ and $g_\theta \geq 2$, one has $c(k + 1) \leq i$, so $\|\Theta\|^2 \ll S^{k+1}$. Fix $D + 1$ coordinates w_{j_a}. The tuple $(\Theta, w_{j_1}^k, \ldots, w_{j_{D+1}}^k)$ of the same weight has a homogeneous relation. Divide this polynomial by the maximum power of its first variable. The resulting relation still holds because Θ is nonzero on a dense open set of the component. Its specialization at first coordinate zero is a nonzero polynomial, also after substituting the k-th powers. In a sequence with $S \to 0$, normalize the w_k in a common frame by $(\sum |w_k|^2)^{1/2}$. The normalized Θ, divided by the k-th power of the same quantity, has absolute value $O(S^{1/2})$. Thus all projective limits satisfy the specialized relation. Imposing these relations for all $D + 1$ choices of coordinates makes the coordinate functions on each affine cone component have transcendence degree at most D. This gives the claimed projective dimension bound. In particular the D generic linear equations $u_i = 0$ miss this set. The marked construction chooses one complex tuple for the finitely many components, algebraically independent over K_0 and in the full-measure set of transverse tuples. The proper incidence loci for avoidance are defined over K_0, so the same tuple satisfies avoidance on every component. The nonzero decaying product on each component · The normalized bound for the extra section · Avoidance by the generic cut · One specialization for all the components.

The subsets with λ_1 bounded below are compact in the flag quotient. Here is the unitary form of the compactness argument. The majorant squared norms on the fixed lattice have constant covolume, since

they are obtained by the determinant-one real changes of orthonormal frame. Mahler compactness gives an integral basis with bounded majorant Gram matrix when the first minimum is bounded below. The determinant is fixed, so its inverse is bounded as well. In such bases the matrices of a fixed set of generators of $\mathcal{O}_E$ have bounded integer entries, because their operator norms in the majorant metrics are bounded. The entries of the original Hermitian form are bounded at every embedding by the majorant Cauchy–Schwarz inequality, and lie in a fixed fractional ideal. There are finitely many possibilities for these two sets of data. Bases with the same data are related by an integral unitary automorphism. Taking finite coset representatives for the marked level in the stabilizer, one gets a compact set of majorants modulo that level. A majorant recovers its maximal positive plane continuously as the positive eigenspace of the corresponding involution. The remaining compact flag fibres preserve compactness. Finiteness of reduced integral types · The criterion on the full marked flag quotient.

Each w_k is cusp and bounded. Its nonzero rank-one frequencies are in the nonnegative cone of the fixed dual period lattice, and their half-traces are bounded below by a positive number. Indeed the lattice meets each bounded trace simplex in finitely many points. The rank-one decay estimate therefore gives $S \to 0$ as $\lambda_1 \to 0$. Thus every sequence escaping compact sets has $S \to 0$. If A/S tended to zero on such a sequence, a subsequential projective limit would be a point of the excluded boundary set with all $u_i = 0$.

All common zeros J consequently lie in a compact part of the marked component. The generic transverse property follows by considering $u_1(x) = \cdots = u_D(x) = 0$ in the product of the flag component with the coefficient space. Variation of the independently chosen coefficients is surjective in these D equations since some $w_k(x)$ is nonzero. The incidence manifold is smooth of the dimension of the coefficient space, and regular values of its projection give transverse zeros. Hence the selected generic J is finite. Off small open neighbourhoods of these points, a sequence with $A/S \to 0$ would either escape compact sets or converge to another common zero. Both are excluded, proving $A \geq b_J S$ there for some $b_J > 0$. The terminal compact component gives this last inequality by compactness alone. The generic transverse cut at the same level.

Value interpolation with the cusp cutoff. On the actual flag component set $\psi = \tau_0 \psi_0$ with $\tau_0 = 0$ or $1 + c_1/a_1$. The potential is invariant and bounded above. Put

$$G = \phi_{w_F} - \log d_\mu - D\phi_{w_L} + \psi.$$

The row gaps give $\partial\bar\partial G \geq M_{\mathrm{int}} \partial\bar\partial \rho_*$ with $M_{\mathrm{int}} > b_{\mathrm{P}} > 4fr^2$. They give also $\partial\bar\partial G \geq 2\eta_1 \partial\bar\partial \phi_{w_L}$ for some $\eta_1 > 0$, where ϕ_{w_L} has strictly positive Levi form. For the second inequality one uses the strictly dominant row $w_F - K_Y - Dw_L$ with the loss $(\tau_0\,\mathrm{rr}/2)w_{\mathrm{de}}$ coming from the lower Levi bound on ψ_0. The chosen cutoff satisfies $\tau_0 \mathrm{rr} \leq 2T$ and the reserve in that row give the claimed positive multiple of the curvature of w_L. At the compact endpoint $\rho_* = \psi = 0$, and the same positive reserve comes from the consecutive gaps on the whole definite row. The marked margins · The curvature reserve and the cutoff · The reserve after inserting the cusp potential · The whole-row definite reserve.

To obtain one coordinate of the prescribed values, centre a big cell at an orthogonal frame above a point $x_j \in J$. Choose a smooth cutoff χ supported in a small ball containing only this lift of a point of J, equal to one near its centre in the holomorphic frame. Write s_i for the holomorphic representatives of u_i on the cell. Solve $\bar\partial v = \bar\partial \chi$ with Euclidean weight

$$G - b_{\mathrm{P}} \rho_* + D \log \sum_{i=1}^{D} |s_i|^2.$$

The first two terms have Levi form at least $(1 - b_{\mathrm{P}}/M_{\mathrm{int}})\partial\bar\partial G$, hence are plurisubharmonic everywhere and strictly so on the cutoff annulus. The logarithm is plurisubharmonic. On the support of the derivative of the cutoff the singular denominator is bounded. The weighted $\bar\partial$ estimate in top holomorphic degree [125, §5.A, Theorem 5.1] gives a correction of finite norm. One may solve off the divisor of a nonzero s_i,

with the strictly plurisubharmonic exhaustion of that open cell, and use square-integrable removability afterward.

In invariant notation the correction satisfies

$$\int_{\mathcal{F}} |v|^2 (W_\psi / k_{b_{\mathrm{p}}})\, A^{-D}\, d\mu < \infty, \qquad W_\psi = e^{-\phi_{w_F} - \psi}, \qquad A = \sum_i \|u_i\|^2.$$

Initially the integral is over the open cell of the solution. The upper bound on A, smoothness of W_ψ, and compact support of the cutoff give the local square-integrability of $P = \chi - v$ in undivided weight. Thus P extends across both deleted analytic sets as a holomorphic section on $\mathcal{F}$, and the integral identity remains true there by setting $v = \chi - P$. At each lifted point of J the correction is holomorphic in a neighbourhood. Since $A \asymp |z|^2$ there, a nonzero value of v would force the integral of $r^{-1}\, dr$. It follows that P has value one at the chosen lift, and zero at all other lifts of points of J. Boundedness above of A also gives $\|P\|_{W_\psi / k_{b_{\mathrm{p}}}} < \infty$. Apply Poincaré summation with W_ψ. The free marked action ensures that exactly the chosen summand contributes at x_j. Rescale in the actual value-line frame at x_j and repeat for every point. This spans all values by bounded cusp sections of weight w_F. The conclusion holds component by component. The cutoff and its singular interpolation norm · Evaluation of the periodizations at the marked points.

In the deep case put $a = \tau_0 a_1 = a_1 + c_1$. The decay bound on ψ at any lattice isotropic direction reads, for every $0 < \epsilon_2 < 1/2$ and $t_\tau(x) \geq 1$,

$$\psi(x) \leq C + B \log(1 + \lambda_1(x)^{-1}) - 8\pi(1 - 2\epsilon_2) a \min_\tau t_\tau(x).$$

In the pointwise periodization estimate (142) the ratio for W_ψ is e^{ψ}. Summing and taking the square root gives $\|\mathcal{P}_{w_F} P(x)\|_{w_F} \leq C_{\epsilon_1} \sup_{B(x,\epsilon_1)} e^{\psi/2}$ on each small orthogonal flag ball. Shrinking the radius compares all neighbouring majorants with the central ones up to a factor arbitrarily close to one. In fact on the reference ball the majorants are a continuous family of positive definite Hermitian forms, and transporting by the central orthogonal frame gives the same comparison on all its translates. This compares λ_1 by taking minima on the lattice, and compares each t_τ by taking reciprocals. For $t_\tau(x) \geq 2$, any $a' < a$ with $a' > 0$ therefore gives

$$\|\mathcal{P}_{w_F} P(x)\|_{w_F} \leq C'(1 + \lambda_1(x)^{-1})^{B'} e^{-4\pi a' \min_\tau t_\tau(x)}$$

by choosing the radius and ϵ_2 small enough.

For a fixed isotropic direction translate the real period parallelepiped to $y_\tau = Y$ at every place. The exact flow formula gives $\min_\tau t_\tau(x_{u+iY}) = Y + \min_\tau t_\tau(x)$. On this translate the majorants are comparable to their initial values by polynomial factors in $1 + Y$, uniformly in u on the parallelepiped. To see this use an initial orthonormal basis of the positive plane. Its flowed columns have size $O(1 + Y)$ and positive Gram matrix I plus a positive semidefinite rank-one term. The formula for the Hermitian projection onto the flowed plane then bounds the majorant above by a polynomial multiple of the initial majorant. The determinant of the majorant is constant under unitary change of the positive plane. This bounds the smallest relative eigenvalue below by a reciprocal polynomial. Thus λ_1^{-1} grows at most polynomially. The same holds for the reciprocal norm factor by the Gram determinant formulas. Applying (139) gives, for a frequency of half-trace $s < a'$, the bound $C''(1 + Y)^{B''} e^{-4\pi(a' - s)Y}$ for its coefficient. Sending Y to infinity kills this coefficient. As a' can approach a, this proves vanishing for all half-traces less than $a_1 + c_1$, in particular at most c_1. The prescribed values can therefore be interpolated in the strengthened cusp space. At the definite endpoint this is ordinary value interpolation. The exact 8π decay of the potential · The Fourier cutoff for the half-potential bound.

Division of the section remaining after the jets. Write $A = \sum_{i=1}^{D} \|u_i\|^2$ and $S = \sum_k \|w_k\|^2$, in the weight-w_L metric. For a section g of weight $w_F + T w_L$ of order at least $a_0 = T - D$ at J, assume $\int_Y \|g\|^2 S^{-T}\, d\mu < \infty$. The weighted quotient norm required for the lift is

$$\int_Y \|g\|^2 S^{-1/2} A^{-(T - 1/2)}\, d\mu. \tag{144}$$

On the complement of small neighbourhoods of J the transverse projective chart gives $A \geq b_J S$, so this is bounded there by a constant times the assumed integral. Near each point of J, $A \asymp |z|^2$ in transverse coordinates and S is bounded below. The radial exponent is $2a_0 - 2(T - 1/2) + 2D - 1 = 0$, which gives finite local integral. At the compact endpoint S has positive minimum and the initial integral is finite for all sections.

Set $W = e^{-\phi_{w_F + T w_L}} S^{-1/2} A^{-(T-1/2)}$ in local frames. It is an invariant squared metric and dominates a positive constant times the corresponding Petersson metric because S and A are bounded above. In big-cell coordinates let s_i and v_k be the holomorphic representatives of u_i and w_k, and put $A_{\mathrm{hol}} = \sum_i |s_i|^2$ and $S_{\mathrm{hol}} = \sum_k |v_k|^2$. The powers of $e^{-\phi_{w_L}}$ cancel and

$$-\log W - \log d_\mu = \phi_{w_F} - \log d_\mu + \tfrac{1}{2} \log S_{\mathrm{hol}} + (T - \tfrac{1}{2}) \log A_{\mathrm{hol}}.$$

The last two logarithms are plurisubharmonic and the metric is smooth off the divisor of any nonzero s_i. The row-gap inequality for $\phi_{w_F} - \log d_\mu$ gives the surjectivity margin $M > b_{\mathrm{P}}^2 f \min(p_0, r - p_0) + b_{\mathrm{P}}$. Using (144) in that surjectivity result provides a holomorphic lift F with $\mathcal{P}_{w_F + T w_L} F = g$. On the pseudoconvex big cell its upstairs norm is

$$\int |F|^2 A_{\mathrm{hol}}^{-(D + a_0 - 1/2)} e^{-\Phi_{\mathrm{red}}} \, dz < \infty, \qquad \Phi_{\mathrm{red}} = \phi_{w_F} - \log d_\mu - b_{\mathrm{P}} \rho_* + \tfrac{1}{2} \log S_{\mathrm{hol}},$$

where dz denotes Lebesgue volume. The reduced weight Φ_{red} is plurisubharmonic by the same row gap.

For $a \geq 1$, apply Skoda's division theorem [124, Theorem 1] on this cell with D generators in dimension D, so the parameter q_{S} of the theorem is $D - 1$. Use $\alpha q_{\mathrm{S}} = D + a - 3/2 > q_{\mathrm{S}}$. The input power of the sum of squares is $-D - a + 1/2$, and the output power is $-D - a + 3/2$. The estimate bounds the sum of squared quotient norms by $\alpha/(\alpha - 1)$ times the input norm squared. Apply the step a_0 times. Grouping ordered products with the same monomial gives $F = \sum_{|\alpha| = a_0} s^\alpha H_\alpha$ with holomorphic H_α and finite norms with reference weight $A_{\mathrm{hol}}^{-(D-1/2)} e^{-\Phi_{\mathrm{red}}}$. This last norm is stronger than the upstairs Petersson norm of weight $w_F + D w_L$ with extra factor $k_{b_{\mathrm{P}}}^{-1}$: the ratio of the Petersson density to the density just used is $A^{D-1/2} S^{1/2}$, which is bounded. Extend each H_α over the analytic complement of the big cell by local square-integrable removability in this metric. The bounded periodizations $h_\alpha = \mathcal{P}_{w_F + D w_L} H_\alpha$ are cusp sections at the same marked level. Periodization commutes with the fixed automorphic monomials, so $g = \sum_{|\alpha| = a_0} u^\alpha h_\alpha$.

For the localized fibre, take a bounded section g_0 of weight $w_F + k w_L$ vanishing on J. Apply the argument to $g_0 u_0^{D+1}$, with base weight $w_F + k w_L$ and $T' = D + 1$. The initial weighted integral is finite by $\|u_0\|^2 \ll S$, boundedness of g_0, and finite invariant volume. The margin for the base weight persists since w_L is semipositive. There is one division step, with exponents $D + 1/2$ and $D - 1/2$. It expresses $g_0 u_0^{D+1} = \sum_i u_i h_i$ with bounded cusp coefficients of weight $w_F + (k + D) w_L$, as required on localizing at u_0. Finiteness of the singular remainder norm · Iteration at the half-integer exponents · Order-one division at the shifted weights.

The ranks and the placed fibre. The marked analytic cusp spaces are obtained by extension of scalars from the rational cusp rules in these weights. The deep condition is the kernel of the low coefficient readings, which commute with scalar extension using the separating CM arcs below. The actual cusp spaces at the shifted weights.

The interpolation and order-one division at the generic fibre identify the weight-w_F section evaluation on the localized algebra with that on the reduced complex intersection. The evaluations of C_{w_F} span this fibre. The short forms spanning C_{w_F} over K_0 therefore give a basis $\theta'_1, \dots, \theta'_h$ by taking independent images over $K_0(z)$. The evaluation matrix extends to every Gauss field, so the basis is fixed for all the places of this run and stage. Order-one division after localization · The fibre basis and the complex values.

The lower rank estimate uses interpolation within the strengthened cusp space. Given any basis of the fibre, interpolate its values by sections v_i with this cusp condition. All the sections

$$v_i u_0^{T - |\alpha|} u_1^{\alpha_1} \cdots u_D^{\alpha_D}, \qquad |\alpha| \leq T,$$

belong to the deep space over $\mathbb{C}$. Multiplication by a bounded section of weight w_L preserves the strengthened cusp condition. At the generic points, divide by u_0^T and take the first nonzero total degree in the parameters u_i/u_0. The leading coefficient in each monomial is the combination of the independent values $v_i|_J$ with the same indices. This proves independence. The same argument applied to θ_i' and $|\alpha| < t_0$ proves independence modulo the span of the products of order t_0. In particular

$$\dim C_{\mathrm{deep}} \geq h\binom{T+D}{D}, \qquad r_Z + h\binom{T-1}{D} \leq n_j.$$

For the upper rank estimate, on an indefinite component use the auxiliary sections $\theta_1, \theta_2, \eta_0$ of weights $2w_L - w_{\mathrm{sc}}$, $(D+4)w_L - w_F$, and $I w_{\mathrm{sc}}$, where $I = i_\eta g_*/2$ and w_{sc} is the scalar weight of the boundary metric, so $w_{\mathrm{sc}} = 2w_{\mathrm{de}}$. The companion and boundary section constructions give nonzero readings on the component for all three. The bound on the boundary norm is $\|\eta_0\|^2 \ll e^{-2i_\eta \phi}$, with ϕ the boundary potential in the estimate for the w_k. For a nonzero F of weight κ, consider the quotient

$$\frac{\|F\|^2 \|\theta_2\|^2 \|\theta_1\|^2 \|\eta_0\|^{2/I}}{\left(\sum_{i=1}^D \|u_i\|^2\right)^{T+D+6+1/2}}.$$

It is positive somewhere. If F vanishes to order $T+D+7$ at every point of J, transversality gives decay to zero there. Off these neighbourhoods the lower bound for the generic sections and the scaled generator estimate $\sum_k \|w_k\|^2 \gg e^{-2c\phi}$ show decay at the cusps, since $4/g_* > 2c(T+D+6+1/2)$ by $Tcg_* \leq 1$ and $T \geq 11D^2$ with $D \geq 2$. At the terminal stage one uses θ_2 and any w_k^2 nonzero on the component in the numerator in place of $\theta_2, \theta_1, \eta_0$. The quotient is then considered on the compact definite flag quotient. Both constructions would give a positive interior maximum. At such a point the logarithms of the absolute values of the local holomorphic numerator factors are pluriharmonic. The contribution from the metrics is $\frac{1}{2}\phi_{w_L}$, since the numerator has total weight $(T+D+6)w_L$. The remaining term is the negative of the pullback metric potential of the map $[u_1 : \cdots : u_D]$, times $T+D+6+1/2$. Its Levi form vanishes in a nonzero direction in the kernel of the differential to $\mathbb{P}^{D-1}$. Strict positivity of w_L gives a positive Levi form in that direction, contradicting the maximum. The order-$(T+D+6)$ jet map at J is therefore injective and

$$n_j \leq h\binom{T+2D+6}{D}.$$

Decay of the norm quotient at the cusps · The flag curvature argument · The same auxiliary weights on the definite cover.

These estimates give the rank shares. Put $U = \binom{T+2D+6}{D}$, $L = \binom{T+D}{D}$, and $R = \binom{T-1}{D}$. Then $Rh \leq n_j$, $r_Z \leq (1 - R/U)n_j$, and $r_Q \leq (1 - L/U)n_j$. Each ratio is a product of D factors $1 - a/(T+2D+6-i)$, for $0 \leq i < D$, with $a = 2D+7$ and $a = D+6$ respectively. Expanding the elementary inequality $1 - \prod_i (1-x_i) \leq \sum_i x_i$ for $0 \leq x_i \leq 1$ gives $r_Z \leq D(2D+7)n_j/T$ and $r_Q \leq D(D+6)n_j/T$. Both bounds are at most $22D^2 n_j/T$.

Containment of the deep sections. Subtract a linear combination of the leading sections $u_0^{T-|\alpha|}u^\alpha \theta_i'$ for $|\alpha| < t_0$ from a deep section. The coefficients can be chosen successively in total degree, since at each degree the starting coefficient is a value on J of weight w_F after division by u_0^T. The remainder vanishes to order t_0 at every point of J. The Fourier estimate for the deep condition and the bounds on the subtracted terms give its finite weighted norm $\int_Y \|g\|^2 S^{-T} d\mu$. The weighted division above then expresses the remainder by order-t_0 products with bounded cusp sections of weight $w_F + D w_L$. The initial weighted integrability and the remainder.

The weight-w_0 short forms span the complete weight-w_0 cusp space. Thus one can take the order-t_0 products from this bounded generating set. Their span and the leading span are both defined by matrices over $K_0[z]$ in the weight-κ space. The inclusion of the deep space after extension to $\mathbb{C}$ implies the inclusion

over $K_0(z)$, since it is equivalent to equality of the ranks before and after adjoining the columns of a K_0-basis of C_{deep}. It holds over every Gauss field by extension. This gives the containment of M_{deep} required in the local index argument.

A finite norming set at the selected weight. At a good complex embedding use the actual marked flag components of either run and write $Z_b = frL$. There is a finite set of framed points, depending on the run, the weight κ, and the embedding, such that every bounded cusp section has at one of these points at least one quarter of its invariant supremum. Its cardinality N_σ satisfies $\log(16N_\sigma) \le 1041Z_b^3$ at a uniform threshold. Here is a counting proof with absolute constants.

Start at a point where a nonzero section has at least one half of its supremum. Apply the transvection argument of the principal cusp norm comparison with the chosen full weight and this run's marking. It uses the quadratic polarization discriminant d, the principal transvection period and the weight spread, with logarithms $O(L)$. In its period-lattice units the first minimum is at least e^{-40L}. In the adapted unit frame the symplectic pairing of the principally polarized period lattice Λ has orthogonal real matrix J. Thus $\Lambda^* = J\Lambda$, up to the sign of J, and its Euclidean covolume is one. Put $q_{\mathrm{real}} = 2fr$. By transference the last successive minimum is at most $q_{\mathrm{real}}e^{40L}$. To pass to a genuine integral basis, project along a shortest primitive vector, take an integral basis in the projected lattice, and lift with the component in the first vector bounded by half its length. Induction bounds the length of each basis vector by q_{real} times the last successive minimum. This gives basis matrices A with entries at most $B = e^{CL}$ for an absolute C, since $\log q_{\mathrm{real}} \le 4L$. The determinant has absolute value one, so the cofactor formula gives $\|A^{-1}\| \le e^{C' q_{\mathrm{real}} L}$ with another absolute constant. The same thick locus, The finite norming estimate.

To enumerate the discrete data of these bases, use one primitive element $\alpha \in \mathcal{O}_E$ with $\log \max_\tau(1 + |\tau\alpha|) \ll fL$. For existence, use the trace dual of the lattice $\mathcal{O}_E$ in the sum of its complex metrics. Multiplication by $|\operatorname{disc}(E)|$ sends the codifferent into $\mathcal{O}_E$, whose first minimum is at least one by the norm. Transference and the preceding basis argument give an integral basis with logarithmic lengths $O(fL)$. An integer combination with coefficients between zero and $(2f)^2$ can separate all embeddings. Indeed each pair of embeddings excludes at most $1/(1 + (2f)^2)$ of that box, by varying a coordinate on which they differ. This gives the asserted bound on α.

For each short basis remember its integral alternating matrix, its integral multiplication matrix for α, and its reduction in the prescribed marking. In the unit frame the alternating pairing has bounded operator norm, and multiplication by α has operator norm at most $\max_\tau |\tau\alpha|$. The bounds for A and A^{-1} make the entries of the two integral matrices at most $e^{C'' q_{\mathrm{real}} L}$ in absolute value. The marking modulo its prescribed integer modulus M' has at most $(M')^{q_{\mathrm{real}}^2}$ values, and $\log M' \le 30L$. For a fixed component there are therefore at most $\exp(O(q_{\mathrm{real}}^3 L))$ discrete records.

Two bases with the same record are identified by an integral lattice map preserving the alternating pairing and commuting with α. It is E-linear since α generates E. Applying the alternating pairing after multiplying either argument by elements of E recovers the full Hermitian pairing by trace nondegeneracy. The marks agree, so this identification is in the marked arithmetic group of the component. It preserves the invariant point norms of the sections.

For each record partition the real matrix box $[-B, B]^{q_{\mathrm{real}}^2}$ into cubes of side $\exp(-C_* q_{\mathrm{real}} L)$ for a sufficiently large absolute C_*. In every occupied cube retain one actual flag frame giving that record. The number of cubes has logarithm $O(q_{\mathrm{real}}^3 L)$. After the marked identifications, a frame and the retained frame in its cube differ by $A'A^{-1}$, within $e^{-C_1 L}$ of the identity for any prescribed absolute C_1 upon choosing C_*. Indeed each matrix entry differs by at most the cube side and the inverse satisfies the cofactor bound.

The last closeness bound controls the values of all sections at once. At a unit frame use lower unipotent holomorphic flag coordinates. The leading Gram determinants stay uniformly away from zero near the origin, and the logarithm of the point metric varies by at most $O(r^2 \sum_{\tau,a} |\kappa_{\tau,a}|)\,\|z\|$. This follows by differentiating the determinants in the product defining the weight metric. Here $\kappa_{\tau,a}$ denotes the corresponding angular row entry. The sum of the absolute weights is bounded by a fixed multiple of $frn_\kappa + \mathrm{DF}(\kappa)$, hence

by $e^{O(L)}$ uniformly in the run and stage. The fixed theta and Lie normalization cancels in the local metric ratio. The coordinate change from frames near the identity to the lower unipotent chart is given by Gaussian elimination with invertible leading blocks. Their inverses are bounded near the identity, so the coordinate difference is at most a polynomial in q_{real} times the frame difference. In a coordinate ball of radius $e^{-C_2 L}$, with absolute C_2 sufficiently large, the metric ratio lies in $[7/8, 8/7]$. Apply one-variable Cauchy estimates along the line from the original frame to the retained frame. The holomorphic reading is bounded in the ball by twice the invariant supremum. By taking C_1 larger than C_2 and the coordinate-change allowance, the two raw readings differ by at most one eighth of that supremum. The retained frame then has invariant value at least one quarter of the supremum.

The number of complex components has logarithm at most $41 Z_b^2$. Take the union of the finite sets, including an arbitrary point if the section space is zero. This gives

$$\log(16 N_\sigma) \le 41 Z_b^2 + C_3 (fr)^3 L + \log 16 \le 1041 Z_b^3.$$

The constant C_3 is absolute from the preceding matrix, marking and Cauchy bounds. The quotient of the upper estimate by Z_b^3 is at most $41/Z_b + C_3/L^2 + (\log 16)/Z_b^3$, which tends to zero uniformly at the classical scales. This gives the numerical slack for 1041 by taking the one common lower threshold. Conjugating the norming set supplies the corresponding set at the conjugate embedding.

Scalar evaluation on the finite set is injective on the complex section space, so its restriction after extension from K_0 bounds the rank by N_σ. The normalized generic fibre rank h is at most that section rank by the interpolation and rank comparison. Thus $\log h \le 1041 Z_b^3$ holds for the Chow multiplier. The square root of the mean squared evaluations also gives its Hermitian comparison norm, with distortion at most $4\sqrt{N_\sigma}$ from the supremum. The rank and Hermitian norm consumers.

Determinant content of the order-t_0 products. Choose a rationally spanning set of short integral sections of weight w_0 from S_0'. The same choice of scaled form and coordinates as in the three products gives this set with $\mathrm{mean}_\sigma \log \|\theta^0\|_\sigma \le G_0 \, \mathrm{S}(w_0)$. Here and below the mean counts each embedding of K_0 into $\mathbb{C}$ once. The weight w_L generators of A_L' have a common logarithmic upper bound m_σ at each embedding with mean $m_\sigma \le G_0 L_0$ and $\log N \le P_w(L_0 - 1)$. These bounds are for the scaled generators, including the own-prime powers on each run. They use the integral Hermite span of the theta series, the scalar norm estimates and the chosen bounds on the companion weights. The corresponding rationally spanning short set in S_K' has mean logarithmic norm at most $G_0 K_*$ per element. The scaled full theta generators · The scaled short generators · Norm and spanning bounds on the marked cover.

Keep independent columns of the form $Z_i(z) = u^\alpha \theta^0$ over $K_0(z)$ until they span the order-t_0 space. There are $s = r_Z$ columns. Since the matrices have entries in $K_0[z]$, the same choice works over every Gauss field. The coefficients lie in M_0 by the multiplicative valuation identity on marked points, since each factor is integral at its own weight. At each finite place a, the minimum valuation of a coefficient of a nonzero maximal minor in a local lattice basis equals G_a. It gives the sum of the Smith exponents of these columns over the Gauss valuation ring. A change of basis over the valuation ring preserves this minimum by Cauchy–Binet in both directions. The case $s = 0$ has determinant one and length zero.

Put $d = [K_0 : \mathbb{Q}]$, $n = n_j$ and suppose $s > 0$. Write H_σ for the conjugation-compatible Hermitian norm given by the finite supremum net. It is the square root of the mean of the squared evaluation norms. The number N_σ of evaluations satisfies $\log(16 N_\sigma) \le 1041 Z^3$, and some evaluation has norm at least one quarter of the supremum. The net gives $H_\sigma(F) \le \|F\|_\sigma \le e^{P_d/2} H_\sigma(F)$, where $P_d = 1041 Z^3 (1 + \log K_*)$ suffices and $Z = fr \log b$. Also write p_κ for the selected archimedean constant in $\mathrm{mean}_\sigma \log \|F\|_\sigma \ge -p_\kappa K_*$ for nonzero $F \in M_0$. The finite net uses the same cusp space and all its original supremum balls at the finite places. The lower bound is obtained by choosing a CM frame near a positive supremum value of F, evaluating its nonzero homogeneous frame polynomial in cubical theta and eigenline units, and applying

the product formula to a nonzero coefficient. For a frame with comparison parameter $\Lambda \geq 1$, the lower bound from that coefficient is

$$-\big(p_{\deg} + \tfrac{1}{2}\log\Lambda\big)(A + frn_\kappa) - A\log(r\Lambda), \qquad p_{\deg} = \tfrac{1}{2}(1 + |\log\Omega|),$$

where A is the sum of the row exponents and $\log\Omega$ is the common CM period parameter. The room estimate on the CM frame and $A + frn_\kappa \leq K_*$ give the required p_κ. The coefficient formula uses the normalized degree of the CM eigenlines and holds also on the definite component. The CM frame coefficient bound · The uniform lower bound on integral forms · The finite net on the selected record.

Use the Euclidean norm $(\sum_\sigma H_\sigma^2)^{1/2}$ on $V \otimes_{\mathbb{Q}} \mathbb{R}$. By the arithmetic–geometric mean inequality, every nonzero vector of the embedded M_0 has length at least $\sqrt{d}\exp(-p_\kappa K_* - P_d/2)$. Euclidean transference in dimension dn then supplies a real spanning family in its trace-dual lattice, all of length at most

$$B_{\mathrm{op}} = n\sqrt{d}\exp(p_\kappa K_* + P_d/2).$$

Via the trace pairing they are K_0-linear forms on V. Their embedded Hermitian operator norms are at most B_{op}, by the Euclidean norm of the vector of all components. They take M_0 into the codifferent $\mathfrak{d}_{K_0}^{-1}$, because M_0 is stable under $\mathcal{O}_{K_0}$ and the trace of each such product is integral. Their K_0-linear span is all of V^*. Consequently some s of these forms give a nonzero determinant polynomial on the columns Z_i. Fix a nonzero coefficient $x \in K_0^\times$ of that polynomial. At a, multiplying each row by a generator of the local different makes the rows integral on M_0. By Cauchy–Binet and the Gauss norm,

$$G_a \leq \mathrm{ord}_a(x) + s\,\mathrm{ord}_a(\mathfrak{d}_{K_0}).$$

This also proves that only finitely many G_a are nonzero.

At an embedding σ, take the sum of the Hermitian norms of the coefficients of each column. By submultiplicativity of the supremum norm in the section algebra and the triangle inequality this sum is at most

$$\big(Ne^{m_\sigma}\big)^{t_0}\|\theta_i^0\|_\sigma.$$

Hadamard's inequality applied coefficient by coefficient bounds $|\sigma x|$ by the product of these column bounds times B_{op}^s. Summing the finite bounds with $\lambda_a = \log Na/d$ and applying the product formula gives

$$\frac{C_Z}{s} \leq t_0(\log N + G_0 L_0) + G_0\mathrm{S}(w_0) + p_\kappa K_* + \tfrac{1}{2}P_d + \log n + \tfrac{1}{2}\log d + \log\mathrm{rd}(K_0).$$

The size is one plus an additive weight size. Hence $t_0L_0 + \mathrm{S}(w_0) = K_* + t_0$ and $t_0(L_0 - 1) = K_* - \mathrm{S}(w_0)$. The selected inequalities give $t_0 \leq K_*$, $\log n \leq P_d$, and $\log d, \log\mathrm{rd}(K_0) \leq P_O$. It follows that

$$C_Z \leq r_Z P_D K_*, \qquad P_D = 2G_0 + P_w + p_\kappa + \tfrac{3}{2}(P_d + P_O).$$

The bound takes the degree and different of K_0, the field of the section system before any extension to period coordinates. Transference with the trace pairing · Determinant content with the determinant bounds.

Marked CM arcs for the low Fourier–Jacobi cut

AI-proposed candidate. Work on one of the two principal-level covers in the index calculation, with its own coefficient field K_0, marking and point norm. Write $M = g_0g_1$, $d_E = |\mathrm{disc}(E)|$, $F^+ = E^+$, $W = K_* + r$, $N_A = d_E$ and $N_c = N\mathfrak{c}$. The auxiliary level parameter is 1, so the level $M' = \mathrm{lcm}(M, 1)$ of the toroidal marking is M. The cutoff is $c_1 = \mathrm{rd}(E)^7\,Tc(2f(m-j))^2(m-j)/2$. Use the actual scalar reading and order

of the sign blocks. In particular, each summand of its analytic weight spread is nonnegative and their sum is at most W. Let P be a complex cusp section with a nonzero coefficient

$$Q = \mathrm{FJ}_{\xi,\nu}\, P, \qquad \xi \in \Lambda \setminus \{0\}, \qquad H(\xi,\xi) = 0,$$
$$\nu_\tau = \tau(\nu_0) \geq 0, \qquad 0 < \sum_\tau \nu_\tau \leq c_1.$$

The τ's are the f embeddings in the actual analytic CM type. Here $\nu_0 \in F^+$ belongs to the trace dual of the period lattice of this direction. This assertion about frequency and the zero constant term on each rational isotropic line are the cusp coefficient conditions of the principal-level space. The same conditions hold for complex linear combinations. In the definite terminal case the spaces have no isotropic directions and the cut has zero quotient. Fourier support for the geometric cusp condition · Rationality and positivity of the actual frequencies.

The following detection statement concerns each specified nonzero $Q = \mathrm{FJ}_{\xi,\nu}\, P$ at any nonterminal stage on either marked cover. Fix its level, host lattice, CM type, weight bound and stage cutoff c_1 before choosing P, ξ, ν. There are uniform constants n_Q, Λ_Q for all these choices. Over a finite extension of K_0 there is a marked CM arc on the cusp of $E\xi$. Its abelian part B has dimension $f(r-2)$ and receives an isogeny from the product of the reference CM axes. The tangent lengths of these maps are at most Λ_Q times the reference lengths at every embedding. The extension class and its polarization transpose are killed by n_Q. There is a coefficient $\ell = [Y^\mu q^{l_0}]$ of the marked arc evaluation that detects the given mode. Its tangent degree $A = |\mu|$ satisfies $A + frn_\kappa \leq K_*$.

The same functional ℓ has a value line with ideal $\mathfrak{I}_\ell$, metric factor $N_\ell(\sigma)$, and normalized degree at most $B_{\deg} = \frac{1}{2}(1 + |\log \Omega| + \log \Lambda_Q)K_*$. At every finite place $w \mid a$, for every integral section F of the cusp lattice, one has

$$\mathrm{ord}_w\, \ell(F) \geq \mathrm{ord}_w\, \mathfrak{I}_\ell - e_{w/a} A \max\{\mathrm{ord}_a\, d_E, \mathrm{ord}_a\, n_Q\}.$$

At each complex embedding σ of the arc field, the functional on the corresponding complex sections is bounded by

$$\big|\ell_\sigma(F)\big| N_\ell(\sigma) \leq e^{C_{\mathrm{arch}}} \|F\|_{\infty,\sigma|K_0},$$
$$C_{\mathrm{arch}} = 4\pi + K_* \log\Big(r\Lambda_Q\, \mathrm{rd}(E) \sqrt{\max(1, c_1)} \Big).$$

For $X \geq 1$ bounding f, r and the upper logarithms of d_E, M, N_A, N_c, c_1, W, one may choose $\log n_Q, \log \Lambda_Q \leq 200X^2$. The construction allows the monodromy direction, l_0, the extension degree and the flag matrix to depend freely on the test. The analytic nonvanishing is asserted at the detecting embedding. The bounds concern the action of the chosen coefficient on the whole section space.

A minimum on the reduced lattice. Use the CM ideal of the marked polarization to write $\Lambda = \mathfrak{A}L'$, with L' Hermitian self-dual. Then $N_A\Lambda \subset L'$ and $N_c L' \subset \Lambda$. The alternating form on Λ is $-\mathrm{Tr}_{E/\mathbb{Q}}(\xi_0 H)$. At each embedding use the positive majorant h_x of H at the sign plane of the frame x. Put $c_\tau = 2|\tau(\xi_0)| \leq 2$ and $\mathrm{tt}_\tau = (c_\tau h_x(\xi,\xi))^{-1}$. Write U_x for the Hermitian complement of the real Witt plane spanned by ξ and its image under the complex structure of the sign plane. Projection along $E_\tau \xi$ identifies the reduction $\xi^\perp / E\xi$ with U_x. It therefore gives the reduced majorant at each τ. The lattice

$$L_B = (L' \cap \xi^\perp)/(L' \cap E\xi)$$

is Hermitian self-dual. Locally the summand on the isotropic line is primitive. In a self-dual module its pairing onto the dual of that summand is surjective, by extension of a functional from a direct summand. Splitting off the resulting hyperbolic pair gives the perfect reduced pairing, also when the lattice on the line is a nonprincipal ideal. The two comparisons with the marked lattice · The self-dual reduction.

Write S for the supremum of the weighted scalar norm of Q over the frames with all $\mathrm{tt}_\tau = 1$. A real central translation preserves the majorants, so the Fourier integral bounds this norm by the supremum of

the norm of P. The complex flow in the isotropic line adds the imaginary translation to tt_τ and multiplies Q by a nonzero scalar. It follows that $S > 0$. Fix a frame x_0 in this slice with weighted norm $> S/2$ for Q. Right triangular orthogonalization preserves that value and the majorant. In these coordinates H has matrix $\mathrm{diag}(1_{p_0}, -1_{q_0})$. Use the constants

$$\begin{aligned} N_0 &= M d_E^2, & H_0 &= N_c^2(2 + 4f^4 \operatorname{rd}(E)^2 N_A^2), \\ K_2 &= 1 + N_0 H_0(2c_1 + 1), & N_2 &= 4f, \\ R_0 &= W + N_2 \log(2K_2 N_2) + N_2 + 2, & \epsilon_0 &= \min\bigl(1, f^2/(N_0^2 N_c^2 c_1 R_0)\bigr). \end{aligned}$$

Suppose some nonzero vector of L_B has squared reduced majorants $< \epsilon_0$ at all embeddings. Lift it to $v \in L' \cap \xi^\perp$. The reduced Hermitian square $H(v, v) \in \mathcal{O}_{F^+}$ has all conjugates of absolute value < 1, so is zero. The quotient could be definite, in which case this is already impossible. In the indefinite case replace v by $v + N_A \alpha \xi$, $\alpha \in \mathcal{O}_E$, so its projection coefficient on ξ has absolute value at most $2f^2 \operatorname{rd}(E) N_A$ at every embedding. Indeed the real trace-dual lattice of the embedded $\mathcal{O}_E$ has minimum at least $2/\operatorname{rd}(E)$ in the Euclidean norm over the f complex factors. This follows from the field-norm lower bound on the codifferent. Successive-minimum transference and rounding in a real basis give the stated covering bound. This adjustment preserves the reduced vector. Orthogonality for h_{x_0} and $c_\tau h_{x_0}(\xi, \xi) = 1$ show that $c_\tau h_{x_0}(N_c v, N_c v) \le H_0$.

Consider the two-dimensional rational isotropic plane with basis $(\xi, N_c v)$. In its two-variable Hermitian Fourier expansion use the character

$$\exp\left(4\pi i \sum_\tau \operatorname{tr}(\chi_\tau B_\tau) \right).$$

With this convention the real period lattice contains the Hermitian matrices with entries in $(M d_E/2)\mathcal{O}_E$. In fact the transvection sends a marked lattice vector u to u plus an integral M multiple of $\xi, N_c v$, since $d_E \xi_0 H(\Lambda, \Lambda) \subset \mathcal{O}_E$ by trace duality. Its determinant and the scalar character on it are one. Testing against the off-diagonal pairs $(b, \bar{b})$ and against $b + \bar{b}$ on a diagonal gives $\chi_\tau = \tau(\beta)$ with β Hermitian and $N_0 \beta_{ik} \in \mathcal{O}_E$. Holomorphicity and polynomial growth in the positive tube kill all indices with a negative direction by an imaginary translation of the coefficient integral. A singular surviving index has a rational kernel line. Averaging along the central transvections of this rational isotropic line kills its coefficient, by the geometric zero constant-term condition. The real integrations can be made on common finite-index period lattices, so rescaling a vector on the line leaves the argument unchanged. Thus every nonzero two-variable coefficient of Q has $\chi_\tau > 0$, $\chi_{\tau,00} = \nu_\tau$, and $\prod_\tau (N_0^2 \det \chi_\tau) \ge 1$. The period lattice and its trace-dual denominator · The first diagonal entry of the two-variable support.

Let e_τ be the U_{x_0}-projection of v normalized by $c_\tau h_{x_0}(e_\tau, e_\tau) = 1$, with coordinates $a_\tau = (a_{0,\tau}, a_{1,\tau})$ in this plane. If all squared reduced majorants of v are at most $\epsilon < \epsilon_0$, then $|a_{1,\tau}|^2 \ge (2N_c^2 \epsilon)^{-1}$. The Schur complement gives $\chi_\tau[a_\tau] \ge |a_{1,\tau}|^2 \det \chi_\tau / \nu_\tau$. Apply the arithmetic–geometric mean inequality separately to the quotients $\det \chi_\tau / \nu_\tau$ and to the ν_τ. It gives

$$R(\chi) = 2 \sum_\tau \chi_\tau[a_\tau] \;\ge\; \frac{f^2}{N_0^2 N_c^2 c_1 \epsilon} > R_0.$$

To count $R(\chi) \le t$, express the second basis vector of the plane in the $c_\tau h_{x_0}$-orthonormal basis (ξ, e_τ). Its squared norm is at most H_0. Cauchy–Schwarz for the positive form χ_τ gives $|\chi_{\tau,11}|, |\chi_{\tau,01}| \le H_0(c_1 + t/2)$. Distinct points of $\mathcal{O}_E$ have embedding sup-distance at least one, by the field norm. Packing sup-balls of radius $1/2$ for the two remaining entries of $N_0 \beta$ yields at most $(K_2 t)^{N_2}$ indices for $t \ge 1$.

Flow by the parameter $-i/2$ along each e_τ. Both this flow and the real period translations of the two-plane preserve $\mathrm{tt}_\tau(\xi) = 1$. The rank-one determinant update puts the ratios of the leading Gram determinants in $[1/2, 1]$, with the full determinant unchanged. The logarithms of these ratios decrease to

their minimum at the sign break and then increase to zero. Thus their consecutive weighted differences cost at most $2^{W/2}$ in the norm. At the translated frame the real torus integral has norm at most S. The character then bounds the coefficient at x_0 by $S\,2^{W/2}\exp(-\pi R(\chi))$. For the shell $\lceil R(\chi)\rceil = l$, use $(K_2 l)^{N_2} \le \exp(2l - W - 2)$, which follows from $\log(l/(2N_2)) \le l/(2N_2) - 1$ and $l \ge R_0$. Here $l \ge 4$. With $a = \exp(2-\pi) < 1/2$, the sum of the coefficient bounds divided by S is at most

$$2^{W/2}\exp(\pi - W - 2)\,\frac{a^4}{1-a} \;\le\; \frac{\exp(6-3\pi)}{1-a} < \frac{1}{2}.$$

This absolute convergence allows summation of the torus Fourier series at the origin. The resulting bound on $\|Q(x_0)\|$ contradicts the choice of x_0. Hence the reduced lattice has minimum $\sqrt{\epsilon_0}$ in the supremum of the reduced majorant norms. The complementary rank-one frequency · Transferring the two-variable coefficient along the normalized flow.

The split point and its torsion denominator. Set $d_B = r - 2$ and

$$\rho_0 = (12\pi(c_1 + W))^{-1}, \qquad \mu_0 = 4(f d_B)^3 d_E/\sqrt{\epsilon_0},$$
$$R_B = 160\sqrt{d_B}\,\mu_0/\rho_0, \qquad \Lambda_Q = (2R_B)^{2p_0-1},$$
$$m_B = f(2f d_B)^3 d_E N_A N_c \sqrt{\mathrm{rd}(E)/\epsilon_0}, \qquad n_Q = 1 + \lceil 1000 m_B/\rho_0\rceil.$$

Write $Y_* = e^X$, where X is as in the detection statement. Here $X \ge r \ge 3$. The formulas for the minimum give $N_0 \le Y_*^3$, $H_0 \le Y_*^{11}$, $K_2 \le Y_*^{16}$, and $R_0 \le Y_*^5$, using $4f\log(8fK_2) \le 76X^2$. Hence $\epsilon_0^{-1} \le Y_*^{14}$. The formulas just above give $\rho_0^{-1} \le Y_*^3$, $\mu_0 \le Y_*^{15}$, $R_B, m_B \le Y_*^{21}$, $n_Q \le Y_*^{28}$, and $\Lambda_Q \le Y_*^{22(2p_0-1)}$. Since $p_0 \le r$, both logarithms are at most $200X^2$. These bounds are fixed by the data preceding the choice of the mode Q. The quadratic bounds on the arc constants.

The absolute trace-dual of L_B for the reduced majorant is obtained from $\mathfrak{D}_{E/\mathbb{Q}}^{-1} L_B$ by the isometric sign reversal and the factor 2. Multiplication by d_E clears the codifferent. Thus its Euclidean minimum is at least $2\sqrt{\epsilon_0}/d_E$. Successive-minimum transference and rounding give a real covering radius at most μ_0 and d_B independent E-vectors of norm at most μ_0 in the reduced lattice. To see the latter assertion, select from $2f d_B$ real independent vectors of that size.

Approximate the first $p_0 - 1$ positive orthonormal axes, scaled to length R_B, by reduced lattice vectors simultaneously at all embeddings. The divided error matrix has operator norm at most $\rho_0/160$. Inverting its positive block puts their span in the graph of a map of norm $< \rho_0/120$. The vectors have majorant length $\le 2R_B$. Project a selection of the short vectors onto the Hermitian orthogonal complement of their span and clear the determinant of the positive Gram matrix. This is the fraction-free projection of the ordinary CM approximation, applied only to the positive block. If that block is empty its determinant is 1. The projection of a vector of squared majorant u has Hermitian square of absolute value $\le 2u$: for a graph Z the complementary graph is $(Z^* y, y)$, and the norm of the inverse of $1 - ZZ^*$ is at most $(1 - \|Z\|^2)^{-1}$. The Gram determinant is bounded by $(2R_B)^{2(p_0-1)}$ by Hadamard in the positive span. Choose the complementary vectors so their projections are independent. The resulting integral vectors λ_i in the reduction therefore form an E-basis with opposite signs orthogonal and $|\tau H(\lambda_i, \lambda_i)| \le \Lambda_Q^2$. Both sign planes are definite for these vectors at every embedding. The unitary mapping to this sign decomposition fixes the Witt plane, and its new sign plane has graph norm at most $\rho_0/100$ in the original coordinates. Projection and determinant clearing on the two sign blocks · Placement inside the actual full frame domain.

Write $I = \Lambda \cap E\xi$ and $\Lambda_p = \Lambda \cap \xi^{\perp}$. The projective module Λ_p/I gives an $\mathcal{O}_E$-linear retraction $m \mapsto f(m)\xi$ of Λ_p to I. After the sign-plane move, write the complex toric projection of m as $\alpha(m)\xi$. On the reduced complex space use the linear functional $n_Q(\alpha - f)$. Approximate it by the Hermitian pairing against a vector of the projection of Λ_p, with factor $d_E(-\xi_0)$. To verify the error, write the functional in the Hermitian form by the isometric sign reversal and use the covering radius of $N_c L_B$, which is contained in

that projection. The squared majorant after the sign move is comparable by the factor $t = (1-\rho_0/100)/(1+\rho_0/100)$ and its reciprocal. The residual functional divided by n_Q, multiplied by ξ, has operator norm

$$\kappa \le \frac{d_E N_c \mu_0}{n_Q \sqrt{t}} \le \frac{m_B}{t n_Q} \le \frac{\rho_0}{1000 t}.$$

In the first bound the factor $|\tau(\xi_0)| h_x(\xi,\xi)^{1/2} \le 1$ uses the height-one slice. Let the approximating lattice vector be m_0. The corrected toric projection is

$$n_Q \alpha'(m)\xi = \big(n_Q f(m) + d_E(-\xi_0) H(m, m_0)\big)\xi \in I. \tag{145}$$

The integrality follows by testing the alternating form against every $\mathcal{O}_E$-multiple of m.

Here the correction can be made in the frame domain. In Witt coordinates $H(\xi,\xi^*) = 1$, use the unitary shear fixing ξ which sends $u \in U_x$ to $u + \delta(u)\xi$ and sends ξ^* to $\xi^* + v - \frac{1}{2}H(v,v)\xi$, where $H(u,v) = -\delta(u)$ and $\delta = \alpha - \alpha'$. The image of the sign plane under this shear has projection given by subtracting δ. The change has majorant operator norm at most 3κ at the intermediate plane, by the two off-diagonal blocks and the quadratic block on the Witt plane. In the original majorant this is at most $3\kappa/t$. The image of a graph Z by a matrix $1 + E$ has graph $(Z + E_{-+} + E_{--}Z)(1 + E_{++} + E_{+-}Z)^{-1}$. Using $\|Z\| \le \rho_0/100$, $\rho_0 < 1/36$ and $\|E\| \le 3\rho_0/(1000 t^2)$, its norm is $< \rho_0/21$. The shear preserves the split sign planes on the reduction. Integral toric projection with the prescribed denominator · Combined graph displacement.

The holomorphic graph chart extends the displaced graph along zZ to a disc of radius > 21. On the graph ball of radius ρ_0, the majorants are comparable by $(1+\rho_0)/(1-\rho_0)$. In the chart's holomorphic frames the Gram determinant factors differ from their values at the origin by at most $\exp(4Wr\rho_0^2) \le \exp(1/4)$: the Gram correction in the off-diagonal graph chart is quadratic in the graph. Central flow back to the height-one slice costs at most $\exp(2)$ in the norm of Q. Indeed the character contributes at most $\exp(8\pi c_1 \rho_0)$, and the rank-one Gram ratio contributes at most $(1 + 3\rho_0)^{W/2}$. In the fixed frame at the origin the norm bound on the disc is therefore $e^{9/4}S$. If the reading vanished at $z = 1$, the Schwarz lemma after a disc automorphism would bound its value at $z = 0$ by $e^{9/4}S/21 < S/2$. We obtain a split point x_* with $Q(x_*) \ne 0$, the bound Λ_Q for its abelian axes and the projection condition (145) with the same integer n_Q. Bounds in the normalized graph chart · Nonvanishing on the displaced split frame.

The marked degeneration and the chosen coefficient. Let $X = \Lambda/\Lambda_p$, of rank one over $\mathcal{O}_E$ and rank $2f$ over $\mathbb{Z}$. The marked alternating pairing identifies I with $X^\vee$ and gives a principal polarization on the torus B of lattice Λ_p/I and reduced complex structure at x_*. Its signature is $(p_0 - 1, q_0 - 1)$. The maps from the two reference CM types are the maps of period lattices $\mathfrak{A}\lambda_i \longrightarrow \Lambda_p/I$. The direct sum has finite cokernel. As in the ordinary CM point construction, these maps give an abelian variety over a number field, a good abelian scheme after finite extension, and maps from the integral CM Lie axes. The pullback of the polarization along the i-th map is scaled by $\pm H(\lambda_i,\lambda_i)$, totally positive for the type in use and at most Λ_Q^2 at every embedding. This follows from the trace pairing on this axis and holds for all embeddings because it is an identity of polarized abelian varieties.

Use the rank-one unitary cusp chart of the marked lattice. In its semiabelian data the extension class $c : X \to B^\vee$ is the class of the toric projection on Λ_p/I. In analytic coordinates this is seen by lifting Λ_p/I to $\xi^\perp$: exponentiating the $I_{\mathbb{C}}/I$ coordinates gives the character of the periods defining the extension. Equation (145) makes c and its polarization transpose c^t n_Q-torsion. Both are algebraic over a finite extension and extend to the good abelian scheme. The monodromy is determined by a symmetric $\mathcal{O}_E$-adjoint bilinear form on X. For the functional $\xi_X : X \to \mathbb{Z}$ corresponding to $\xi \in I$, let $p_\xi : X \to E$ be its trace representative, so $\mathrm{Tr}_{E/\mathbb{Q}}(a p_\xi(x)) = \xi_X(ax)$ for $a \in \mathcal{O}_E$. In the convention of this chart the form is

$$b_s(x,y) = \mathrm{Tr}_{E/\mathbb{Q}}(2s\, p_\xi(x)\overline{p_\xi(y)}), \qquad s \in F^+, \quad \tau(s) > 0.$$

Choose s so b_s takes values in $2M\mathbb{Z}$ and $\mathrm{Tr}(\nu' s) = \mathrm{Tr}(\nu_0 s)$ for a nonnegative mode ν' implies $\nu' = \nu_0$. For the latter property, first vary s in a small positive open box around the all-ones vector. Only finitely many

points of the trace-dual frequency lattice can give an equality there, by nonnegativity and comparability of the coordinates of s. Delete their finitely many rational hyperplanes. Take a rational F^{+}-point in the remaining box and multiply it by a positive integer. The integral multiple can be chosen with b_s divisible by $2M$ since $p_{\xi}(X) \subset \mathfrak{D}_{E/\mathbb{Q}}^{-1}$. Then adding one to the complex parameter of the arc effects a principal-level transvection. The exponent $l_0 = \mathrm{Tr}_{E/\mathbb{Q}}(v_0 s)$ lies in $\mathbb{Z}$, since this transvection is a period. Isolation in the actual period lattice · The rank-one period chart with the original marking.

For the constant term of the toric period block, use

$$T_0(x, y) = -H_B(c^t x, c^t y)/(2i) + \Im H_B(u\, c^t x, c^t y), \qquad u + \bar{u} = 1$$

with the host element u and the lifted vectors $c^t x$. The convention is H_B linear in the first variable. This form is symmetric and adjoint because the two Hermitian terms of $u + \bar{u}$ sum to H_B. At $n_Q x, n_Q y$ the factor $\exp(2\pi i T_0 + \pi H_B(c^t x, c^t y))$ is one, with c^t and T_0 evaluated at these multiples. Indeed the H_B terms cancel and $\Im H_B(u\, n_Q c^t x, n_Q c^t y)$ is the integral Riemann pairing of two abelian periods. Changing the constant block to T_0 changes the periods by a complex central rank-one transvection. The marking is retained and Q changes by a nonzero character. The monodromy block is $w b_s$. The same period and rigidification formulas apply when the lift of the marked lattice is chosen by the symmetric cubical line. The normalized toric constant block · Changing the block inside the marked central orbit.

To pass to algebraic data use the toroidal construction from a polarized semiabelian extension, its symmetric cubical trivializations and the positive monodromy form. Here the normalized Poincaré trivialization on $n_Q X \times n_Q X$ is the rigidified unit. This normalization also gives descent of the full constant block. The bilinear comparison of two trivializations has values on $X \times X$ determined up to roots of unity by its values on $n_Q X \times n_Q X$. The defect in the symmetric theta trivialization is an even quadratic map $q_0 : X \to \mathbb{C}^{\times}$. For its polarization b_{q_0} one has $q_0(x)^2 = b_{q_0}(x, x)$ and $b_{q_0}(x, y)^{n_Q^2} = 1$ for a defect with trivial restriction. It suffices to adjoin the values on a $\mathbb{Z}$-basis and the corresponding $2n_Q^2$-th roots of unity. More explicitly one first chooses algebraic generators of the one-dimensional fibres at the finitely many torsion points. The equations for the bilinear and quadratic comparisons then say that a positive power of each constant lies in the enlarged number field. This realizes the normalized complex data over a finite field extension, with compatibility of the level and the cubical rigidification. The marked level is obtained in the semiabelian uniformization by adjoining the indicated finite torsion points and their lifts. The divisibility $2M \mid b_s$ gives integral exponents for the marked q-data. The generic toroidal point has parameter $q = \exp(2\pi i w)$ and the abelian model stays the same B. Root-of-unity bound for the quadratic comparison · Algebraic Poincaré and cubical descent · Descent of the marked arc · Simultaneous descent of the normalized torsion lifts.

The generic toric tangent letter in the τ-slot has d log-coordinates from the τ-component of the trace functional p_{ξ}, with the conjugation prescribed by the marked action. Use as abelian letters the images of the chosen generators of the reference CM Lie axes, lifted through the splitting of the extension after $[n_Q]$ on B and divided on tangents by n_Q. This gives one toric letter for each of the $2f$ embeddings and $f(r-2)$ abelian letters. The exact tangent sequence and the basis on the abelian part give a basis at the generic point. The toroidal coordinate frame comes from these tangents and a nonzero t_B of the rigidified theta line of B. Allow arbitrary block matrices Y in the letters. The function on the frame torsor extends as a polynomial in Y by the polynomial representation of the tangent weight, and is homogeneous of total degree A. Let $G_{l_0}(Y)$ be its q^{l_0}-coefficient. Eigenspace labels are transported by the embedding.

The analytic toroidal period chart and the marked isomorphism identify the q-family, at all invertible block matrices, with the flow of the same rank-one central direction. This is seen from the triangular period matrix of the extension: the toric period entries alone acquire $w b_s$, and the matrices of the tangent letters are constant in the logarithmic torus and abelian coordinates. The identified framed evaluation differs from its scalar reading by the theta normalization. Along the holomorphic flow the absolute value of this comparison is constant, as verified below. The multiplier is then constant by the maximum principle on the upper half-plane, by taking the ratio on the nonvanishing locus of a reading and using holomorphic

continuation. Thus taking the Laurent coefficient isolates exactly Q, since the frequency with trace l_0 is unique. In particular some block frame over the detecting embedding gives $G_{l_0}(Y) \neq 0$: at x_* all analytic tangent bases differ from the arc letters by such invertible block matrices, and a central shift preserves $Q \neq 0$. Consequently a monomial coefficient $\ell(F) = [Y^{\mu}]G_{l_0}(Y)$ detects the given section. For every allowed arc, vanishing of the low mode gives zero on the dense invertible-block locus, so the whole polynomial is zero. This proves the kernel comparison over the complex field as well as on rational sections. Matching the toroidal letters with a detecting frame · The Laurent comparison on the marked flow.

Integral comparison and value-line degree. Fix a finite place w of the field of definition of the normalized arc, and let R be its completed valuation ring with uniformizer ϖ. Give $A = R[[q]]$ the (q)-adic topology. This is a normal noetherian domain, complete for the radical ideal (q). The abelian scheme B_R, the split Raynaud extension $G_R^{\natural}$, its two constant lattices and the rigidified ample cubical line on B_R come from the CM data. Here $G_R^{\natural}$ is the extension of B_R by the split torus with character group X. The two period lattices are identified by the principal polarization. The constant parts of the period trivializations are generators of their integral fibres. This follows by raising the bilinear comparison to the n_Q^2-th power and the symmetric theta comparison to the $2n_Q^2$-th power, using the rigidified units at the trivial torsion multiples. The full trivializations have these constant parts times $q^{b_s(y,x)}$ and $q^{b_s(y,y)/2}$, respectively. The evenness and positivity of b_s give the integral vanishing and convergence conditions for an ample degeneration datum.

Apply the unmarked Mumford equivalence to this datum over A [127, § 4.1, Definitions 4.2.1.10–4.2.1.13, and Theorem 4.4.16]. It gives a semiabelian scheme G with its cubical line $\mathcal{L}$. The formal completion along (q) identifies its Lie algebra with the constant $\mathrm{Lie}(G_R^{\natural})\otimes_R A$. This is an identification over A since these Lie modules are finite projective and A is complete. The abelian locus is $D(q)$, since the monodromy ideal $q^{b_s(y,y)}A$ is the unit ideal at a point of $D(q)$ for every y, and is nonunit at every point with $q = 0$ for $y \neq 0$ [127, Proposition 4.5.3.11 and Corollary 4.5.3.12]. The $\mathcal{O}_E$-action is induced by the equivariant degeneration data. Symmetry, the adjoint identity for the action, and the host rigidified polarization identity extend with the cubical structures by their uniqueness on the normal base [127, Proposition 3.3.1.5, Proposition 3.3.2.2 and Theorem 3.3.2.3]. All these constructions allow arbitrary residual characteristic, including 2, 5, and primes dividing n_Q.

The full level of the prescribed cover is used over $D(\varpi q)$. In the Raynaud description an M-division of a period indexed by y has the specified constant division points and torus characters with monomials $q^{b_s(y,x)/M}$. Indeed the M-th multiple in the extension recovers the prescribed period, and different choices of the divisions differ by $G^{\natural}[M]$. The M-torsion description in the Mumford quotient gives the period cosets indexed by X/MX, together with the M-torsion of $G^{\natural}$ [127, Theorem 4.5.3.10 and § 5.2]. All constant divisions, and the required roots of unity, are available after a finite extension of the field of definition. The exponents are integral because $2M \mid b_s$. Choosing them with their original labels at the complex embedding specifies the original generic marking over $D(\varpi q)$, including its Weil pairing. The $\mathcal{O}_E$-equivariance and pairing equalities are algebraic identities and agree with the marked lattice after the faithful complex base change of the Laurent field. The valuation test extends G, its action, cubical data and frames over the trait and uses the marked point over its characteristic-zero field. At a prime dividing M this definition still uses the same generic level. It imposes no étaleness condition on torsion over the valuation ring. The integral frame test with level on the generic fibre · The polarized semiabelian structure used in that test.

To check primitivity of the theta section, write t_B for a generator of the theta line over R. The toric Fourier expansion on the formal Raynaud extension gives a map from sections of $\mathcal{L}$ to their zero toric coefficient in At_B. Every other coefficient is determined by that one, since the principal polarization identifies the period and character lattices. For the character corresponding to y the coefficient is a translate of the zero one, multiplied by its integral cubical trivialization and $q^{b_s(y,y)/2}$. The coefficient construction

uses the split-torus character summands and the formal pullback to the Raynaud extension. It is therefore compatible with continuous specialization of A, including reduction modulo ϖ. The ample datum associated to $(G, \mathcal{L})$ is the given one under the Mumford equivalence, so the Fourier recurrence uses these same trivializations [127, Theorem 4.2.1.14 and (4.3.1.7)].

On $D(q)$ the principal theta pushforward is invertible, by abelian cohomology and base change. It is free there because $A[1/q]$ is factorial. Choose a generator ω and an integer $N \geq 0$ such that $q^N\omega$ extends to G. Write $a(q)t_B$, $a(q) \in A$, for its zero toric coefficient. Suppose a height-one prime other than (q) divides $a(q)$. For (ϖ) specialize to $(R/\varpi)[[q]]$. Every other such prime is generated by a distinguished irreducible polynomial up to a unit, by Weierstrass preparation. In that case specialize to the completed valuation ring of the finite extension in which the polynomial has the chosen root. In each case q remains nonzero and topologically nilpotent. The specialized zero coefficient vanishes, so all specialized toric coefficients vanish by the recurrence. The section on the formal completion is consequently zero. For an invertible sheaf on the smooth group over this complete domain, a section vanishing on the completion vanishes near the special fibre by Krull intersection in a local trivialization. It then vanishes on the generic fibre, by integrality of the total space. This contradicts the generator over $D(q)$. Hence $a(q) = q^k u(q)$ for some unit $u(q)$ of A.

Normalize the section over $D(q)$ by its zero coefficient t_B. It extends over A: if q^m times this section extends with $m > 0$, all Fourier coefficients reduce to zero modulo q. Its restriction to $G_{R[[q]]/(q)}$ is therefore zero. Multiplication by q on the invertible sheaf is injective and its cokernel is the restriction to $q = 0$, so exactness on global sections allows division by q. Repeating gives the extension. Its reduction on $q = 0$ is the pullback of t_B, which is nonzero on each fibre.

Now specialize q to any nonzero element of positive valuation in a finite extension of R. The pullback theta section is integral and its reduction is nonzero on the special fibre. It generates the sections on the semiabelian trait. Indeed any competing section is a scalar multiple over the fraction field, and a negative valuation of the scalar would force the pullback generator to vanish on the special fibre. Base change of the rigidified normalized theta section gives the same section in the Laurent reading, by its zero toric coefficient. This proves the comparison of the theta frame in the finite test at every place. The nonzero reduction test for the theta generator.

On $D(\varpi q)$ this construction gives the characteristic-zero marked family with its chosen theta frame. Choose integral bases of the saturated $\mathcal{O}_E$-eigenspaces of the constant Raynaud Lie algebra. Each is free over R, and its formation commutes with the finite flat extensions used for traits, by the kernel description of the eigenspace. These bases on the generic marked family give the regular evaluation on the block frame torsor. Its ring of letter coordinates is

$$A[1/(\varpi q)][Y]\left[\left(\textstyle\prod_\tau \det Y_\tau\right)^{-1}\right].$$

The evaluation, viewed over the Laurent field, is a polynomial in the letters. Division by the product of the block determinants, whose leading monomial has coefficient 1, shows that this polynomial lies in $A[1/(\varpi q)][Y]$. In particular a single power of ϖq clears its denominators. Over a specialized trait the theta generator and integral tangent bases just proved give the required integral frame when the block determinants are units. These are the two facts needed for the Gauss estimate as the positive valuation of q tends to zero.

The multiplicative characters of the divided toric letter have coordinates in the codifferent, so multiplication by d_E makes its coordinates integral. For an abelian letter use a local generator of its reference CM line. Its image in the Lie algebra of B is integral. The normalized $[n_Q]$-splitting of the semiabelian extension is defined over the DVR. Hence multiplying this letter by n_Q makes it an integral tangent letter of the same model. The theta contribution is entirely the valuation of the cubical line.

Fix $0 < \eta \in \mathbb{Q}$ and test at $v_w(q) = \eta$ over a ramified finite extension, with unit block determinants. At that radius a maximal Gauss term among the letter and Laurent coefficients is attained, and only finitely

many terms have maximal norm. Reduction after dividing by this norm gives a nonzero Laurent polynomial in the unit of q and the residue matrices. Over a further finite residue extension one can arrange that this polynomial and the block determinants are nonzero. Lifting gives tests in the punctured disc of the original marked arc with integral frames, so the supremum-integral rule has norm at most one there. It follows that each coefficient has norm at most $|\varpi|^{-\eta m}$ for its Laurent exponent m. Letting $\eta \downarrow 0$ proves the integral coefficient bound. The matrices of the scaled letters can now be substituted into this polynomial over R, even if their determinants reduce to zero. Now substitute the integral letter matrices after the two scalings, using homogeneity in each set of columns. The trace representative and the $[n_Q]$-splitting give respectively the costs $v_w(d_E)$ and $v_w(n_Q)$ per column exponent. Thus the whole degree-A coefficient costs at most $A\max(v_w(d_E), v_w(n_Q))$. The value line is

$$\mathcal{N}_\ell = (T_B^\vee)^{\otimes n_\kappa} \otimes \bigotimes_i (\ell_i^\vee)^{\otimes \mu_i},$$

where T_B is the cubical theta line, ℓ_i the reference Lie line (unit line for a toric letter), and μ_i sums the monomial powers of entries attached to the same letter. With $\mathfrak{I}_\ell$ its value ideal this gives

$$v_w(\ell(F)) \geq v_w(\mathfrak{I}_\ell) - e_{w/a} A \max\big(v_a(d_E), v_a(n_Q)\big) \qquad (F \in M_0).$$

Compatibility of toroidal models and the fixed good abelian scheme · The signed valuation of the theta normalization · The abelian letter lift with denominator n_Q · Bounded denominators in the marked arc chart · Integral framed traits on the punctured disc · The Gauss coefficient estimate on the integral frames.

With the probability Haar Gaussian metric on T_B and the π-scaled Lie metric, each abelian reference line has normalized degree $\frac{1}{2}\log\Omega$, by the CM eigenline calculation for the same host. The actual abelian dimension is $g_B = f(r-2)$. The integral tangent maps from the reference axes are generically an isomorphism, and the image of each has length at most Λ_Q times its reference length. Hadamard on their wedge therefore gives normalized degree at least $g_B(\frac{1}{2}\log\Omega - \log\Lambda_Q)$ for $\det \mathrm{Lie}\, B$. The principal cubical key formula in these metrics is

$$\widehat{\deg}_{\mathrm{norm}}\, T_B^\vee = -\tfrac{1}{2}\widehat{\deg}_{\mathrm{norm}} \det \mathrm{Lie}\, B + \tfrac{g_B}{4}\log(2\pi).$$

It follows that, with $p_{\deg} = (1 + |\log\Omega|)/2$, $\widehat{\deg}_{\mathrm{norm}}\mathcal{N}_\ell \leq (p_{\deg} + \frac{1}{2}\log\Lambda_Q)K_*$, because $\log(2\pi) < 2$ and $A + g_B n_\kappa \leq K_*$. The value-line degrees and the dimension of the abelian part.

Archimedean comparison at the same arc. At an embedding of the arc field in the given CM orientation, the comparison of the toroidal and scalar readings on the orbit has the form

$$\begin{aligned} G_{l_0}(Y)e^{2\pi i w l_0} &= \gamma\, Q(x_w), \\ |\gamma| &= \left(C\prod_\tau \mathrm{tt}_\tau(x_w)^{-1}\right)^{n_\kappa/2} N_{0,n_\kappa}(x_w). \end{aligned} \tag{146}$$

Here N_{0,n_κ} is the scalar point-norm factor for the theta weight alone. For the metric identity, choose a symplectic period basis of the arc extending a basis of its toric periods I. If Π_w is the associated Siegel matrix, the probability L^2 squared norm of its standard theta section is $(\det(2\Im\Pi_w))^{-1/2}$, by integrating the characters in the real period coordinates and then the remaining Gaussian over the real vector space. The abelian coefficient has the identical formula using the abelian block of $\Im\Pi_w$. The ratio of these Gaussian factors is $2^{-f}\sqrt{\det G_I}$, where G_I is the real Gram on I of the polarization metric. This follows from the complementary principal minor of $(\Im\Pi_w)^{-1}$, whose toric block is G_I. If $b_1, \dots, b_{2f}$ is the trace basis of $\{a \in E : a\xi \in I\}$, the Gram in this basis has entries $\sum_\tau \Re(\tau(b_i)\overline{\tau(b_j)})/\mathrm{tt}_\tau(x_w)$. Its determinant is

$$|\mathrm{disc}(b)| \prod_\tau \big(2\mathrm{tt}_\tau(x_w)\big)^{-2}.$$

The fractional ideal contains $\mathcal{O}_E$, so its trace discriminant is at most d_E in absolute value. With the abelian generator t_B the metric identity therefore uses $0 < C \le \|t_B\|^2\sqrt{d_E}$. The theta-only norm factor is the product over the embeddings of $D_{p_0}^{n_\kappa/2} D_r^{-n_\kappa/4}$ in absolute value, for the leading Gram determinants D_i. The Gram update gives that D_{p_0} and tt_τ scale by the same factor along the flow while D_r is constant. Thus $N_{0,n_\kappa}\prod_\tau \mathrm{tt}_\tau^{-n_\kappa/2}$ is constant along the flow. This proves the absolute-value assertion used in taking the Laurent coefficient. Gaussian norm of the full theta generator · Gaussian norm of the abelian coefficient · The toric Gram determinant and its discriminant bound.

For the normalized constant block, the height $\mathrm{tt}_\tau(x_w)$ is exactly $\tau(s)\Im w$. Indeed the imaginary part of T_0 cancels the abelian term in the Schur complement for the toric periods. The canonical trivialization at the multiples of n_Q makes this identity compatible with all embeddings. Flow now from x_w to the point of parallel height $\epsilon_c = 1/\max(1, c_1)$. The change in the Fourier character has modulus

$$\exp\left(4\pi \sum_\tau \nu_\tau(\epsilon_c - \mathrm{tt}_\tau(x_w))\right) = e^{4\pi\epsilon_c \sum_\tau \nu_\tau}|e^{2\pi i w l_0}|.$$

In (146) this cancels the Laurent factor exactly. The remaining exponential is at most $e^{4\pi}$. In the new frame the theta metric is evaluated at height ϵ_c, and its comparison factor is at most

$$\|t_B\|^{n_\kappa}\big(\mathrm{rd}(E)\max(1, c_1)\big)^{f n_\kappa/2}.$$

The normalized toric period block · Exact height without a constant offset · The theta and character along the parallel flow.

For the tangent factor, the logarithmic derivative on a toric character introduces $(2\pi i)^{-1}$. The norm of the corresponding toric Lie letter in the π-scaled metric is at most $\epsilon_c^{-1/2}$. On the abelian Lie directions, the lift through the normalized $[n_Q]$-splitting divided by n_Q is isometric for the abelian metric. In period coordinates it lands in the complex subspace of the real span of the zero-monodromy periods. There the Riemann pairing, hence its Hermitian metric, is the abelian one, and the central flow acts trivially. Thus the lifted abelian letter of generator u_i has length at most $\Lambda_Q\|u_i\|$. These computations give a common bound $\Lambda_Q\epsilon_c^{-1/2}\|u_i\|$, with $\|u_i\| = 1$ for the toric letter. Use the polydisc with radii $\|u_i\|^{-1}$ for entries multiplying letter i. Each resulting tangent column has length at most $r\Lambda_Q\epsilon_c^{-1/2}$, by the triangle inequality. The scalar point-norm formula bounds its tangent weight factor by the product of these column lengths to the prescribed powers, using Hadamard in each block. On invertible blocks the Fourier integral at height ϵ_c is bounded by the supremum of F in the given point norm. The polynomial extends the bound to singular blocks. Cauchy's coefficient estimate in Y, followed by the metric divisions defining the value line, yields

$$|\sigma(\ell(F))|\, N_\ell(\sigma) \le \|F\|_{\infty,\sigma|K_0}\exp(4\pi + P_{\mathrm{arc}}K_*),$$
$$P_{\mathrm{arc}} = \tfrac{1}{2}\log\max(1, c_1) + \log\Lambda_Q + \log\mathrm{rd}(E) + \log r.$$

At the opposite CM orientation use the paired complex-conjugate embedding. The absolute coefficient, the value-line metric, and the supremum norm agree under this conjugation and eigenspace relabelling. Together these estimates give the required common family for detection and the quotient height, on the same cover with the same coefficient weights. The positive direction s used to isolate l_0 contributes no height factor. Isometry for the normalized abelian lifts · Toric and abelian lengths at parallel height · The archimedean coefficient estimate on the actual value line.

The quotient by the low Fourier–Jacobi cut. Give the quotient by C_{deep} the infimum of the supremum norms of the complex lifts at each embedding. Each detecting row kills C_{deep} by its Fourier–Jacobi coefficient, so its archimedean bound holds with this quotient norm. For each nonzero integral quotient

class x, choose a row with $\ell(x) \neq 0$. Apply the product formula to its value, using the finite-place inequalities and the degree of its value line. Then

$$\frac{1}{[K_0 : \mathbf{Q}]} \sum_{\tau : K_0 \hookrightarrow \mathbf{C}} \log \|x\|_{\mathrm{quot},\tau} \geq -C_{\mathrm{arch}} - B_{\deg} - A(\log d_E + \log n_Q).$$

In this lower bound the archimedean norm uses all complex embeddings over K_0, with their equal multiplicities. The finite sum of the maxima of the two orders is bounded by the sum of both logarithms. Thus the same test controls the first minimum of the integral low-mode quotient.

Take $Q = M_0/(M_0 \cap C_{\mathrm{deep}})$. This is torsion free of rank $s = r_Q$, and after localization or completion it is the quotient by the corresponding saturated cut. Choose s short sections $\theta_1, \dots, \theta_s$ from S'_K whose images are a rational basis of Q. Their existence follows by projecting the rationally spanning short set. The resulting cokernel is torsion. Its local length over $\mathcal{O}_{K_0,a}$ is q_a, since passage to the Gauss valuation ring preserves the uniformizer and Smith exponents. At the definite stage the quotient is zero. We now treat the indefinite stages with $s > 0$.

The marked CM arcs give separating rows ℓ on this quotient, with their normalized fractional ideals $\mathfrak{I}_\ell$ and absolute values $|x|_\tau N_\ell(\tau)$ on the CM value lines. For each nonzero row let A be its letter degree, so $A + frn_\kappa \leq K_*$ by (117). The preceding arc proof gives the common allowances

$$\begin{aligned} b_a &= A \max\bigl(\mathrm{ord}_a \,|\, \mathrm{disc}(E)|, \mathrm{ord}_a\, n_Q\bigr), \\ B_{\deg} &= \bigl(p_{\deg} + \tfrac{1}{2}\log \Lambda_Q\bigr) K_*, \\ C_{\mathrm{arch}} &= 4\pi + K_* \log\bigl(r\Lambda_Q \,\mathrm{rd}(E)\sqrt{\max(1, c_1)}\bigr). \end{aligned}$$

The first measures the finite losses in the scaled frame, the second bounds the normalized degree of the value line, and the third bounds the increase over the logarithmic supremum.

Choose finitely many separating arc functionals on V/C_{deep}, successively reducing their common kernel over K_0 to zero. Pass to a finite normal field extension L_1/K_0 containing all their coefficient fields and adjoin their Galois conjugates. The common kernel over L_1 of the resulting rows is Galois stable, so descends to K_0 and is zero. Hence one can select s rows forming a basis of the dual quotient over L_1. Conjugation carries each line and bound to the corresponding line and bound at the conjugate places, and extension multiplies the finite-place orders by the ramification index. These choices work at all places at once.

Let $\Delta \neq 0$ be the determinant of the chosen rows on the short quotient basis (θ_i). At a place w of L_1 over a, the map given by the rows sends the localized quotient lattice into the direct sum of fractional ideals $\varpi_w^{-e_{w/a}b_a}\mathfrak{I}_{ell_i,w}$. It is injective on that lattice. The cokernel of the short basis in the quotient lattice therefore injects into the cokernel of its image in this larger lattice. More precisely, the smaller length is the length in this larger lattice minus the length of the larger lattice modulo the image of the quotient lattice. The larger length is a determinant valuation, giving

$$e_{w/a} q_a \leq \mathrm{ord}_w(\Delta) - \sum_{i=1}^{s} \mathrm{ord}_w(\mathfrak{I}_{ell_i}) + s e_{w/a} b_a.$$

Here localization of the saturated cut commutes with the field extension. The equality of base-changed length with $e_{w/a}q_a$ follows from the elementary divisors over the original valuation ring.

Use $\sum_{w|a} e_{w/a}\lambda_w = \lambda_a$ with the weights on L_1. The product formula and Hadamard's inequality on Δ bound the sum by the following expression, where H_{sh} bounds the mean logarithmic norm of each short form θ_i.

$$C_Q \leq s\left(H_{\mathrm{sh}} + C_{\mathrm{arch}} + B_{\deg} + \tfrac{1}{2}\log s + \sum_a \lambda_a b_a\right).$$

The short form bound may be taken as $H_{\mathrm{sh}} = Z^4 K_*$. Indeed the scaled short bound $G_0 K_*$ is dominated by this budget. To see the determinant estimate directly, multiply row i at each embedding by $N_{\ell_i}(\tau)$

before applying Hadamard to the columns. The mean of the s line degree corrections is then at most sB_{deg}. The mean of the logarithmic norms of the columns is taken over the embeddings of K_0, because every such embedding has the same number of extensions to L_1. This argument is unchanged by real or complex embeddings when each embedding is counted once. The common field and the row lines · The local determinant and its normalized descent.

For the numerical bound let C_A be the coefficient of the bound for the chosen numerical parameters and C_K the coefficient of the weight bound. The splitting and approximation constants can be taken with

$$K_b = 20 + |C_A| + 7|C_A| + |C_K|, \qquad C_0 = 17 + 2K_b + 200K_b^4.$$

In fact $Z \geq 1$ and the logarithms of the inputs to the construction of n_Q, Λ_Q are at most $K_b Z$. These inputs include the marked level and the two lattice norms. For c_1 this follows from $Tc \leq 1$, and for the weight input from $K_* \leq C_K k f b$. The approximation uses the first minimum cutoff and the radius $1/(12\pi(c_1 + K_* + r))$. The explicit bounds for its index and its comparison of frames give $\log n_Q, \log \Lambda_Q \leq 200(K_b Z)^4$. Substitution in the three estimates gives

$$C_{\mathrm{arch}},\ B_{\mathrm{deg}},\ \sum_a \lambda_a b_a\ \leq C_0 Z^4 K_*.$$

The rank bound for the marked section system gives $\log s \leq \log n_j \leq 1041 Z^3$ for this calculation. Thus we may use

$$C_Q \leq r_Q P_Q K_*, \qquad P_Q = 2(1 + 3C_0 + 1041)Z^4.$$

All estimates use the run's marked cover and its geometric field. Both own-prime choices meet the same inequalities before passage to the period coordinates. The cutoff and frame denominator bounds · The quotient constants on a marked run · Domination by the chosen numerical budget.

For the common index budget one can use $P_O \leq 217Z^2$, $G_0 \leq 2000Z^2$, $P_w \leq 256Z^2$, $p_\kappa \leq 3600Z^4$, $P_d \leq 4164Z^4$, and $P_H \leq 5 \cdot 10^6 Z^{10}$. In particular $P_D \leq 14500Z^4$. The threshold for $Z = frL$ gives $1 + 3C_0 + 1041 \leq Z$, so $P_Q \leq 2Z^5$. The selected bound satisfies $P \geq 1000Z^{12}$ and $Z \geq 1000^{20}$. The four terms in $60G_0 + 180P_H + 44P_D + 22P_Q$ therefore sum to less than $10^9 Z^{10} \leq P^2 D$, as required for the combined index estimate.

The determinant and the boundary on the generic intersection. Use $F = K_0$ and the scaled sections $w_0, \dots, w_N$ generating A'_L on this run. Here the subscript labels a generator of weight w_L, and $w_F + Dw_L$ continues to denote the weight of the second short module. Set $L_0 = \mathrm{S}(w_L)$, and write $\overline{f}$ for the average of $f(\sigma)$ over the embeddings of F in $\mathbf{C}$. In the generic intersection put D independent rows $U = (U_{ik})$ for the vanishing equations and set $B = F(U)$. The additional row used for dividing the sections is a purely transcendental extension of B. It has Gauss order zero on its coefficients and leaves the local lengths and norm orders below unchanged.

For each component C let X_C be the closure of the moving image in $\mathbf{P}^N_F$, and use its normalization in the original component field Ω_C. The ratios of the w_k have transcendence degree D. The universal incidence is, over this normalization, the vector bundle of D row vectors annihilating the tautological point. Its function field L_C is finite separable over B, of degree h_C. It has the form $\Omega_C(U^{\mathrm{free}})$ with DN free variables, obtained by eliminating one coefficient per row. Set $L = \prod_C L_C$ and $h = \sum_C h_C$. These are the finite algebra and rank used for the local section fibre. A normalized branch above a height-one place of the base has the marked order of its restricted component valuation. For a Gauss place above a, the total of $e_\zeta f_\zeta$ is h. At a horizontal place, any branch with nontrivial component valuation is the Gauss extension of that valuation in the DN remaining variables. Indeed the base residue field has transcendence degree $D(N+1)-1$ over F, and the restricted residue field has transcendence degree at most $D-1$. Eliminate in a column of minimal order so all the substitution ratios are integral. Equality in the residue transcendence degree bound then makes the remaining variables independent over the restricted residue field. The same

calculation over $\kappa(a)$ at a Gauss place uses $D(N+1)$ and D. In particular the dimensions, ramification, residue degrees, and shifts in the finite norm are those of the local section calculation. The degree over a Gauss place and The normalized branch correspondence.

For a nonzero polynomial P in row variables, let $H_0(P)$ be the sum of the normalized finite log Gauss norms and of the archimedean means of $\log|P|$ on independent complex unit spheres, one per row. Let $H_{\mathrm{tor}}(P)$ use instead independent unit circles for each variable. Finite norms and both archimedean means are additive on products in logarithmic notation. Constants have height zero by the number-field product formula. For a block homogeneous P of degrees d_i one has

$$-\tfrac{1}{2}H_N \sum_i d_i \ \le H_0(P) \le H_{\mathrm{tor}}(P), \qquad H_N = \sum_{k=1}^{N} k^{-1}.$$

For the lower bound choose a nonzero lexicographically leading coefficient. Apply the one-variable Jensen inequality to this coefficient successively at the coordinate radii of the spheres. Each coordinate has expected log radius $-H_N/2$, as follows from the beta distribution of its squared magnitude. The finite Gauss bounds on the same coefficient cancel its archimedean contribution. The upper bound follows by increasing each coordinate radius to one in Jensen's circular mean. The same argument proves log integrability for polynomials, and hence for the logarithms of the absolute values of rational functions.

Write $\operatorname{ord}_v s$ for the order of a section in its marked weight frame on a branch. Orders and norms add under tensor product in logarithmic notation. For a horizontal irreducible polynomial $P \in F[U]$, set

$$n_P(s) = \sum_{\zeta|P} f_\zeta \operatorname{ord}_\zeta s, \qquad a(s) = \sum_P n_P(s) \deg_{U_1} P, \qquad N_a(s) = \sum_{\zeta|a} f_\zeta \operatorname{ord}_\zeta s.$$

These definitions include all components, and s is nonzero on every component. A horizontal branch trivial on the component field has order zero. Every horizontal P with a nonzero component order is block homogeneous, with equal degree in the D rows. To see this, its centre in the normalized moving image is a prime divisor by the residue dimension calculation. Its image in X_C is a divisor since the normalization map is finite. The incidence image of that divisor is the hypersurface of tuples cutting it, defined by the Chow polynomial in D rows. There are only finitely many nonzero $n_P(s)$ and $N_a(s)$. Take the graph of the rational map from the normalized moving image to the marked projective model of the weight lines and normalize the graph. It is an isomorphism above the generic point of each divisor of the normalized moving image. Thus the orders at the relevant component centres are the orders of the rational sections of the pulled back weight lines there. Their divisors have finite support. The incidence image of each such divisor determines at most one height-one place of the base. The horizontal norm factors and their support.

For any weight G used below, choose a nonzero rational section s on every component and define

$$\begin{aligned}\mathcal{H}_G = \sum_a \lambda_a N_a(s) + \int \sum_{\zeta \text{ in the complex fibre}} -\log\|s(\zeta)\|\, dU \\ + \sum_P n_P(s) H_0(P), \qquad a_G = a(s).\end{aligned} \tag{H1}$$

The archimedean integral uses the D independent spheres and counts the h geometric points of the generic fibre. Both $\mathcal{H}_G$ and a_G are independent of the rational frames s and additive in G. In fact, replacing s by sr, for $r = (r_C) \in \prod_C \Omega_C^\times$, changes the first two terms by the negative of the finite and archimedean log norms of $\mathrm{Norm}_{L/B}\, r$. The divisor norm of r has order $\sum_{\zeta|P} f_\zeta \operatorname{ord}_\zeta r$ at P, so its factorization cancels the change in the last term. The total degree of these factors in any row is zero, since scaling the row preserves the fibre and leaves the norm of r fixed. This proves the assertion for a_G as well. All archimedean norms here are the marked weight norms. Their log integrability follows from the same pointwise estimates used below. In fact for $r \in \Omega_C^\times$, the characteristic polynomials of r and r^{-1} over B bound both absolute values

by the Cauchy root bound on every point of the generic complex fibre. The coefficients are rational in U, so these bounds make $\sum_\zeta |\log|r(\zeta)||$ integrable. In particular $U_b = \log(\beta/\|w_b\|)$ is integrable for a section w_b nonzero on C, since it is bounded above by a constant plus the sum of the positive logarithms of the nonzero ratios w_k/w_b.

At a nonterminal stage let t_1 be the auxiliary section of weight $2(w_L - w_{\mathrm{de}})$ and set $r = t_1^{n_\eta}\eta_0^2/w_b^{2n_\eta}$. The archimedean bound on β and the upper bound for $\|t_1\|$ give

$$(1 - cg_*)\log\|\eta_0\| \geq \tfrac{1}{2}\log|r| - n_\eta\left(U_b + C_5\,\mathrm{DF}(w_L) + cC_2 + \tfrac{1}{2}\log\|t_1\|_{\sigma,\sup}\right).$$

Since $1 - cg_* \geq 1/2$ and η_0 is bounded above in norm, its logarithm is integrable. The bound on β then makes $\log\|w_b\|$ integrable, and the weight identity with r gives this for t_1. For w_F, choose a short section θ_i' nonzero on C. The logarithms of $\|\theta_i'\|$ and of $\|t_2\|$ are bounded above and their sum is $(D+4)\log\|w_b\| + \log|\theta_i' t_2/w_b^{D+4}|$. Both are therefore integrable, and so are the corresponding rational frames by comparison with their ratios. In the terminal case the bound on $-\log\beta$ directly treats $\log\|w_b\|$, after which the same second weight identity applies. These calculations concern the measurable generic complex fibres of the original marked section system. Frame independence and the two weight identities.

The normalized moving height. For a branch set $b_\zeta = \min_k \mathrm{ord}_\zeta\, w_k$ and put $\beta = (\sum_k \|w_k\|^2)^{1/2}$ at a complex point. Use a test section $s = \sum_k x_k w_k$, $x_k \in F$, nonzero on each component. Form $\mathcal{H}_{\min}$ by replacing $\mathrm{ord}_\zeta\, s$ by $\mathrm{ord}_\zeta\, s - b_\zeta$ and $\|s\|$ by $\|s\|/\beta$ in (H1). Let $F_C(X,U)$ be the Chow form of X_C in $D+1$ rows, raised to $[\Omega_C : F(X_C)]$. Its degree in each row is h_C. As a polynomial in X, it is the norm of $\sum X_k w_k$ on the D-row incidence fibre, with the weight frame removed and with its $F[U]$-content removed. This follows by factoring over the geometric generic fibre. At any P the Gauss order of that norm in the new row is the sum over the branches of the minimal coefficient orders times their residue degrees. Consequently the polynomial $F_C(x,U)$ has horizontal factor exponents $\sum_{\zeta|P,C} f_\zeta(\mathrm{ord}_\zeta\, s - b_\zeta)$. The difference between the Gauss orders of $F_C(x,U)$ and $F_C(X,U)$ at a gives the analogous vertical sum. At infinity the new unit-sphere row averages the logarithm of a linear form of length β to $\log\beta - H_N/2$, in any unit weight frame. Add these formulas and use the number-field product formula for the constant in $F_C(x,U)$. This gives

$$\mathcal{H}_{\min} = \sum_C H_0(F_C) + \tfrac{1}{2}hH_N, \qquad a_{w_L} = h + \sum_P \deg_{U_1}(P)\sum_{\zeta|P} f_\zeta b_\zeta. \tag{H2}$$

In particular this calculation counts the actual generic degree of the moving system, including any degree of its function-field extension. The generic norm form and its specialization and The height with the sphere correction.

The two global estimates before inserting the constants. Choose rational weight-w_F frames e_C, and let $b_i = (\theta_i'/e_C)_C \in L$ for the selected short basis. The trace discriminant $\Delta = \det(\mathrm{Tr}_{L/B}(b_i b_k))$ is nonzero. Over the Gauss DVR at a the normalized weight lattice consists of $z \in L$ with $\mathrm{ord}_\zeta\, z + \mathrm{ord}_\zeta\, e_C \geq 0$ on every branch. Its quotient by the basis columns has length $I_a^0 \geq I_a$. The comparison follows by placing the local section fibre in L after adjoining the free row and completing at a. Its generator in the normalized fibre has precisely the same shift, and an inclusion of normalized lattices preserves the indicated direction of length. The added row and completion preserve lengths, since each base uniformizer still has order one. The required fractional ideal can be changed to the integral closure by multiplication by an element whose branch orders are $\mathrm{ord}_\zeta\, e_C$, using weak approximation in the semilocal normalization. Its determinant changes by its field norm, of order $N_a(e)$. The discriminant of the integral closure has nonnegative order, by integrality of trace. The change of basis for the b_i therefore gives

$$2I_a \leq \mathrm{ord}_a^{\mathrm{Gauss}}\,\Delta + 2N_a(e). \tag{H3}$$

The identical trace argument at every horizontal P applies because the θ_i' have nonnegative marked orders. Thus $\Xi = \Delta \prod_P P^{2n_P(e)}$ is a nonzero polynomial in $F[U]$, block homogeneous of degree $2a_{w_F}$ in every row. Here Δ is homogeneous of degree zero under row scalings, since the b_i come from the component sections.

Apply the product formula for the height of Ξ, using the finite Gauss norm to sum the $\operatorname{ord}_a^{\mathrm{Gauss}} \Delta$ in (H3). The horizontal factor terms combine with the last term of (H1). On the complex fibre, Δ is the square of the determinant of the evaluation matrix $(b_i(\zeta))_{\zeta,i}$, since the trace pairs via the sum over embeddings. Hadamard bounds the square of this determinant by $h^h \prod_i \|\theta_i'\|_{\sigma,\sup}^2$ times the product of the inverse squared norms of the e_C over the h points. The latter product combines with the archimedean part of (H1). With $\ell_N = \log(N+1) + 1/2$, the lower bound on $H_0(\Xi)$ is at least $-2Da_{w_F}\ell_N$. This proves

$$2\sum_a \lambda_a I_a \le h\log h + 2\sum_i \overline{\log \|\theta_i'\|_{\sigma,\sup}} + 2\mathcal{H}_{w_F} + 2Da_{w_F}\ell_N. \tag{H4}$$

The right hand side of (H3) also shows finite support, since the trace discriminant is a fixed rational function and the shift $N_a(e)$ has finite support. Local discriminant and index and Global trace bound and horizontal correction.

At $j < m$, take $s = \eta_0$, of weight $n_\eta w_{\mathrm{de}}$ with $n_\eta = i_\eta g_*$. The local section calculation measures $v_\zeta(\eta_0) = \operatorname{ord}_\zeta(\eta_0)/e_\zeta$ in base units, so $N_a(\eta_0) = \mathcal{N}_a$ exactly. Its horizontal orders are nonnegative. In (H1) the horizontal contribution is at least $-Da(\eta_0)\ell_N$, and the archimedean contribution is at least $-h\overline{\log\|\eta_0\|_{\sigma,\sup}}$. Put $H_{\mathrm{de}} = \mathcal{H}_{n_\eta w_{\mathrm{de}}}/n_\eta$ and $a_{\mathrm{de}} = a(\eta_0)/n_\eta$. The boundary estimate is

$$\frac{1}{i_\eta}\sum_a \lambda_a \mathcal{N}_a \le g_*(H_{\mathrm{de}} + Da_{\mathrm{de}}\ell_N) + hP_\eta. \tag{H5}$$

The factor g_* comes from the weight of η_0; the residue degrees in N_a count the normalization points once. There is no further conversion on the geometric field. The boundary-height estimate and The equality of local readings.

Inserting the run bounds. Let t_1, t_2 be the actual integral auxiliary sections of weights $2(w_L - w_{\mathrm{de}})$ and $(D+4)w_L - w_F$, nonzero on each component. Their mean log sup norms are at most $2P_I DL_0$ and $P_I DL_0$. Write

$$H_L = \mathcal{H}_{w_L}, \quad H_F = \mathcal{H}_{w_F}, \quad a_L = a_{w_L}, \quad a_F = a_{w_F},$$
$$a_1 = \tfrac{1}{2}a(t_1), \quad a_2 = a(t_2).$$

The effective orders and the weight identities give

$$0 \le a_{\mathrm{de}}, a_1 \le a_L, \quad 0 \le a_F, a_2 \le (D+4)a_L, \qquad a_L \le h + \tfrac{1}{2}cg_*a_{\mathrm{de}} \le \tfrac{4}{3}h. \tag{H6}$$

For the penultimate inequality use $b_\zeta \le c\,\operatorname{ord}_\zeta(\eta_0)/(2i_\eta)$ on horizontal branches in (H2). The branch depth bound and $\sum_{\zeta|a} e_\zeta f_\zeta = h$ give the vertical shift of $\mathcal{H}_{\min}$ at most

$$h\sum_a \lambda_a \mathrm{nus}(a) + \frac{c}{i_\eta}\sum_a \lambda_a N_a(\eta_0).$$

The archimedean shift uses the bound $-\log\beta \le C_5\,\mathrm{DF}(w_L) + (c/i_\eta)(-\log\|\eta_0\|) + cC_2$. In the horizontal shift, replace b_ζ by $(c/i_\eta)\operatorname{ord}_\zeta(\eta_0)$, allowing the possible negative height of P by $H_0(P) \ge -D\deg_{U_1} P\,\ell_N$. Adding gives

$$\begin{aligned} H_L &\le \mathcal{H}_{\min} + h(P_N + C_5 + C_2)L_0 + Dcg_*a_{\mathrm{de}}\ell_N + cg_*H_{\mathrm{de}},\\ H_{\mathrm{de}} &\le H_L + hP_I DL_0 + Da_1\ell_N,\\ H_F &\le (D+4)H_L + hP_I DL_0 + Da_2\ell_N. \end{aligned} \tag{H7}$$

The last two inequalities follow by taking the heights of t_1, t_2 in (H1). The finite terms there are nonnegative and the horizontal terms have the sphere lower bound. The first inequality uses $0 \le c \le 1$,

$\mathrm{DF}(w_L) \le L_0$, $L_0 \ge 1$, and $\sum_a \lambda_a \mathrm{nus}(a) \le P_N L_0$. Degree comparison of the auxiliary weights and The finite, horizontal and archimedean shifts.

The coefficient estimate below gives

$$\sum_C H_0(F_C) \le (D+2)h(P_0' L_0 + G_0 L_0 + \log(N+1) + 1). \tag{H8}$$

The run has $\ell_N \le P_w L_0$, $\log(N+1)+1 \le P_w L_0$, $H_N/2 \le P_w L_0$, and $P_H = 2(P_0' + P_N + C_5 + C_2 + P_I + P_w) + G_0$. Set

$$X = \mathcal{H}_{\min} + h(P_N + C_5 + C_2)L_0 + \tfrac{2}{3}DhP_w L_0, \qquad Y = Dh(P_I + \tfrac{4}{3}P_w)L_0.$$

Since $cg_* \le 1/2$, (H6)–(H7) imply $H_L \le X + cg_* H_{\mathrm{de}}$, $H_{\mathrm{de}} \le H_L + Y$, and $H_F \le (D+4)(H_L + Y)$. Use the nonnegative upper bound X_0 on X obtained from (H2) and (H8). Solving the two inequalities yields

$$H_{\mathrm{de}} \le 2X_0 + 2Y, \qquad H_F \le (D+4)(2X_0 + 2Y),$$
$$2X_0 + 2Y \le (D+1)h(2G_0 + 6P_H)L_0.$$

For the last inequality, the largest coefficient among the six summands in the parentheses defining P_H is $6(D+1)$, the coefficient of P_w. The coefficient of G_0 is $2(D+2)$. Equations (H4)–(H6) and the short-basis norm $\sum_i \overline{\log \|\theta_i'\|_{\sigma,\sup}} \le hG_0F_0$ give the two fibre terms in the index estimate below. At $j = m$, use $c = 0$ and $b_\zeta \le e_\zeta \mathrm{nus}(a)$ at vertical branches and $b_\zeta = 0$ at horizontal branches. The archimedean shift has no boundary term, and the same t_2 identity gives the first inequality. The boundary term B is zero there. If the fibre rank is zero, the lengths and boundary sums in this calculation are zero. Combining the numerical bounds and The nonterminal and terminal substitutions.

The coefficient estimate for the moving Chow forms. Use $D+2$ rows V of length $N+1$, and $D+2$ variables $x = (x_0, \dots, x_{D+1})$. Set

$$K_1 = D+1, \qquad M_1 = (D+2)(N+1),$$
$$n = 8M_1^2 K_1 h(1 + \lceil P_O \rceil), \qquad n' = 8nM_1 K_1 h.$$

Here $P_O \ge 1$ bounds the logarithm of the root discriminant of F. The lower norm estimate needed at this point is on the actual products of n of the scaled w_k. For an $\mathcal{O}_F$-linear combination s of these products, nonzero on C, it gives

$$-P_0' n L_0 \le \overline{\log \|s|_C\|_{\sigma,\sup}}. \tag{H9}$$

Use the principal-level estimate (118) at the weight nw_L. Each w_k is an integral cusp section on this same geometric cover. Multiplication of their rules multiplies each framed evaluation, so integral linear combinations of n such products lie in the integral supremum lattice at nw_L. Products of the bounded cusp readings still have zero central coefficient, since all Fourier indices are nonnegative and the zero index is absent.

Here C is an irreducible component over the coefficient field K_0 of the original run, in the smooth generic cover $X_o \otimes K_0$. Such components are open and closed, since regular local rings are domains. Its constituent complex components over any embedding of K_0 may be conjugate. Extend the section on C by zero on the other K_0-components. The resulting section is still cuspidal on all complex components, and on a generic framed trait it has either the original value or zero. Thus it is in the same integral lattice, also when special fibres meet. Components over the coefficient field · Component restriction and integral evaluations.

At a fixed embedding σ of K_0 the component norm is the supremum on the readings from embeddings $\Omega_C \hookrightarrow \mathbf{C}$ extending σ. They are dense in each complex flag chart of C. For an integral affine chart over K_0, remove the zero sets of the countably many nonzero chart functions. Each is nowhere dense on every complex component, because multiplication by that function remains injective after the scalar extension

and the complex chart is smooth. Points in the complement give embeddings of the fraction field, and Baire's theorem gives density. The same argument on the flag charts uses the independent flag variables of the generic reading. The actual Lie and theta point norm is continuous in local frames. Hence the supremum on the zero extension equals this component supremum at each σ. The suprema in the actual point metric.

For the $N + 1$ chosen generators the multiplier is $n = 8((D + 2)(N + 1))^2(D + 1)h(1 + \lceil P_O \rceil)$. For $L_0 = \mathrm{S}(w_L)$ one has $\mathrm{S}(nw_L) \le nL_0$. Use the selected bound B_{gen} for $\log(N + 1)$. With $\log h \le 1041Z_b^3$ and $1 + \lceil P_O \rceil \le 2 + P_O$ gives

$$1 + \log \mathrm{S}(nw_L) \le 1 + \log L_0 + \log\bigl(8(D + 2)^3(2 + P_O)\bigr) + 2B_{\mathrm{gen}} + 1041Z_b^3.$$

The selected degree constant P_0' is $C_L Y^2$ times this factor. Choosing the fixed constant $C_L \ge C_*$ gives $-P_0' nL_0 \le \overline{\log \|s|_C\|_{\sigma,\sup}}$, as in (H9). The product lattice at the chosen weight · The constants at this multiple.

Let Q_C be the irreducible Chow polynomial of X_C; $F_C = Q_C^{m_C}$ up to an element of $F^\times$, where $m_C = [\Omega_C : F(X_C)]$. Substitute the $D + 1$ rows

$$U_i = x_0 V_i - x_i V_0, \qquad 1 \le i \le D + 1.$$

After inverting x_0 this change of row coordinates is an automorphism, so the resulting Q_C is $x_0^b G_C$ up to a constant, for an irreducible polynomial G_C coprime to x_0. Write $s_C = \deg_x G_C$, $u_C = \deg_V G_C$. Evaluation at $x_i = V_i \cdot y$, with $y_k = w_k/w_b$ on C, annihilates the substituted Chow polynomial, because those rows share the point y. The universal scaled evaluation $x_i = Z(V_i \cdot y)$ has a prime kernel of height one. Indeed, the image field is contained in $\Omega_C(V, Z)$, of transcendence degree $M_1 + D + 1$ over F. Over $F(V)$ the general linear projection by the $D + 2$ rows maps the projective D-dimensional image generically finitely onto its image. This follows by choosing independent general hyperplanes through the projection centre, or by Noether normalization. The image after adjoining the scale of one nonzero x_i thus has transcendence degree $D + 1$ over $F(V)$. The kernel is (G_C). This also shows $s_C \ge 1$, since no nonzero polynomial in the free V variables vanishes under evaluation alone. The change of rows gives

$$m_C s_C \le h_C, \qquad u_C \le (D + 1)h_C.$$

Here the x-degree bound uses that the normalized residual incidence equation defines the general projection. More explicitly, fix general V. Its $D + 2$ coordinates give a morphism from X_C to $\mathbf{P}^{D+1}$, since the centre avoids X_C. The image is a hypersurface, whose degree times the generic degree of the projection is $\deg X_C$, by the projection formula for the hyperplane class. The polynomial $G_C(V, x)$ is its irreducible equation over $F(V)$, with that degree in x. This yields the first bound, including when F is not algebraically closed. The V-degree bound follows from the total degree of the original Chow polynomial in its rows. The incidence relation and Its degrees and evaluation kernel.

The coefficient of highest x_0-degree in the substituted F_C is the original F_C evaluated at $V_1, \dots, V_{D+1}$. Jensen in x_0 and the coefficient Gauss inequality give $H_{\mathrm{tor}}(F_C) \le m_C H_{\mathrm{tor}}(G_C)$. In combination with $H_0(F_C) \le H_{\mathrm{tor}}(F_C)$ this reduces (H8) to bounding the torus height of G_C. Only $H_0(F_C) > 0$ needs a bound, in which case $H_{\mathrm{tor}}(G_C) > 0$.

The monomial coefficient lattice of polynomials of homogeneous x-degree n and V-degree at most n' has rank

$$A(n, n') = \binom{n + K_1}{K_1}\binom{n' + M_1}{M_1}.$$

Give it the Euclidean coefficient metric at each embedding. For the following calculation abbreviate $A = A(n, n')$ and $k = A(n - s_C, n' - u_C)$. Multiply this latter monomial space by G_C and saturate its image in the integral coefficient lattice. Its normalized arithmetic degree is at most $-kH_{\mathrm{tor}}(G_C)$. At a finite place, divide G_C by its content in that DVR. The shifted copies of the primitive polynomial are independent modulo the uniformizer, as multiplication by its reduction is injective in the residue polynomial ring. This computes

the saturation contribution exactly. At an archimedean place, the volume of the k vectors G_C times the monomials is the product of their Gram–Schmidt distances. In lexicographic order every such distance is the infimum of $\|G_C P\|_2$, where P has leading coefficient one and preceding terms from the allowed monomials. Parseval and Jensen on the torus give $\log\|G_C P\|_2 \geq \int \log|G_C| + \int \log|P| \geq \int \log|G_C|$. Combining these determinant contributions gives the degree inequality.

Apply Minkowski's convex body theorem to the saturated quotient of rank $A-k$, with its quotient Euclidean metrics. Its normalized degree is at least $kH_{\mathrm{tor}}(G_C)$, by additivity and the degree zero of the ambient standard lattice. A nonzero integral quotient vector has a representative Q whose quotient norm $q_\sigma(Q)$ satisfies

$$\overline{\log q_\sigma(Q)} \leq -\frac{k}{A-k} H_{\mathrm{tor}}(G_C) + \tfrac{1}{2}\log A + P_O. \tag{H10}$$

For example, apply Minkowski to boxes of real intervals and complex discs of coordinate radius $R/\sqrt{A-k}$ in orthonormal quotient frames. The covolume is the usual number-field discriminant factor times the exponential of minus the unnormalized degree. The box estimate yields the half logarithm of the rank and half the logarithm of the root discriminant, which are bounded by the displayed allowance. Degree of the saturated relation lattice and Minkowski's bound on the quotient.

The evaluation of Q at $x_i = \sum_k V_{ik} w_k$ has a coefficient of V nonzero on C, since G_C generates the rational relation ideal. Each coefficient is an integral combination of n-fold products of w_k. Fix such a coefficient s nonzero on C. At any embedding, the corresponding coefficient functional annihilates the relation subspace and has operator norm at most $\sqrt{A}((N+1)e^{S_\sigma})^n$, where $S_\sigma = \max_k \log\|w_k\|_{\sigma,\sup}$. Indeed each x-degree-n monomial contributes a coefficient of absolute value at most the total absolute sum of its terms, at most $(\sum_k \|w_k\|_{\sigma,\sup})^n$. Cauchy–Schwarz over the A monomials gives the asserted bound. Taking the infimum over a representative at each embedding and applying (H9)–(H10) yields

$$\frac{k}{A-k} H_{\mathrm{tor}}(G_C) \leq n(P_0' L_0 + \overline{S_\sigma} + \log(N+1)) + \log A + P_O. \tag{H11}$$

Coefficient evaluation and its quotient norm.

For clarity, the binomial ratio gives $k/(A-k) \geq n/((D+2)s_C)$. Indeed set $s = s_C, u = u_C$, $W = n + (K_1+1)s$, and $z_1 = K_1 s/(n-s+1)$, $z_2 = M_1 u/(n'-u+1)$. The product formula for the two binomial ratios gives $A/k \leq \exp(z_1+z_2)$. Using $s \leq h$, $u \leq K_1 h$, $n \geq 8M_1^2 K_1 h$ and $M_1 \geq K_1 + 1$ gives

$$z_1 W \leq K_1 s + s/4, \qquad z_2 W \leq 2/7 \leq s/3.$$

For the first inequality expand $W = (n-s) + (K_1+2)s$ and use $s/(n-s) \leq 1/(8M_1^2 K_1 - 1)$. For the second use $W \leq 2n$ and $n' - u + 1 \geq 7nM_1 K_1 h$. Hence $1 - z_1 - z_2 \geq n/W$, and $\exp(z_1+z_2) \leq (1-z_1-z_2)^{-1} \leq W/n$. Also $\log A + P_O \leq n$. In fact $n' M_1 \leq n^2$ and $D \geq 3$ give $A \leq n^{4M_1}$; $\log n \leq \sqrt{n}$, $n \geq 32M_1^2$, and $P_O \leq n/32$ give the desired bound. Insert these in (H11), use $\overline{S_\sigma} \leq G_0 L_0$, and multiply by m_C. The result is

$$H_0(F_C) \leq (D+2)h_C(P_0' L_0 + G_0 L_0 + \log(N+1) + 1).$$

Summing over C proves (H8). Binomial ratio bound and The component sum.

At a coefficient place a, use the Gauss extension with ring O_a', of uniformizer ϖ_a, and write $M = M_0 \otimes O_a'$. All the modules are now extended to O_a'. Write N_1 for the span of the products $u_0^{T-|\beta|} u_1^{\beta_1} \cdots u_D^{\beta_D} \theta_i'$ for $|\beta| < T - D$, with the short θ_i' forming a basis of the rational weight-w_F fibre. Write Z_0 for the lattice spanned by a rational basis selected from the products $u_1^{\beta_1} \cdots u_D^{\beta_D} \theta^0$ for $|\beta| = T - D$ and short $\theta^0 \in S_0'$. Its saturation in M is Z_{sat}. The rank of Z_0 is r_Z. The length of Z_{sat}/Z_0 equals the minimum Gauss valuation G_a of the nonzero r_Z-minors of the column matrix, by Smith form. For rank zero the empty minor gives $G_a = 0$. Use $S = \mathrm{span}_{\mathrm{Frac}\, O_a'}(N_1 + Z_0)$. The deep space lies in S by section division. Intersect it with M to obtain M_{deep}. The rational quotient by the deep space has a basis given by the images of r_Q elements of S_K'. Let q_a be the length of the quotient of M by M_{deep} and the lattice generated by these r_Q elements. At the terminal stage the deep space is the whole cusp space, so $q_a = 0$.

Let E_a be a nonnegative integer with $E_a \geq T\mathrm{nus}(a)$ and $X = \varpi_a^{E_a} M$. The image of X modulo $\mathcal{T}$, modulo the image of $X \cap S$, has length at most q_a. This follows by multiplying the deep lattice and its short complement by $\varpi_a^{E_a}$. Put $J_0 = N_1 + Z_0$. The image of $X \cap S$ is a quotient of $((X \cap S) + J_0)/J_0$. Reduction after division by u_0^T through degree $T - D - 1$ has kernel $\mathrm{span}(Z_0)$ on S. On $(X \cap S) + J_0$ this kernel is contained in Z_{sat}, and it contains Z_0 on J_0, so the kernel contributes at most G_a. The fibre length and pole allowances from the normalized jet calculation contribute at most

$$\binom{T-1}{D}\left(I_a + \sum_{\zeta} f_\zeta m_\zeta\right).$$

Here I_a is the length of the normalized order-zero fibre modulo the short basis. At each branch with ramification and residue degrees e_ζ, f_ζ, the pole allowance is $m_\zeta = \max(0, e_\zeta(Tv_\zeta - E_a))$. Both entries of this maximum are integers. Put $\mathcal{N}_a = \sum_\zeta e_\zeta f_\zeta v_\zeta(\eta_0)$, with v_ζ in the units of the coefficient place. The point valuation bound gives $\sum_\zeta f_\zeta m_\zeta \leq Tc\mathcal{N}_a/(2i_\eta)$ for $j < m$, and all $m_\zeta = 0$ for $j = m$. With the middle term read as zero at $j = m$, we have

$$\mathrm{length}\big((X + \mathcal{T})/\mathcal{T}\big) \;\leq\; r_a := \binom{T-1}{D}\left(I_a + \mathbf{1}_{j<m}\frac{Tc\mathcal{N}_a}{2i_\eta}\right) + G_a + q_a.$$

The right side is independent of the choice of E_a within the stated range. Local index from the four terms.

The weighted sum and arbitrary subspaces. For a place a over p, put $\lambda_a = [\kappa(a) : \mathbf{F}_p]\log p/[K_0 : \mathbf{Q}]$ and $R = \binom{T-1}{D}$. Use the weighted sums $A = \sum_a \lambda_a I_a$, $B = \mathbf{1}_{j<m}(Tc/(2i_\eta))\sum_a \lambda_a \mathcal{N}_a$, $C_Z = \sum_a \lambda_a G_a$, and $C_Q = \sum_a \lambda_a q_a$. The simultaneous rank and height estimates give

$$\begin{aligned} R(A + B) &\leq n_j\left((60G_0 + 180P_H)D^2K_*/T + (5G_0 f r^2 + P_d)T + TcP_\eta/2\right), \\ C_Z + C_Q &\leq n_j(44P_D + 22P_Q)D^2K_*/T. \end{aligned}$$

To obtain the first line, divide (H4) by two and use $a_F \leq (4/3)(D+4)h$, $H_F \leq (D+4)(D+1)h(2G_0 + 6P_H)L_0$, and the short-basis norm $hG_0\mathrm{S}(w_F)$. In (H5), multiplication by $Tc/2$ makes the coefficient of $H_{\mathrm{de}} + Da_{\mathrm{de}}\ell_N$ at most $1/2$. Together, the terms proportional to hL_0 in $A + B$ are at most $(30G_0 + 90P_H)D^2hL_0$: expand the bounds and use $D + 4 \leq 5D$, $D + 1 \leq 2D$, $P_w \leq P_H$, and $D \geq 1$. The other terms are bounded by $h[(5G_0 f r^2 + P_d)T + TcP_\eta/2]$. For the first line use $Rh \leq n_j$, $P_w \leq P_H$, $Tcg_* \leq 1$, $\log h \leq P_d$, $\mathrm{S}(w_F) \leq (D + 2)\mathrm{S}(w_L) + 5fr^2T$, $T\mathrm{S}(w_L) \leq 2K_*$, and $(D + 4)(D + 1) \leq 10D^2$. Here $K_* = \mathrm{S}(\kappa)$. For the second line use $r_Z \leq r_Q + 22D^2n_j/T$, $r_Q \leq 22D^2n_j/T$, and the nonnegative height bounds P_DK_* and P_QK_* per rank. The selected inequalities for the run give

$$\sum_a \lambda_a r_a \leq n_j\big(P^2D^3K_*/T + P^3T + P^2Tc\big).$$

In fact the bound on the first coefficient is

$$60G_0 + 180P_H + 44P_D + 22P_Q \leq P^2D.$$

The other two bounds are $5G_0 f r^2 + P_d \leq P^3$ and $P_\eta/2 \leq P^2$. Since $K_* \leq C_K kfb$, $4C_K \leq P^2 f$, and $T \geq P^4D^3kf/(md_2)$, the first term is at most $fmbd_2/4$. The same lower bound for T gives $Tfmd_2 \geq 4$. Thus $P^3T^2 \leq b$ gives the same bound for P^3T. Lastly $Tcg_* \leq 1$ gives $P^2Tc \leq P^2 \leq P^3T$. Consequently $\sum_a \lambda_a r_a \leq n_j fmbd_2$. Numerical estimate on the four weighted terms · The height bounds of the selected family.

Completion and the Gauss extension preserve the length in (4.3.7). Indeed, in a Smith basis for $M_0, \mathcal{T}$ the summands of the length are $(b_i - E_a)_+$, where b_i are the Smith exponents, and each of these extensions preserves the uniformizer. Passing to the period field at a place $v \mid a$ multiplies the exponents by $e(v/a)$,

including $E = e(v/a)E_a$. Write M_v, L_v for the placed lattices. On a subspace W of the period coefficient space, $M_v \cap W$ is a lattice of rank $\dim W$. Its image in $M_v/(M_v \cap \varpi_v^{-E} L_v)$ has length at most $e(v/a)r_a$. Scaling its kernel by ϖ_v^E gives a sublattice of $L_v \cap W$. Hence

$$\operatorname{length}\big((M_v \cap W)/(L_v \cap W)\big) \le E \dim W + e(v/a)r_a.$$

This holds even for subspaces defined only over the period field or its local completion. The containment $L_v \subseteq \sigma L_{K,v}$ then bounds the signed determinant difference on W by

$$\operatorname{ord}_v \det(L_{K,v} \cap W) - \operatorname{ord}_v \det(M_v \cap W) \le (E - \operatorname{ord}_v \sigma) \dim W + e(v/a)r_a,$$

where both determinants are measured in the same determinant line. The containment $L_v \subseteq \sigma L_{K,v}$ also bounds the signed determinant difference between $L_{K,v} \cap W$ and $M_v \cap W$ by

$$(E - \operatorname{ord}_v \sigma) \dim W + e(v/a)r_a.$$

Indeed the same determinant line measures both lattices, and scalar multiplication adds its order times $\dim W$. With the analogous place weight λ_v in the period field, $\sum_{v|a} e(v/a)\lambda_v = \lambda_a$. The two runs each have this full bound. Use the run of own prime p'_{de} at places over p_{de} and the run of own prime p_{de} elsewhere. The weighted residual at the eligible places in each run is at most $n_j f m b d_2$. The residual applies to every subspace in that run with the same choice of numbers. Descent and the residual bounds on both runs.

The primewise heights of the short weight

The harmonic depth on the selected runs. At stage j, use the chart on the rank-r target with signature $(P_j, Q_j) = (k + j - 1, m - j)$ and level $M_0 = g_0 g_1$. The two section-product runs have the same marked moduli datum. At a place of residue characteristic p_{de} choose the run of own prime p'_{de}. At all other places choose the run of own prime p_{de}. This rule applies to both members of a conjugate pair, so the own-prime factors in the section products are units at the two places used. The marked run and its alternate prime.

Write $a_{\tau,i}$, $1 \le i \le d_\tau$, for the nonnegative decreasing rows of the actual short weight w_L, where $d_\tau = P_j$ on Φ and $d_\tau = Q_j$ on its conjugate. The positive row at $\tau \in \Phi$ is read from w_L at the prescribed real label underlying τ. The row at $\bar\tau$ is the negative-block row of that same label, with its prescribed order. Put $D_L = \sum_{\tau,i} a_{\tau,i}$. The degree, ordering and copy count in the local infimum are those of w_L. In particular the selected $n = n_{w_L}$ satisfies $n \ge fr(r+1) \ge 2fr$, and $\sum_\tau d_\tau = fr$. Assign an injection $(\tau, i) \mapsto b(\tau, i)$ into the n copies. The two rows at the recorded embeddings and The copy count of the short weight.

For each point in the depth supremum, let V be its complete discrete valuation ring and a the place of the geometric coefficient field K_0 that it extends. If the relative ramification is e', then $\operatorname{ord}_V(p) = e' e_a$, where $e_a = \operatorname{ord}_a(p)$. We calculate first with the absolute normalization $v(p) = 1$. The supremum uses exactly the integral framed points with finite residue field. At such a point the integral chart is a formal chart of the smooth semiabelian model with the tangent directions of the given marking. In the chart, the τ-vectors $u_{\tau,i}$ are a basis of the integral τ-eigenvectors. Their span is saturated in the entire Lie module: an integral vector with a nonzero multiple in the span satisfies the same eigenvector equations, since one can cancel the nonzero scalar in the torsion-free Lie module. The given basis of the integral eigenspace then contains the vector in its span. Saturation over the discrete valuation ring gives an integral right inverse to the matrix of these vectors. This right inverse persists under valued extension. Its span stays saturated and has the full generic eigenrank, so it stays the integral eigenlattice. Thus each full typed flag individually admits integral dual linear forms, including at the ramified primes. The finite residue field of an algebraic test and Saturation in the marked chart.

For the character calculations, split the torus of the special fibre after a finite residue field extension and lift that extension to a complete trait. Raynaud uniformization of the polarized degeneration identifies the

formal group at the identity with that of a split torus extension of an abelian scheme. The finite extension in this construction is allowed for each point of the supremum because its residue field is finite. The group law and the $\mathcal{O}_E$-action in the transported chart are the base changes of the original marked chart. The character differentials below are therefore evaluated on the base changes of the original $u_{\tau,i}$, with their given order. The original integral Poincaré frame and the canonical gauge are also used after base change. Uniformization after trait extension and Compatibility of the transported group law and action.

The harmonic polynomial on the prescribed copies. With $a_{\tau,d_\tau+1} = 0$, use the polynomial

$$\mathcal{P}(Z) = \prod_{\tau} \prod_{i=1}^{d_\tau} \left[\det(Z_{\tau,s,b(\tau,t)})_{1\le s,t\le i}\right]^{a_{\tau,i}-a_{\tau,i+1}}.$$

Its coefficients are integral. Each leading minor is unchanged when an earlier row is added to a later one. The degree in row (τ, s) is $\sum_{i\ge s}(a_{\tau,i} - a_{\tau,i+1}) = a_{\tau,s}$. The variables of conjugate types use disjoint copies, so every term in each conjugate contraction annihilates $\mathcal{P}$. These are the row and harmonic conditions in the depth infimum. In particular every differential in a used column acts in one typed tangent space. The admissible harmonic operators and The leading-minor polynomial.

Let $\mathcal{H}(x) = h(x, w_0 x)$ be the single-copy canonical gauge in the original chart. Its value at the identity is zero. On n copies the value of the harmonic germ test at level M_0 is

$$\mathcal{P}(D)\left(F \exp\left(M_0 \sum_b \mathcal{H}(X_b)\right)\right)\Big|_{X=0}.$$

The marked gauge in the depth definition and The integral biextension after scalar extension.

Character columns at the small exceptional primes. For $p \mid Q_S$, the imaginary quadratic subfield $K_{\mathrm{im}} = \mathbf{Q}(\sqrt{-d})$ splits at p. At odd p this is the host's splitting condition. At 2, $w_0 = (1 + \sqrt{-d})/2$ has polynomial $X^2 - X + (d + 1)/4$, with two distinct roots modulo 2 since $d \equiv 7 \bmod 8$. The two roots lift over $\mathbf{Z}_2$. In either case the factors of E above the distinguished quadratic place have labels in Φ, and their conjugate factors have the other labels. Use R for the valuation ring in a complete algebraically closed extension of the uniformizing trait's fraction field. The character lattice in this formal group is the valuation-ring span of the differentials of a finite list of integral characters from the dual Tate module.

On the large types denote its projected colengths by ℓ_τ. The abelian dual period argument on each whole local factor w gives colength at most $[E_w : \mathbf{Q}_p]Q_j/(p-1)$ in its dual Lie module. The Tate Hodge quotient is ω_B on the original w-side, and its period determinant uses the Lie ranks of $B^\vee$ on that side, hence the smaller ranks of B at the conjugate labels. Lifting the split torus characters and then projecting to the separate saturated eigenflags gives

$$\sum_{\tau\in\Phi} \ell_\tau \le f Q_j(1 + \delta_p/2).$$

Here δ_p is the host's mean local discriminant depth on these factors and $\delta_p \le \log_2 Q_S$. The projection cost is at most $Q_j/2$ times the discriminant valuation on each whole factor by the Vandermonde lattice calculation. This includes the dyadic case, using the determinant comparison over the common quotient modulo 2. The host's split-place conditions and The large flag colength on the transported chart.

For the small types use the inclusion of $w\omega_B$ in the span of the abelian character differentials, where $w = \zeta_p - 1$. It follows from the inclusion of tM in the Tate lattice of the integral period module and its Hodge quotient. The torus characters lift through the finite flat Cartier-dual torsion of the Raynaud extension and span the torus cotangent lattice. Consequently the cokernel on the semiabelian cotangent lattice is a quotient of the abelian cokernel and is killed by w. Since $p \in wR$, the full span contains p times the cotangent lattice. Restrict to the saturated τ-flag using its integral right inverse. The span of

the projected character differentials contains pR^{d_τ} and its colength is at most $d_\tau = Q_j$. The Hodge–Tate inclusion for the small flags.

For each type one can choose a single ordered list $\chi_{\tau,1}, \dots, \chi_{\tau,d_\tau}$ of actual integral characters with all leading differential minors bounded by that full colength. To see this, let $L_i \subset R^i$ be the span of the projections of the finite generating list to the first i rows. Inductively choose the first $i-1$ columns to form a basis of L_{i-1}. The map $L_i \to L_{i-1}$ is surjective. Subtract from each candidate column an integral combination of the previously chosen columns that matches its first $i-1$ rows. The resulting residuals generate the rank-one kernel. Choose a nonzero residual of least valuation in the last coordinate, and use its original character as the next column. This gives a basis of L_i. The leading determinant at size i has valuation $\operatorname{length}_v(R^i/L_i)$, which is at most the full colength by projection. All determinants are nonzero. For the small types we can thus take $\ell_\tau = Q_j$. The simultaneous leading character minors.

Eliminating the gauge on each used column. In $\mathcal{O}_{K_{\mathrm{im}}} \otimes \mathbf{Z}_p$ choose the idempotent e_τ that acts as the identity on the τ-directions. Conjugation exchanges the two idempotents, so $e_\tau \bar{e}_\tau = 0$. The coefficient action on the formal group extends to the completion. Indeed $[p](x) = px + O(x^2)$, so the coefficients of any fixed jet of $[p^v](x)$ tend to zero as v increases. Comparing two representatives modulo p^v by the group law proves convergence of their endomorphisms on each jet.

For $a \in \mathcal{O}_{K_{\mathrm{im}}}$, write $N_a = a\bar{a} \in \mathbf{Z}_{\geq 0}$. The adjoint identity moves $[a]$ from the first argument of h to the second as $[\bar{a}]$. The second argument is then $[N_a]w_0 x$. Repeated addition in that argument gives N_a times the original gauge. The corrective terms for both operations are logarithms of integral units, so

$$\exp\bigl(\mathcal{H}([a]x) - N_a \mathcal{H}(x)\bigr) \in R[[x]]^\times$$

with constant term one. Approximate e_τ by such a, so N_a tends to zero. All these power series operations are continuous coefficient by coefficient. In particular the exponential of a series with constant term zero uses a finite polynomial in its coefficients at every fixed degree. The limit gives

$$W_\tau(x) = \exp\bigl(\mathcal{H}([e_\tau]x)\bigr) \in R[[x]]^\times.$$

This uses the integral biextension on the given chart. The integral multiplier for a single copy.

In the copy $b = b(\tau, i)$ take the integral germ

$$F_b(x) = \chi_{\tau,i}([e_\tau]x)\, W_\tau(x)^{-M_0}.$$

Use $F_b = 1$ in the unused copies, and $F = \prod_b F_b$. If u is one of the directions in a τ-copy, the invariant derivative in direction u commutes with pullback along $[e_\tau]$, by the homomorphism property and $d[e_\tau]u = u$. At the identity the derivative therefore has the same value after the entire gauged germ is restricted along $[e_\tau]$ in that copy. This applies at every order. On this restriction, idempotency gives

$$F_b([e_\tau]x) \exp(M_0 \mathcal{H}([e_\tau]x)) = \chi_{\tau,i}([e_\tau]x).$$

Each remaining invariant derivative on a character has eigenvalue $d\chi_{\tau,i}(u)$. The value of $\mathcal{P}(D)$ is exactly the product of leading character minors to the indicated row-gap powers. It is nonzero and has valuation at most

$$\sum_\tau \ell_\tau a_{\tau,1} \leq f Q_j (1 + \delta_p/2) \max_{\tau \in \Phi} a_{\tau,1} + Q_j \sum_{\tau \notin \Phi} a_{\tau,1}. \tag{147}$$

An empty row contributes zero. The construction uses one copy per typed column and the original connection at M_0. The harmonic value from typed character columns.

Descent of the bound to the given test. The polynomial $\mathcal{P}$ is integral over the original trait. Its differential evaluation with the given gauge is linear on the degree-D_L jet of the germ. Every value on an integral jet over R has valuation at least the least valuation of the values on monomial jets in the original integral chart. This follows by expanding in the finite monomial basis and using the valuation inequality for sums. Hence the infimum over integral germs in the original chart is bounded above by (147), normalized at $v(p) = 1$. The relative valuation used in the unscaled depth is $e_a v$ since $\operatorname{ord}_V(p) = e' e_a$. Taking the supremum over the tests, taking the maximum with zero, and dividing by e_a gives (147) for α_a. The valued-extension comparison.

Use the selected $\kappa = T w_L + w_F$ on this same record. Multiplying the entries of w_L by T gives bounds by the entries of κ because the rows of w_F are nonnegative. In the κ-row at $\tau \in \Phi$, the first $k-1$ entries are at most $2b + S_c$ and the remaining j are at most η. All Q_j entries of its conjugate row are at most η. These bounds give

$$\begin{gathered} T a_{\tau,1} \le 2b + S_c \quad (\tau \in \Phi), \qquad T a_{\tau,1} \le \eta \quad (\tau \notin \Phi), \\ T D_L \le D_K \le f\big((k-1)(2b + S_c) + m\eta\big) \le k s_{B,j}, \qquad s_{B,j} = \mathrm{wt}_j + f(S_c + \eta), \quad \mathrm{wt}_j = 2fb. \end{gathered}$$

Inserting the first-row inequalities in (147) and using $Q_j \le m$ and $\delta_p \le \log_2 Q_S$ yields

$$T\alpha_a \le f Q_j\big((1 + \delta_p/2)(2b + S_c) + \eta\big) \le s_{C,j}, \qquad s_{C,j} = Cm(2 + \log_2 Q_S) s_{B,j}. \tag{148}$$

Here $C \ge 1$ is the fixed constant of the weight selection. The first-row allowances of the actual short weight and The selected total degree bound.

A principal-symbol test at q and d. The same marked flags and the same polynomial $\mathcal{P}$ are used at the two residual host primes. For each assigned copy $b(\tau, i)$ take an integral linear form $L_{\tau,i}$ that is dual to $u_{\tau,i}$ in the τ-flag. Use the integral germ

$$\prod_{\tau,i} L_{\tau,i}(X_{b(\tau,i)})^{a_{\tau,i}}.$$

It is homogeneous of degree D_L. The harmonic operator applied to it at the identity only uses the order-D_L principal symbol and the constant term one of the gauge exponential. Within a τ-block the symbol evaluates only the diagonal monomial of the leading minors. Its coefficient is one, so the value is $\prod_{\tau,i} a_{\tau,i}!$. This uses the integral dual for each type separately and applies even if different types have the same residue label. The depth in units with $v(p) = 1$ is therefore at most $D_L/(p-1)$. The host sizes $p > 4t_{\mathrm{cut}} r + 1$ at $p = q, d$ imply $k \le p - 1$. It follows from the total degree bound that

$$T\alpha_a \le s_{B,j} \quad (p = q, d). \tag{149}$$

The integral diagonal symbol calculation and The size check at both residual host primes.

Insertion into the paired local bounds. For $p \mid Q_S$ set $c_p(D) = D \operatorname{ord}_p Q_S$. For $p = q, d$ set $c_p(D) = \lceil D/2 \rceil$. The three types of bad prime under consideration are disjoint from the type primes g_0, g_1. The factors from the CM norms and the level in the common scalar are therefore units at these primes, leaving

$$\frac{\operatorname{ord}_v \sigma}{e_v} = T c_p(D_L) + c_p(D_F)$$

on each selected run. This uses the scalar of its actual weights, whose full factor has $8fr + 4$ copies. At such a prime the ceiling in the marked coefficient field is $E_a = \lceil e_a T(\alpha_a + c_p(D_L) + 1) \rceil$. Its extension to the period coefficient field is $E = e_{v/a} E_a$ with $e_v = e_{v/a} e_a$. Consequently the normalized difference is at most

$$\frac{E - \operatorname{ord}_v \sigma}{e_v} \le T\alpha_a + T + 1/e_a - c_p(D_F) \le T\alpha_a + T + 1.$$

The bounds (148) and (149) are uniform at all coefficient primes above the rational prime. Apply the same reasoning at the place below $\bar{v}$. The chosen own-prime rule and rational-prime classification apply there as well. This gives the two $s_{C,j}+T+1$ bounds on $p \mid Q_S$ and the two $s_{B,j}+T+1$ bounds at q, d. The reference value is zero at q and at $p \mid Q_S$. At d the reference has the form $-\frac{k-1}{2}\sum_{t\in A}\min(w_j(0,t), w_j(1,t))$. Since $|A| = f/2$ and each minimum lies in $[0, 2b']$, this lies between $-fkb'/2$ and zero, giving the prescribed bound $-2fkb' \le F_j^0(d) \le 0$. The paired small-prime rounding and The paired bound at q, d. The residual-prime reference.

The marked integral chart at an ordinary prime. Write $n = k-1$, $j' = m-j$, and $r = n+j+j'$. Let p be an ordinary prime of the chosen run. The level scalar G of that run is the positive integer $g_0 g_1$, a p-unit, and E is unramified at p. At an algebraic framed point, pass to a finite extension of its trait on which the semistable torus splits and use the formal completion of the Raynaud extension

$$0 \longrightarrow T_{\mathrm{tor}} \longrightarrow \mathcal{G} \longrightarrow B \longrightarrow 0.$$

The polarized Raynaud construction identifies the pointed completion with the completion of this extension, pulls the rigidified biextension from the principally polarized abelian part B, and transports the $\mathcal{O}_E$-action. Use the chart coordinates to pull back this pointed identification. Each letter of the row operator is the invariant vector with initial value given by the corresponding framed direction. It therefore goes to the invariant vector of the transported direction. These identifications apply also when the torus rank is zero. SGA 7 I, Exposé IX, §7. Pass further to the completion of an algebraic closure C, with valuation $v(p) = 1$, ring R, and residue field κ. The $\mathcal{O}_E$-action on the formal group extends to $\mathcal{O}_E \otimes \mathbf{Z}_p$: the iterates of $[p](Z) = pZ + O(Z^2)$ tend to zero on every finite nilpotent jet over R/p^e, for every e. Thus local-factor idempotents split the group. On Lie algebras the unramified embeddings split each factor integrally, because distinct embeddings are separated by units. In each type the marked frame is an integral basis of its typed tangent summand; its leading directions span a direct summand.

Let $h(X, Y)$ be the additive gauge of the rigidified polarized biextension, zero on the two axes. In the two addition laws, the symmetry law, and the CM adjoint law, the defects of h are logarithms of integral units with constant term one. The CM adjoint of a is $\bar{a}$. These laws can be used also for the completed coefficient action, by taking coefficientwise limits on each finite jet. In formal logarithm coordinates D_u for a marked tangent vector u becomes the constant directional derivative. The evaluated harmonic operator is

$$P(D)\left[f \exp\left(G\sum_b h(X_b, w_0 X_b)\right)\right]_{X=0}, \qquad w_0 + \overline{w}_0 = 1,$$

where f is an integral germ in the copies of the chart. Multiplication by an integral unit in the gauge can therefore be absorbed by using its inverse in f. If two additive gauges have the same integral multiplication laws, after the unit correction their difference is bilinear in formal logarithm coordinates. The diagonal operator associated to this bilinear difference acts on P by conjugate-type contractions. A harmonic P is killed by every such contraction, so the value is unchanged.

The two orbit arrangements. For the alternating inert and signal classes the host gives $p \ge 5$. In the distinct-factor arrangement label an orbit by (g, s), $g \in \{0, 1\}$, $s \in \mathbf{Z}/D\mathbf{Z}$, with D even, conjugation $(g, s) \mapsto (1-g, s)$, and residue Frobenius $(g, s) \mapsto (g, s+1)$. The large Lie rank $n+j$ is on $g \equiv s \bmod 2$, and the other rank is j'. On one physical copy the two orthogonal idempotent factors give a pair of variables X, Y. The sum law, adjoint law, and symmetry replace the gauge of this pair by $h(X, Y)$ modulo the logarithm of an integral unit. Indeed the two self terms are orthogonal and the two cross terms combine by $w_0 + \overline{w}_0 = 1$.

In the self-conjugate arrangement the real orbit has odd length d, and the local CM orbit has length $D = 2d$. Index the full orbit by τ_s with $\bar{\tau}_s = \tau_{s+d}$ and large rank on the even indices. On two physical

copies of the local factor choose u with $u\bar{u} = -1$. The norm on the quadratic residue-field extension is surjective, and a Teichmüller lift of a residue solution has exactly this norm. Put

$$M = \begin{pmatrix} 1 & 1/2 \\ u & -u/2 \end{pmatrix}, \qquad M^*M = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad M^* = \overline{M}^{\,t}.$$

This is an integral automorphism of the formal group on the two copies. After M, label type τ_s by s on side zero and by $s + d$ on side one. The displayed cross matrix gives precisely the same cross gauge modulo an integral unit, because distributing the two diagonal gauges gives $\sum_{g,g'} h(X^g, (M^*M)_{gg'} w_0 X^{g'})$ and the two remaining terms combine by symmetry and the adjoint law. The same matrix transforms the harmonic contractions. A block joining (g, s) to $(1 - g, s + i)$, i odd and $0 < i < D$, has no contracted pair of variables. Distinct blocks use disjoint copies. Projection modulo d gives the real target $s + i$, and, if $d = 1$, the block at $i = 1$ uses large directions at the same embedding on the two virtual sides. As the matrix acts only on the copies, it commutes with every typed change of marked rows.

Polarization and first Cartier coefficients. We use the cross chart on the two sides, denoting the base group law on each side additively. Let $B_2(Z; X, Y)$ be the integral multiplication unit whose logarithm is $h(Z, X + Y) - h(Z, X) - h(Z, Y)$. Suppose an α_p in the residue formal group on the first side has tangent e. If the first variation $c(X, Y) = D_e^{(1)} B_2(0; X, Y)$ is an additive coboundary on the opposite residue formal group, this α_p is zero. To see this, use the p-Weil pairing defined by the polarized biextension, evaluated on εe, $\varepsilon^2 = 0$, and formal p-torsion Y. In the first argument every unit in the p-fold multiplication has value one, since both summands of the first argument are multiples of ε and the unit is rigidified in both. In the second argument the product is

$$1 + \varepsilon \sum_{a=0}^{p-1} c(aY, Y) = 1$$

by the coboundary equation and $pY = 0$. This argument is on torsion schemes, so it applies also with nilpotent test coefficients. The formal p-torsion of B lifts to that of the Raynaud extension faithfully flat locally: lift through the smooth split torus and correct the resulting p-multiple by adjoining torus p-th roots, which reduce to one at the identity. The étale part of the residue $B[p]$ pairs trivially with an α_p, as $\mathrm{Hom}(\alpha_p, \mathbf{G}_m) = 0$. Perfection of the polarized pairing on $B[p]$ then forces the image of this α_p in B to be zero. The same vanishing of characters rules out an α_p in the split torus. For that vanishing, differentiate $F(x + y) = F(x)F(y)$ on $\kappa[x, y]/(x^p, y^p)$: its coefficients through x^{p-1} are $b^q/q!$ if the linear coefficient is b. The coefficient of $x^{p-1}y$ gives $b^p = 0$. The argument on the hyperbolic product uses the same perfect pairing, transported by the integral matrix M.

For a tangent vector v_0 on the opposite side set $s_{v_0}(X) = -D_{v_0}^{(2)} h(X, 0)$. Its addition defect is integral, by differentiating the first multiplication unit on the second axis. The second multiplication law also gives

$$D_{v_0}^{(2)} h(X, Y) - D_{v_0}^{(2)} h(X, 0) \in R[[X, Y]]. \tag{A}$$

Here and below derivatives without evaluation coordinates are invariant derivatives. The mixed linear coefficient of h between types other than conjugate partners is integral: the adjoint defect multiplies it by $\tau(a) - \bar{\tau}'(a)$, a unit for a suitable a. On the conjugate partners we may subtract the adjoint-compatible bilinear form with those mixed coefficients in logarithm coordinates. In the cross gauge symmetry supplies the transpose on reversing the sides. This changes no harmonic evaluation. After this normalization $d s_{v_0}(0)$ is integral.

Use an additive central coordinate Z_{v_0} for each member of a basis of opposite tangent vectors. There is an integral commutative formal extension with logarithm

$$(\log X,\ (Z_{v_0} - s_{v_0}(X))_{v_0}). \tag{B}$$

Indeed its multiplication defect in Z_{v_0} is the integral addition defect of s_{v_0}; the CM action on these logarithm coordinates is the action on $\log X$ and the dual of the conjugate action on the Z-space. In the original coordinates the extra central term for a is $s_{v_0}(aX) - s_{\bar{a}v_0}(X)$, integral by the differentiated adjoint law. The differential of (B) has an integral inverse, so a tangent vector can be prescribed with zero central logarithm component.

An integral curve with this tangent can be chosen with logarithm $\sum_{\ell\geq 0}(A_\ell, B_\ell)t^{p^\ell}$. For clarity, p-typify any integral curve by subtracting its averages over ℓ-th roots of unity for primes $\ell \neq p$, using formal sums and division by the unit ℓ. These commuting projections stabilize at every degree in logarithm coordinates. For a tangent of type a in an unramified factor, average also the substitutions $[u]^{-1}\gamma(\tau_a(u)t)$ over its Teichmüller group. The coefficient at t^{p^h} then has type $a + h$, in residue Frobenius order. Components in the central space of that type pair with the opposite tangent of conjugate type. Use the same definition of typical curves on both sides.

The Frobenius operation on curves takes the formal sum over p-th roots of unity and replaces t^p by t. It is integral: in integral coordinates the sum is invariant under those substitutions, so all its powers are multiples of p. In logarithm coordinates it sends the coefficient sequence to $p(A_{h+1}, B_{h+1})_{h\geq 0}$. Verschiebung is substitution of t^p; its inverse on curves with zero first logarithm coefficient is again integral. Iterating Frobenius shows $v(A_h), v(B_h) \geq -h$, where the valuation of a vector is the minimum over an integral basis. Call a curve good if $A_1 = 0$ and $B_0 = 0$. For such a curve, subtract the integral central linear curve $pB_1 t$ from its Frobenius and divide by Verschiebung. Applying the same bounds to the result gives

$$v(A_h), v(B_h) \geq 1 - h \qquad (h \geq 2). \tag{C}$$

For a good curve $\gamma(t)$, let $\Gamma(t) = \sum_{\zeta^p=1}^{\mathcal{G}} \gamma(\zeta t)$. Its base logarithm starts with $pA_2 t^{p^2}$. The addition identity in the first argument gives

$$\sum_{\zeta^p=1} h(\gamma(\zeta t), X) - h(\Gamma(t), X) = \log U(t, X), \qquad U \in 1 + t^p R[[t, X]]. \tag{D}$$

In fact the left side has no lower powers of t, and U is integral as a product of multiplication units. In particular $b(X) = p[t^p]h(\gamma(t), X)$ is integral. By (A), every invariant derivative of b, modulo the maximal ideal, equals its value at zero. Modulo that ideal these constants are the components of pB_1, by (B) and integrality of the central coordinate. If pB_1 vanishes in the residue field, then $\bar{b}$ is a p-th power. The second multiplication identity gives its addition defect modulo the maximal ideal as

$$\overline{p[t^p] \log B_2(\gamma(t); X, Y)} = \bar{c}(X, Y)^p.$$

Only the p-th term in the logarithm survives here. Taking Frobenius roots in the perfect residue field shows that $\bar{c}$ is a coboundary. The curve γ reduced modulo t^p defines an α_p: compare the two addition series in logarithm coordinates, where the first nonlinear term has degree at least p^2, and then reduce this integral identity in bidegrees $< p, < p$. The preceding pairing argument now gives

$$pB_1 \equiv 0 \quad \Longrightarrow \quad A_0 \equiv 0. \tag{E}$$

Congruences in (E) and the next argument are in κ.

There is a saturated family of good tangents of codimension at most j' in any given integral plane in a large type. Begin with typed typical curves for a basis of that plane. The A_1 coefficients land in the next, small type. Choose p-th roots of the entries of the matrix of pA_1 and take the saturated kernel of that matrix of roots over R. Such a kernel has an integral basis extendible to a basis of the given plane: row reduction over a valuation ring uses the smallest nonzero valuation as pivot in each row or column elimination. For a kernel basis vector with coefficients c_b, use the formal sum of the curves $\gamma_b(c_b t)$. Its pA_1 is in pR by the multinomial formula. Subtract the Verschiebung of a typed integral typical curve with linear coefficient equal to this integral A_1, to make A_1 zero. The first tangent and $B_0 = 0$ are preserved.

On the residue span of any saturated family of good tangents of large type a, the semilinear map given by pA_2 has kernel of dimension at most j'. Here semilinearity uses the p^2-th power on scalars, from substitution in the curves. To prove the assertion, take another good curve $\gamma'(s)$ on the opposite side of large type conjugate to $a+1$, and put $H(t,s) = h(\gamma(t), \gamma'(s))$. In (D) the unit with $X = \gamma'(s)$ now satisfies $U - 1 \in t^p sR[[t,s]]$. The coefficient of $t^{p^2} s^p$ in $p \log U$ gives

$$(p[t^p s]H)^p \equiv p^2[t^{p^2} s^p]H - p[t^{p^2} s^p]h(\Gamma(t), \gamma'(s)). \tag{F}$$

Again only the p-th logarithm term survives at that bidegree. Applying (D) in the s-variable shows $p[t^{p^2} s^p]H \in R$, so the first term on the right of (F) is zero in κ. The last term pairs pA_2 with the integral covector $p[s^p]D^{(1)}h(0, \gamma'(s))$. Its integrality follows from (D) with the sides exchanged and from the differentiated symmetry unit. Thus $pA_2 \equiv 0$ forces pB_1 to annihilate the good tangents in the opposite large type. These tangents have codimension at most j' by the preceding construction. By (E), pB_1 is injective on the kernel under consideration, with scalar action by the p-th power. This proves the rank bound, since κ is perfect. Both rank arguments apply to the transported hyperbolic polarization and to the distinct-factor chart.

AI-proposed candidate (the marked Frobenius pairing). Let R be the valuation ring of the completed algebraic closure of the marked trait, with $v(p) = 1$ and $p \geq 5$. Use its polarized Raynaud completion and its unramified CM action, with large and small Lie ranks $n + j$ and j' alternating in Frobenius order in each side of the paired chart. The two sides come from distinct conjugate factors or from the integral two-copy change of the self-conjugate factor. Let D be the full even CM period. Take i odd, $1 \leq i < D$, a saturated marked n-plane in a large type a on the first side, and a saturated marked n-plane in the type conjugate to $a + i$ on the second side. Let G be a p-unit, taken to be $g_0 g_1$ in the application. Put $q = p^i$ and

$$\rho = \frac{1}{p^2 - 1}, \quad \epsilon = \frac{1}{p-1}, \quad \alpha = \rho + \frac{3}{8q}, \qquad \beta = \rho + \frac{1}{2} - \frac{1}{8q}, \qquad E_i = 1 + \frac{q-p}{p^2-1}.$$

There are saturated subplanes of common rank $n_0 \geq \max(0, n - ij' - j)$ in the two given planes. In their formal logarithm coordinates x, y, with integral bases $(u_c), (v_b)$, the cross gauge $\widetilde{h} = Gh$ has the form

$$\widetilde{h} = -p^{-E_i} \sum_{c=1}^{n_0} x_c^q y_c + H.$$

The expansion uses the normalized gauge of the chosen rigidification. The series H vanishes on $y = 0$. Every coefficient $H_{\mu,\nu}$ satisfies

$$v(H_{\mu,\nu}) + \alpha|\mu| + \beta|\nu| \geq \epsilon.$$

Each first y-derivative of H has nonnegative coefficient weight for (α, ϵ). At $i = 1$ the remainder is integral in bidegrees of left degree $< p^2$ and right degree $< p$. The series converges at the two closed radii $p^{-\alpha}$ and $p^{-\beta}$.

On these selected planes the exact bidegree $(q, 1)$ of $p^{E_i}\widetilde{h}$ lies over R and reduces to the perfect Frobenius-semilinear pairing $-\sum_c x_c^q y_c$. This bidegree is unchanged by integral rigidification units and conjugate-type bilinear shifts. The coefficient uses the formal logarithm of the chart, which is transported by the tangent map of any pointed chart change. In particular, if $\mathfrak{h}_{\mathrm{act}}$ is the actual scaled diagonal gauge on the one or two physical copies of a logical row, then for the transported directions

$$\frac{p^{E_i}}{q!} D_{u_c}^q D_{v_b} \mathfrak{h}_{\mathrm{act}}\big|_0 \equiv -\delta_{bc} \pmod{\mathfrak{m}_R}.$$

These mixed jets are defined by the invariant directional derivatives of the actual chart. For $b = c$ their unscaled valuation is $(q-1)/(p-1) - E_i$. For example, at $i = 1$ the actual mixed derivative of order $(p, 1)$ is a unit in every surviving matched pair. At the terminal stage $j = m$, where $j' = 0$ and $n > m$, the rank bound guarantees such pairs on an alternating chart.

The release's normal-form statement uses $n_0 \ge n - (i+1)j' - j$, and the prime-table estimate uses the coarser row loss $(i+1)m$. The proof below gives $n_0 \ge n - ij' - j$ by using $K = (i+1)/2$ backward restrictions and $K-1$ forward steps. At its final level the A_2 term already belongs to the weighted error, so only the first central coefficient is needed there.

Proof of the marked normal form. The curve and polarization relations are the ones developed in the theory of Cartier curves and of polarized Raynaud extensions. The display pairing construction gives the residue identities [131, Propositions 15 and 90 and §4.3]. The prescribed-plane rank bound, the logarithmic coefficient and the two weighted radii are the quantitative assertions here.

Fix an odd $0 < i < D$, a large type a on one side, and the paired large type of the opposite side conjugate to $a + i$. Write $K = (i+1)/2$, $Q = p^2$, $\rho = (p^2-1)^{-1}$, $\gamma_p = p\rho$, and

$$\alpha = \rho + \frac{3}{8p^i}, \qquad \beta = \rho + \frac{1}{2} - \frac{1}{8p^i}, \qquad e_s = (Q^s - 1)/(Q-1).$$

On the first side use the actual marked leading n-plane at type a, and then the full tangent spaces at the distinct types $a + 2s$, $1 \le s < K$. Proceed backwards starting at $s = K-1$. At level s first choose a good family of codimension at most j' by the preceding construction. If $s < K-1$, look at the matrix of the pA_2 coefficients from that family into the complement of the already retained family at level $s+1$. Take the saturated kernel of the matrix of its p^2-th roots. Formal sums of the corresponding rescaled good curves remain good. Their pA_2 has complement coordinates in pR, by the multinomial formula. Thus each retained family has codimension at most $(K-s)j'$ at its level, and for its curves $A_2 = p^{-1}g + \lambda$ with g in the retained integral span at the next level and λ integral. Make the Verschiebung correction to achieve $A_1 = 0$ before imposing the condition on A_2.

Iterate the residue maps pA_2 through these retained spaces. By (F) each map drops rank by at most j'. At the final type, (E) makes the residue map pB_1 injective on the image. Restrict its covectors to the marked leading n-plane on the opposite side, which drops rank by at most j. Hence there are

$$n_0 \ge \max\{0, n - ij' - j\} \tag{G}$$

independent surviving directions in the first retained family. Take them from the original retained basis, extend to a basis of the actual leading n-plane, and make an integral change of basis within the opposite marked n-plane to normalize the surviving pairing. The residue matrix on these directions has a unit n_0-minor, so the normalization may be made over R, with an invertible change-of-basis matrix. Each change replaces the leading determinant by a unit multiple of itself. In particular it preserves the specified leading plane of w_L.

We verify the error estimates in the formal logarithm variables of these two planes. Complete the base tangents of the retained curves to a typed integral basis and choose integral typed typical curves for the additional tangents. The formal sum of all the base curves, with independent one-variable parameters t_b, is an integral coordinate system by its invertible linear term. In logarithm coordinates x on the first side, its inverse solves

$$x_b = t_b + \sum_{c,\, h \ge 1} A_{c,h,b}\, t_c^{p^h}. \tag{H}$$

Restrict x to the n_0 directions just selected. The weighted order of a monomial is its coefficient valuation plus α times its degree. At level $s < K$, write x_l for the selected input coordinates and let $P_{s,bl}$ be the integral matrix of the iterated g-maps with p^2-power action on coefficients, starting at the coordinate inclusion P_0 and using the g's chosen above. The solution of (H) is

$$t_b = (-1)^s p^{-e_s} \sum_l P_{s,bl} x_l^{Q^s} + O_\alpha(1+\rho) \tag{I}$$

in the retained slots of level s. In all other slots $t_b = O_\alpha(1+\rho)$. Here $O_\alpha(u)$ means that every monomial has weighted order at least u.

Here is a direct check of (I). Iterate (H) in fixed-point form by degree, starting with the displayed pure terms. Their weighted orders are at least $w_s = \rho + 3Q^s/(8p^i)$. On a good input, terms with $h \ge 3$ in (H) have weighted order at least $1 - h + p^h\rho \ge 1 + \rho$ by (C), and the integral part of the $h = 2$ term has order at least $Qw_s \ge 1 + \rho$. On all other inputs, the terms have order at least $-h + p^h(1+\rho) \ge 1 + \rho$. The remaining terms are $p^{-1} g\, t_b^Q$. The mixed multinomial coefficients are divisible by p, and replacing an input by its error also has order at least $1 + \rho$ after division by p. The pure powers have the matrix P_{s+1}, using $e_{s+1} = 1 + Qe_s$ and odd Q. At the final level their order is at least $\rho + 3p/8 \ge 1 + \rho$. The invariant is preserved under each iteration and hence in the degreewise limit.

In the integral t-coordinates, equation (B) writes $D^{(2)}_{v_0} h(X, 0)$ as the sum of the central-log terms $B_{c,h,v_0} t_c^{p^h}$ and an integral germ. Indeed the sum of the lifted curves is an integral section above X in these coordinates, and its central coordinate Z_{v_0} is integral. By (A) one may further replace $D^{(2)}_{v_0} h(X, 0)$ by $D^{(2)}_{v_0} h(X, Y)$, at the cost of another integral germ. On the retained inputs, the terms $h \ge 2$ have nonnegative weight by (C) and (I). On the other inputs this follows from $v(B_h) \ge -h$ and (I). The term $h = 1$ can pair with the chosen opposite type only at $s = K - 1$. Its pure contribution is

$$p^{-E_i} \sum_l \Phi_{v_0 l} x_l^{p^i}, \qquad E_i = 1 + pe_{K-1} = 1 + \frac{p^i - p}{p^2 - 1}.$$

Replacing an input by its error in (I) has nonnegative order here, as do all mixed multinomial terms. In the derivative of Gh, the matrix Φ is the final pB_1 pairing of the iterated matrices P_s, times a sign and G. Its indicated residue minor is a unit by (E)–(G). Normalize it to $-I$ as in (G).

On the second side put $\epsilon = 1/(p-1)$. The usual inverse logarithm in integral typical coordinates has weighted order $\ge \epsilon$ on the ϵ-disc because every nonlinear term of the fixed-point calculation has order $\ge -h + p^h\epsilon \ge \epsilon$. Consequently, in the pair of selected logarithm planes, the phase itself has the form

$$Gh(x, y) = -p^{-E_i} \sum_{l=1}^{n_0} x_l^{p^i} y_l + H_1(x, y). \tag{J}$$

Every y-derivative of H_1 has nonnegative order for the weights (α, ϵ). The phase and the displayed main term vanish for $y = 0$. For a coefficient $c_{\mu,\nu}$ of the remainder and $\nu_l > 0$, integration of its y_l-derivative gives

$$v(c_{\mu,\nu}) + \alpha|\mu| + \beta|\nu| \;\ge\; \epsilon + (\beta - \epsilon)|\nu| - v_p(\nu_l) \;\ge\; \epsilon.$$

The last step uses $v_p(b) \le b/p$ for $b > 0$ and $\beta - \epsilon \ge 1/p$ for $p \ge 5$. This proves both coefficient conditions in the short-window jet theorem. At $i = 1$, the inverse in (H) is linear modulo degree p^2 on the first plane and linear modulo degree p on the second. Thus the remainder is also integral in bidegrees with left degree $< p^2$ and right degree $< p$, since integration in these right degrees divides by a unit.

On the selected planes, the exact bidegree $(p^i, 1)$ of $p^{E_i} Gh$ is integral and has residue $-\sum_l x_l^{p^i} y_l$. Indeed every remainder coefficient of that bidegree has valuation at least $E_i - p^i\alpha = 5/8 - \gamma_p > 0$ after scaling, by the derivative estimate. This defines a Frobenius-semilinear pairing in the formal logarithm planes. An integral change of pointed chart carries its formal logarithm by the induced linear tangent map. An integral rigidification unit also has logarithmic derivative of nonnegative order in the first-disc weight, while a bilinear gauge change contributes only in bidegree $(1, 1)$. Thus the scaled bidegree on these planes is independent of these choices. The proof works also with the smaller weights $\rho + 1/(4p^i)$ and $\epsilon + 1/p$. The final fixed-point bound needs only $p/4 \ge 1$, and the integration bound uses the difference $1/p$. This strict slack gives convergence at the closed radii $p^{-\alpha}$ and $p^{-\beta}$.

The test of a prescribed block of w_L. Use n_0 disjoint rows of the two-side chart from (J). On these rows apply the short-window determinant jet theorem to the phase $\sum_\rho Gh(x_\rho, y_\rho)$. For an integer $y \ge 0$

it evaluates the leading determinants to the powers $p^i y$ and y on the exponential of this phase. For $n_0 > 0$ every forward interval of

$$S(p,i,n) = \left\lceil 64n^2\big(i+3+3\log_p(n+1)\big)\right\rceil$$

consecutive integers contains an admissible y. One may always use zero. The normalized valuation of the nonzero evaluation at an admissible y is at most $n_0 y((p^i+1)\gamma_p - 1) + n(\log_p n + 1)$. If $n_0 = 0$, use the empty test for every y.

For each of the remaining $n - n_0$ rows use a germ whose principal monomial is the specified diagonal coordinate to the power $p^i y$ on the first side and the specified diagonal coordinate to the power y on the second. Use integral coordinates dual to the chosen marked vectors. The degree in each physical row forces every derivative in that row to hit this principal monomial. Thus only the surviving n_0-square determinant is left on the other rows, and the extra cost is at most $(n - n_0)(p^i + 1)y/(p - 1)$ by the factorial bound. Pull back the product of these germs, using the integral inverse of M in the self-conjugate arrangement and the inverse of the integral unit from the gauge comparison. This is an integral germ for the actual marked chart. The resulting nonzero value has normalized valuation at most

$$ny\big((p^i+1)\gamma_p - 1\big) + \frac{(i+1)m(p^i+1)y}{p-1} + n(\log_p n + 1). \tag{K}$$

Here $n - n_0 \le 2Kj' + j \le (i+1)m$ and the leading coefficient in (K) is nonnegative. The differential polynomial is the product of the two specified leading minors, on the typed columns determined by the two-side chart. It remains harmonic on the corresponding physical copies. This realizes the block on the chart of w_L.

Character and principal-symbol tests. For a leading $n \times n$ minor on one large type use n characters, one on each of n disjoint copies. Their determinant of invariant differentials, restricted to the marked n-plane, can be chosen of valuation at most the typed colength of the character differentials: expand a full-size minor of minimum valuation along the marked rows. The value of the determinant power u is then the u-th power of the minor of these character differentials, because the characters are invariant-derivative eigenfunctions. All characters have value one at the identity. On a single side the gauge is an integral unit on restriction to that side, and its inverse is absorbed in the test. In the self-conjugate case this uses the two-copy change M.

For the unramified character estimate use the Cartier-dual Tate module of the Raynaud p-divisible group. A compatible sequence of dual torsion points defines a formal character: modulo any finite nilpotent jet and p^e, some power of $[p]$ kills the input, so evaluation takes place in a finite torsion group and is independent of further choices. The exact dual torsion sequence and surjectivity of the transition maps on abelian dual torsion lift every compatible sequence on the Cartier dual of the torus torsion. The restricted characters then span the dual torus Lie lattice, so their colength on each type is bounded by the colength of the abelian Hodge–Tate map. In the latter comparison, if D_s is the determinant valuation of the Tate inclusion on embedding s, the unramified Frobenius relation is

$$pD_{s-1} - D_s = \dim \mathrm{Lie}(B^\vee)_s.$$

One obtains this by taking determinants in $\beta_s X_s = \varphi(X_{s-1})$ and reducing modulo p: the cokernel of β_s is the indicated conjugate Lie quotient over A_{inf}/ξ, and $v^\flat(\xi \bmod p) = 1$. The Tate inclusion lies between tM and M, with $t = \varphi^{-1}([\epsilon_{\mathrm{cyc}}] - 1)$; this follows by evaluating the integral duality against the dual Tate frame, whose pairing on Tate frames is t times a unit matrix. Thus the Smith valuations after either specialization, $A_{\mathrm{inf}}/(\xi)$ or $A_{\mathrm{inf}}/(p)$, are at most $1/(p-1) < 1$ and agree by reduction to the common quotient. The Hodge target is a quotient. Howard, Theorem 2.2.2. Iterating the relation gives $D_s = \sum_{h \ge 1} p^{-h} \dim \mathrm{Lie}(B^\vee)_{s+h}$. Hence on the large types the cost of the character test is at most $uj'/(p-1)$ in the split arrangement and

$$u\left(\frac{(n+j)p + j'}{p^2 - 1}\right) \le u\left(n\gamma_p + \frac{m}{p-1}\right) \tag{L}$$

on an alternating orbit, including the full self-conjugate CM orbit.

For any other flag gap use its leading minor on separate columns, with the integral germ cut out by the integral dual diagonal coordinates to those powers. One set of columns for each type suffices for all these minors: the product of their diagonal monomials has coefficient one, and every other term includes an off-diagonal variable. The germ has the full order of the operator. Only its principal symbol contributes and the gauge has constant coefficient one. The value is a product of factorials, of valuation at most the total degree divided by $p-1$. Use disjoint columns for conjugate types.

The capacities and the signed plan. Write $a_{\tau,b}$ for the rows of w_L, with a terminal zero, and use its actual flag gaps $\omega_{\tau,h} = a_{\tau,h-1} - a_{\tau,h}$. The core gap at the large type above a real label s is $W_s = \omega_{\tau,n}$; write $w_s = w_j(z_s)$. The choice of the short weight gives a common $0 \le \theta_0 \le 1$, with $n\theta_0 \le 5m$, and, writing $S = S_c + \eta$,

$$(1-\theta_0)w_s \le TW_s \le w_s + C_1 S, \qquad T\,\mathrm{rest} \le C_1 fkS, \tag{M}$$

where rest is the total degree of all gaps off the core. The principal-symbol tests cover this remaining degree at normalized cost at most rest $/(p-1)$.

For an alternating real orbit of length $h \ge 2$, let I_0 be the least odd integer ≥ 3 with $p^{I_0} > k^{3/2}$. Start with a maximizing nonnegative real plan with capacities $W'_s = (1-\theta_0)w_s/T$ and constraints

$$\sum_{\substack{0<i<h \\ i \text{ odd}}} (p^i x_{s,i} + x_{s-i,i}) \le W'_s. \tag{N}$$

Discard the edges $i \ge I_0$; their total contribution to the maximized sum is at most $p^{-I_0}\sum_s W'_s$. For each remaining edge choose the largest admissible integer $y_{s,i} \le x_{s,i}$ from the chart block test. The deficit is at most $S(p,i,n)+1$, by the forward windows and the admissibility of zero. Allow for signed thresholds by using the actual integer remainders

$$u_s = W_s - \sum_i (p^i y_{s,i} + y_{s-i,i}) \ge 0.$$

Each outgoing edge spends $p^i y_{s,i}$ at its source, and $y_{s,i}$ at its target. Use (K) on these two specified large types and (L) on the pure u_s block. In the odd real orbit case the two virtual types over a real edge are the ones from the full $2h$-cycle, and spend exactly these same two capacities.

The terms from (K) and (L) proportional to $n\gamma_p$ sum to $n\gamma_p \sum_s W_s$, and each edge subtracts $ny_{s,i}$. In the row-loss term of (K), the actual loss is bounded by $\min(n,(i+1)m)$. Thus the remaining degree-dependent errors are bounded by a constant times

$$\frac{\min(n,(I_0+1)m)}{p-1}\sum_s W_s + \frac{\mathrm{rest}}{p-1}.$$

The bound by n here controls the excess $TW_s - w_s$ in (M), even when $mI_0 > n$. There are at most $f(I_0-1)/2$ retained edges in all orbits. The window and per-edge errors are in total $O(fn^3(\log k)^2)$ before multiplying by T, since $I_0 \le 3 + 2\log k$ for $p \ge 5$.

Let F_{orb} be the untruncated reference in the orbit, with the capacities nw_s. By homogeneity, the plan for (N) would give exactly $(1-\theta_0)F_{\mathrm{orb}}$ after multiplying by T, if one used W'_s in the leading γ_p term and discarded and rounded no edges. The replacement of W'_s by W_s costs at most

$$n\gamma_p \sum_s (\theta_0 w_s + C_1 S),$$

and the discarded-edge loss is at most $np^{-I_0}\sum_s w_s$. The reference is nonnegative because summing the constraints bounds the maximizing sum by $\sum_s nw_s/(p+1)$. Thus we may bound $(1-\theta_0)F_{\mathrm{orb}}$ by F_{orb}. Combining the errors, the total, after multiplying by T, is

$$O\left(\frac{m\log k\sum_s w_s + fkS}{p} + kp^{-I_0}\sum_s w_s + Tfn^3(\log k)^2\right). \tag{O}$$

For pure conjugation use the block at $i = 1$ on the full two-cycle and choose the largest admissible $y \le W_s/(p+1)$ at each real label. Here the unused gap is $u_s = W_s-(p+1)y$. The leading cost $n\gamma_p W_s - ny$ is at most $nW_s/(p^2-1) + n(S(p,1,n)+1)$. This gives the reference $n\sum_s w_s/(p^2-1)$ with errors bounded by (O), using (M). In the split case there are no edges: the characters alone give the core with cost at most $j'\sum_s W_s/(p-1)$, and the off-core costs are as before. The reference is zero and (O) bounds the errors in this case as well.

Copies and the large-exponent threshold. Each edge and each character remainder requires at most n logical rows. There are at most $(I_0-1)/2$ relevant odd levels per real label, and one such remainder. Doubling the rows where the integral change M is used and supplying all off-core directions costs at most

$$f\big((I_0+1)n + 4r\big)$$

physical copies. The pure-conjugation choice also fits this bound with $I_0 \ge 3$. At split primes the fn character rows and at most $2fr$ principal-symbol rows suffice. The selected budget $n_{w_L} \ge 10r(f + f^2n^3L^2/m)$ covers these numbers, since $m \le n \le r$, $f \ge 2$, $L \ge 3$, $\log k \le L$, and $I_0 \le 3 + 2\log k$. Assign disjoint physical positions to different blocks and to off-core types. The product of all tests is row-unipotent invariant, since it uses only leading minors, and harmonic, since each block is harmonic and different blocks have disjoint support in the physical copies. Its row degree at (τ, b) is $\sum_{h>b}\omega_{\tau,h} = a_{\tau,b}$. This is the required weight w_L on the given number of copies.

When $p \ge mb/(Tfn^3L)$, use principal symbols for every gap, with small exponents. The selected inequalities give $Ta_{\tau,0} \le 3b$. Take $s = \lceil 3b/(T(p-1))\rceil$, so $a_{\tau,0} \le s(p-1)$. Divide the multiset of leading minors at each type into s groups with at most $p-1$ factors per group. Give each group its own columns and diagonal principal-symbol germ. Every factorial argument is $< p$, so the evaluation is a unit. There are at most $4frs$ required copies. The threshold gives $s \le 1 + (9/2)fn^3L/m$ for $p \ge 3$, so these copies fit the same budget. Below the threshold, the term $Tfn^3(\log k)^2$ in (O) is at most $mb\log k/p$.

Descent and the bound on the chosen run. The finite free module of row-degree polynomials on the marked copies contains the integral harmonic module as the intersection of the kernels of the contraction maps and the coefficient maps for row-unipotent invariance. For the latter, expand the row shift as a polynomial in its parameter and take its nonconstant coefficients. These finite kernel conditions commute with the flat base change to R. The invariant differential evaluation with its canonical gauge is bilinear in the polynomial and the germ, and uses only the finite germ jet through the weight degree. Over R the resulting value is consequently a finite integral combination of evaluations of integral polynomials and integral germ monomials over the original trait. The valuation of at least one original evaluation is at most that of the constructed one. This gives the same upper bound for the infimum defining the depth on the original pointed chart, with valuation normalized by $v(p) = 1$.

The construction applies to every algebraic framed point. Thus after the supremum and the maximum with zero, the inequalities still bound $T\alpha_a$. In (O), $p^{-I_0} = \mathrm{th}_k(p)$ and $\sum_s w_s \le \mathrm{wt}_j$. Above X_2, the profile ratio is $< p$, and the plan using only length-one edges with $px_s + x_{s-1} = nw_s$ is nonnegative. Indeed its cyclic alternating sum can be paired in positive consecutive pairs, with a positive extra term on an odd cycle. Its sum is the largest possible sum in (N) with capacities nw_s. Thus the untruncated reference above X_2 is bounded by $2k\,\mathrm{wt}_j/p^2$, also for pure conjugation. This is exactly the loss from replacing it by its truncated value zero. We obtain, with an absolute constant enlarged for the chosen weight constant C_1,

$$T\alpha_a \le F_j^0(p) + C\Bigg(\frac{m\log k\,\mathrm{wt}_j + fk(S_c+\eta)}{p} + k\,\mathrm{wt}_j\,\mathrm{th}_k(p) + \mathbf{1}_{p>X_2}\frac{k\,\mathrm{wt}_j}{p^2}\Bigg). \tag{P}$$

When $p > Y$, one can also use the full principal symbol directly, as all factorial arguments are at most the degree of w_L, strictly less than p. Then $\alpha_a = 0$.

The two runs have the same marked chart data, and the ordinary test above uses the run whose own prime is different from p. At the selected place the scalar is a unit and the actual normalized lattice exponent is $\lceil e_a T \alpha_a \rceil / e_a$. Thus (P) gives the ordinary entry of the prime table, with rounding loss at most $1/e_a$ for $p \le Y$, and gives zero for $p > Y$.

The ceiling and the rational-prime reference. Write $D_L = \mathrm{DF}(w_L)$, $D_F = \mathrm{DF}(w_F)$, and $D_K = \mathrm{DF}(\kappa) = TD_L + D_F$. For $v \mid a \mid p$ write $e_a = \mathrm{ord}_a\, p$, $e_v = \mathrm{ord}_v\, p$, and $e(v/a) = e_v/e_a$. Let α_a denote the unscaled harmonic-germ depth of w_L, measured in units with $v(p) = 1$. This is the nonnegative supremum over the algebraic integral framed points of the infimum of the admissible harmonic evaluations on the n_{w_L} copies with the marked biextension gauge. Put

$$\mathcal{B} = \{g_0, g_1, q, d\} \cup \{p \,:\, p \mid Q_S\}, \qquad c_p(D) = \begin{cases} D \,\mathrm{ord}_p\, Q_S & p \mid Q_S, \\ \lceil D/2 \rceil & p = q, d, \\ 0 & \text{otherwise}, \end{cases}$$

$$\nu_*(a) = \begin{cases} D_L/(g_i - 1) & p = g_i, \\ \alpha_a + c_p(D_L) & p \ne g_0, g_1, \end{cases} \qquad \mathrm{nus}(a) = e_a \nu_*(a),$$

$$E_a = \left\lceil e_a T \left(\nu_*(a) + \mathbf{1}_{p \in \mathcal{B}} \right) \right\rceil, \qquad E = e(v/a) E_a.$$

In the norm factors of the common scalar σ, write $\nu_{jv} = \mathrm{ord}_v\, n(\kappa)/e_v$. The prescribed halves $\pi_i, \bar{\pi}_i$ of g_i fix these orders. The normalized order of σ is

$$\frac{\mathrm{ord}_v\, \sigma}{e_v} = \nu_{jv} + f r n_{w_F} \mathbf{1}_{p \in \{g_0, g_1\}} + T c_p(D_L) + c_p(D_F).$$

Here $n_{w_F} = 8fr + 4$. The remaining own-prime powers in the three section products are units at the eligible places. The ceiling in the marked coefficient field · Equality with the local section parameter · Normalized order of the common scalar.

The reference $F_j^0(p)$ is specified at each rational prime. At $v \mid p$ use $F_{jv}^0 = F_j^0(p)$. At either g_i it is $-D_K/2$. At d it is $-\frac{k-1}{2} \sum_{t \in A} \min(w_j(0,t), w_j(1,t))$ for the stage profile w_j on the numerical labels $\{0,1\} \times A$. At q, the primes dividing Q_S, the ordinary split primes, and the ordinary primes above X_2 it is zero. At an inert prime $p \le X_2$ acting by pure conjugation it is $(k-1) \sum_z w_j(z)/(p^2 - 1)$. At every other inert or signal prime $p \le X_2$, let z_s, s mod h, be an orbit of the Artin permutation of the numerical labels and put $D_s = (k-1) w_j(z_s)$. The orbit contributes

$$\frac{p \sum_s D_s}{p^2 - 1} - \max \left\{ \sum_{s,i} x_{s,i} \;:\; \sum_{\substack{0 < i < h \\ i \text{ odd}}} (p^i x_{s,i} + x_{s-i,i}) \le D_s \quad \text{for all } s \right\} \text{ over } x_{s,i} \ge 0,$$

Summing the constraints shows this is nonnegative. For $p > X_2$ the demand ratios are at most p. The feasible solution of $p x_s + x_{s-1} = D_s$ has nonnegative coordinates, as is seen by pairing successive terms of its alternating geometric numerator $\sum_{u=0}^{h-1} (-1/p)^u D_{s-u}$. It attains the upper bound $\sum_s D_s/(p+1)$ for the maximum. Thus truncating the ordinary reference above X_2 removes at most $2k \sum_z w_j(z)/p^2$ from the untruncated bound. At d the minima are at most $2b'$, and $|A| = f/2$, so $-2fkb' \le F_j^0(d) \le 0$. Reference at each prime · Demand ratios and the reference cutoff.

At $p = g_i$, the cyclotomic subfield gives $e_a \ge g_i - 1$. The chosen weights satisfy $TD_L \le D_K < (g_i - 2)/2$, so $T\nu_*(a) < 1/2$. Consequently

$$(E - \mathrm{ord}_v\, \sigma)/e_v \le 2T - \nu_{jv}.$$

The equality $n(\kappa)\overline{n(\kappa)} = (g_0 g_1)^{D_K}$ gives $\nu_{jv} + \nu_{j,\bar{v}} = D_K$ separately at each of the two primes. At $p = q, d$ or $p \mid Q_S$, the term $Tc_p(D_L)$ cancels between the ceiling and the scalar, leaving the upper bound $T\alpha_a + T + 1$. The marked principal-symbol test at q, d gives $T\alpha_a \le s_{B,j}$. The integral character test for the marked flag at $p \mid Q_S$ gives $T\alpha_a \le s_{C,j}$.

At the remaining primes the scalar is a unit and the difference is $\lceil e_a T\alpha_a \rceil / e_a$. On the selected stages $Y \ge D_L$, so for $p > Y$ the diagonal factorial test is a unit and the depth is zero. For $p \le Y$ put

$$a_j(p) = C\left(\frac{m \log k \,\mathrm{wt}_j + f k(S_c + \eta)}{p} + k\,\mathrm{wt}_j\,\mathrm{th}_k(p) + \mathbf{1}_{p>X_2}\frac{k\,\mathrm{wt}_j}{p^2}\right).$$

The ordinary integral harmonic tests give $T\alpha_a \le F_j^0(p) + a_j(p)$. The chart-degree term $\mathrm{ad}(p)$ of the general estimate is zero on the selected record. The normalized difference is therefore at most $F_j^0(p) + a_j(p) + 1$. All these inequalities apply also at $\bar{v}$, because the eligible run and prime classification depend on the rational prime. Type-prime weight and ramification bounds · Small-prime cancellation and rounding · The two residual host primes · Ordinary depth and large-prime vanishing · Zero chart-degree term.

The two transport indices. The geometric indices used here are those of the two runs on the common tied record, under the identification in Section 6.1. In particular the comparison uses their own-prime halves and their shared geometric coefficient field in the same row coordinates as the period maps.

Work at $v \mid p$ and a fixed stage. In the common compact-target coordinates let M be the supremum lattice, $L \subseteq M$ the scaled geometric section lattice, and L_K the short-section lattice at the same weight. The finite mirabolic group H preserves M, L. Let U be its upper unipotent subgroup, ψ the prescribed generic character, $u = |U|$, and $\delta = [H : U]$. Write V for the rational target, $n = \dim V$, V_ψ for the character eigenspace, and $Z = V_\psi \cap \mathrm{im}\, e$ for the good cut of dimension $d = d_{\mathrm{row}}$. The unnormalized trace projects into V_ψ. The type projector e commutes with the U-action.

The geometric comparison is $L \subseteq \sigma L_K$, where $\sigma \ne 0$ is the common scalar of the three scaled section products. With E the geometric ceiling exponent, choose a with $\mathrm{ord}_v(a) = E$. The section index on the chosen run gives

$$i_v(M \cap W, L \cap W) \le E \dim W + r_{vj}$$

on every rational cut W, with $r_{vj} \ge 0$. Thus $M' = M \cap a^{-1}L$ satisfies $\mathrm{length}(M/M') \le r_{vj}$. Indeed in a local Smith basis the excess is $\sum_i (b_i - E)_+$, obtained by taking the cut through the entries with $b_i > E$. This also works if the section index bounds a smaller placed lattice. These are the section-product comparison, the stable supremum lattice in the period coordinates, and the section index on each of the two actual runs. The three placed lattices and the scalar · The two ceiling indices.

Put $\mathcal{A} = \sum_{h \in U} \psi(h)^{-1}\rho(h)$. The Mackey condition for the prescribed finite-field mirabolic character gives $\mathcal{A}\rho(g)V_\psi = 0$ for $g \notin U$ and $\mathcal{A}|_{V_\psi} = u$. In fact a conjugate $g^{-1}hg \in U$ with $\psi(g^{-1}hg) \ne \psi(h)$ gives the first vanishing. Applying $\mathcal{A}\rho(g_i^{-1})$ for representatives of H/U shows that the $\rho(g_i)V_\psi$ sum directly. Let B, B' be the sums of these translates of $M \cap V_\psi$ and $M' \cap V_\psi$. The same operators on $B \cap M'$ give $u(B \cap M') \subseteq B'$, since M' is H-stable. Moreover $B/(B \cap M')$ injects into M/M'. It follows that

$$\delta\, i_v(M \cap Z, M' \cap Z) \le r_{vj} + n\,\mathrm{ord}_v\, u.$$

Here the Z-cut quotient injects into the V_ψ-cut quotient.

The trace lattice contains $u(L_K \cap V_\psi)$, and locally the off-own-prime trace comparison puts it inside the reference theta lattice Θ. Because $aM' \subseteq L \subseteq \sigma L_K$, this gives $(ua/\sigma)(M' \cap Z) \subseteq e\Theta$. For the reference lattice Θ use the fixed own-prime character δ_{p_o}. The actual tuple at v uses $\delta_{p_o}\chi_v$, with the extra signed ray character χ_v. In the common faithful coefficient reading inducing v, the image of $e\Theta$ is contained in the projected base test lattice, and the projected specialized test columns lie in the image of $A_j(t(v))$. These are the normalized base-class comparison and integral twisted-column placement on this same one-step model. The latter uses $\theta(\pi e_X x) = e\theta(x)$ on the integral source level lattice, where π projects

along the theta kernel. The test determinant gives a normalized signed index at most ϵ_{vj}. Each outer inclusion contributes a nonpositive signed index. Dividing the extended coefficient length by its relative ramification over O_v gives $i_v(e\Theta, A_j(t(v))) \le e_v\epsilon_{vj}$. The common cusp map identifies J_j^{row} with $M \cap Z$. Hence

$$i_{jv}(t(v)) \le d\big(E - \operatorname{ord}_v \sigma + \operatorname{ord}_v u\big) + \frac{r_{vj} + n \operatorname{ord}_v u}{\delta} + e_v\epsilon_{vj}.$$

The trace comparison at the chosen place · The projected class map · The simultaneous signed reading.

For the good-type share, the normalized stage action identifies the actual V_ψ with the P_0-type above and carries the central selector to e. This identification holds in the placed coordinates of each marked run. The half-level embeddings for these runs use the upper shifted rank at most twice the lower shifted rank. At the terminal stage the free action on fibres gives the regular model directly. Thus $\dim(V_\psi \cap \ker e) \le \exp(-x^{e_z/7}) \dim V_\psi$ by (114), and $(3/4)\dim V_\psi \le d$ at a common threshold. The run-level regular bounds and The exact good rank in the placed coordinates.

For the character share itself, write T_{O} for the chosen Hodge-line multiple on the marked run and put $D = f\binom{r}{2}$. At the nonterminal stage the two-level cusp rank bounds at the lower determinant shift and at the unshifted weight have common component degree factors with binomial terms $\binom{T_{\mathrm{O}}+D}{D}$ and $\binom{T_{\mathrm{O}}+2D+6}{D}$, respectively. The lower binomial divided by the upper is at least $1 - D(D+6)/T_{\mathrm{O}}$ by multiplying the D factors $(T_{\mathrm{O}} + a)/(T_{\mathrm{O}} + D + 6 + a)$. Put $y = 1/(4(fk)^3)$. The selected size of T_{O} allows the determinant shift and gives $D(D+6)/T_{\mathrm{O}} \le y$. Multiplication by the translates of the shift section followed by the ψ-average injects the lower-weight cusp space into u copies of V_ψ, by the translate determinant. Hence $(1-y)n \le u \dim V_\psi$. At the terminal stage free orbits give $u \dim V_\psi = n$. The same rational dimensions apply to the actual placed lattices because the coordinate isomorphism respects the character equations. The two binomial rank bounds, The character projection from the lower weight, The selected shift and size, and The actual placed character dimensions.

The finite mirabolic index satisfies $\delta/u = \prod_{k'=0}^{1}\prod_{a=1}^{r-1}(1 - g_{k'}^{-a})^f$, by the order of GL_{r-1} and the unipotent order. Hence $1 - \delta/u \le f/(g_0 - 1) + f/(g_1 - 1) \le y$ at the same threshold. The three comparisons give $n/\delta \le 4d/(3(1-y)^2) \le 64d/27 < 3d$. The index in the actual product of mirabolic groups.

The p-normalized exponent of u is $fr(r-1)/2$ at g_0, g_1 and zero elsewhere. Indeed u is the product of the orders of the upper unipotent groups over the split residue fields, of order $g_i^{r(r-1)/2}$ per factor. The two unipotent terms at one place are therefore bounded by $q_v e_v d$, where $q_v = 2fr^2$ at the type primes and zero elsewhere.

Put $c_{jv} = (E - \operatorname{ord}_v \sigma)/e_v$. The depth estimates on the same run and weight give $c_{jv} + c_{j,\bar v} - 2F_{jv}^0 + 2q_v \le D_j(p)$. The four applications are as follows. At g_0, g_1 each depth is bounded by $2T - \nu_{jv}$ on its side, with $\nu_{jv} + \nu_{j,\bar v} = -2F_{jv}^0$, giving $D_j(p) = 4T + 4fr^2$. At the small exceptional primes both depths are at most $s_{C,j} + T + 1$ and $F_{jv}^0 = 0$. At q, d both are at most $s_{B,j} + T + 1$, with $F_{jv}^0 \ge -2fkb'$. At all other primes up to Y, each depth is at most $F_{jv}^0 + a_j(p) + 1$, giving $D_j(p) = 2a_j(p) + 2$. In the remaining range both depths vanish and $F_{jv}^0 \ge 0$, so $D_j(p) = 0$. The depth estimates on the marked chart · The four numerical bounds.

At conjugate places $t(v) = t(\bar v)$, $e_v = e_{\bar v}$, and $\lambda_v = \lambda_{\bar v}$. Put $F_{jv} = e_v F_{jv}^0$. Since $\sum_j \epsilon_{vj} \le p^{-2}$ at either read, addition gives

$$\frac{1}{d}\sum_j \big(i_{jv}(t(v)) + i_{j,\bar v}(t(v))\big) \le \sum_j \left(2F_{jv} + e_v D_j(p) + \frac{r_{vj} + r_{\bar v j}}{\delta d}\right) + \frac{2e_v}{p^2}.$$

When $v = \bar v$, the same term is added twice as required by the Hermitian change of basis. At p one has $\sum_{v|p} \lambda_v e_v = \log p$. Under extension from a geometric field, $\sum_{v|v_0} \lambda_v e(v/v_0) = \lambda_{v_0}$. The geometric residual sum on either run is at most $nfmb\,d_2 \le nfmb$. Thus both terms of the weighted residual contribute at most $4nfmb/(\delta d) \le 12fmb$ per stage.

For $L = \log b$, the table has $\sum_p D_j(p) \log p = O(fmbL^2)$. At ordinary $p \le Y$, $a_j(p) = A_1/p + A_2\,\mathrm{th}(p) + \mathbf{1}_{p>X_2}A_2/p^2$, with $A_1 \ll fmb \log k + fk(S_c + \eta_{\mathrm{sz}})$, $A_2 \ll fkb$, and $\sum_p \mathrm{th}(p) \log p \ll 1/k$. Apply the prime

sums $\sum_{p\le Y}\log(p)/p \ll 1+\log Y$, $\sum_{p\le Y}\log p \ll Y$, and $\sum_{p>X_2}\log(p)/p^2 \ll X_2^{-1}$, using $\log Y \le 2L$, $Y \le mb$, $k \le X_2$, and $k(S_c+\eta_{\mathrm{sz}})L \le mb$. The small exceptional set, q, d, and the two type primes have total $O(fmbL)$ by the chosen size bounds. Thus the weighted term from $2F_{jv}$ is $\mathcal{F}$, and the separate place remainder is as bounded below.

Assembly of the finite determinant bound. Use the support consisting of the twisted places and the places where $G_v(0)$ is not a unit. The stage indices, stage Gram valuations and reference exponents vanish outside the twisted places by the prescribed base support. On that twisted support apply the paired transport inequality above. In terms of the unweighted F_{jv}, with $2\sum_{j,v}\lambda_v F_{jv} = \mathcal{F}$, the required local upper bound for their sum per row is $\sum_j(2F_{jv}+e_{jv})+r_v$, where $r_v = 2e_v/p^2$ at $v \mid p$. This is a comparison of the transported lattices $A_j(t)$ with J_j^{row}. The Gram inclusions above apply before this transport. For the place remainder use $\sum_{v\mid p}\lambda_v e_v = \log p$. Then $\sum_v \lambda_v r_v < 1$. For instance $\sum_{p<240}\log(p)/p^2 < 0.4891$, by trial division and the first twelve terms of $\log y = 2\sum_{i\ge0}((y-1)/(y+1))^{2i+1}/(2i+1)$ after reduction to $1 \le y \le 2$, with the geometric bound on the tails. In intervals $[6a, 6a+6)$ beginning at 240 there are at most two primes, and $(\log y)/y^2$ decreases. Their contribution is at most $\frac{1}{3}\int_{234}^{\infty} y^{-2}\log y\,dy < 0.0092$. Base support and Weighted prime remainder.

Let $E_{1,j}$ bound the weighted sum of $e_{jv}+J^*_{jv}$, and $E_{2,j}$ that of $\eta^{\mathrm{cut}}_{jv}/d_{\mathrm{row}} + \mathrm{ord}_v(sg_{jt})$. With $E_3 = 1$ and the bound for the quotient Gram term this gives

$$\mu(C^{(0)}) - J_\kappa \le \mathcal{F} + \sum_j L_{3,j} + L_G + \sum_j (E_{1,j}+E_{2,j}) + E_3.$$

Here $L_{3,j}$ is the stage bound from (99). The use of the weighted inequality uses the Weyl estimate above and the paired transport indices on the prescribed stage lattices. Division by the common row rank.

The compensated mean at the two row levels. Use the actual principal-level lattice and its two character restrictions from (156). By (157), the principal norm squared on j_iR_i is h_i times the level norm squared, for $R_0 = R^{(0)}$ and $R_1 = R$. Divide each squared metric by the same positive κ_R at the appropriate embedding. Its effect on primal slopes is to add J_κ. Write $\widehat{\mu}$ for these compensated slopes in this paragraph, and let $\lambda = \log|G| \le \lambda_0$.

Over a finite splitting extension of K, let V_s be a retained irreducible isotypic space of the principal module, of the form $W_s \otimes U_s$ with W_s irreducible. Normalized degrees and ranks are unchanged by this extension. The base Whittaker line in W_s has dimension one. Put $a_s = \dim U_s$, $n_s = \dim e_1W_s$, and $L_s = L_{\mathrm{amb}} \cap e_0V_s$ with the principal metric. Then $\mathrm{rk}\,L_s = a_s$. The isotypic projectors are orthogonal in every transported principal metric. The central character formula makes $|G|$ times each projector integral on L_{amb}. Thus the direct sum of the L_s lies in S_0 and contains $|G|S_0$. By the level sandwich, $j_0R_0 \subseteq S_0$ and $h_0L_s \subseteq j_0R_0$. Using the metric factor h_0 and squared-slope units gives

$$\frac{\sum_s a_s\widehat{\mu}(L_s)}{\sum_s a_s} \ge \widehat{\mu}(R_0) - 3\lambda, \qquad \widehat{\mu}(L_s) \le \widehat{\mu}_{\max}(R_0) + \lambda.$$

For the second inequality, multiplication by h_0 is an integral injection of squared norm at most h_0 from L_s to R_0. The first interface uses only the index killed by $|G|$ and the comparison of means under inclusion of full-rank lattices.

In W_s the translates of a nonzero vector on the e_0 line span W_s. Applying e_1 gives a spanning family of e_1W_s. Choose $g_{s,1},\dots,g_{s,n_s} \in G$ whose projected translates are independent. The same algebraic rank holds at each embedding. The map

$$(v_{s,j})_{s,j} \longmapsto j_1^{-1}\left(h_1^2\sum_{s,j} e_1\rho(g_{s,j})v_{s,j}\right) \quad \text{from} \quad \bigoplus_s L_s^{\oplus n_s} \text{ to } R_1$$

is integral, since $h_1 e_1$ preserves the principal lattice and the second factor h_1 puts the output in $j_1 R_1$. It is injective and has full rank, by the independent projected translates on W_s tensored with the identity on U_s. Within one isotypic space its squared norm is at most $h_1^3 n_s$, using Cauchy–Schwarz, unitarity and the factor h_1^{-1} of the level metric. Orthogonality of the isotypic spaces gives the same maximum as the bound on the whole map. It is at most $|G|^4$. Taking determinants yields

$$\widehat{\mu}(R_1) \geq \frac{\sum_s n_s a_s \widehat{\mu}(L_s)}{\sum_s n_s a_s} - 4\lambda.$$

The determinant proof works for any actual common multiplicity bound $(1-\theta)\bar{n} \leq n_s \leq (1+\theta)\bar{n}$ with $\bar{n} > 0$, $0 \leq \theta \leq 1/3$. Let H bound $\widehat{\mu}_{\max}(R_0)$ from above. The nonnegative deficits $(H+\lambda) - \widehat{\mu}(L_s)$ have sharp weighted average at most $(1+\theta)/(1-\theta)$ times their base weighted average, by these multiplicity inequalities. Since this ratio is at most $1+3\theta$, the estimates just proved give

$$\widehat{\mu}(R_1) \geq \widehat{\mu}(R_0) - 7\lambda - 3\theta\big(H - \widehat{\mu}(R_0) + 4\lambda\big).$$

By duality this is the desired upper mean comparison for $\mu(C) - J_\kappa$, with base mean $\mu(C^{(0)}) - J_\kappa$. The relative multiplicity tolerance θ concerns each retained simple type separately. The actual level metric and lattice sandwiches · The base line and sharp multiplicities.

The multiplicities in the two character spaces. Use the residue frames of (41) and the common principal action of (156). Each retained simple W_s, after extension to a splitting field, is a tensor of local irreducibles over the $2f$ chosen split halves. The central good projector has squarefree labels at every one of these places. For a squarefree label its local irreducible is the parabolic induction of the cuspidal representations of the distinct factors, as in the squarefree calculation of Section 3. The character levels are the full upper unipotent with its generic character at the base and the symmetric subgroup J in the distinguished row count at the sharp level. The prescribed permutation of the sharp frame conjugates the subgroup and its character together. In the conjugate row realization the characters are inverted together with the action. The calculation below uses any nontrivial additive character, so it applies with either of the prescribed residue readings. The two level subgroups and characters · The tensor of squarefree types · The retained simple types in the mean passage.

Write q for the residue order at one place, ψ for its nontrivial additive character and U for the upper unitriangular subgroup of $\mathrm{GL}_r(\mathbf{F}_q)$. Number the coordinates from zero. For $0 \leq a \leq h < r$ let $U(a,h) \subset U$ have $u_{ih} = 0$ for $a \leq i < h$, and put

$$c_h(u) = \sum_{i+1 \neq h} u_{i,i+1}.$$

For a representation V of the finite general linear group, let D_h be the dimension of the $\psi(c_h)$-eigenspace on U. Let $T_{a,h}$ be the same dimension on $U(a,h)$ and $W_h = T_{0,h}$. The sharp subgroup in position a with block length c is $U(a, a+c)$. Its character is generic on each of the blocks of ordered coordinates $[a, a+c)$ and its complement. Let $S_{a,c}$ be that character dimension. Thus $S_{0,c} = W_c$.

Here is the finite Fourier calculation for these dimensions. Fix $a < h$ and define

$$L_e = \{u \in U \,:\, u_{aj} = 0\ (a < j < e),\ u_{ih} = 0\ (e \leq i < h)\}, \qquad a+1 \leq e \leq h.$$

For $e < h$, the two groups L_e, L_{e+1} are generated by $L_e \cap L_{e+1}$ and respectively the root groups $1 + xE_{ae}$, $1 + yE_{eh}$. The intersection is the kernel of $u \mapsto u_{ae}$ on L_e and of $u \mapsto u_{eh}$ on L_{e+1}, as follows from the missing row and column entries. If the common character is $\psi(\tau(u) + u_{ah})$, with τ a sum of adjacent entries, it extends by the same formula to both groups. In their product law for u_{ah} every intermediate summand vanishes. On the common character space the commutator of the two root groups is the scalar $\psi(xy)$. The second root group therefore permutes all the character spaces of the first transitively. The same holds with their roles interchanged. Each decomposition of the common space has q summands

of equal dimension, so the two extensions have the same multiplicity. Iterating exchanges the column interval in L_{a+1} for the row interval in L_h.

For $c > 0$ and $a + c + 1 < r$, move the first coordinate after the block $[a, a + c)$ to its front in $S_{a,c}$. This makes a missing row $u_{aj} = 0$ for $a < j < h = a + c + 1$. The complementary chain uses u_{ah}, and the other character entries are the adjacent ones with $u_{a,a+1}$ and $u_{h-1,h}$ omitted. The exchanges just proved give the missing column of $S_{a+1,c}$ with this same character. Hence $S_{a,c} = S_{0,c}$ for all positions. For $c = 0$ this is the full generic character at every position.

Now the quotient $U(a+1, h)/U(a, h)$ is the additive group via u_{ah}. The characters in its decomposition are the zero coefficient and the $q - 1$ nonzero coefficients of u_{ah} in $\psi(c_h(u) + tu_{ah})$. The latter have equal dimensions, by conjugation with the diagonal scalar on coordinates $< h$ and identity on the others. For coefficient 1 use the exchanges from L_{a+1} to L_h with $\tau = c_h$. Moving index a to the end of $[a, h)$ then gives the subgroup and character of $S_{a,h-a-1}$. In detail, the block $a + 1, \dots, h - 1$ keeps its successive edges, the complementary chain follows the old index a through h, and every entry from that block to the old index a is zero. Consequently

$$T_{a,h} = T_{a+1,h} + (q - 1)W_{h-a-1}, \qquad W_h = D_h + (q - 1)\sum_{j<h} W_j.$$

The second equation uses $T_{h,h} = D_h$. Setting $\gamma_q(0) = 1$ and $\gamma_q(l) = (q - 1)q^{l-1}$ for $l > 0$, the recurrence gives

$$S_{a,k} = W_k = \sum_{h=0}^{k} \gamma_q(k - h)D_h. \tag{150}$$

Indeed the partial sums satisfy $\sum_{j=0}^{h} W_j = D_h + q\sum_{j=0}^{h-1} W_j$. The change of position · The root exchanges · The Fourier recurrence.

Apply this to V induced from the inflated tensor of irreducible cuspidal representations of the standard Levi. By Bruhat decomposition the parabolic–U double cosets are words in the Levi block labels with their prescribed multiplicities. Mackey's formula tests $\psi(c_h)$ against the inducing tensor at the intersection attached to each word. If two successive labels ascend on either side of the cut at h, the character is nontrivial on a root subgroup of the inducing parabolic's unipotent radical, so that contribution vanishes. If a label occurs on both sides of the cut, the cross-cut roots within that Levi block form the unipotent radical of a proper parabolic in the block. The character is trivial on this radical, and cuspidality again gives zero. Each remaining word places whole blocks on each side in decreasing block order. Its same-block roots are maximal upper unipotent groups with generic characters. The interblock roots in the intersection cross the cut, where the character is trivial. Thus there is one contribution of dimension one per subset of whole blocks of total size h. Gel'fand's existence and uniqueness of maximal-unipotent character lines apply to each inducing cuspidal representation. Cuspidality excludes a character trivial on any adjacent root, since it would be trivial on the radical of the corresponding maximal parabolic. The diagonal torus conjugates the remaining generic characters. [45, Theorems 0.4'–0.5', pp. 17–18] Double cosets and the cut character.

In particular $D_0 = 1$, giving the base rank one. Also $D_h \le \binom{r}{h}$, by mapping a block subset to its union of coordinates. Since $1 \le k < r \le q$,

$$\begin{aligned} \frac{W_k}{q^k} &\ge 1 - \frac{1}{q}, \\ \frac{W_k}{q^k} &\le 1 - \frac{1}{q} + \sum_{h=1}^{k} \binom{r}{h} q^{-h} \le 1 - \frac{1}{q} + e^{r/q} - 1 \le 1 + \frac{2r - 1}{q}. \end{aligned}$$

The last inequality uses $e^z \le 1 + 2z$ on $[0, 1]$. The pointwise local bound.

Return to the selected product G. Write $q_z = |k_z|$ for its residue orders, $g_{\min} = \min(g_0, g_1)$ and

$$\bar{n} = \prod_{z \in P_g} q_z^k, \qquad \theta = \frac{8fr}{g_{\min}}.$$

There is one chosen split half over each of the f places of F above g_i, with $q_z = g_i$ there. The rank of the sharp projector on W_s is the product of its local ranks. The product of the lower factors is at least $1 - \sum_z q_z^{-1}$, and the product of the upper factors is at most $\exp((2r-1)\sum_z q_z^{-1})$. Since $4fr \le g_{\min}$, the same exponential inequality gives

$$(1-\theta)\bar{n} \;\le\; \dim e_1 W_s \;\le\; (1+\theta)\bar{n} \qquad \text{for every retained simple } W_s. \tag{151}$$

This uses the irreducible factor before taking the tensor with its multiplicity space. The dimensions are also the dimensions of the character projectors over a splitting field of K, by invariance of rank under extension. The bound is pointwise on the retained types at every index. The residue group on the chosen record · Product of the local multiplicities.

The size of the compensated spread. Write $\lambda = \log|G|$ and put

$$B_{\text{base}} = \max\left(0, \mathcal{F} + \sum_j L_{3,j} + L_G + \sum_j (E_{1,j} + E_{2,j}) + E_3\right),$$
$$Q_{\text{sp}} = \max\bigl(1, B_{\text{base}} + H + 28\lambda\bigr).$$

Here H is a uniform majorant of $\mu_{\max}(R^{(0)}) + J_\kappa$ from the base part of the row norm estimate. One can use $H = x^{C_H} b$ with the fixed exponent in the family application. The base determinant inequality gives $\mu(C^{(0)}) - J_\kappa \le B_{\text{base}}$. Since $\lambda \ge 0$, Q_{sp} bounds the expression $H + \mu(C^{(0)}) - J_\kappa + 4\lambda$ in the mean comparison. It also bounds the corresponding expression with 28λ in the division inequality. The compensation and the required error factor.

All the terms in B_{base} have an upper bound of the form $Ax^c b$ with fixed constants. Indeed $L_{3,j}$ in (99) has period term $(p_j - q_j)\sum_{\rho \in P} w_{u,j}(\rho)\log\Omega$ and error $O(B_0 \log\log x)$. The stage weights are at most $2b$, $|\log\Omega| \ll L$, and the number of stages, embeddings and the ranks grow at most polynomially. The reference estimate gives $\mathcal{F} \ll x^c b$, by summing the ordinary $O(x^c b/p)$ exponent against $\log p$ up to $X'' \le a_E$; the exceptional reference contributions are nonpositive. The weighted paired depth table and the Weyl lattice comparison give $E_{1,j} = O(fmbL^2)$ as a majorant. For $E_{2,j}$ use (121) and (91). The normalized Weyl and dual Weyl pairings on this base face are integrally dual, so their common coefficient duality integer is 1. The remaining stage denominators divide $|G|D_\Gamma$: the level index divides $|G|$, and the class stabilizer orders divide the common multiple D_Γ on the selected record. Thus one can take $L_G = \lambda$, and

$$E_{2,j} \le \eta_j / D_0 + \log(|G|D_\Gamma).$$

The selected size $\log D_\Gamma \ll f\ \mathrm{rank}(X_m)\log(f\ \mathrm{rank}(X_m))$ and $N_c \le N_2 kb\, x^{2\varepsilon}/n^2$ prove the polynomial times b bound on these terms as well. The remainder is $E_3 = 1$. The unit coefficient duality allowance · The prescribed class stabilizers · The common stage denominator · The upper bound on the base mean.

Finally $\lambda \le fr^2(\log g_0 + \log g_1)$ by the finite general linear group order. Fix the accuracy and all family constants, including the exponent in H, before increasing x. Since $b = \exp(x^\varepsilon)$ with $\varepsilon > 0$, beyond one threshold we have $B_{\text{base}} \le b^2$, $H \le b^2$, $28\lambda \le b^2$ and $72fr \le b^6$ for every selected index, hence $Q_{\text{sp}} \le 3b^2$. The same threshold gives $4fr \le g_{\min}$, using the selected $g_{\min} \ge (kb)^8$. Consequently

$$3\theta Q_{\text{sp}} \le \frac{72fr\,b^2}{g_{\min}} \le 1, \qquad 0 \le \theta \le \tfrac{1}{3}. \tag{152}$$

The implication for θ uses $Q_{\mathrm{sp}} \geq 1$. The upper bounds on B_{base} and H use the base lattice throughout. The threshold applies simultaneously to all selected translates and words. Each use of the mean comparison keeps the prime pair, coefficient realization and both projectors of its own selected record. The numerical threshold for the division inequality.

Apply (151) and (152) in the determinant comparison. Its added cost is at most $1 + 7\lambda \leq 1 + C_{\mathrm{lev}} m \lambda_0$ for the fixed $C_{\mathrm{lev}} \geq 28$. Adding the height estimate cancels the row normalization J_k.

For the weight calculation let $b_j = p_j - q_j$, write $w_{j\rho} = w_{u,j}(\rho)$, and use ν_i for the shifted row coordinates at (j, ρ). Write $H_{j\rho} > 0$ for the corresponding hook product. Put $\phi(z) = z \log z - z$ with $\phi(0) = 0$, and $\Phi(z) = \phi(z) - z \log \Omega$. The weight comparison on the prescribed arrays is

$$\sum_i \nu_i = \mathbf{1}_{\rho \in T} a_\rho + k w_{j\rho} + \Delta_{j\rho}, \ \sum_i \phi(\nu_i) - \log H_{j\rho} \quad \leq \mathbf{1}_{\rho \in T} \phi(a_\rho) + E_{j\rho}.$$

The sampled coordinate is used only when $\rho \in T$. Consequently

$$\sum_i \Phi(\nu_i) + b_j w_{j\rho} \log \Omega - \log H_{j\rho} \leq \mathbf{1}_{\rho \in T} \Phi(a_\rho) + E_{j\rho} + \big(|b_j - k| w_{j\rho} + |\Delta_{j\rho}|\big) |\log \Omega|.$$

This is the cancellation between the stage covolume main term and the height. We give the row and hook comparison used here. Put $w' = w_{j\rho} - (k + (2j-1)(t_{\mathrm{cut}} + 1))/2$, where t_{cut} is the twist cutoff. It is a nonnegative integer in the chosen weights. Let u_i, $0 \leq i < k$, be the first source block and v_i, $0 \leq i < k-1$, the interlacing column weights. Subtracting the half-sum coordinates gives

$$\beta_i = u_i - (k-1)/2 + i, qquad \gamma_i = v_i - (k-2)/2 + i.$$

The first block of the ν_i is $w' + \beta_i$. The corresponding hook row lengths are $w' + \gamma_i + (j-1)t_{\mathrm{cut}}$. Both are integers. Since $u_i > v_i > u_{i+1}$, the shifted differences give $\beta_i \geq \gamma_i \geq \beta_{i+1}$, since an integer strictly less than $1/2$ is nonpositive. Outside T, the endpoint bounds give $0 \leq \beta_i \leq c_*$, where c_* is the cloud width. On T, the γ_i and interior β_i lie in this interval. The two source endpoints there are $a_\rho + w'$ and a number in $[0, w']$. The remaining source block has sum at most $m(t_{\mathrm{cut}} + 1)\eta_*$, where η_* is the auxiliary weight displacement. Shifted weight data. Half-sum shifts and interlacing.

The hook product is the factorial quotient in the explicit height: in row i of a partition with p rows of lengths a_i, the hooks run from 1 to $a_i + p - i - 1$ with $a_i - a_l + l - i$ omitted for each $l > i$, numbering rows from zero. This follows by counting the columns extending to row l when the row index in the hook changes. Each hook is at least the number of boxes remaining in its row, including its own box. Hence $H_{j\rho}$ is at least the product of the row factorials. For integers $a, d \geq 0$, with $a \leq A$ and $A \geq 1$, one has $\phi(a) - \log(d!) \leq (a-d)_+ \log A$. For $d \geq a$ this follows by integrating $\log y$ up to a to bound $\phi(a)$ by $\log(a!)$. For $d < a$, integrate from d to a and use the upper bound $\log A$ on the integrand. Apply this with $A = a_E \operatorname{rd}(E)$. Outside T, pair the first $k-1$ source rows with the hook rows. Their total unpaired positive difference is at most $\sum_i (\beta_i - \gamma_i) \leq c_*$ by telescoping $\beta_i - \beta_{i+1}$. The last source row costs at most $(w' + c_*)L_a$. On T, pair the interior source rows with the preceding hook rows, which are at least as long. The first endpoint costs at most $\phi(a_\rho) + w' L_a$ and the last at most $w' L_a$. The remaining source block costs at most $m(t_{\mathrm{cut}} + 1)\eta_* L_a$. This proves the asserted bound for $E_{j\rho}$. The sum of the first block differs from $k w' + \mathbf{1}_{\rho \in T} a_\rho$ by at most $k c_* + w'$, and the remaining block has the same size bound as above. Finally $|w_{j\rho} - w'| \leq n_*$ and $k n_*, k c_*, m(t_{\mathrm{cut}} + 1)\eta_* \ll b$ in the size range, where n_* is the largest column rank. Thus $E_{j\rho} \ll b L_a$ and $|\Delta_{j\rho}| \ll b$. With $|b_j - k| \leq m$, $w_{j\rho} \leq 2b$, and $|\log \Omega| \leq L_a \ll L$, their sum is $O(mB_0)$. Cancellation at each embedding.

The non-hook analytic errors in the explicit height contribute $O(N_T a_E L_a + mB_0)$. The terminal evaluation error has order $(N_T b' n_* + f(m\eta_* + c_*)n_* + f n_*^2)L_a$. The stage errors have order $f n_*^2 (L_a + m^3 L)$, with the total lower-hook allowance $O(mf n_*^2 L_a)$. The size inequalities used here are $b' n_* \leq a_E$, $m\eta_* n_* L \leq b$, $c_* n_* \leq b$, $n_*^2 \leq b$, and $m^3 n_*^2 L \leq b$. Summing therefore gives the stated scale. The finite determinant and stage covolume errors require the separate bound $O(m(B_0 + N_c)x^{e_D \varepsilon})$. For the additional twisted length,

the normalized local expression is $e_v\eta_{jv}/D_0$ plus valuations of the stage group order and the chosen depth s_j. If g_{gr} is a common multiple of the stage group orders, summing with λ_v gives

$$E_{2,j} \le \eta_j/D_0 + \log g_{\mathrm{gr}} + \log s_j.$$

Here $\sum_v \lambda_v \operatorname{ord}_v(g) = \log g$ for a positive rational integer g. Thus the first term is bounded by (91) on the actual twists. Identification of the twisted length. Group-order and depth allowances. With that bound the result is

$$\mu(C) + h_w - \mathcal{F} \le P_{\mathrm{div}} - S_a \log T_w + O\big(N_T a_E L_a + m(B_0 + N_c)x^{e_D\varepsilon}\big), \qquad S_a = \sum_{t\in T} a_t. \tag{153}$$

Here P_{div} comes from factorials and the normalized CM metric. The component height estimate enters this calculation through $L_{3,j}$ in the bound on the mean of the row dual. The local determinant bounds, the metric estimates on the actual recursive maps and the comparison of the shifted rows with the hook product are required in this application. The row-dual mean comparison. Numerical error absorption.

The explicit height used for the lower bound at this scale is

$$h_w = -S_a \log T_w \; + \sum_{j=1}^{m} \sum_{\tau} \sum_{i<n_{j-1}} \Phi(\nu_{w,j,\tau,i}) \; + H_0^{\mathrm{arch}} - J_\kappa.$$

Here n_j is the stage column rank, $\nu_{w,j,\tau,i}$ is the shifted source weight, and H_0^{arch} consists of the positive error allowances minus the lower hook terms. This formula also gives direct control of height variation by the interior weight differences. Explicit height expression and variation between words.

4.3.8 The least target slope

The frame bound uses $\mu_{\min}(C) = -\mu_{\max}(R)$. From the explicit height formula we obtain $-(h_w + J_\kappa) \ll x^{c_1} b$ for a fixed c_1. The number of shifted coordinates and hook factors is at most a fixed power of x. The shifted weights and positive integer arguments of the hook factorials are bounded by $x^c b$ for a fixed c. For $z \ge 1$ one has $\Phi(z) \ge -(1 + |\log \Omega|)z$, and $\log((d-1)!) \le d \log d$ for each positive integer d. The hook denominators are integers at least one. All the other logarithms in $h_w + J_\kappa$, including $\log T_w$, are bounded by fixed powers of x. The sum thus has the stated lower bound. The lower bound on the explicit height.

The donor estimate above gives $\mu(N'_w) \le \mu_{\max}(N'_w) \ll x^{c_2} b$. Here is the coarse upper bound on the gains used in it. On each contributing orbit the shifted source degrees lie in $[0, 2a_E]$. The receiving degrees and the subtracted error terms in the two gain expressions are nonnegative. Each expression is therefore bounded by $p/(p^2 - 1)$ times the sum of the shifted source degrees on the orbit, using the sum of all positive odd powers of $1/p$. Maximizing with zero and retaining the specified orbits preserves this bound. The orbits partition the sample, so $g_{jv} \le 4N_T a_E/p$. Gains vanish above X''. Using $\sum_{v|p} e_v\lambda_v = \log p$ gives $\mathcal{G}(u) \ll mN_T a_E(1 + \log X'')^2$. Here the sum of $(\log p)/p$ is bounded by the sum over all integers, which is $O((1 + \log X'')^2)$ by integration. For $\mathcal{F}$, at an ordinary prime the reference exponent on an orbit is at most the baseline $p\sum D_s/(p^2 - 1)$, because the allocation zero is feasible. After division by the orbit length and summation over embeddings this bounds it by a fixed polynomial in the rank parameters times $\max_{j,\rho} w_{u,j}(\rho)/p$. At the auxiliary primes and at d the reference terms are nonpositive. Since the weights are $O(x^c b)$, the same integer-sum bound applies. Finally $-Q_w \ll mB_0$ follows from $Q_w = mS_a\lambda_g + O(mB_0)$ with $\lambda_g \ge 0$. Local gain definitions. Nonnegative corrections. Reference exponents. Stage weight and local reference bounds.

The remaining bound is on the normalized row lattice itself.

$$\mu_{\max}(R) + J_\kappa \le x^{c_3} b. \tag{154}$$

At the base, use the common injective test Ψ above. Give the base row lattice $R^{(0)}$ the squared norms $\kappa_{R,\sigma}^{-1}\|\cdot\|^2_{R^{(0)},\sigma}$, and call the resulting lattice $\widetilde{R}^{(0)}$. Its maximal slope is $\mu_{\max}(R^{(0)}) + J_\kappa$. Let M_{ev} be the lattice

of all fitted evaluations. If the source test values before the square normalization are in $\mathfrak{p}_v^{-d_v}(M_{\mathrm{ev}})_v$, the fitted common map shows that

$$\Psi(R_v^{(0)}) \subseteq \mathfrak{p}_v^{-a_v}(M_{\mathrm{ev}})_v, qquad a_v = d_v + \mathrm{ord}_v(q_{\boldsymbol{\mu}(v)}).$$

The root of unity has valuation zero. When these exponents have finite support the local inclusions give $\Psi(R^{(0)}) \subseteq IM_{\mathrm{ev}}$, where $I = \prod_v \mathfrak{p}_v^{-a_v}$. On each rational subspace of $R^{(0)}$, apply the exterior-power degree inequality to the injective map to its image and saturate that image in IM_{ev}. This yields

$$\mu_{\max}(R^{(0)}) + J_\kappa \le \mu_{\max}(M_{\mathrm{ev}}) + 2\sum_v \lambda_v a_v + h_\infty(\Psi : \widetilde{R}^{(0)} \to M_{\mathrm{ev}}).$$

For the first term, M_{ev} is the orthogonal sum of coordinate duals of the prescribed integral Weyl lattices on the definite vertices. The underlying Hermitian lattices are self-dual, since at each finite place they are unitary translates of the self-dual prescribed lattices. The normalized contravariant pairing identifies each coordinate dual with the integral dual Weyl module, up to conjugation. The divided-power highest weight module has the induced module of the same weight as its contravariant dual, realized by the integral bideterminants. In the polynomial model Fischer pairing is contravariant under transpose, and normalizing by the squared norm of the highest polynomial makes the extremal pairing one. Thus this is the integral duality in the definite bases. Evaluation lattices and metrics and integral Weyl duality. Let s_0 be the product of the highest polynomial's Fischer squared norms over the host embeddings. At a finite place and its conjugate, the cross Gram matrices in the self-dual Lie bases are unimodular over the local ring. Multiplication by s_0 restores the Fischer pairing with monomial entries $\alpha!$, integral substitution and determinant twists of units. It follows that s_0 times the rational Hermitian pairing of the coordinate dual is integral. Integral comparison at paired places.

For a rank-a sublattice with pseudo-basis matrix D and product ideal I_D, this says $s_0^a \det(D^* H D) I_D^* I_D$ is integral, where H is the rational Gram matrix. Its determinant is nonzero at the coefficient embedding where the involution acts by complex conjugation. At other embeddings σ, put $\sigma^{\#} = \overline{\sigma \circ *}$. The embedded Gram matrices agree at σ and $\sigma^{\#}$ by Hermitian symmetry. Determinant Cauchy–Schwarz bounds the embedded absolute value of $\det(D^* H D)$ by the geometric mean of the two positive Gram determinants of the frames $\sigma(D)$ and $\sigma^{\#}(D)$. Multiplying over all embeddings and using the ideal norm gives a squared logarithmic covolume at least $-a \log s_0$. This applies on every rational subspace, hence $\mu_{\max}(M_{\mathrm{ev}}) \le \log s_0$. If $R_* \ge 1$ bounds the weight spread at rank n_*, each highest polynomial is a product of at most R_* minors, of total degree at most $n_* R_*$ and coefficient absolute sum at most $(n_*!)^{R_*}$. The monomial factorials bound the logarithm of its Fischer squared norm by $3 n_* R_* \log(n_* R_*)$. In the actual grid the spread is bounded by the cloud width plus at most $O((m+1)(g_{\mathrm{sep}}+1) + n_*)$, with g_{sep} the separation of the auxiliary weight blocks. The cloud width is at most b and the separation is bounded by a fixed polynomial in the ranks times one plus the square root of the cloud width. Consequently $\log s_0 = O(x^c b)$ for a fixed c. Conjugating coefficient embedding and location of the definite grid weights.

The two remaining terms in the fitted estimate can be evaluated before the passage to the sharp level. At a finite place use the chosen tuple $\boldsymbol{\mu}(v)$. The denominator loss d_v is bounded by the sum of the two chains' per-step losses, the final polynomial pairing and initial row losses, and the losses of the common scalar, tensor coefficients and rational stabilizers. Test against integral column vectors supported at one definite class. The rational transporter carries the indicated class lattices into one another, and these tests compute the coordinates of the dual lattice M_{ev}. The ultrametric inequality gives the same loss bound for the full sum as for each term.

All rational stabilizers have a common denominator of logarithm $O(D_{\mathrm{ev}}^2)$, where $D_{\mathrm{ev}} = [E : \mathbf{Q}] \max_i \dim_E V_i$ for the fitted vertices V_i. A finite stabilizer preserves the saturated rational representation's integral lattice obtained by summing the translates of any full lattice. It injects into $\mathrm{GL}_{D_{\mathrm{ev}}}(\mathbf{F}_3)$ after padding the rank. Indeed a prime-order element in the principal kernel would have form $1 + 3^a B$ with $a \ge 1$ and B nonzero modulo three. In the binomial expansion of its order-th power the first term has

strictly smallest valuation, also at order three. This contradicts prime order. Termwise losses and Passage to the fitted sum.

Initial tests and the corner pairing. For the column extension the compact group is the principal terminal level K^+ used in the principal test column span. It has trivial right character, and its two distinct type primes make the lattice stabilizers at that level trivial. Thus the covariance extends arbitrary integral coefficient values in a single terminal class. Intermediate class stabilizers are retained with their reciprocal weights in the unfolded sum. The terminal level and its full coefficient space.

The two chains start at the definite test and at the value of the full good row. At a definite test class c, take a vector in the prescribed Weyl lattice in that class, and extend it by the class covariance and by zero elsewhere. If a finite argument g lies in that double class, write $g = \gamma c k$ with rational γ and k in the compact level. Right covariance acts by the prescribed level unit. The left action of γ transports the class lattice to the lattice attached to g, since k fixes the reference lattice there. Inverting the chosen frame gives the lattice in the initial test coordinates. Applying tensor Weyl functors to the transport preserves this lattice inclusion. The same transport is an isometry of the definite coefficient form at each conjugate embedding. Thus the initial definite test has zero finite loss and its local squared norm is bounded by that of the coefficient vector at c. The test class is one of the summands of the evaluation dual. The transported test evaluation.

For the initial row use the actual presentation $\mathrm{formC}(t)(j_X(t)\,\mathrm{formOf}(t,x))(g)$. The good basis includes the integral cut in the full class lattice, before taking the row quotient. In that class presentation the value at $g = \gamma c k$ is its value at c transformed by γ and multiplied by the level character at k. The character is in the coefficient field and has finite order, so is an integral unit even at a coefficient prime dividing its order. The rational transporter identifies the actual E-lattice at g with the class lattice transported from c. After applying the fixed contragredient polynomial substitution and the inverse chosen frame, the finite row value is in the prescribed dual-frame coefficient lattice. In the standard local bases the contragredient substitution replaces minors by signed complementary minors, with the determinant shift accounted for by the dual weight. It takes integral polynomials to integral polynomials. Its covariance sends the transported lattice to the lattice in the new frame. This proves zero finite loss for the initial row at every tested twist. At the base twist the rational transport preserves the definite form at every embedding. The contragredient map preserves the normalized Fischer metrics, since it intertwines the corresponding unitary actions and identifies their unit extremal lines up to sign. Hence the squared norm of the initial row value is at most

$$|P|\, D_\Gamma \|x\|^2_{\text{full row},\sigma}, \qquad D_\Gamma = \operatorname{lcm}_c |\Gamma_c|.$$

Indeed the class norm of the row basis is $|P|$ times the row form and each summand of the class norm has weight $|\Gamma_c|^{-1}$. This is the bound on the full good row, so for the quotient use the least-norm lift. A finite lift of an integral quotient row may be taken in the original integral cut by the definition of the quotient lattice. Class covariance and the placed finite row and Class norm on the full good row.

The corner pairs the outputs of the two fitted chains. In the corner reading, over the tested finite class, the upper Lie lattice is the last-coordinate extension $M \oplus \mathcal{O}_E$ of the lower lattice M. With a local basis B of M, the corresponding dual frames are $B^{-\mathrm{t}}$ and $(B \oplus 1)^{-\mathrm{t}}$. The branching form in these frames uses the last-coordinate parameter 1, as follows by equivariance under $B \oplus 1$. Write N_Λ, D_Λ for the standard dual-Weyl and Weyl lattices at the upper weight, and N_α, D_α for the lower weight. The normalized bilinear Fischer form makes these two lattices dual at each weight. Because its coefficients before division by the squared extremal norm are integral,

$$s_\Lambda N_\Lambda \subseteq D_\Lambda, \qquad s_\alpha N_\alpha \subseteq D_\alpha,$$

where s_Λ, s_α are the products of the squared extremal Fischer norms over the host embeddings.

The adjoint of the actual corner branching operator is the lower-group map taking the lower highest vector to $c_\rho \xi_\rho$ at each host embedding, with ξ_ρ and c_ρ as in the terminal branching formula above. The map preserves integral polynomials on generators of D_α: the coefficients of the generic opposite unipotent orbit of its highest vector map to the corresponding coefficients of $c_\rho \xi_\rho$, an integral product of flag minors with integer scalar. These polynomials satisfy the prescribed upper column degree and column-addition equations, so they are in N_Λ. This proves that the adjoint sends $s_\alpha N_\alpha$ into N_Λ. Taking the integral bilinear dual shows that the branching operator sends D_Λ into $s_\alpha^{-1} D_\alpha$. The lower input to the pairing is in the contragredient polynomial coordinates. Reversing the same complementary-minor substitution in the lower frame gives an element of N_α. The paired scalar therefore has finite loss at most $\operatorname{ord}_v(s_\Lambda s_\alpha)$ on the unscaled corner lattices, in addition to the two input losses. Integral column equations, The actual branching operator, and The split corner and its finite frames.

At infinity the same last-coordinate parameter is 1. The normalized Fischer Cauchy–Schwarz inequality bounds the squared norm of the pairing by that of its two operands times $\prod_\rho c_\rho^2 k_\rho$. For both weights the shifted spreads at the corner are $O(b)$ in the fitted grid with gap parameter 1. Put N_{g} for the bound on grid ranks used here. The highest polynomial has degree $O(N_{\mathrm{g}} b)$ and the sum of absolute coefficients is at most $(N_{\mathrm{g}}!)^{O(b)}$. The monomial factorials and $N_{\mathrm{g}} = O(b)$ give $\log(s_\Lambda s_\alpha) = O(f N_{\mathrm{g}} b \log b)$. In the archimedean bound there are at most $(N_{\mathrm{g}} - 1)\sum_i (\lambda_i - \lambda_{i+1}) = O(N_{\mathrm{g}} b)$ factors in each c_ρ, each $O(b)$, and k_ρ is at most the upper Weyl dimension by the Gaussian frame calculation above. Thus the log squared norm multiplier in the corner is $O(f N_{\mathrm{g}} b \log b)$ as well. These are the corner bounds for the two fitted chains, with inputs in the fixed corner frames. The fitted corner weight range and Fischer pairing in the corner.

Bounds in the placed bottom coordinates. For the elliptic lattice $\mathbf{Z} + \mathbf{Z}\tau$, take $\Omega^{12} = \Delta_{\mathbf{Z}+\mathbf{Z}\tau}$. Let f be the odd reduced sigma function with $f'(0) = 1$. For $u = a\tau + b$ of order dividing N, set

$$S_u(\tau, z) = \eta(\tau)^{-1} \theta_1(\tau, z + u) e^{2\pi i a z + \pi i a u}.$$

The shifted f multiplied by Ω equals S_u times $e^{\pi z^2/(2 \operatorname{Im} \tau)}$ and a root of unity. This follows from the normalized odd theta formula and $\theta_1'(\tau, 0) = 2\pi\eta(\tau)^3$. The normalized order-$n$ derivatives of $S_u(\tau, z) e^{G_2(\tau) z^2/2}$ form a finite modular family of weight n. Under a modular change of basis the linear terms in the Jacobi Gaussian cancel the marking phases, and the remaining quadratic factor cancels the anomaly of G_2. Changes of lifts of the torsion labels give only finite unit phases. In every member's Tate expansion the z-frequencies are in $(2N)^{-1}\mathbf{Z}$. Each normalized derivative costs at most $v_p(2N)$, and the Gaussian derivatives at order $2j$ contribute $(2j-1)!!(G_2/(2\pi i)^2)^j$. Their total additional loss is at most $(n/2) v_p(12)$. After clearing these losses, the symmetric functions of the twelfth powers divided by Δ_{mod}^n have integral Laurent coefficients and are polynomials in j. Thus $j(\tau)$ integral gives the same loss at the CM point.

The reduced jet corrects the holomorphic jet by $\exp(-s_2 z^2/2)$ with $s_2 = G_2 - \pi/\operatorname{Im}\tau$. On the discriminant-one good Weierstrass model the integral CM eigenprojection used below gives $v_p(s_2/\Omega^2) \geq -\max(v_p(|\operatorname{disc} K_d|)/2, v_p(12))$. For a nonzero half-period the reciprocal of its constant $S_u(\tau, 0)$ loses at most $v_p(2)$ by the same weight-zero polynomial argument. Its reciprocal Tate product has integral coefficients apart from a possible $(1 - \zeta_2)^{-1}$ in its constant coefficient. The normalized even section is the half-period translate of f divided by its value at zero. Composing translations changes only a unit phase and the new point has order dividing $2N$. Consequently in the discriminant-one differential the lower valuation for its translated order-n derivative is

$$-n\big(v_p(N) + C v_p(6|\operatorname{disc} K_d|)\big) - v_p(2).$$

This bound holds at every algebraic embedding over p.

At a tested prime take representatives of the definite classes with local integral frames. Rational weak approximation is available from the Cayley chart of the unitary group, including for the finer open level

cosets. The lattice intersections with the standard frames then have index prime to p. Use the placed finite Fourier expansion for this choice of representatives. In dimension (a, d_{bot}) at modulus Q' it uses translates of order dividing MQ', with $p \nmid M$, and coefficients of loss at most $ad_{\mathrm{bot}}v_p(Q')$. This is the finite presentation to check in the moved frames. A harmonic letter at a host embedding uses the trace-dual basis of $\mathcal{O}_E$ over the fixed imaginary quadratic subfield. The idempotent for one embedding is $\sum_i \tau(e_i^\vee)e_i$; thus each letter loses at most $\delta_p^{\mathrm{let}} = v_p(\operatorname{disc} E)/[E : \mathbf{Q}]$ on the Galois host. Expand the minor polynomials and use Leibniz's rule on the fad_{bot} elliptic factors. With $D_{\mathrm{jet}} = a\sum_\tau n_\tau$ the coordinate loss is at most

$$D_{\mathrm{jet}}\big(\delta_p^{\mathrm{let}} + v_p(Q') + Cv_p(6|\operatorname{disc} K_d|)\big) + fad_{\mathrm{bot}}v_p(2) + ad_{\mathrm{bot}}v_p(Q').$$

Include also the common stabilizer denominator obtained by reduction modulo three. In the selected tuple at $v \mid p$ one has $Q' = QQ_n(\boldsymbol{\mu}(v))$ with $p \nmid Q_n(\boldsymbol{\mu}(v))$, so the fixed modulus Q suffices for the valuation. The sum weighted by $\log p$ uses the discriminant, fixed conductor and stabilizer budgets. It is $O(x^c b)$ in the admitted derivative range. Placed Fourier expansion and Determinants and trace letters.

For the fixed archimedean norms, use the same expansion at each conjugate. Its coefficient ℓ^1-norm is bounded by $Q^{ad_{\mathrm{bot}}/2}$ in the normalized finite test. At infinity the elliptic vacuum means multiply, and the Gaussian and Cauchy estimates below cost at most $O(D_{\mathrm{jet}}(\log(D_{\mathrm{jet}} + fad_{\mathrm{bot}} + 2) + \log^+(\operatorname{area}\Lambda)^{-1}))$ in the log squared norm. Products of minors have at most $(a!)^{\sum_\tau n_\tau}$ monomial terms. The number of minor tuples is at most $\binom{d_{\mathrm{bot}}}{a}^{\sum_\tau n_\tau}$. The class weights sum to the definite mass, with a finite level index at most $Q^{d_{\mathrm{bot}}^2+4}$. The mass formula bounds its log above by a fixed polynomial in the degrees, ranks, and logarithmic discriminant. For a conjugate of a coordinate times its conjugate, Cauchy–Schwarz bounds the class sum by the geometric mean of the two corresponding conjugate squared norms. Averaging and applying the CM elliptic means below gives the $O(x^c b)$ upper bound for the fixed base norms. This uses the placed finite presentation and its conjugate kernel in the same class frames at every embedding. Conjugate coordinates and Mean archimedean bound.

Placement of the definite coordinates. Fix a bottom cell of Section 2.3.22 and write t for its auxiliary tuple. Use the standard Hermitian matrices of ranks $a, a', d = a + a'$, with the curve of lattice $\delta\mathcal{O}_{K_d}$ and the fixed auxiliary half-period. For each tuple t, write $\mathfrak{m} = \mathfrak{m}(t)$ for the selected conductor ideal in E. It is prime to its conjugate and to the half-different and 2. At each exceptional place it selects one of the two split factors. Use e for the corresponding idempotent, so the opposite polarization uses $\bar{e}$. The rank-a arithmetic pair uses $(\mathfrak{m}, 0)$ and the rank-a' pair uses $(\bar{\mathfrak{m}}, a)$. The second coordinate denotes the starting row of the identity block of the translation. The target level of the two pairs is the same block subgroup J with lower-left (a', a) block in $\mathfrak{m}$. The finite characters of the two pairs are prescribed by (β^{-1}, ν_t^{-1}) and (α_t, ν_t^{-1}), respectively, in the arithmetic action. The analytic norm calculation for the larger pair takes the target splitting character $\chi_B = \nu_t$ by the inverse in the arithmetic-to-analytic character convention. The standard factor for doubling uses χ_B^{-1}. Since ν_t is trivial on the base ideles, $\nu_t^{-1}(z) = \nu_t(\bar{z})$, so the primitive $L(\nu_t^{-1}, \cdot)$ equals $L(\nu_t, \cdot)$ by permuting the places of E under conjugation. Ranks, characters, translates and target levels · Character in the norm formula.

Here is the finite comparison, with both normalizations. For an ideal $\mathfrak{a}$ in the chosen polarization and a source of rank b, put $n = (\mathrm{N}\mathfrak{a})^{bd}$. Let L_0 be the standard theta lattice in the matrix space and L_1 the lattice obtained by replacing the target lattice by $(\mathfrak{a}/\bar{\mathfrak{a}})^d$. The two indices over $L_0 \cap L_1$ are n, as is seen entry by entry at the selected split places and their conjugates. The indicator of the zero coset modulo $\mathfrak{a}$ is the average $n^{-1}\sum_\xi \xi$ of the additive characters of the matrix residue group. In the lifted theta model the characters are the Heisenberg translations in the dual finite group. Indeed multiplication by such a character and translation by its dual element satisfy the same commutation law $M_\xi T_z = \xi(z)T_z M_\xi$, and both fixed vacuum sections correspond to the integral indicator in the Schrödinger coordinates. Use the trivialization on the integral isotropic subgroups. Locally the character of the prescribed additive pairing vanishes there by the half-different condition. The character average therefore reads as $n^{-1/2}I_{01}e_0$, since

the common-cover intertwiner I_{01} already includes $n^{-1/2}$. Replacing the zero coset by the identity-block coset applies the same Heisenberg translation in both descriptions. The scalars introduced by choosing torsion lifts are roots of unity. Normalized arithmetic coset vector · Character orthogonality with the two square roots · The conjugate polarization.

In the larger leg the finite section is supported on integral d-by-a matrices whose lower block is zero modulo $\mathfrak{m}$ and whose upper block is invertible there. Write $Y = (Y_1, Y_2)^t$ in row blocks. The congruences are $Y_1 \bmod \mathfrak{m} \in \mathrm{GL}_a(\mathcal{O}_E/\mathfrak{m})$ and $Y_2 = 0 \bmod \mathfrak{m}$. The section is weighted by the evaluated finite unit determinant character on Y_1. Carrying the identity-block coset by representatives of this finite general linear group partitions the support exactly. The coefficient on the coset of r equals the coefficient at the identity-block coset times the determinant character of r in the right substitution $Y \mapsto Yr^{-1}$. The Weil action on the same support has this same factor. The prescribed negative leg conjugates the reading and hence conjugates these coefficients on both sides. In the smaller leg the section uses the dual matrix $\xi_0 Y^*$, in the opposite polarization, with its first a columns in $\mathfrak{m}$ and an invertible right block on the support. Passing back to the d-by-a' matrix gives the $(\bar{\mathfrak{m}}, a)$ cosets above. Here ξ_0 is a unit at these places. In each leg the initial coefficient is the prescribed character at an integral unit idele. The conductor kills principal congruence, so the coefficient is independent of the representative of the coset. All varying factors in these comparisons, apart from the common $n^{-1/2}$ already present in each arithmetic coset vector, are values of the finite-unit character or Heisenberg phases. In particular their ratios are algebraic units at every tested prime. The character and the exact polarized cosets specify the relative phases of the summands. Both determinant-weighted sections · Larger leg finite reading · Smaller leg finite reading.

In changing the target frame to the standard identity frame, take the split local stabilizer carrying the old basis to the new one at the selected factor and its adjoint inverse at the other. The pullback takes each of the two prescribed block supports to the re-framed one. The two determinants upon substitution contribute the same target action in both legs and possibly differing finite-unit character values. Thus I_t and I'_t acquire the same target right translate and their scalar of proportionality acquires only a root of unity. The exact original proportionality is recovered by including this factor. Common target translation and quotient of finite-unit phases.

The scalar comparing normalized vacua. There is one analytic vacuum on each of the two standard matrix spaces before imposing the tuple conductor. In theta coordinates it spans the same one-dimensional line as the normalized arithmetic vacuum. The arithmetic normalization at a place uses an orthonormal frame of the K_d-Hermitian matrix space, locally a basis of the underlying integral lattice, and a scalar unit there. For clarity, the change between two such normalizations has a numerical test independent of algebraicity of their complex norm. Let G be the unitary matrix of transition between the frames over K_d. For the two pulled back theta products in the frames, the proportionality factor c after normalized common-cover transport has $|c| = 1$ and $c^2 = \zeta \det G$ for a root of unity ζ, in the orientation from the first to the second frame.

One way to calculate the last assertion is to use the reproducing kernel of the polarized torus, obtained by summing the Bargmann Gaussian over its lattice Λ_0 with the prescribed semicharacter α. Write A for the complex coordinate change on points (matrix G^*). Assume first that $1 - A$ is invertible. On the common period lattice Λ' the semicharacter agrees under A. In the Gaussian trace on the graph of A, one unfolds on the cosets of $(1 - A)\Lambda'$ in Λ_0. For the term with representative x, translate by $y = (1-A)^{-1}x$. Completing the square, the integral of this term is the fixed real Gaussian integral divided by $\det(1 - A)$, times $\alpha(x) \exp(-\pi i \operatorname{Im} H(y, x))$, where H is the normalized positive Hermitian form. The phases give a finite quadratic Gauss sum. Its nonzero phase is a root of unity, by reducing by its radical and evaluating the primary quadratic factors. The same trace using the common-cover average for the frame-changed product keeps only the zero coset, by orthogonality of the translated sections. It is c times a positive real factor. Hence $c \det(1 - G^*)/|\det(1 - G^*)|$ is a root of unity. The identity $\overline{\det(1 - G^*)} = (-1)^{\dim G} \det G \det(1 - G^*)$ now gives $c^2 = \zeta \det G$. If $1 - A$ is singular, insert an intermediate frame

obtained by a norm-one scalar in K_d avoiding the two finite spectra. The product of the two normalized transports differs from the direct transport by the root of unity of the finite quadratic pairing on the common cover. The determinants multiply, which gives the same assertion. Unfolding the Gaussian trace · Reading the same trace in the theta line · General frame transition.

At a finite embedding over p, each of two normalizations may use the stated locally integral frame. The transition and its inverse preserve that same localized matrix lattice, so $\det G$ is a p-unit in the specified embedding of K_d. Then $c^2 = \zeta \det G$ implies that c is algebraic and a p-unit. The remaining local normalization scalars and the common-cover composition phase are units as well. Consequently any of the permitted arithmetic vacua differs from a single chosen one by an algebraic unit at the required places. Take the analytic normalization constant once relative to this chosen vacuum. This proves uniformity across t, even if the arithmetic vacuum presentations depend on t or on the place. Local frames in the normalized vacuum ratio.

Class sums and the Petersson normalization. The differentiated kernels at finite source and target representatives can now be read on the arithmetic theta functions. The finite Weil transport is the same for both readings. This follows by carrying the fixed vacuum with the common-cover averaging operators, then using covariance for each torsion translate in the displayed finite section. A phase in the lift of a finite isometry is fixed by the prescribed determinant characters. At τ, changing from the Fock creation operator to the curve derivative multiplies the latter by $c_\tau = \delta/\sqrt{2\pi\xi_R}$ in the standard frames, with $\xi_R = |\tau(\xi_0)|$. Indeed the Gaussian restriction of the vacuum has annihilation relation, and in the holomorphic chart $a^+(Z)$ is $(2\pi\xi_R)^{-1/2}\partial_Z$ at the origin. The further substitution $z = \delta Z$ gives the stated multiplier. Applying a minor product multiplies the arithmetic jet by $C_b = \prod_\tau c_\tau^{b\,n_\tau}$ for source rank b. The reflected larger leg conjugates this formula. The twist t has trivial infinite component, so the degrees n_τ are fixed in the cell, and the complementary products have the same n_τ on the two legs. The compact source real action on the minors is the determinant to these powers, which is canceled by the vacuum exponent and splitting character. Negative source and conjugate jets · Fock-to-curve period multiplier · Same complementary degrees.

The source maximal compact groups at finite places fix these two sections in their respective split models. On a selected factor the determinant factor of the Schrödinger action on the section cancels the determinant weight just computed. Off the selected factors the sections are the spherical vacua and the same claim holds with unramified characters. The two definite class sum formulas therefore use the maximal source class groups with the original stabilizer weights. Integration includes the fixed source Haar volumes, independent of t. It gives

$$\overline{p_t(T,h)} = \Omega_A u_A(t)\, I_t(T,h), \qquad I'_t(T',h) = \Omega_{A'} u_{A'}(t)\, p'_t(T',h),$$

for all the finite target representatives and all complementary minor tests, with one nonzero $\Omega_A, \Omega_{A'}$ each and algebraic units $u_A(t), u_{A'}(t)$ at the places under consideration. The constants include the indicated Fock-to-curve periods, the choice of the original standard vacua and the Haar volumes. When the original target frame is used, include the finite-unit factor from the common right translate in their product. The common quotient comparison and the signed complementary minors then give the two inverse coordinate relations with $B_t = \kappa u_t c_0(t)$. In particular u_t has valuation zero at every w with $w \nmid Q_n(t)$. Smaller source class formula · Larger source class formula · Both integral coordinate relations.

For the norm use the positive definite analytic realization reading the arithmetic rank-a vector itself, with source maximal level of volume one. Write $F_{t,T}$ for the theta lift of the constant source vector with the same finite section and the minor Fock vector T. The same pointwise reading and integration give $F_{t,T}(h) = C_a p_t(T,h)$ at each target class representative. In this realization there is no extra vacuum ratio, since the analytic vacuum is chosen by the arithmetic one. The target maximal compact has volume one for the reference measure. At level J the measure with level volume one is $[U_d(\widehat{\mathcal{O}}_F) : J]$ times that measure. The squared sum over all T is invariant under the real target unitary group: on the tuple of

minors the action is the tensor of the exterior powers of a unitary matrix with a determinant phase. It is also invariant under J, since the same finite eigencharacter multiplies every minor coefficient. Thus the sum can be integrated by the same stabilizer-weighted target class formula defining the arithmetic norm, giving exactly

$$\|P_t\|^2 = |C_a|^{-2}\,[U_d(\widehat{\mathcal{O}}_F) \,:\, J] \sum_T \|F_{t,T}\|^2_{\mathrm{Pet,max}}.$$

The period correction is $|C_a|^{-2} = \prod_\tau (\pi/\operatorname{area}\Lambda)^{an_\tau}$, using the standard scale $|c_\tau|^2 = \operatorname{area}\Lambda/\pi$. Exact norm class sum · Period correction · Passage to maximal target volume.

In the theta-space norm the distinct congruence cosets in the matrix support are orthogonal, since their indicator functions on the split factor have disjoint supports. Each coset vector has squared norm $1/n$ of the normalized vacuum squared norm, and the finite unit weights have complex absolute value one. At depth e_v the number of cosets is $q_v^{e_v a^2}\prod_{j=1}^{a}(1-q_v^{-j})$. Thus the local factor is $q_v^{-e_v aa'}\prod_{j=1}^{a}(1-q_v^{-j})$. Section 2.3.22 gives the compact support and primitive L-function on these block conditions.

Finally the target index at v is the number of rank-a summands in the free residue module:

$$q_v^{e_v aa'}\,\frac{\prod_{j=1}^{d}(1-q_v^{-j})}{\prod_{j=1}^{a}(1-q_v^{-j})\prod_{j=1}^{a'}(1-q_v^{-j})}.$$

Multiplying by the local section norm just computed gives $\prod_{j=a'+1}^{d}(1-q_v^{-j})$, the ratio of the full and the partial doubling factors in their required orientation. Consequently dividing by the partial factor and multiplying by the level index and test norm leaves exactly the reciprocal of the full factor. The source and target masses, the normalized vacuum norm, the period correction and the sum of Fock norms are fixed within the cell. Applying convergent doubling with the source exponent string gives the claimed one nonzero common norm scalar times $\prod_{s=1-e_0}^{h_0} L(\nu_t, s)$. Petersson norm in convergent doubling · Level cancellation in the arithmetic norm.

The complementary minor matrices. Let $a + a' = d$ in the bottom triangle. For a unitary matrix written in column blocks $U = [Z\ Z']$ the complementary minor identity is

$$\det Z_I = \varepsilon_I \det(U)\,\overline{\det Z'_{I^c}}.$$

Products of the minors obey the identity with $\det(U)^{n_\tau}$ and the product sign. Thus passage to the complementary products is the signed permutation of the minor tuples. It induces the intertwiner on the corresponding rectangular Schur spaces. For instance evaluate any relation among the products on the first block of $\bar{U}$ and use the formula. The resulting relation on the complementary block holds on all unitary frames and hence as a polynomial identity. The right determinant character singles out the rectangular highest weight in the polynomial model, and the other multi-degrees map to zero. The common real quotient, with the determinant twist in its complementary action, then makes the two projections proportional by Schur's lemma. The induced two matrices on the actual minor coordinates are a signed permutation and its inverse. This construction requires the compact bottom real frames and the single common quotient of the two matched theta problems in Section 2.3.22. Complementary projections · Integral permutation in the triangle.

The bottom norms in the square ratio. Write P_μ, P'_μ for the two bottom vectors at the tuple selected at v, and P_0, P'_0 for the base pair in the same cell. Let L_μ, L_0 be their placed Hecke functions and $\rho_s = L_\mu(s)/L_0(s)$. With the same class norm and lowest-type normalization, the bottom norm formula has the common scalar A_{cell}

$$\|P_\mu\|^2 = A_{\mathrm{cell}} \prod_{1-e_0\le s\le h_0} L_\mu(s), \qquad \|P_0\|^2 = A_{\mathrm{cell}} \prod_{1-e_0\le s\le h_0} L_0(s).$$

Thus its value as a field element is $A\prod_s \rho_s$ for the nonzero base norm A. Let A' be the other base norm, that of P_0'. The range is $s = h_0 - k$ for $0 \le k < a$, as in Section 2.3.22. The placed formula gives the scalar in exactly this coordinate norm. The common bottom norm scalar and Use for the matched bottom vectors.

Let $\gamma \in K^\times$ be the ratio of bottom proportionality constants on the placed triangle, as in (19). The two frame identities comparing the conjugate coordinates have the form $\overline{p_\mu} = B_\mu M_\mu p_\mu'$ and $B_0 p_0' = M_0' \overline{p_0}$, with integral matrices and B_μ/B_0 equal to γ times an algebraic unit at v. In the class sum divide the first norm by B_μ and multiply the second norm $\|P_0'\|^2$ by B_0. The resulting terms are products of integral matrix entries, two coordinates and a reciprocal stabilizer. Use class representatives integral at v and the uniform coordinate and stabilizer lower bound $-\delta_v$. The finite level exponents of both pairs are within the local bound for that comparison. The same bounds hold for the conjugated factors at every embedding over the rational prime below v. Hence each normalized norm has valuation at least $-3\delta_v$. Multiplying them and using the unit relating B_μ/B_0 to γ gives

$$\operatorname{ord}_v \gamma \le \operatorname{ord}_v A + \operatorname{ord}_v A' + 6\delta_v + \sum_{1-e_0 \le s \le h_0} \operatorname{ord}_v \rho_s.$$

The integral comparison matrices and The termwise valuation.

In the square identity the two tower products have ranges $[e_0 + 1, e + q_*]$ and $[e + 1, e + q_*]$, with $e_0 + p_* = e + q_*$, $e_0 \le h_0$. Moving them across the identity and including the bottom norm gives coefficient

$$\mathbf{1}_{[1-e_0,h_0]}(s) + \mathbf{1}_{[e_0+1,e+q_*]}(s) - \mathbf{1}_{[e+1,e+q_*]}(s)$$

on $\operatorname{ord}_v \rho_s$. If $e \le e_0$ the last two terms remove only $[e+1, e_0]$ from the bottom interval. If $e \ge e_0$ they add $[e_0 + 1, e]$. All coefficients are nonnegative. The away-unit comparison at the tested prime replaces $s \le 0$ by $1 - s$ in the valuations. The reflected indices lie in $[1, e_0]$, disjoint from the possible addition. Thus the positive indices have multiplicity at most two and are at most $h_{\mathrm{cell}} = h_0 + p_* + q_*$. For the fixed toric quantities Z_s, Z_s^0, Z_s^1, the reference depletion index $c(v) \in \{0, 1\}$ and the common normalization exponent g_{sv}, put

$$T_{s,v} = |\operatorname{ord}_v Z_s + g_{sv}| + |\operatorname{ord}_v Z_s^{c(v)} + g_{sv}|.$$

The selected toric bound (155) at the actual ratio gives $\operatorname{ord}_v \rho_s \le T_{s,v} + e_v/p^2$ for $v \mid p$. The resulting upper valuation of the cell ratio is at most

$$\operatorname{ord}_v A + \operatorname{ord}_v A' + 2\sum_s T_{s,v} + 2h_{\mathrm{cell}} e_v/p^2 + 6\delta_v + 2m \sum_{q \in S_0} c_q \operatorname{ord}_v(q).$$

The last sum uses the base conductor budget. Include the supports of A, A' when summing the local inequality. At places assigned the base tuple the square ratio is one, and the same bottom estimate gives $0 \le \operatorname{ord}_v A + \operatorname{ord}_v A' + 6\delta_v$. The product formula sums the signed bottom terms to the archimedean mean of $\log|\sigma(AA')|$. The fixed-height bound for each depletion controls the weighted sum of $T_{s,v}$. The sums of $(\log p)/p^2$ and of the conductor terms use the given support and weights. With the $O(x^c b)$ bounds on these data and the uniform coordinate allowance, the polynomial number of indices and cells gives the required upper bound on $\sum_v \lambda_v \operatorname{ord}_v q_{\boldsymbol{\mu}(v)}$. Square identity and signed bottom terms and Sum at the selected tuples.

The toric denominator N_s is used for the integral polynomials before the character is selected. For its partial-zeta part, smoothing at a positive principal ideal (a) prime to p multiplies the depleted class vector of negative special values by $1 - a^{2sf}$, which gives an integral vector by Deligne–Ribet. The omitted Euler operators have form $1 - (\mathrm{N}\mathfrak{p})^{2s-1} P_{\mathfrak{p}}$ with $P_{\mathfrak{p}}$ a class permutation and $\mathfrak{p} \mid p$. They are invertible over $\mathbf{Z}_p$ by their geometric series, so the smoothing also clears the undepleted vector. If p is odd, choose a primitive modulo p^2. For $J = 2sf$ this costs zero unless $p - 1 \mid J$, in which case the valuation is $1 + \operatorname{ord}_p(J)$. At two, $a = 3$ gives $2 + \operatorname{ord}_2(J)$, with the power of two in the normalization adding at most f. The total

weighted cost is $O(J\log(J+1)+f)$. The remaining terms of the prescribed profile are supported on the base conductor and the discriminant, with coefficients bounded by a fixed multiple of one plus the weight size. They have logarithmic cost $O(x^c b)$. The integrality for the local toric value can be seen directly in its normalized jet. Let $F = E^+$, $f = [F : \mathbf{Q}]$, with parallel weight $2s$ and total derivative order h_{jet}, so that the total weight is $w_{\mathrm{tor}} = 2sf + 2h_{\mathrm{jet}}$. Write $\delta = v_\ell(|\operatorname{disc} K_d|)$, $a = \max_{v|\ell} \operatorname{ord}_v(\mathfrak{f})$ for the base conductor $\mathfrak{f}$, and put $c = a$ on its support and $c = v_\ell(\operatorname{disc} E)/[E : \mathbf{Q}]$ elsewhere. A class representative can be taken equal to one over ℓ by weak approximation, so its finite character factor is a unit. Indeed a power of the representative is principal modulo order units and the generator is an ℓ-unit. Let e be the maximum valuation loss of the normalized partial zeta constants just bounded.

In coordinates $1, \theta$ for $\mathcal{O}_{K_d}$, the scaled lattice at every cusp has integral coordinates of valuation at least $\delta/2$ and contains $\ell^{\delta+a}\mathcal{O}_{F,\ell}^2$. This follows from $(\theta-\bar\theta)b_1 = z-\bar z$ and its analogue for the constant coordinate, after scaling by the discriminant and an ℓ-unit integer to clear the other primes. The Lipschitz expansion in the second coordinate gives a trace dual ideal, so the overall coefficient loss is at most $e + sv_\ell(\operatorname{disc} F) + 2sf(\delta+a)$. Each frequency has loss at most $2c+2a+\delta/2$ since $c \ge v_\ell(\operatorname{disc} F)/f$. The h_{jet} derivatives add at most that loss per derivative. The holomorphic torsion correction $(G_2 + \wp(P))/(2\pi i)^2$ at prime-to-ℓ order is integral at every cusp by its Tate expansion. Its transformation cancels the inhomogeneous term of the derivatives, so the jets form a finite level family. After clearing the coefficient bound, the elementary symmetric functions of the twelfth powers divided by $\Delta^{w_{\mathrm{tor}}}$ are polynomials in j integral at ℓ, by their Laurent expansions. Evaluating at integral $j(\tau)$ proves the bound for the normalized holomorphic jets, also at weight two with the convergent Fourier expansion of the regularized series.

For the CM correction, scale the elliptic discriminant to one and take a good integral Weierstrass model after finite extension. If η is its invariant differential, the class of $X\eta$ is integral since its principal part at infinity is $u^{-2}du = d(-u^{-1})$. The conjugate CM eigenclass differs from it by $c_1\eta/(\theta-\bar\theta)$ with c_1 integral, by projecting the integral de Rham lattice along the Hodge line. Adding the value at prime-to-ℓ torsion gives valuation at least $-\delta/2$. The correction therefore fits the same derivative budget. Undoing the dilation recovers $2sf\delta$ in valuation, and converting to the conjugate derivative recovers $h_{\mathrm{jet}}(\delta - \mathbf{1}_{\ell=d})/2$. Consequently the normalized value loses at most

$$e + sv_\ell(\operatorname{disc} F) + w_{\mathrm{tor}}a + h_{\mathrm{jet}}(2c + \mathbf{1}_{\ell=d}/2).$$

This fits the two prescribed conductor and discriminant cases, with just e at an unramified place prime to the conductor. Summing class terms with their unit factors preserves the bound. Taking twelfth powers gives the local integrality used for the two fixed depletions. All-cusp coefficient loss, CM specialization, and The normalized derivative. Partial zeta denominator and the toric profile.

Toric specialization on the same ray quotient. At one stage and side, the toric label is the character $\mu_{0,j}\,\mathrm{del}_\ell\,\chi_{m-i-1}^{-1}$. The untwisted grid character has trivial finite ray twist, so the label is the same fixed depletion used above. For the auxiliary set Q, let Δ_Q be the finite quotient used for the Selmer test. Reciprocity gives a map $\pi_Q : \mathrm{Pic}(\mathcal{O}_{\mathfrak{c}}) \to \Delta_Q$ for $\mathfrak{c} = \mathfrak{f}\prod_{q\in\mathrm{below}(Q)} q$. At a split q the two tame characters are inverse on the diagonal units of the order and trivial on the principal units. Away from Q they are unramified. Thus the map kills order units, including those over ℓ, and takes a Frobenius prime to the inverse arithmetic Frobenius class.

Write $T_Q(C)$ for the weighted value in Picard class C and $T_0 = \sum_C T_\varnothing(C)$. The latter is nonzero by the base Hecke L-value, with the nontrivial auxiliary character in the case $s = 1$. The coefficient at $\delta \in \Delta_Q$ is

$$A_Q(\delta) = T_0^{-1} \sum_{\pi_Q(C^{-1})=\delta} T_Q(C).$$

The CM isogeny for a unit idele u over ℓ changes the lattice of class C to that of $C[u_E]^{-1}$, where u_E is the lift to E. The holomorphic jet transforms by that change of level and torsion labels and differential.

Its coefficients transform this way by the q-expansion principle, and the weight-two correction uses the same conjugate CM eigenline. The period, area and fixed character factors occur also in T_0. Therefore an inertia automorphism g transports the ratio for a Picard character ξ to the ratio for $g\xi$, times $(g\xi)([u_E])^{-1}$. For ξ coming from Δ_Q this factor is one. Fourier inversion gives descent of $A_Q(\delta)$ to the prescribed local coefficient field, which contains the host embeddings and the label's fixed finite values. The preceding denominator bound clears these coefficients uniformly in Q, since T_0 is fixed for this label. CM transport of the jets and Descent on the ray quotient.

If an auxiliary split pair becomes unramified under a specialization, compare the order with and without the pair. At that prime the order consists of pairs congruent modulo q. There is one unit class and for each pair of positive valuations there are $\mathrm{N}q - 1$ leading-unit ratios. A ramified tame character sums to zero on the latter. For an unramified character with value Y at one prime, put $x_q = (\mathrm{N}q)^{-s}$. The local factor in the sum for the smaller order changes under addition of the pair by

$$e_q(Y) = (1 - Yx_q)(1 - Y^{-1}x_q) + (\mathrm{N}q - 1)x_q^2.$$

Indeed the two independent positive valuations give geometric series with ratios Yx_q, $Y^{-1}x_q$, each starting at degree one. Dividing the resulting order factor by the unrestricted Euler factor gives $e_q(Y)$. The fixed base conductor has the exact character exponent, so nonunit ratio cosets there cancel against its own character. The class regrouping is valid by absolute convergence to the right of one and extends to $s = 1$ by the Hecke L-functions. The local identity at all characters gives the deletion identity in the finite group ring by Fourier inversion. The order Euler factors and Deletion in the ray group ring.

The toric test on the paired ray coordinates. Fix a record, a cell (i, j), and a marked coefficient place $w \mid \ell$. Set $a = m - i - 1$ with zero-based stage indices and

$$\vartheta = \mu_{0,j}\chi_a^{-1}, \qquad \chi = \vartheta\,\mathrm{del}_\ell\,.$$

The reference character of the tuple is trivial. The character of the marked tuple is $\mathrm{del}_\ell\,\psi_w^{\mathrm{c}}$, where ψ_w^{c} is the complex reading of the chosen moving character. Thus the cell characters are ϑ at the base, χ at the fixed depletion, and $\chi\psi_w^{\mathrm{c}}$ at the marked tuple. The depletion is supported at p_{de} if $\ell \neq p_{\mathrm{de}}$ and at p'_{de} otherwise. It has already been included in χ. The character identities on the chosen tuples.

Let $\mathfrak{f}$ be the base ideal of the real field. The auxiliary χ_a has exact conductor exponent two over the four rational type and scalar-character primes and exponent one over the extra base split place. The numerator characters $\mu_{0,j}$, del_ℓ and ψ_w^{c} are tame at the former places and unramified at the latter. They cannot change these exponents. The paired moving set Q is split and avoids the base conductor, the fixed levels and the places above ℓ. It is the full conductor support of the moving character, so the character $\chi\psi_w^{\mathrm{c}}$ has exact exponent one above each pair in Q. The same support holds for every primitive power of the moving character. Each coordinate is a ray character minus its conjugate. Its value on any real idele is zero additively, since the reciprocity image of such an idele is fixed under conjugation. Thus the moving character is trivial on real ideles, also when $\ell = 2$. The base conductor exponents and The exact moving support.

For the order of modulus $\mathfrak{f}\prod_{q\in\mathrm{below}(Q)} q$, let $T_Q(\eta) = \sum_C \eta(C)^{-1}T_Q(C)$ be the weighted class sum with a Picard character η. Write T_0 for the untwisted class sum at the base order. The local Euler calculation gives, for $s \geq 1$,

$$T_Q(\eta) = L_E(s, \chi\eta) \prod_{\substack{q\in\mathrm{below}(Q)\\ (\chi\eta)_P \text{ unramified},\ P\mid q}} e_q\big((\chi\eta)_P(\varpi_P)\big), \qquad T_0 = L_E(s,\chi) \neq 0,$$

where P is one of the two primes above q and e_q is the split order factor displayed above. The character in this formula is the idele character obtained from η. The identity follows by grouping integral ideles by

order class in the half-plane of absolute convergence and continuing to s. The global units of the order are the real units. For a unit v, the ratio $v/\bar{v}$ is a root of unity by the absolute values of its conjugates. This ratio is congruent to one at the odd unramified extra base prime by the order congruence. The principal units there have no nontrivial torsion, so $v = \bar{v}$. Thus the grouping uses the same unit multiplicity for both orders. At a prime of $\mathfrak{f}$, positive-valuation cosets cancel against the nontrivial character of the ratio of leading units. In fact, at exponent h, valuations below h must agree. If their common value is $n > 0$, the leading-unit ratio runs through $(1 + \mathfrak{q}^{h-n})/(1 + \mathfrak{q}^{h})$. At valuations at least h it runs through the full unit ratio group modulo $1 + \mathfrak{q}^{h}$. The exact exponent h makes each of these character sums zero. The remaining term is the unique unit class, as in the primitive Euler product at a ramified place. The primitive product and the unramified order factors.

In the paired prime construction the first Frobenius readings agree modulo the kernels of the residual toric characters with a fixed tame inertia element at the scalar-character prime of del_{ℓ}. This element can be taken in the prime-to-ℓ part of tame inertia with depletion value of order three, or of order five when $\ell = 3$. The base $\mu_{0,j}$ is unramified there and the auxiliary χ_a is trivial on the Teichmüller units there. Thus the residual toric value on the inertia element differs from one, for either orientation. These kernels can be included for the finite list of labels in the shared Frobenius prescription. Since $\mathrm{N}q \equiv 1 \pmod{\ell}$, the reduction of $e_q(Y)$ at a first pair with all coordinates trivial is $(1 - \overline{Y})(1 - \overline{Y}^{-1}) \neq 0$. Deleting those pairs keeps the valuation of the empty-order test unchanged.

At the u-th second pair set the second coordinates to one multiplicatively. In the group-ring Euler factor the remaining Frobenius value is $1+X_u$ or its inverse, by the elementary first coordinate prescription. With $U = 1 + X_u$, the reduction of the factor is $(1 - \alpha U^{\pm 1})(1 - \alpha^{-1} U^{\mp 1})$, where α is the residual value of the fixed toric character at the second Frobenius. This is a unit at $X_u = 0$ if $\alpha \neq 1$ and has order exactly two if $\alpha = 1$, including at $\ell = 2$. The inverse powers are taken modulo $U^{\ell^M} - 1$. For large M that relation does not affect degree two in the residue ring. The product has a unit leading monomial of degree at most twice the number of second pairs. The integral toric polynomials therefore satisfy the same bounded-order hypothesis as the lattice determinants. The first-pair computation identifies their reference valuation with the cleared empty-order sum. The first residual toric value and The two deletion calculations.

Apply simultaneous specialization to the finite signature list, these toric polynomials and the projected stage determinants on the same prime array. Both sides read powers of the same root, with the side + reading the character and the side − its inverse. Use total tolerance $\epsilon_{\ell} = \ell^{-2}$ for the signature losses beyond their cyclotomic baselines, total tolerance ϵ_{ℓ} for the stage and side lattice losses, ϵ_{ℓ} for each toric test, and ϵ_{ℓ} for the sum of the tame costs weighted by t. These follow by taking each polynomial and tame tolerance at most $\epsilon_{\ell}/(\#\{\text{signatures}\} + 2m + 4Rt + 1)$ for R pairs of each kind. All four conclusions hold for every power prime to ℓ of the chosen character, with exact conductor support and trivial decomposition over ℓ. At a given index the finitely many marked primes of the base coefficient field can be grouped into conjugate pairs. Choose for each such pair once. In successive choices include all primes over the rational supports already used in the avoidance set. This gives disjoint rational supports for the different pairs. Extending a base place to the larger coefficient field changes the root reading by a unit power and can change the side. The estimates persist on these readings. One specialization and its error bounds and The paired place choices.

Reading the primitive ratio at the marked place. On an installed record, let ψ'_w be the ray reading from which the signed local Selmer and stage lattice tests are formed. The datum shared by the two coefficient places of a conjugate pair records $(\psi'_w)^{\mathrm{sgn}(w)}$, with the chosen place orientation $\mathrm{sgn}(w) \in \{1, -1\}$. Each test then uses the orientation of its signature or stage. Both signs are absorbed by the estimates for

all primitive powers. Write Art for its map on Galois elements and π_Q for its map from the Picard group. At an unramified prime P away from the finite exceptional set they satisfy

$$\pi_Q([\varpi_P]) = \mathrm{Art}(\mathrm{Frob}_P)^{-1}, \qquad \iota_w\big(\psi_w^{\mathrm{c}}(P)^{-1}\big) = \psi'_w\big(\mathrm{Art}(\mathrm{Frob}_P)\big),$$

where the second equality uses the arithmetic coefficient embedding and takes place after adjoining the torsion roots. The finite idelic part of ψ_w^{c} descends to a Picard character ψ_c. Ideals outside the finite exceptional set generate the finite Picard group. Thus the first identity and the second identity match the ray character with ψ_c there and then on the whole group. They require the Frobenius reading for the actual tuple on the record. The arithmetic reading on Frobenius.

Let I be the CM reading in a common valued extension of the local fraction fields and $I_+ = I$, $I_- = I \circ \mathrm{conj}$. Comparison with the arithmetic coefficient embedding on the ℓ-power roots gives an exponent c prime to ℓ such that $I(\psi_c(C)^c) = \psi'_w(\pi_Q(C))$. Take any primitive power a'. On the side b group the toric coefficients by $\pi_Q(C^{-1})$, and denote evaluation of these coefficients by $D_{Q,b}$. The empty-order evaluation is $D_{\varnothing,b}(1)$. In the same field one has

$$\frac{D_{Q,b}\big((\psi'_w)^{b\,a'}\big)}{D_{\varnothing,b}(1)} = I_b\left(\frac{T_Q(\psi_c^{c\,a'})}{T_0}\right).$$

Here the two sides of a toric label have the same class weighting under their respective readings. The inverse in the fibre supplies the inverse of the Picard character in the sum. Complex conjugation in I_- inverts its roots of unity, which gives the negative orientation. Any scalar used to normalize the class coefficients is common to the two levels $Q, \varnothing$ and cancels. Comparison on the cyclic roots and The stage, weight and signed class evaluation.

Use the side whose place lies under w, taking the conjugate reading if needed. Over the common small coefficient field, Galois transitivity on the extensions of that place gives an automorphism g aligning the CM reading with the valuation at w on the ratio. The small field contains the host embeddings and the fixed finite values of χ. Thus g fixes the labels, the fixed character weights and the finite idelic moduli. Choose a reciprocity idele of the quadratic CM field for its action on the normalized CM values. The same change of CM lattice and level as in the inertia descent reindexes T_Q and T_0. Its period, modulus and fixed character factor cancels, leaving a Picard character value as multiplier of $T_Q(g\psi_c)/T_0$. That multiplier is a root of unity. The action on the cyclic character has the form $g\psi_c = \psi_c^u$ with u prime to ℓ. Taking $a' \equiv c^{-1}u$ modulo the order of ψ_c reads the actual valuation using the displayed evaluation quotient. The embeddings can be fixed for the entire label list at w and then c, u, a' depend only on the place and the chosen ray reading. For the trivial moving character one can take $a' = 1$. The toric estimate holds on every primitive reading with the same conductor, Selmer and stage lattice bounds. Thus the local valuation can be read without altering the twist installed in the tuple. Cancellation under CM reciprocity and The actual place and the unit factor.

All moving primes are ramified for this character. The order product in the formula for $T_Q(\psi_c)$ is empty. At the fixed conductor the primitive factors are one, by the exact exponents above. At other finite places, including bad host primes and word-level primes outside the conductor, retain the usual primitive Euler factors. At primes over ℓ the moving character is trivial on the decomposition group, so the selected and depleted local factors agree there. The partial numerator of the one-step theta comparison omits the fixed conductor. The bottom norm comparison omits the fixed conductor and the tuple's ramified places. At each omitted place the primitive factor on the corresponding numerator is one. For the bottom comparison the inverse character gives the same global L-function, since conjugation permutes the primes and inverts the character. These equalities hold in the convergent half-plane and continue to the required integers. The bottom doubling denominator is restored by the level index and local section norm already computed for that comparison. The omitted fixed places and The bottom numerator.

Consequently the ratio controlled by the toric test is

$$\mathcal{R}_{s,w} = \frac{L_E(s, \chi\psi_w^{\mathrm{c}})}{L_E(s, \chi)}.$$

The quotient on the right of the signed evaluation has normalized valuation at most ℓ^{-2}. Both numerator and denominator are nonzero. With $e_w = \operatorname{ord}_w(\ell)$ the normalized valuation on the coefficient field is $\operatorname{ord}_w / e_w$, also after a finite local extension. Therefore

$$\operatorname{ord}_w \mathcal{R}_{s,w} \le e_w/\ell^2. \tag{155}$$

The square and bottom norm use $\rho_s = \mathcal{R}_{s,w}\, L_E(s,\chi)/L_E(s,\vartheta)$. With the common normalization of $Z_s^{c(w)}$ and Z_s, the last factor has valuation $\operatorname{ord}_w Z_s^{c(w)} - \operatorname{ord}_w Z_s$, which is bounded above by $T_{s,w}$. Thus the two fixed depletions supply exactly the stated reference budget. No error from the split order factors is added to this budget. The normalization of the valuation and The depleted ratio in the square.

At the family application this calculation uses the particular ray datum chosen on the preliminary record and carried through the two passes of the selection. The character reading, the marked place and the transported stage lattice are required on each actual index. The raw size and translate conditions describe the family to which that selection is applied. The preliminary choice and its realization.

The finite kernel comparison in the two fitted chains. In (15), retain one tensor occurring in the common coset expansion. Its two chains use the same class arguments and the same expansion depth $k_{\mathrm{jet}} > 0$. This depth is an integer modulus of the finite test functions. Let $\partial(p)$ be the maximum prescribed local depth at host places over p in the fixed support S_0, and let Σ be the sum over $w \in S_0$ of the local depth times $\log \mathrm{N}w$. At a coefficient place $v \mid p$ the selected tuple is away from the added conductor over p except at the places in S_0. In the local integral frames of its tensor the test is constant on translates of $p^{\partial(p)}$ times the arrow lattice at p. This last property concerns the actual level and the elementary staircase cut of the test. Once it holds, taking the least coset expansion modulus preserves the bound on its p-part. Indeed the subgroups of periods of a finite product add, so that the greatest common divisor of any two admissible integer moduli is again admissible. The least one divides each admissible one. In particular

$$v_p(k_{\mathrm{jet}}) \le \partial(p), \qquad \sum_p \partial(p) \log p \le \Sigma.$$

For the sum choose, above each p, a place attaining the maximum and use $\log \mathrm{N}w \ge \log p$. Periodic test in the completed frames and The least expansion depth.

For a step of ranks n_s, n_t, the polarization lattice in the finite Schrödinger model has unramified position rank $f n_s n_t$ over the rational finite adeles. On an occurring coset translate the position lattice to the required coset representative and select its finer residue class by the character sum on $L_+/k_{\mathrm{jet}}L_+$. The characters are the modulations by the opposite half, divided by k_{jet}, by perfect pairing of the two halves. On a quotient of this group the same formula uses its quotient order. Thus it represents the test as a linear combination of Heisenberg translates of its vacuum, with translation parameters of order dividing the prescribed k_{jet}-torsion in the pair lattice. The coefficients are roots of unity divided by an order dividing $k_{\mathrm{jet}}^{f n_s n_t}$, with total absolute value at most one at every complex embedding. If the coset has empty support the combination is zero. The finite model with its complementary half and The two specific test factors.

To read a matrix at v, one can first move the two definite vertices to lattices standard at p. Weak approximation in the rational unitary group gives the local move. In the fixed definite frames it is a rational unitary matrix at each end, and it identifies the old vertex lattice with its moved one. Simultaneous rational invariance of the theta sum and the linear change of variables on the Fock polynomials give, if U_s, U_t are these coordinate changes,

$$K_{\mathrm{moved}}\, \rho_s(U_s) = \rho_t(U_t) K_{\mathrm{original}}.$$

The character of the real compact change is included on both sides as a determinant power. For the opposite path the physical source and target are reversed, and the table is transposed in the two tangent indices and conjugated. The substitution by the inverse change of variables gives the same covariance identity in the row-chain direction. As the changes identify the transported coefficient lattices, their use for this comparison costs no denominator. The move to standard lattices at a rational prime, Covariance on the definite path, and Covariance on the opposite path.

The common analytic table in these frames is read by removing the product of CM periods of the jet monomial and the step constant. The individual normalization needed here is an identity of entire functions in the moving theta-line frame. Write F_0 for the normalized self-dual comparison vacuum in the standard frame, $\mathcal{I}_{st}F_0$ for its transfer by the common self-dual refinement of the two lattices, and U_{x_a} for a reduced torsion translation. The table to which the polynomial operator is applied consists of the derivatives of

$$\kappa_{\mathrm{vac}} \sum_a c_a \zeta_a U_{x_a} \mathcal{I}_{st} F_0.$$

Here c_a are the same Fourier coefficients and $k_{\mathrm{jet}} x_a$ belongs to the actual transported period lattice. The roots of unity ζ_a come from the Heisenberg cocycle after moving a finite point to its rational lift. The rational lift may be chosen by the lattice decomposition modulo k_{jet}. A difference in its choice is an integral period, whose theta covariance leaves the reduced translate with the prescribed phase. The finite and real pieces act with their product of central characters. In the opposite direction the same identity is read after conjugating it and transposing the two sets of tangent indices. The placed normalized table and The finite-to-rational torsion transport.

The scalar for the moved vacuum is determined before the monomials are evaluated. In the standard lattices their compact stabilizers act on the normalized vacuum line by a character of finite image, since the lattice vector is smooth and fixed by an open level. This gives a root of unity in the identity with the normalized standard jet table. Decompose each remaining finite argument into a rational move, a split-supported move and a standard lattice stabilizer. For the valuation at p use split support away from p on the lattices standard there. The split move multiplies a normalized vacuum by the source and inverse target splitting characters of the two determinants, times a root of unity. Indeed the local Schrödinger change of variables gives these characters and a positive absolute determinant factor. It is the unitary comparison of the two vacuum lines, so the positive factor is one after the two vacuum normalizations. The character values and their inverses have absolute value one at every complex embedding and are algebraic units above p, since their finite support is away from p. This follows by taking a power in the finite ray group and evaluating the algebraic infinity type on a principal generator, including the unitary normalization by the norm.

For the rational move the transformation of the real Fock frame contributes the two determinant characters of the unitary coordinate changes. In the theta-line frame apply the change of variables $B \mapsto U_t B U_s^{-1}$ to the Hermitian Hom lattice. Comparing the elliptic products in the two bases gives a square of the scalar equal to its determinant times a root of unity. Its determinant is $\det(U_t)^{n_s} \det(U_s^{-1})^{n_t}$ in each host embedding. The transports to and from the common self-dual refinement may add eighth roots of unity. The compact frame scalars have modulus one under all conjugates. If the before and after lattices are standard at p, the matrices U_s, U_t and their inverses are integral at every host place over p, so each of the determinants is a p-unit. This proves the assertion for the rational move as well. Finally every comparison gives the same scalar because $\mathcal{I}_{st}F_0 \neq 0$ and some Taylor coefficient at zero is nonzero. Hence one scalar κ_{vac} represents the moved vacuum at every embedding, and its valuation is zero in the frames standard at p. The scalar identity requires the normalized standard vacua and the transported analytic and algebraic frames on the same class arguments. The standard vacuum comparison, The split movement scalar, Determinant in the rationally moved theta line, and Uniqueness and the paired finite-place comparison.

In standard local frames split the self-dual comparison vacuum as a prime-to-p finite average of products of normalized elliptic vacua. For the average, choose a locally integral basis of the Hermitian Hom lattice for the trace form. The product of the corresponding elliptic polarizations has index prime to p in the finite common refinement. Pullback to the common refinement and the average over the kernel give the comparison operator. All new translation orders divide a prime-to-p multiple of k_{jet}; its additional coefficients have valuations at least zero. The square-root factor of the normalized intertwiner is a p-unit as well. In these Hom coordinates, each normalized elliptic derivative of order i has loss $O(i(v_p(k_{\mathrm{jet}}) + v_p(6d)) + v_p(2))$ by the normalized torsion-translate estimate above. Changing a direction to a single embedding direction uses a trace-dual letter and costs at most $\delta_p = v_p(\operatorname{disc} E)/[E : \mathbf{Q}]$. The derivative in a product uses only the multinomial Leibniz coefficients, which are integral. Together with the coefficient losses in the Fourier formula, this gives the following normalized loss per table entry:

$$O\left((D' + fn_s n_t)(\delta_p + v_p(6d) + v_p(k_{\mathrm{jet}}))\right).$$

The valuations are tested at each algebraic embedding over p, so conjugating the opposite-path table preserves the bound. The arithmetic matrix entries belong to the tied coefficient field. The loss in the integral ord_v convention is therefore this bound times e_v. The prime-to-p product frame, The product of elliptic derivatives, and Normalized table bounds.

The fitted polynomial step first substitutes the bilinear variables and the determinant powers, then evaluates the jet monomials. In the standard Weyl polynomial bases these operations express its matrix entries as integral combinations of the indicated table entries. One may read coordinates in the standard bideterminant basis by integral straightening; the highest-weight conditions give the same column degrees in every nonzero term. Thus $D' \le \sum_{\rho,i} \nu_{\rho,i}$, with $\nu_{\rho,i}$ the shifted polynomial weights of the larger vertex. In the fitted grid these weights are at most $12b$. The bound is uniform over the cosets and classes of the actual two chains, since $n_s, n_t \le N_{\mathrm{g}} \le b$. Consequently the weighted finite loss per step is

$$O\big(fN_{\mathrm{g}} b(\log \mathrm{rd}\, E + \log(6d) + \Sigma)\big) \;\le\; O\big(fN_{\mathrm{g}} b(\log b + \Sigma)\big).$$

Here each δ_p contributes its normalized different, and both $\log \mathrm{rd}\, E$ and $\log(6d)$ are $O(\log b)$ in the family. Including the released small-prime allowance $O(fN_{\mathrm{g}}^2 \log_p(Cb))$ at $p \le Cb$ adds at most $O(fN_{\mathrm{g}}^2 b \log b)$ to a weighted step loss. It is harmless for the same budget. The argument gives a placewise simultaneous bound, so that one may sum over steps and then apply the ultrametric inequality to all alternative terms of the class expansion. It uses the actual normalized table identity and the staircase depth estimate in the prescribed frames. Integral conversion of the jet table to a matrix and The shift range in the two fitted chains.

At infinity, a fitted step multiplies a source polynomial by the prescribed determinant powers, then substitutes bilinear expressions in the target and jet variables. Each jet monomial is evaluated by the corresponding derivative of the step kernel at zero. If the target degree is D', homogeneity gives jet orders $|\alpha| = D'$. For a kernel with $|F(z)|^2 \le A \exp(\pi \|z\|^2/\mathcal{A})$ and d'_Z jet variables, Cauchy's inequality on the polydisc of radii $\sqrt{\mathcal{A}}$ bounds the squared derivative by $Ae^{\pi d'_Z}(\alpha!)^2 \mathcal{A}^{-D'}$. If the two ranks are a, a', determinant multiplication and substitution cost at most $(aa')^{D'}$ in the coefficient absolute sum. In a homogeneous degree D'_0 in q' variables this sum is at most $(1+D'_0)^{q'/2}$ times the Fischer norm by Cauchy–Schwarz. In the target the squared Fischer norm is at most $D'!$ times the square of the coefficient sum. The normalizing Fischer extrema cost at most the input extremal squared norm. On the two fitted chains the total input and output degrees are $O(fN_{\mathrm{g}} b)$ and their numbers of polynomial variables are $O(fN_{\mathrm{g}}^2)$. Since $f, N_{\mathrm{g}} = x^{O(1)}$ and $N_{\mathrm{g}} \le b$, the factorial, substitution and polynomial counting costs have logarithm $O(fN_{\mathrm{g}} b \log b)$. Thus the squared operator norm is at most $A \max(1, \mathcal{A}^{-1})^{D'} \exp(O(fN_{\mathrm{g}} b \log b))$. The fitted jet operator.

The mean bounds in its Gaussian estimate follow from the normalized elliptic lattices and theta values. In a reduced basis write the lattice as $\omega(\mathbf{Z} + \tau\mathbf{Z})$ and $y = \Im\tau$. The discriminant-one normalization gives

$|\omega| = 2\pi|\eta(\tau)|^2$ for Dedekind's eta function. Thus $\mathcal{A} = |\omega|^2 y \ll y\exp(-\pi y/3) \ll 1$. The conjugates are indexed with equal multiplicity by the CM ideal classes. The area identity from Chowla–Selberg in this normalization is

$$\overline{\log(\mathcal{A}/\pi)} = \gamma + \log(2\pi/\sqrt{d}) - \frac{L'(1,\eta_d)}{L(1,\eta_d)}.$$

The field and character estimates give $\overline{\log^+(\mathcal{A}^{-1})} = O(L)$, since the area has a uniform upper bound and the right side is $O(L)$. Area normalization.

For the derivative-one odd theta section the absolute scalar in unitary coordinates $z = u + v\tau$ is $C_\tau = |\omega|/(2\pi|\eta(\tau)|^3)$. Integrating its Fourier expansion first over u eliminates cross terms. The Gaussian integral over the v intervals gives mean squared norm $S = C_\tau^2/\sqrt{2y}$, hence $S^2 = 2\pi^2/\mathcal{A}$. The weighted Fourier sum bounds every squared point norm by $O(C_\tau^2)$, since the Gaussian sum is uniformly bounded for $y \geq \sqrt{3}/2$. For the three nonzero half-periods the squared norms a_t at zero after unitary translation have product four, by $\vartheta_2(0)\vartheta_3(0)\vartheta_4(0) = 2\eta(\tau)^3$. Each is at most $O(\sqrt{y}\,S)$, so the translated section normalized to value one has mean squared norm $S/a_t \ll yS^3$. Since $y \leq \sqrt{d}/2$ in a reduced CM basis, the mean positive log of this normalized norm is $O(L)$ for the selected half-periods. Theta and translation convention.

Conjugation of the fitted theta kernels. On the normalized CM elliptic curve, write the zeta quasiperiod as $s_2\gamma + \pi\bar{\gamma}/\mathcal{A}$. The class $[x\,dx/y] + s_2[dx/y]$ has periods $-\pi\bar{\gamma}/\mathcal{A}$. A nonreal CM endomorphism acts by its conjugate eigenvalue on this class, and by its own eigenvalue on $[dx/y]$. The algebraic de Rham splitting therefore determines s_2 and transports it to the conjugate elliptic lattice. The normalized odd theta series $f(z)$ is the Weierstrass sigma series multiplied by $\exp(-s_2 z^2/2)$. It transports coefficientwise.

For a torsion translation x, choose $m_0 \geq 2$ with $m_0 x/2$ a period. Put $u = x/(2m_0)$ and $M = m_0^2$. The function $D_N(z) = f(Nz)/f(z)^{N^2}$ is elliptic and algebraic, with the divisor given by the nonzero N-torsion and the pole of order $N^2 - 1$ at zero. The leading coefficient is N in the invariant coordinate, so its normalized Taylor series transports with that coordinate. For u nonzero modulo periods put $B_u(z) = (U_u f(z))^{4M}$, where U_u is reduced translation. The multipliers and oddness give

$$U_x f(2m_0 z) = B_u(z)D_{2m_0}(z+u), \quad B_u(z)^M = f(2Mz)/D_{2M}(z+u), \quad B_u(0) = -D_{M-1}(u)/D_{M+1}(u).$$

The quotient in the middle is a unit series. Its M-th root is determined by the given constant, so the phase transports along with the deeper torsion point. The doubled-period case acts trivially. Normalizing the translated even test by its value at zero preserves this conclusion. The signs and cocycles in the finite averages are powers of Weil pairings on the deeper torsion table, and therefore also conjugate with the table. Torsion translation and normalized jets.

For the higher-dimensional transports, pull the polarizing line bundles back to a common finite cover on which the translations belong to their theta groups. A reduced translation has finite order since a suitable multiple is an even period. Over its algebraic torsion point the finite-order condition on a theta group lift is an algebraic equation for its scalar, so the lift conjugates algebraically. The common formal frame is obtained from the invariant coordinate and the reduced elliptic series. The rigidified isomorphisms of the pullback bundles respect it. The finite subgroup projections conjugate as sums with the same indices, and a square root of an index can change only its common sign. The rational isometries in the fourfold identity also transport, since an integer multiple induces a homomorphism of the CM tori. Taking the fourth root changes a vacuum by at most a common unit phase. One can choose a finite extension of the tied field containing the translated germs and all marked transport constants in the fixed expansion. The jet operator is linear in its table, uses integer coefficient substitutions, and is equivariant for the weight-relabeling at an embedding. Thus an equality of the normalized table at the chosen embedding is an equality of algebraic coefficients. For the direct path scalar extension along σ gives the jet operator with entries read

in the conjugated CM frame. For the opposite path complex conjugation at the chosen embedding realizes the coefficient involution. Entrywise conjugation on the two source and target polynomial models gives that path by applying the direct jet construction at $\sigma\circ *$ and conjugating the coefficient reading. The transposed jet variables correspond to interchanging the two ends of the opposite physical arrow. This preserves the Fischer norms under conjugation of the weight and of the definite Hermitian form. These are the two reads on the original coefficient matrices, with the same determinant characters as in their original finite movement. The statement uses the algebraically normalized theta-germ identity on the actual marked frames. Scalar extension of the jet operator, The conjugate table in the opposite order, and The opposite coefficient involution.

For the theta vacua occurring in this comparison, the finite Heisenberg averages are orthogonal projections onto the two section spaces in the common refinement. Their transports are the restrictions of these projections multiplied by the square root of the common index. Since the two transports are inverse adjoints, they preserve the mean norm. The fourfold identity then expresses the mean norm of one vacuum through the elliptic factors. On a self-dual theta lattice in complex dimension g', a ball of radius $r_0 = \min(1, \ell_{min}/3)$ embeds in the quotient. After a unitary translation, the mean-value inequality on the ball bounds its pointwise squared norm by $e^{\pi} g'!/(\pi^{g'} r_0^{2g'})$ times the mean squared norm. The alternating pairing gives covolume one in the Hermitian coordinates. In the prescribed Hom lattices the nonzero Hermitian lengths lie in the inverse relative different after the elliptic area is factored out. Taking the ideal norm and using the arithmetic–geometric mean inequality bounds $\ell_{min}^{-2} \ll \sqrt{d}\,\mathrm{rd}(E)$. Thus the log cost of passing from the mean norm to the envelope, and of the finite combinations of translated vacua, is at most $O(fN_{\mathrm{g}}^2 \log b) + O(fN_{\mathrm{g}}^2)\log^+ \|F_{\mathrm{ell}}\|^2$ in the envelope. Here F_{ell} is the normalized elliptic vacuum on that conjugate curve. The fourfold norm calculation uses at most $O(fN_{\mathrm{g}}^2)$ elliptic factors, and the square of the finite Fourier coefficient sum is at most one. The $\mathcal{A}^{-D'}$ term and the preceding Cauchy bound therefore give, at any extending embedding ρ over the coefficient embedding σ,

$$\log^+ \|K_\sigma\|^2 \le C\left(fN_{\mathrm{g}} b \log b + fN_{\mathrm{g}} b \log^+ \mathcal{A}_\rho^{-1} + fN_{\mathrm{g}}^2 \log^+ \|F_{\mathrm{ell},\rho}\|^2\right).$$

One may average the right-hand side over the extensions of each σ, since it bounds the same operator in the same prescribed metrics at every extension. All embeddings have the same number of extensions. Thus the mean step allowance is $O(fN_{\mathrm{g}} b \log b)$ on each chain. This estimate uses only the normalized Heisenberg coefficients and the elliptic mean bounds above. Conjugate step kernels, Both paths.

Assembly in the quotient norm. For the fitted map the same term estimate applies to every definite test class and every input vector. In the σ-norm let B_σ bound the product of the endpoint squared operator norms and those of the $2m$ kernel operators. For one term its bilinear pairing has absolute value at most $\sqrt{B_\sigma}\|u\|\|x\|$. Multiplication by the finite tensor and scalar prefactors and taking the sum bounds its operator norm by $\sqrt{B_\sigma}\sum_\beta |c_{\beta,\sigma}|$, where the sum ranges over the supported terms, with the class and expansion indices included. The stabilizer denominators have absolute values at most one. For the direct sum of definite classes, use Cauchy–Schwarz in their coefficient norms before taking the dual. In coordinates its Gram is the inverse transpose of the Gram on the direct sum of the tests. Thus a bound on the squared pairing against all test vectors is exactly the same bound on the output squared norm.

Finite tensors, measures and support. Consider a fixed split place of norm Q and depth a. The split character in a flip is a base splitting character times the finite twist times the inverse grid character. Its values on units are roots of unity. For a uniformizer and its inverse, each of the two algebraic infinity types has absolute exponents at most $\lfloor 8b \rfloor$ in the unitary Hecke character convention. Raise the uniformizer idele to a power in its finite ray group. The resulting principal generator has ideal a power of the uniformizer ideal. The character at this generator is the product of its conjugates raised to the respective exponents. The sum of their nonnegative valuations is the valuation of its norm. Each type therefore

costs at most $\lfloor 8b \rfloor v(Q)$, in either direction. The finite twist costs zero. All complex readings of these values have modulus one, since the unitary infinity types have opposite exponents at the two embeddings of a CM pair. The two types and the finite twist.

In one square flip the normalized Fourier sum on the changed columns has denominator Q^{ah} with $h \le 2N_g^2$. The inverse Fourier sum is unnormalized. The intervening integral uses a compact open J of volume one. Completing a primitive scaled row frame to an integral invertible frame introduces a determinant with at most $2N_g a$ uniformizer powers. In the remaining rank $d' \le N_g$ the integral is supported on matrices in $\pi^{-a}\mathcal{O}^{d'\times d'}$, periodic modulo integral matrices before the determinant character is inserted. A Smith shell with an elementary exponent ≥ 0 has zero integral. For fixed compact left and right frames, multiplying the corresponding diagonal coordinate by a unit changes the argument by an integral matrix and multiplies the integrand by the ramified determinant character of that unit. Averaging over that unit kills the shell. All remaining exponents are in $[-a,-1]$. Each such shell is compact and is a finite union of left J-cosets, on which the integrand is constant by the matching character covariance. The shell integral is therefore a sum with no volume denominator. The ad' uniformizer powers and the frame determinants, including the determinant half-power, give a total loss $O(aN_g(b+N_g)v(Q))$ per flip.

For the complex bound change variables by $\pi^a I_{d'}$ in the remaining integral. Now the support is integral and the positive power of the determinant modulus is at most one. The determinant factors from this change are bounded by $Q^{O(aN_g^2)}$. Partition the integral support into matrix congruence classes modulo π^a. If a class contains a singular matrix, put it in Smith form by integral changes of frame. Multiplying the null coordinate by units leaves the congruence class invariant and kills the integral by the same character averaging. In a class containing no singular matrix choose a lift M in $\mathrm{GL}_{d'}$. Smith reduction gives $M^{-1}\pi^a\mathcal{O}^{d'\times d'} \subset \pi\mathcal{O}^{d'\times d'}$, so the class is a translate of a subgroup of $\mathrm{GL}_{d'}(\mathcal{O})$. Its mass is at most $[\mathrm{GL}_{d'}(\mathcal{O}) : J] \le Q^{ad'^2}$. The sum over at most $Q^{ad'^2}$ classes and the two Fourier sums hence cost $O(aN_g^2 \log Q)$ in the complex logarithm per flip. This applies at every embedding of the base data. Fourier, primitive-frame and shell factors, The unit-volume shell sum, and The conjugate complex bound.

There are m^2 flips. At a fixed exceptional place the paired finite support in the principal opposite-path calculation has the single staircase representative, with the completed-frame period of the chosen depth on both halves. Its diagonal exponents are between 0 and ma. In the movement of the rectangular variables these exponents enter with the ranks of the two ends divided by two. Hence the $2m$ edges add $O(m^2N_g^2a)$ powers of $Q^{1/2}$. Each internal inverse measure has the form Q^z, where $|z| \le m^2N_g a$ is the absolute Jacobian exponent of the diagonal change on its free unipotent entries. There are at most m of these factors and $m \le N_g$.

The filled local integral has at most mN_g^2 free additive variables. On a coset modulo $\pi^{ma}\mathcal{O}$ in these variables the moved tensor is constant. In the oscillatory phase the difference induced by a translation is linear modulo the additive character conductor. Each product of two increments contains π^{2ma}, which clears the two staircase scales in the phase. Thus an integral on a coset is zero or its value at a representative times $Q^{-man_{\mathrm{var}}}$, where n_{var} is the number of variables and the integral coordinate box has volume one. The archimedean estimate on the base uses the support in a translate of $\pi^{-ma}\mathcal{O}^{n_{\mathrm{var}}}$, whose volume is $Q^{man_{\mathrm{var}}}$. At an added twist place the prescribed coefficient place has another residue characteristic. The same finite integration can use any sufficiently deep cosets. Their volumes are powers of the other residue cardinality, and the flip loss there is zero at the coefficient place. Completed-frame periods at fixed places and Integral periods away from the twist. The inverse unipotent volume is a rational integer divided by a power of the corresponding residue prime: use the index formula with the integral compact open and its intersection with the moved open. The index from the moved open to the intersection is a power of that prime, by its filtration by integral root groups. It contributes no loss at the coefficient place. Staircase determinant and inverse volumes, Volumes at a different coefficient prime, and Filled coordinates and their character sums.

For the common scalar s_{AA} the oscillator measure changes by a product $\prod_w(\sqrt{\mathrm{N}w})^{h_w}$. In each of the m^2 square comparisons the measure is a product of changes of lattice measure on rectangular blocks, each

with at most N_{g}^2 entries. The scale of each coordinate is a product of the conductor scales of the additive character, corresponding to $\mathfrak{A}, \overline{\mathfrak{A}}$ on the two sides. The bound on their exponent, counting the repetition along the staircase and both signs, is

$$|h_w| \le 16m^3 N_{\mathrm{g}}^2 \big(|v_w(\mathfrak{A})| + |v_w(\overline{\mathfrak{A}})|\big).$$

The identity $\xi_0 \mathfrak{A}\overline{\mathfrak{A}}\mathfrak{D}_E = \mathcal{O}_E$, with nonpositive valuations of $\xi_0, \mathfrak{A}, \overline{\mathfrak{A}}$, bounds the norm log of these ideal counts by $\log|\operatorname{disc} E|$. Every complex reading of each square root has the same modulus. The product formula for each $\mathrm{N}w$ then bounds the weighted finite loss and the complex absolute log by $O(m^3 N_{\mathrm{g}}^2 f \log \operatorname{rd} E)$. This is $O(mfN_{\mathrm{g}} b \log b)$, since $m^2 N_{\mathrm{g}} \le b$ and $\log \operatorname{rd} E \ll \log b$. The host ideal counts and Square roots in the common scalar.

The finite coset choices are specified by residue coordinates on the $2m$ steps. At a place of norm Q and depth a there are at most $Q^{2n_s n_t a}$ choices for the two residue blocks of a step. Summing $a \log Q$ in Σ bounds their log count by $O(mN_{\mathrm{g}}^2 \Sigma)$. For the class count the base levels contain the principal congruence of the grid modulus $\mathfrak{m}$, of log norm at most 2Σ by extension of the depth ideals to E. At a row vertex this follows from the defining integral block congruences. At a top vertex conjugation by the staircase fixes the original rank block and scales each added off-block coordinate by its designated power. The principal congruence meets those scaled congruences and preserves the split determinant character. Residue coordinates, Top staircase level, Row level, and The grid modulus.

A rank-n vertex has at most $(\mathrm{N}\mathfrak{m})^{n^2}$ refinements of each full-level class, by reduction inside $\mathrm{Mat}_n(\mathcal{O}_E/\mathfrak{m})$. In the standard definite genus the Hermitian lattice is self-dual over $\mathcal{O}_E$. Its integral trace lattice has dual $\mathfrak{D}_E^{-1}$ times the lattice and determinant $|\operatorname{disc} E|^n$. Its squared minimum is at least one. In dimension $D = 2fn$, Minkowski's second theorem bounds all its successive minima by $(2D)^D |\operatorname{disc} E|^{n/2}$. One obtains a basis with lengths at most 2^D times this bound by extending primitive bases through the spans of independent vectors realizing successive minima, in each extension subtracting a multiple of the preceding basis to reduce the component in its span. Apply the same argument to $\mathcal{O}_E$ itself. In the reduced basis record the integral trace Gram matrix and the matrices for multiplication by this bounded field basis. Their entries are bounded by the basis lengths and the norm of the dual basis, obtained from the Gram cofactors. These matrices determine the Hermitian form, since the forms $\mathrm{Tr}_{E/\mathbb{Q}} H(\alpha x, y)$ for $\alpha \in E$ identify H by trace duality. Thus the log number of classes at full level is bounded by a fixed polynomial in $f, n, \log|\operatorname{disc} E|$. These parameters are at most fixed powers of x. The full-level count therefore costs $O(fN_{\mathrm{g}} b \log b)$ in the logarithm, since $b = \exp(x^{\varepsilon})$. The product over at most $2m+1$ class vertices costs $O(mfN_{\mathrm{g}} b \log b + mN_{\mathrm{g}}^2 \Sigma)$. Refinement within the standard definite genus.

The term bounds apply on the support of (15). The finite tensor and scalar prefactors and the class count contribute at most

$$O\big(mfN_{\mathrm{g}} b \log b + m^2 N_{\mathrm{g}}^2 \Sigma\big)$$

to the archimedean log on the base. The weighted finite loss, including the stabilizer denominator and the two step chains, is $O(fm^2 N_{\mathrm{g}}^2 b(\log b + \Sigma))$. The scalar, tensor and class contributions, and Using the test dual.

On the quotient Q_0 the bound before the operator ι_0 follows by taking the least-norm lift in the full source form. At each embedding the compensated target Gram on a block is the quotient Gram multiplied by $\prod_j \lambda_{j,\alpha,0,\sigma}$. Changing to the compensated norm therefore adds at most $-\min_\alpha \sum_j \log \lambda_{j,\alpha,0,\sigma}$ to the log squared operator norm. The minimum here is over the actual base type decomposition at the given embedding. By (44) its averaged loss is $O(mfn^2 \log(b \operatorname{rd}(E) g_0 g_1 n))$, since all hook products are at least one. Taking the mean proves the bound for $h_\infty(\Psi)$ on the placed table and the compensated quotient Gram. Compensated quotient comparison. Passage to an injective lattice map. Weighted denominators in the fitted tests.

The depth sum in these estimates has polynomial size in x. At the auxiliary place q_0, the depth is one plus the integer part of

$$\frac{Cf^2(r+1)^6 \log\bigl(b(r+1)g_0 g_1 p_{\mathrm{de}} p'_{\mathrm{de}} |\operatorname{disc} E|\bigr)}{\log q_0}.$$

This is one specified split place of norm q_0. At the other fixed places the depth is two. The sum of their log norms is at most $f \log(g_0 g_1 p_{\mathrm{de}} p'_{\mathrm{de}})$. For the classical host, $q_0, p_{\mathrm{de}}, p'_{\mathrm{de}} < X'$, $\log g_i = O(L)$, and $\log|\operatorname{disc} E| = 2f \log \operatorname{rd}(E)$. Together with the polynomial ranks this gives $\Sigma \le x^{c_\Sigma}$ with a fixed exponent and one threshold for the selected family. Thus the weighted finite test loss, the signed square budget on the chosen readings, the mean squared operator log, and the evaluation lattice's maximal slope are each $O(x^c b)$. The constants depend on the admitted row parameters and the fixed accuracy range. They are uniform in the words and translates. The integral row generators are those of (92), and the square-index comparisons use their selected normalized coefficient readings. Substitution in the exterior-power estimate proves the base version of (154). The depth parameter, The fixed-prime size sum.

The two actual row levels. Use the group $G = \prod_{v \in P_g} \mathrm{GL}_r(k_v)$ and the normalized right action of the level-change calculation. Set $R_0 = R^{(0)}$, $R_1 = R$, and $h_i = |P_i|$, where P_1 is the sharp subgroup P_s of (41). In the adapted residue frames P_0 is the product of the full upper unipotent subgroups with generic characters. The prescribed symmetric ordering of the frame coordinates carries P_1 and its character to the product of the subgroups J and characters in the distinguished row count. The reordering is inner, by a permutation matrix at each factor. Write χ_i for the level eigencharacter in the normalized action. Use the same character transport on the subgroup and on its projector. Exchanging the conjugate realizations at the distinguished complex place conjugates both the action and the characters and preserves norms. The actual base and sharp characters and The finite inner conjugation.

Construct the principal lattice L_{amb} by running the gluing in (52) at the trivial subgroup, with the same twists, common row readings and central good projector e_X. At tuple t write V_t for the principal class-value space and $\mathscr{L}_t$ for its normalized value lattice. The normalized finite action ρ_t preserves $\mathscr{L}_t$. Indeed it translates the framed classes, preserves their value lattices on representatives, and preserves the class pairing. Write $V_{\mathrm{p},t} = e_X V_t$ and $V_{i,t} = e_i V_{\mathrm{p},t}$, with $e_i = h_i^{-1} \sum_{p \in P_i} \chi_i(p)^{-1} \rho_t(p)$ in this chart. Let Z_t be the selected adjoint cut from (41). The corresponding projected lattices are

$$\Gamma_{a,t} = \pi_{a,t}(\mathscr{L}_t \cap V_{a,t}), \qquad \pi_{a,t} : V_{a,t} \longrightarrow Z_t \cap V_{a,t}, \quad a \in \{\mathrm{p}, 0, 1\},$$

where $\pi_{a,t}$ is the rational orthogonal projection in the class Gram. These are the images of the integral values before projection, as in (52). The descent is equivariant, so its kernel, its adjoint complement, and the principal good cut are G-stable. It follows by unitarity that $\pi_{\mathrm{p},t}$ commutes with G. Its restriction to $V_{i,t}$ is $\pi_{i,t}$, since the level Gram is h_i^{-1} times the restricted Gram. Every $\chi_i(p)$ is a root of unity, so $h_i e_i$ carries $\mathscr{L}_t \cap V_{\mathrm{p},t}$ into $\mathscr{L}_t \cap V_{i,t}$. Applying the projections gives

$$h_i e_i \Gamma_{\mathrm{p},t} \subseteq \Gamma_{i,t} \subseteq \Gamma_{\mathrm{p},t} \cap \operatorname{im} e_i.$$

The integral action on the normalized values, The projected lattice, and Compatibility with the projector.

Let j_i be inclusion in the principal row coordinates at the base tuple, and $J_{i,t}$ the inclusion at tuple t. Let $B_{a,t}$ be the downward transition to the base tuple at level a. The transitions satisfy $j_i B_{i,t} = B_{\mathrm{p},t} J_{i,t}$ because both sides realize the same descent in the common principal row chart. For the same reason $B_{\mathrm{p},t}$ intertwines ρ_t with the action ρ on the principal base row space. These equalities hold over K by the faithful rational readings in (41). In particular j_i identifies $R_i \otimes K$ with the e_i-image on that row space. At each coefficient prime use $t = \boldsymbol{\mu}(v)$ on all three levels. Transporting the inclusions and localizing, then intersecting over the primes, gives a G-stable principal lattice and

$$h_i S_i \subseteq h_i e_i L_{\mathrm{amb}} \subseteq j_i R_i \subseteq S_i, \qquad S_i = L_{\mathrm{amb}} \cap \operatorname{im} e_i. \tag{156}$$

At primes outside $|G|$ these inclusions are equalities. The central cut is imposed by intersection with its image, and the adjoint projection is applied to the resulting integral vectors at all three levels. The comparison uses the integral operator $h_i e_i$ on this common good type space. The lattice at P_1 is the fitted word-period lattice with its prescribed metric, and the lattice at P_0 is the base lattice of the test Ψ. The same rational transitions, The gluing of the action and the cuts, and The two row lattice and metric identifications.

At the distinguished complex place the Petersson normalization in (41) divides the principal Gram by h_i. At every complex embedding the transported class action permutes the framed classes and their stabilizer weights. Its coefficient actions are unitary for the transported fibre metric because the embedding of the CM host respects its conjugation. The finite character factors have absolute value one at every embedding. Consequently the class action is unitary there and the character and central projectors are orthogonal. In the principal class quotient the shortest lift of an e_i-vector belongs to the character space. Indeed applying the orthogonal projector e_i to any lift preserves its image and can only decrease its norm. The quotient metric for the character space thus restricts the principal quotient metric and contributes the same factor h_i^{-1} from the class Gram. The G-action is unitary on the principal quotient by the corresponding least-norm argument. Both inclusions and the G-action preserve the joint Hecke summands. On each summand the final metrics multiply the quotient Gram by the same positive $\kappa_R \prod_j \lambda_{j,\alpha,0,\sigma}$. Thus ρ is unitary for the final principal metric and

$$\|r\|_{R_i,\sigma}^2 = h_i^{-1} \|j_i r\|_{\mathrm{amb},\sigma}^2. \tag{157}$$

The scalar κ_R is the common row scalar in the fitted height, so the same comparison holds in the compensated metrics defining the shift J_κ. The quotient Gram and final metric transport and The exact fitted row and scalar.

The principal space lies in the actual good central cut. Each irreducible in this cut has squarefree local labels at all type places. The base record calculation gives one generic Whittaker line on each local parabolic induction of distinct cuspidal factors. It applies to the character of P_0 in the adapted frame and its inverse, which are conjugate by an alternating diagonal matrix. Taking the product of the local lines shows $e_0 \neq 0$ on each of the retained simple representations. This holds in the principal rational space by scalar extension. The good character product and base multiplicity one.

Use the map

$$r \longmapsto \left(j_0^{-1} \left(h_0^2 e_0 \rho(g) j_1 r \right) \right)_{g \in G} \quad \text{from } R_1 \text{ to } \bigoplus_{g \in G} R_0.$$

It is integral because $h_0 e_0$ preserves the ambient lattice and the second factor h_0 puts its output in $j_0 R_0$. Its kernel over K is zero. The common kernel of the maps $e_0 \rho(g)$ on the ambient space is G-stable, and a nonzero simple submodule of that kernel would have zero e_0-component. The squared operator norm of this map is at most $|G| h_0^3 h_1 \leq |G|^5$ in the stated metrics, using the unitarity of $\rho(g)$ and orthogonal projection. The same bound holds when the norms on both sides are divided by $\kappa_{R,\sigma}$. The maximal slope of an orthogonal sum is the maximum of the summand maximal slopes. The exterior-power slope inequality for the displayed integral injection gives

$$\mu_{\max}(R) + J_\kappa \leq \mu_{\max}(R^{(0)}) + J_\kappa + 5 \log |G|.$$

With $\log |G| = \lambda_0 \leq 4b$, this has the required size.

Let s_{raw} be the rank-weighted mean of $\mu(N'_w) - h_w$. For each word the paired maximal slope bound gives

$$\mu(N'_w) - h_w + \mu_{\max}(R) \leq \mu_{\max}(M_w) + \mathcal{G}(u) + \mathcal{F} - Q_w + C_1 m B_0 - (h_w + J_\kappa) + (\mu_{\max}(R) + J_\kappa).$$

Write $P_t = \max(C_t, 5) + 3$, where C_t is the fixed exponent in the twist-cutoff size bound. The first term and $C_1 m B_0$ are at most $x^{4P_t+1} b$ each. In the gain and reference term, the bound $g_{jv} \leq 4 N_T a_E / p$ on the

ordinary support, the analogous reference bound above, and $-Q_w \ll mB_0$ give the upper exponent $9P_t+2$. For the negative compensated height, use the preceding explicit bounds for the factorials and the weights. With $S = 1+f+k+m+t+N(m)+L$, the sum is at most $(500+2|C_{T_u}|+30|C_{\mathrm{sz}}|)S^9 b$, where C_{T_u} and C_{sz} are the fixed level and height constants. Since $S \le x^{P_t}$, this is bounded by $x^{9P_t+1}b$ at a common threshold. If the base row estimate is $Ax^{c_F}b$, take $c_F \ge 0$ and set $C_{F,0} = c_F + 1$. The finite level comparison adds at most $20b$, so the sharp compensated bound is at most $x^{C_{F,0}}b$ for $x \ge A+20$. The exponent c_F depends on the earlier row parameters at fixed precision, and the threshold is uniform in the translate. Hence $C_{F,0}$ enters the staged lower bound for C_H before choosing the common family. These estimates justify the fixed requirement

$$C_H \ge 1 + \max(4P_t+1, 9P_t+2, 9P_t+1, C_{F,0})$$

in the family application. It is chosen before the scale and translate. Averaging over the words and using $x \ge 9$ to absorb the five terms gives

$$s_{\mathrm{raw}} + \mu_{\max}(R) \le x^{C_H} b. \tag{158}$$

The frame estimate uses this same C_H in (62). The reference-lattice exponent, The gain and reference exponent, The lower compensated height, and Combination at a common exponent.

4.4 The CM period and the centering term

Put $v_0 = \mathrm{rd}(E)/2$ and let $c(i_1,i_2,h) = (-1)^{i_1}$, so $P = \ker c$. For the classical target, the Artin pullback of c on units is $(\frac{\cdot}{d})$ by (3). Since d is a prime congruent to 7 mod 16, quadratic reciprocity identifies the primitive character as $\chi_d = (\frac{-d}{\cdot})$ of conductor d. Put $\ell_c = L'(1,\chi_d)/L(1,\chi_d)$, which is real.

The elliptic period. Set $K_d = \mathbf{Q}(\sqrt{-d})$. The elliptic datum in the host construction uses a lattice homothetic to an ideal of $\mathcal{O}_{K_d}$, scaled to have $g_2^3 - 27g_3^2 = 1$. Choose a number field F' containing its Weierstrass invariants and the images of the embeddings of the host. Denote the normalized lattice at an embedding $\sigma : F' \hookrightarrow \mathbf{C}$ by Λ_σ. Its area is determined by the ideal class. Indeed, if $\Lambda_\sigma = \omega_\sigma(\mathbf{Z}+\mathbf{Z}\tau_\sigma)$, with $y_\sigma = \operatorname{Im}\tau_\sigma > 0$, the discriminant normalization and the discriminant identity on $\mathbf{Z}+\mathbf{Z}\tau_\sigma$ give

$$|\omega_\sigma|^{12} = (2\pi)^{12}|\eta(\tau_\sigma)|^{24}, \qquad \log\frac{\mathrm{area}(\Lambda_\sigma)}{\pi} = \log(4\pi) + 2\log(\sqrt{y_\sigma}|\eta(\tau_\sigma)|^2).$$

The maximal-order theorem of complex multiplication identifies the conjugates of the j-invariant with the ideal classes of K_d, each once. Every embedding of the j-field has the same number of extensions to F'. Thus the average over σ is the ideal class average. [69, Proof of Proposition 6.1, ideal-class orbit of the CM invariant]

For an integral ideal $\mathfrak{a}$ of shape $\tau = x_0 + iy$, its covolume is $\sqrt{d}\,\mathrm{N}\,\mathfrak{a}/2$. The map $\alpha \mapsto (\alpha)\mathfrak{a}^{-1}$ takes nonzero $\alpha \in \mathfrak{a}$ onto the integral ideals in $[\mathfrak{a}]^{-1}$, with fibres of size $w_d = |\mathcal{O}_{K_d}^\times|$. The partial zeta function of $[\mathfrak{a}]^{-1}$ is consequently

$$\frac{(2/\sqrt{d})^s}{w_d} \sum_{(a,b)\in\mathbf{Z}^2\setminus\{(0,0)\}} \frac{y^s}{|a\tau+b|^{2s}} \qquad (\Re s > 1).$$

Kronecker's first limit formula gives the residue π of the sum at 1 and its constant term $2\pi(\gamma - \log 2 - \log(\sqrt{y}|\eta(\tau)|^2))$. [70, Ch. I, § 1, Theorem 1] Hence all partial zeta functions have the same residue, and the ratio of constant to residue for the ideal of shape τ is

$$\log(2/\sqrt{d}) + 2\gamma - 2\log 2 - 2\log(\sqrt{y}|\eta(\tau)|^2).$$

Their sum is $\zeta_{K_d}(s) = \zeta(s)L(s,\chi_d)$, whose ratio is $\gamma + \ell_c$. Averaging the ratios and using the formula for the normalized area gives

$$\bar{P} = \frac{1}{[F':\mathbf{Q}]}\sum_\sigma \log\frac{\mathrm{area}(\Lambda_\sigma)}{\pi} = \gamma + \log(2\pi/\sqrt{d}) - \ell_c. \tag{159}$$

The host metric. Write $\mathfrak{D} = \mathfrak{D}_{E/K_d}$. The host specifies $I = \mathfrak{b}^{-1}$ with $\mathfrak{b}\bar{\mathfrak{b}} = \mathfrak{D}$. For the classical host, $I_d \cap P = \{0\}$, so E/K_d is unramified at d by (4). The relative different is supported over $2, q$, both split in K_d. Conjugation pairs its prime factors with distinct primes of the same exponents, which allows the prescribed factorization. Use the embeddings $\tau : E \hookrightarrow F'$ extending the chosen embedding of K_d. The integral coordinate module in the theta metric is the image of $I \otimes_{\mathcal{O}_{K_d}} \mathcal{O}_{F'}$ in $\prod_\tau F'$ under the embedding map. Its intersection with the τ-axis is the fractional ideal $J_\tau = \tau(I\mathfrak{D})\mathcal{O}_{F'}$ on that axis. To see this, the trace-dual of I is $I^{-1}\mathfrak{D}^{-1}$. The trace pairing becomes the dot product under $E \otimes_{K_d} F' \simeq \prod_\tau F'$. The integral trace dual commutes with extension to $\mathcal{O}_{F'}$ because I is finite projective over $\mathcal{O}_{K_d}$. Thus the axis vector with τ-coordinate s belongs to the integral module exactly when it pairs integrally with the dual, or equivalently when $s\,\tau(I^{-1}\mathfrak{D}^{-1})\mathcal{O}_{F'} \subseteq \mathcal{O}_{F'}$. Inverting this fractional ideal gives the stated intersection. Coordinate line in the release.

For $0 \neq s \in J_\tau$ the host's mean logarithmic period in the squared norm convention is

$$\frac{2\log[J_\tau : s\mathcal{O}_{F'}] - \sum_\sigma \log\bigl(\pi|\sigma(s)|^2/\operatorname{area}(\Lambda_\sigma)\bigr)}{[F' : \mathbb{Q}]}.$$

Metric and index conventions. The index is $|\mathrm{N}_{F'/\mathbb{Q}}\, s|/\mathrm{N}_{F'/\mathbb{Q}}\, J_\tau$. Consequently the terms involving s cancel, and the remaining finite-ideal correction to $\bar{P}$ is $-2\log \mathrm{N}_{E/\mathbb{Q}}(I\mathfrak{D})/[E : \mathbb{Q}]$ by extension of the ideal norm. Since $\mathrm{N}(I\mathfrak{D})^2 = \mathrm{N}\,\mathfrak{D}$ and $|\operatorname{disc} E| = d^{[E:K_d]}\,\mathrm{N}\,\mathfrak{D}$, this correction equals $\log(\sqrt{d}/\operatorname{rd}(E))$. Together with (159) this gives the common period used in the height estimate,

$$\log \Omega = \gamma + \log(\pi/v_0) - \ell_c. \tag{160}$$

Only d and $\operatorname{rd}(E)$ enter this expression, so the Artin character of the base host and the height normalization at the tied hosts give the same ℓ_c and v_0.

The auxiliary level. At a fixed translate put $F = E^+$, $r_{\mathrm{row}} = k + m - 1$, and $\lambda_g = \log(g_0 g_1)$ for the two fixed auxiliary primes. Write $D_\star$ for the row reference dimension of Section 3, and d_w^* for the coefficient of $(T_w^k - 1)^f$ in the full column dimension. The chosen auxiliary level satisfies

$$X_w = \left(\frac{D_\star}{d_w^*}\right)^{1/(fk)}, \qquad 1 \le \frac{T_w}{X_w \Gamma^{-2/(fk)}} \le 1 + \frac{1}{fk}, \qquad 0 \le \frac{\log \Gamma}{fk} \le 1.$$

The interval and the common Γ are part of the level choice. The chosen interval and its logarithm.

For a decreasing tuple $\mathbf{z} = (z_i)_{0 \le i < h}$, set $V(\mathbf{z}) = \prod_{0 \le i < j < h}(z_i - z_j)$ and $J_h = \prod_{i=0}^{h-1} i!$. Denote the chosen row and compact column Harish–Chandra parameters at ρ by $\mathbf{y}_\rho$ and $\mathbf{q}_{\rho,w}$. The Weyl dimension of the column is $V(\mathbf{q}_{\rho,w})/J_k$. For the row, Harish–Chandra's formal-degree formula in compact-dual Haar normalization gives $d_\rho = V(\mathbf{y}_\rho)/J_{r_{\mathrm{row}}}$. In this normalization an invariant top density is divided by the volume of its transfer to the compact real form. The denominator is the product of the positive-root values of the half-sum of positive roots, which for $U(h)$ are $j - i$ for $i < j$. For the compact central factor, push probability Haar measure on $U(1)$ times the normalized semisimple Haar measure through $U(1) \times SU(p,q) \to U(p,q)$. In the matrix-coefficient orthogonality formula the central character has modulus 1, so the central integral leaves the formal degree unchanged. Rescaling the real Haar at a place multiplies the adelic volume by the scale factor and divides d_ρ by it. [62, (3.10), (3.12), Theorem (8.1)] [61, § 41, Theorem 16] Identification in the chosen compact-dual measure.

Let $\mathfrak{M}_h$ be the mass of the standard positive definite self-dual genus of rank h over E/F, i.e. the sum of the reciprocals of its lattice automorphism orders. The quotient volume of the row group in compact-dual real measures and with the full self-dual finite stabilizers of volume 1 is $\mathfrak{M}_{r_{\mathrm{row}}}$. In fact the row form and the definite form have isometric self-dual lattices at the finite places, since E/F is split or unramified there. In the inert case one splits off unit-length vectors by reduction to a nondegenerate Hermitian form over the residue field and by surjectivity of the norm on units. In the split case one takes dual bases of the two

factors. The quotient volumes agree by the Tamagawa number 2 of these unitary groups. The invariant top forms on the two inner forms correspond because inner conjugation has determinant 1 on the Lie algebra. The finite-volume normalizations agree by local isometry, and the real ones by transfer of the invariant density to the compact dual. The convergence factors in the Tamagawa measure agree because the inner forms have the same motive. For the definite form the double-coset decomposition weights each lattice class by the inverse order of its stabilizer, proving the volume equality.[63, (2.2)–(2.6), § 4]

Passing from a full finite stabilizer to a subgroup multiplies the quotient volume divided by level volume by the subgroup index. At g_i there are f places of F of residue cardinality g_i, all split in E. The column reduction is upper unitriangular, with $\binom{k}{2}$ free entries. On the row side the reduction is upper unitriangular in a flag with central block of size $k+1$, in which only the upper pairs among the first k entries are free. This excludes k upper entries and leaves $\binom{r_{\mathrm{row}}}{2} - k$ free entries. Each choice gives an invertible matrix. Dividing the order $g^{h^2}\prod_{s=1}^{h}(1-g^{-s})$ of $\mathrm{GL}_h(\mathbf{F}_g)$ by the two subgroup orders gives the ratio of indices

$$\frac{I_{\mathrm{row}}}{I_{\mathrm{col}}} = (g_0 g_1)^{fm(2k+m-1)/2} \prod_{i=0}^{1} \prod_{s=k+1}^{r_{\mathrm{row}}} (1-g_i^{-s})^f.$$

Row reduction and free entries. At T_w the compact column level fixes a nonzero covector. The orbit of that covector has $T_w^k - 1$ elements at each split place. With the specified trivial lattice stabilizers, counting its orbits over the definite genus gives the class number $\mathfrak{M}_k I_{\mathrm{col}}(T_w^k-1)^f$. In counting columns the trivial stabilizers let every double coset contribute its full compact representation dimension, including with the imposed local characters. Together with the row formula this gives

$$d_w^* = \mathfrak{M}_k I_{\mathrm{col}} \prod_{\rho} \frac{V(\mathbf{q}_{\rho,w})}{J_k}, \qquad D_\star = \mathfrak{M}_{r_{\mathrm{row}}} I_{\mathrm{row}} \prod_{\rho} \frac{V(\mathbf{y}_\rho)}{J_{r_{\mathrm{row}}}}.$$

The same full column coefficient and row main term. The freeness needed for the column count can also be seen from the large split prime in the two fixed levels. A finite lattice automorphism group intersected with the lifted unipotent reduction subgroup is a g_i-group, since the principal congruence kernel is pro-g_i. An element of order g_i would require $g_i - 1 \le [E : \mathbf{Q}]k = 2fk$, by restriction of scalars and the cyclotomic minimal polynomial. The chosen $g_i > 2fk+1$ rules this out.

For $k < s \le r_{\mathrm{row}}$ the successive mass quotient is

$$\frac{\mathfrak{M}_s}{\mathfrak{M}_{s-1}} = \frac{\big((s-1)!\,(v_0/\pi)^s\big)^f}{(2v_0)^{f/2}}\, L_F(s, \eta_{E/F}^s), \tag{161}$$

where $\eta_{E/F}$ is the quadratic character of E/F and the even power is interpreted as the trivial character. In the unramified self-dual case, the mass formula gives

$$\mathfrak{M}_h = 2\,2^{-hf} \prod_{s=1}^{h} |L_F(1-s, \eta_{E/F}^s)|.$$

The local correction factors are 1 because the lattice stabilizers are hyperspecial. [63, Proposition 2.13, (4.2)–(4.3), Propositions 4.4–4.5] The character $\eta_{E/F}^s$ is unramified at finite places and has parity s mod 2 at every real place. Its functional equation gives

$$|L_F(1-s, \eta_{E/F}^s)| = |\operatorname{disc} F|^{s-1/2} \left(\frac{2(s-1)!}{(2\pi)^s}\right)^f L_F(s, \eta_{E/F}^s).$$

Here the real gamma ratio is $\Gamma_{\mathbf{R}}(s+a)/|\Gamma_{\mathbf{R}}(1-s+a)| = 2(s-1)!/(2\pi)^s$, with $a = s \bmod 2$. Taking the quotient of consecutive masses and using $|\operatorname{disc} F| = (2v_0)^f$ gives (161). For $s \ge 2$ the Euler factors give $|\log L_F(s, \eta_{E/F}^s)| \le f \log \zeta(s)$.

In $D_\star/d_w^*$ the factorials telescope against the Weyl denominators. Taking logarithms gives, uniformly in w,

$$\begin{aligned}\log X_w = &\left(m-1+\frac{m(m-1)}{2k}\right)\log(v_0/\pi) - \frac{m-1}{2k}\log(2v_0) + \left(m+\frac{m(m-1)}{2k}\right)\lambda_g \\ &+ \frac{1}{fk}\sum_\rho \log\frac{V(\mathbf{y}_\rho)}{V(\mathbf{q}_{\rho,w})} + O(1).\end{aligned} \tag{162}$$

Indeed the mass contribution from the L-values is bounded by $(m-1)\log\zeta(2)/k$, and the finite-level Euler factors contribute at most $2k^{-1}\sum_{i=0}^{1}\sum_{s>k} g_i^{-s}$ in absolute value. Both contributions are independent of the word. Identity before bounding the products.

The Harish–Chandra arrays at a place ρ have the following form. Write $v_i = v_{\rho,i}$ for the $k-1$ nonnegative entries in the row and $-\beta_j = -\beta_{\rho,j}$ for its m negative entries. The v_i decrease with step 4Δ, where $\Delta \geq 1$. The β_j increase with gaps at least 1. For $0 \leq j < m$, the stage weights $w_j = w_{u,j+1}(\rho)$ satisfy

$$4k\Delta \leq \beta_0/x, \qquad 0 < \beta_j < w_j \leq 2b, \qquad w_j \leq (1+1/x)\beta_j.$$

The interior column entries are $v_{i-1} - 1/2$ or $v_{i-1} - 1/2 - \Delta$ for $1 \leq i \leq k-2$. The two outer entries are independent of the word. All column differences in increasing index order are between 1 and $2a_E$. These are the chosen grid and order conditions at the selected translate. The embeddings ρ correspond to P, and $v_i \leq 4k\Delta$. Grid conditions and matching of the weights.

The interior log Vandermondes for $v_0, \dots, v_{k-3}$ and for the interior column differ by $O(k\log k)$. At index distance $h \geq 1$, the initial difference is $4\Delta h$ and the perturbation has magnitude at most Δ, so the change in its logarithm is at most $1/(3h)$. Summing over the pairs proves the bound. Adding the missing entry v_{k-2} on the row or the two endpoints on the column contributes $O(kL)$ in either case, since the added differences lie between 1 and $\max(4k\Delta, 2a_E)$. Thus

$$\log\frac{V(\mathbf{y}_\rho)}{V(\mathbf{q}_{\rho,w})} = (k-1)\sum_{j=0}^{m-1}\log\beta_j + \log V(-\boldsymbol{\beta}_\rho) + O(kL + km/x).$$

Here the cross differences are $v_i + \beta_j$, and $0 \leq \log(1+v_i/\beta_j) \leq 1/x$. The negative block contributes between 0 and $\binom{m}{2}\log(2b)$. Replacing the first term by $k\sum_j \log w_j$ costs at most $m\log(2b) + km/x$. Since $m \leq k$ and $m \leq x$, the last display yields the lower bound $k\sum_j \log w_j - O(kL)$ and the upper bound $O(kmL)$.

For the variation between two words at the same translate, the only word-dependent Vandermonde in (162) is the column one. Interior differences of distance h have logarithmic variation at most $2/(3h)$, by comparing both with the same grid. For an endpoint and the i-th interior entry counting from that endpoint, with $i > 1$, the two distances are at least $3\Delta(i-1)$, and their difference is at most Δ. The nearest entry at either end costs at most $\log(1+\Delta)$, because both distances are at least 1. This gives $O(k\log k + \log(1+\Delta)) = O(k\log k)$ at each place. Here $4k\Delta \leq 2b/x$ and $k \geq xm$ imply $\log(1+\Delta) \ll k$. The endpoint-to-endpoint factor is common. Division by fk gives $O(\log k)$ in $\log X_w$. The common Γ cancels from the difference of the logarithmic prime windows, whose residual width is at most $1/(fk)$.

Finally $0 \leq \log(v_0/\pi) \leq \log(2v_0) \leq L$ and $0 \leq \lambda_g \ll L$ on the chosen family. The coefficients of $\log(v_0/\pi)$ and λ_g in (162) are at least $m-1$ and m, respectively. The prime window changes the lower bound by at most 2. Since $L_a \geq L$, we obtain a uniform constant C_{count} for

$$\begin{aligned}&\Lambda_{\mathrm{count}}(u) \leq \log T_w \leq C_{\mathrm{count}} mL, \qquad |\log T_w - \log T_{w'}| \leq C_{\mathrm{count}}\log k, \\ &\Lambda_{\mathrm{count}}(u) = \sum_{j=1}^{m}\frac{1}{f}\sum_{r\in P}\log(w_{u,j}(r)v_0/\pi) + m\lambda_g - E_{\mathrm{count}}, \qquad E_{\mathrm{count}} = C_{\mathrm{count}}L_a(1+m^2/k).\end{aligned} \tag{163}$$

The calculation gives this statement at the chosen levels and with the actual stage weights. The corresponding arithmetic transport sums over the same stages and embeddings after the tied identifications. Count at an actual index. Transport to the tied translates.

The finite allowance and the word exponent. Let $E_{w,j}(v)$ be the prescribed signed local allowance entering $Q_{w,j} = 2\sum_v e_v\lambda_v E_{w,j}(v)$. Write

$$d_{w,j,\rho} = \sum_{a<N(j-1)} \nu_{w,j,\rho,a}, \qquad C_{i,w,j} = \prod_{\rho} \rho(\overline{\pi_{g_i}})^{d_{w,j,\rho}}.$$

Here ρ in the product is taken in the common coefficient field with its involution. The prime generator in the split table at g_i is precisely π_{g_i}, with $\pi_{g_i}\overline{\pi_{g_i}} = g_i$. Each coefficient of the fixed rectangular jet is homogeneous of degree $d_{w,j,\rho}$ in the block at ρ. In the split table the integral derivative is computed after multiplication of its variable by $\rho(\overline{\pi_{g_i}})$. Thus the derivative term in a finite bound has denominator $C_{i,w,j}$. A separate additive Fourier average on the selected split half contributes at most $g_i^{fN(j-1)N(j)}$ to the denominator. To see this power, coarsen the modulus to $g_i c'$ with $g_i \nmid c'$ using the periodicity at g_i. The prime-to-g_i Fourier index is a unit. The g_i-primary lattice quotient in the average is

$$\overline{\pi_{g_i}}^{-1}\Lambda\big/\big(\pi_{g_i}/\overline{\pi_{g_i}}\big)\Lambda,$$

up to scalar rescaling, where Λ is the period lattice in the Hom space. Its order is $g_i^{fN(j-1)N(j)}$, since the complex determinant of the norm-g_i generator has squared absolute value $g_i^{fN(j-1)N(j)}$. The split theta table is evaluated in the scaled integral tangent coordinates. If $\overline{\pi_{g_i}}$ is small, the prime-primary translates are the étale half, and the integral Taylor estimate uses the tangent matrix with all entries divisible by $\overline{\pi_{g_i}}$. If it is a unit, the connected-half projector has integral Taylor coefficients. Both branches give the bound with the same split-table factor and with unit prime-to-g_i phases. Applying the integral Weyl jet projection and the integral terminal branching map preserves it. In the class sum the source stabilizer denominators are cleared by D_Γ. The actual prime generators and entry loss, Passage to the coefficient-field valuation, The normalized split Fourier index, The two split orientations, and The integral projections and the terminal evaluation.

For the allowance used in the division calculation one also subtracts the reference contribution $D_{\mathrm{Fock},j}/2$ in g_i-valuation units. Here $D_{\mathrm{Fock},j}$ sums the positive Fock rows and the absolute values of the negative Fock rows. In squared norms the exponent is $-D_{\mathrm{Fock},j}$. The definition of the allowance gives, at $v \mid g_i$,

$$E_{w,j}(v) = \operatorname{ord}_{g_i} D_\Gamma - \tfrac{1}{2}D_{\mathrm{Fock},j} + \operatorname{ord}_v(C_{i,w,j})/e_v + c_g f(n_*^2 + m^2\eta_{\mathrm{sz}}).$$

The last term is the chosen comparison allowance. It covers the rectangular Fourier loss and the pooled chain losses in the common local norm bounds. This signed allowance enters the word exponent. Its sum can be computed directly from the prime generators. Indeed $C_{i,w,j}\overline{C_{i,w,j}} = g_i^{\sum_\rho d_{w,j,\rho}}$, and the involution permutes the places above g_i, preserving $e_v\lambda_v$. Since $\sum_{v\mid g_i} e_v\lambda_v = \log g_i$, their contribution is exactly

$$\left(\sum_\rho d_{w,j,\rho} - D_{\mathrm{Fock},j} + 2\operatorname{ord}_{g_i} D_\Gamma + 2c_g f(n_*^2 + m^2\eta_{\mathrm{sz}})\right)\log g_i.$$

The definition of the signed allowance, The signed Fock degree, and The conjugate-place sum.

For the size of the degree, the first k rows have entries $w_{j\rho} + u_a$ up to half-sum shifts of size $O(n_*)$. At an unsampled embedding the cloud and both endpoints cost $O(k(S_c + n_*))$ in that block. At a sampled embedding the sum of the two endpoints is $a_\rho + w_{j\rho} + O(S_c + n_* + b'/x + \eta_{\mathrm{sz}})$. This follows by adding their shifts and using the window bound $w_{j\rho} = b_\rho + O(\eta_{\mathrm{sz}} + b'/x)$ at the lower endpoint. The interior rows at that embedding have the same cloud bound. In the appended blocks each row is bounded by $2\eta_{\mathrm{sz}}$ by the level-node differences. Consequently

$$d_{w,j,\rho} = (k-1)w_{j\rho} + \begin{cases} a_\rho, & \rho \in T, \\ w_{j\rho}, & \rho \notin T \end{cases} + O\big(k(S_c + n_*) + m(t_{\mathrm{cut}} + 1)\eta_{\mathrm{sz}} + \mathbf{1}_{\rho\in T}b'/x\big).$$

For the Fock degree put $p = k + j - 1$, $q = m - j$ and $e = w_{j\rho}$. The positive partition is the shifted cloud followed by the first j nodes, and the negative partition in reverse order uses the remaining q nodes. Their degree at ρ, minus $(k-1)e$, is

$$\sum_{a<k-1} v_a + \sum_{B<j}(e - b_B) + \sum_{j\le B<m}(b_B - e) - \tfrac{1}{2}\big(p(p + t_{\mathrm{cut}}) + q(q + t_{\mathrm{cut}})\big).$$

In particular its absolute value is $O(kS_c + m\eta_{\mathrm{sz}} + kn_*)$. The sum of the two degree calculations, the size of the level nodes, and the bound $\log g_i \ll L$ give a deviation $O(B_0)$ from $S_a \log g_i$. Here $f\eta_{\mathrm{sz}} \le b$, $m \le f$ and $n_*^2 \le \eta_{\mathrm{sz}}$ give $f(n_*^2 + m^2\eta_{\mathrm{sz}}) \ll fmb$. The two endpoints of the actual jet, The shifted rows and appended blocks, and The sum of the paired Fock degrees.

For the places away from g_0, g_1, group the allowances by their rational prime. At T_w the nonzero contribution is bounded by

$$2\left(2\sum_{\rho} W_{w,\rho} + fn_* + N_T(t_{\mathrm{cut}} + 1)(m+1)/4\right)\log T_w.$$

Here $W_{w,\rho}$ is the maximum over the column chain of the absolute shifted row weights with the sampled first coordinate omitted. The bound uses $\sum_{v|T_w} e_v\lambda_v = \log T_w$; the coefficient of two can be improved using the restriction to the seen places. The actual $W_{w,\rho} \le 2b$ and the level bound $\log T_w \ll mL$ already give $O(B_0)$. The total stabilizer allowance is at most $2\log D_\Gamma = O(B_0)$. The chosen stabilizer multiple divides the Minkowski integer for rank $2fn_*$ over $\mathbf{Z}$, so $\log D_\Gamma = O(fn_*\log(2fn_*))$.

At the discriminant prime d, pair the two places ρ_0, ρ_1 of each host label r, with the first one the possible sampled place. The discriminant allowance apart from the stabilizer is

$$-\frac{k-1}{2}\sum_r \min(w_{j\rho_0}, w_{j\rho_1}) + \frac{1}{2}\max_w \sum_r c_{w,j,r},$$

$$c_{w,j,r} = d_{w,j,\rho_0} - \mathbf{1}_{r\in T}(\nu_{w,j,\rho_0,0} - \nu_{w,j,\rho_0,1}) + \mathbf{1}_{r\in T}\nu_{w,j,\rho_1,0}.$$

This removes the large first coordinate at the sampled place and uses the first coordinate at the other place. In fact $c_{w,j,r} \ge 0$. The first-place weight is $O(b')$. At a sampled place the two replacements cost $O(b + kb')$ by the cloud and the companion weight bound, and elsewhere $c_{w,j,r} = O(kb')$. Thus the absolute discriminant allowance is $O(N_T b + fkb')$, and its weighted cost is $O(B_0)$ since $kb' \ll b$ and $\log d \ll L$. The exact level and ramified allowances, The stabilizer order on the same family, and The two discriminant weights.

The remaining nonzero allowance, apart from stabilizers, is on the specified inert or signal primes $p \le X''$. For the branch with a single trivial orbit its magnitude times $\log p$ is bounded by $O(fn_*^2 L + fb(\log p)/(p^2-1))$. The other branch has weighted bound

$$O\left(fmb\frac{\log p}{p} + f\eta_{\mathrm{sz}}L + \mathbf{1}_{p\in C}fkb\frac{\log p}{p^3}\right),$$

where C is the cubic correction class among the prescribed primes. In the orbit term this follows from $\gamma(p) \le 2/p$ times the input weight sum $O(fmb)$, because its first-type orbit only uses thin weights $O(b' + \eta_{\mathrm{sz}})$. The other nonlogarithmic terms use $\sum_\rho w_{j\rho} = O(fb)$. The dimension terms are bounded by $O(f\eta_{\mathrm{sz}}L)$, using the size inequalities on the two logarithmic derivative losses. Summation uses $X''\eta_{\mathrm{sz}} \le b$, $n_*^2 \le \eta_{\mathrm{sz}}$, and $\sum_{p\le X''}(\log p)/p = O(1 + \log X'') = O(L)$. In C one has $p \equiv \pm 1 \pmod{2^e}$, with $2^e \ge x^2$ and $k \le x^5$. Summation along these two progressions bounds $k\sum_{p\in C}(\log p)/p^3$ by $O(k\log(2^e)/2^{3e}) = O(1)$. Thus the entire remaining sum is $O(B_0)$. Combining these bounds on the actual support yields, uniformly,

$$Q_{w,j} = S_a\log(g_0g_1) + O(B_0), \qquad Q_w = mS_a\log(g_0g_1) + O(mB_0).$$

The trivial orbit branch, The nontrivial orbit branch, The sparse cubic correction, and The weighted sum at all off-type places.

Cancellation. The division height has main term $P_{\mathrm{div}} = m\sum_{t\in T}(a_t \log a_t - a_t - a_t \log \Omega)$. Substitution of (160) and (163) cancels the factors $\log(\pi/v_0)$ and $m\lambda_g$. With

$$M(a, w) = a\Big(\log a - 1 - \gamma + \ell_c - \frac{1}{f}\sum_{r\in P} \log w_r \Big), \tag{164}$$

the exact identity for the leading terms is

$$P_{\mathrm{div}} - S_a\Lambda_{\mathrm{count}}(u) + mS_a\lambda_g = \sum_{t\in T}\sum_{j=1}^{m} \mathfrak{R}M(a_t, w_{u,j}) + S_aE_{\mathrm{count}}. \tag{165}$$

4.5 The upper bound at each translate

Apply the two-run determinant bounds under the common-record identification and the source slope and height bounds under the ramified donor identification in Section 6.1. The resulting upper bound is conditional on both identifications and the stated candidate conclusions.

Use the centered statistic

$$D(u) = \sum_{t\in T}\sum_{j=1}^{m}\big(A(a_t, t, w_{u,j}) - \mathfrak{R}M(a_t, w_{u,j})\big). \tag{166}$$

Its prime-power cutoff is (96) at every translate and coincides with the cutoff in Section 5.

Take a nonnegative word-uniform bound $E_{\mathrm{div}}(u)$ for the remainder in (153). The lower level bound and $S_a \ge 0$ give the same allowable height for all words at this translate,

$$h^*(u) = \mathcal{F} - \mu(C) + P_{\mathrm{div}} - S_a\Lambda_{\mathrm{count}}(u) + E_{\mathrm{div}}(u) \ \ge h_w.$$

Use $h^*(u)$ in the frame estimate and let s be the rank-weighted mean of the adjusted source slopes $\mu(N'_w) - h^*(u)$. Write $R_*(u) = \sup_w R_{\mathrm{loc}}(w)$. Subtracting $h^*(u)$ from the lower bound in (119) cancels $\mathcal{F}$. Using (165) and taking the rank-weighted mean gives

$$\begin{aligned} s - \mu(C) \ \ge \mathcal{G}(u) &- \sum_{t\in T}\sum_{j=1}^{m} \mathfrak{R}M(a_t, w_{u,j}) \\ &- O\big(N_T a_E L_a(1 + m^2/k) + m(B_0 + N_c)x^{e_D\varepsilon} + R_*(u)\big)\,. \end{aligned}$$

The height, count and source slope errors combine here since $S_a \le 2N_T a_E$, $m \ge 1$ and $e_D \ge 0$. For the upper spread, subtract the lower adjusted mean slope of the donor at w' from the upper adjusted maximal slope at w. The common height and the term $\mathcal{G}(u) + \mathcal{F}$ cancel. The difference between $Q_{w'}$ and Q_w is $O(mB_0)$. Taking the rank-weighted mean over w' and the maximum over w gives

$$U \ll N_T a_E L_a(2 + m/k) + mB_0 + R_*(u).$$

Increasing each height to $h^*(u)$ gives $s \le s_{\mathrm{raw}}$, so (158) bounds $s + \mu_{\max}(R) \le x^{C_H}b$ at the same prescribed exponent. Since $C = R^\vee$, it follows that $H_* = \max(0, s - \mu_{\min}(C)) \le x^{C_H}b$. With the trace and row fraction bounds in the actual metrics, (69) gives $s - \mu(C) \le 8U + 12b$.

Put $B = N_T a_E(1 + m^2/k) + m(B_0 + N_c)$. The size inequalities $kb' \le a_E$, $X'' \le a_E$, and $\log T_w \ll mL$, together with the index part of the gain estimate, give

$$R_*(u) \ll N_T mb x^{e_L\varepsilon} + N_T a_E(m^2/k)L + a_E + 1 \ll Bx^{\max(1,e_L)\varepsilon}.$$

Indeed $mb \le a_E$, $a_E \ge 1$ and $N_T \ge 1$ by (70). The other size inequalities give $L_a \le 2L$ and $m/k \le 1$. Thus $U \ll Bx^{\max(1,e_L)\varepsilon}$. The same upper bound holds for $s - \mu(C)$ since $b \le B$. The prime-power part of (97)

gives $\sum_{t,j} A(a_t, t, w_{u,j}) - \mathcal{G}(u) \ll N_T m b x^{e_L \varepsilon}$. Combining the gain estimate with the lower adjusted slope comparison bounds $D(u)$ above by $s - \mu(C)$ plus their errors. Consequently a uniform constant C_2 gives

$$D(u) \le C_2 B x^{\max\{e_D, e_L, 1\}\varepsilon} \le C_2 B x^{(e_D + e_L + 2)\varepsilon}.$$

All implied constants depend only on the fixed comparison data. For fixed $\varepsilon > 0$ the factor x^{ε} absorbs them past a common threshold. Inserting $B_0 = 2nmbx^{\varepsilon}$ and $N_c \le N_2 k b x^{2\varepsilon}/n^2$ gives the bound

$$\begin{aligned} &D(u) \le E_x \qquad (u \in H), \\ &E_x = x^{(e_D + e_L + 3)\varepsilon} \left[N_T a_E (1 + m^2/k) + m \left(2nmbx^{\varepsilon} + \frac{N_2 k b x^{2\varepsilon}}{n^2} \right) \right]. \end{aligned} \tag{167}$$

The factor is $x^{5\varepsilon}$ for $(e_D, e_L) = (2, 0)$. Section 5 uses (167) at every translate in H_0, with the same base field, samples and cutoff.

5 Fourier inversion and the pointwise bound

Fix the target and $\varepsilon > 0$ as in §1. At a real scale $x \ge 3$, use the Artin data and character estimates at the base index i_0. Keep the sampling sets, degrees and cutoff T, T_2, a_t, N of Section 4, and the weights $w_{u,j}$ from the selected indices $i(u)$ in (166). Use the same family at every translate, as in Section 6.1, and assume the bound $D(u) \le E_x$ of (167) for every translate in the average.

5.1 Sampling degrees and weights at the tied indices

Write $n = |H|$, $f = 2n$, $N_T = |T|$, $N_2 = |T_2|$, and $a_E = xb$ with $b = \exp(x^{\varepsilon})$. In the chosen records the other scales are

$$X' = \exp(x^{\varepsilon/3}), \qquad b' = \frac{a_E}{X'}, \qquad w'' = \frac{b'}{xn}, \qquad \delta_1 = x^{-\varepsilon/8}.$$

The sampling integers a_t, b_t are read at i_0 and carried by the same $t \in T$ to each $i(u)$. Their prescribed intervals are

$$a_E \le a_t \le a_E + \frac{a_E}{k x^{13}}, \qquad b' \le b_t \le b' \left(1 + \frac{2n\delta_1 x^{\varepsilon} L_{\mathrm{a},G}}{k} \right), \qquad L_{\mathrm{a},G} \le x^{\varepsilon + \varepsilon/100}.$$

Here $L_{\mathrm{a},G}$ is the parameter in the upper endpoint of the interval for b_t. The degree window gives $0 < a_t \le 2a_E$ and, with $y_t = a_t/b$, gives $0 \le y_t - x \le 1/(kx^{12}) \le 1$. The intervals for the sampling integers and their transport are recorded in the size data and the family data.

Use the coordinates $r \in P = \{(0, i, s) : i \in \mathbf{Z}/2\mathbf{Z},\ s \in H\}$ for the stage weights. In the original q, d sign order, the CM embedding at $(0, i, s)$ has signs $((-1)^i, 1)$ and last component s. The prescribed CM type has exactly one embedding with each such label. At $i(u)$ the actual weight is its Harish–Chandra entry at stage j. With the common $\eta_{\mathrm{wt}} > 0$ of the tied records, its prescribed intervals are

position r	interval for $w_{u,j}(r)$
$(0, 1, s + u),\ s \in T_2$	$[b, b + \eta_{\mathrm{wt}} + 1]$
$(0, 0, t),\ t \in T$	$[b_t + b'/x, b_t + b'/x + \eta_{\mathrm{wt}}]$
all other positions in P	$[w'', w'' + \eta_{\mathrm{wt}}]$.

The first two rows are disjoint in their second coordinates. The intervals use the selected embeddings and weights, as in the three weight ranges. The rank and size ranges give $n \le x^{9/4}$, $x^4 \le k \le x^5$, and $\log \eta_{\mathrm{wt}} \le x^{\varepsilon/4}$ at large x. In particular,

$$\frac{2n\delta_1 x^{\varepsilon} L_{\mathrm{a},G}}{k} \le 2x^{-7/4 + (201/100)\varepsilon} \longrightarrow 0, \qquad w'' = \frac{\exp(x^{\varepsilon} - x^{\varepsilon/3})}{n} \longrightarrow \infty, \qquad \frac{\eta_{\mathrm{wt}}}{b'} \longrightarrow 0.$$

Thus all weights are ≥ 1 at large x. Also $w''/b' = 1/(xn) \leq 1/x$, and $b_t/b' \leq 1 + o(1)$, so both of the last two rows have upper bound $2b'$ at a common threshold. Since $a_t/w_{u,j}(r) \leq 2a_E$, the cutoff $N = \lceil 2a_E \rceil$ in (96) applies to the actual degrees and weights. These same scales give $kb' \leq a_E$ since $k \leq x^5 \leq X'$ for large x, and $X'' = a_E/w'' \leq a_E$.

Put $\vartheta = \exp(-c_6 x^{\varepsilon/8})$ with the host constant $c_6 > 0$ from §1. Since $f = 2n$ is polynomially bounded and $\eta_{\mathrm{wt}} \leq \exp(x^{\varepsilon/4})$, both $6f(x+1)(\eta_{\mathrm{wt}}+1)/b$ and $2f(x+1)\vartheta$ are $o(x^{-1})$. Thus at a common threshold

$$6f(x+1)(\eta_{\mathrm{wt}}+1) \leq b/x, \qquad 2f(x+1)\vartheta \leq 1/x. \tag{168}$$

Each threshold here follows from the fixed precision and the uniform record bounds. In particular it applies to every j and u at one scale, using the actual a_t, b_t from the common seed.

5.2 Finite Fourier inversion

Use the group Γ, the subgroup $P = \ker c$ of order $f = 2|H|$, and the bijection $p_t : \Gamma \setminus P \to P$ in (95). The base field is a subfield of $\mathbf{Q}(\zeta_{Q_S qd})$, and its Artin map is the quotient of the cyclotomic action. For $\xi \in \widehat{\Gamma}$, let χ_ξ be the unique primitive Dirichlet character inducing $\xi \circ \mathrm{Art}$ on units modulo $Q_S qd$. Its conductor divides $Q_S qd$, and for $\xi = 1$ its conductor is 1. Write $G_\xi = G_{\chi_\xi}$ for the centered sum of (2), with (1) for the trivial character. For $\xi \neq 1$ put $\ell_\xi = L'(1, \chi_\xi)/L(1, \chi_\xi)$, using the nonvanishing at 1 in §1.

For a one-dimensional character ξ of the Galois group of a finite abelian extension of $\mathbf{Q}$, let $V_\xi = \mathbf{C}$ with the action by ξ. At each prime p, let I_p be inertia and let F_p lift the arithmetic Frobenius, which sends z to z^p on the residue field. To compare the two Euler factors for our field, write $Q_S qd = p^\alpha N_p$ with $p \nmid N_p$. By the cyclotomic decomposition of Section 1, inertia is the image of $(\mathbf{Z}/p^\alpha\mathbf{Z})^\times$ in Γ, and a Frobenius lift can be represented by a unit whose component modulo N_p is p mod N_p. Thus ξ is trivial on I_p exactly when its pullback factors through N_p, or equivalently $p \nmid \mathrm{cond}(\chi_\xi)$. In that case every lift has $\xi(F_p) = \chi_\xi(p)$. When $p \mid \mathrm{cond}(\chi_\xi)$, the invariant space is zero and $\chi_\xi(p) = 0$. Both cases give

$$L_p(s, \xi) = \det\Bigl(1 - p^{-s} F_p \mid V_\xi^{I_p}\Bigr)^{-1} = \bigl(1 - \chi_\xi(p) p^{-s}\bigr)^{-1} \qquad (\Re s > 1).$$

The logarithm given by the series in p^{-s} is

$$\log L_p(s, \xi) = \sum_{h \geq 1} \frac{v_\xi(p^h)}{h\, p^{hs}}, \qquad v_\xi(p^h) = \frac{1}{|I_p|} \sum_{i \in I_p} \xi(hF_p + i) = \xi(F_p)^h \frac{1}{|I_p|} \sum_{i \in I_p} \xi(i) = \chi_\xi(p^h).$$

The inertia average equals 1 on characters trivial on I_p and vanishes on the others. This uses the primitive character even at primes dividing $Q_S qd$ whose local p-part is trivial. The coset $hF_p + I_p$ is independent of the chosen lift because changing F_p by an element of I_p changes hF_p by an element of I_p. [71, (3.2)–(3.5), (3.15), and §4.2]

For a finite abelian group B written additively, the dual group $\widehat{B}$ has $|B|$ characters. For every $z \in B$ and $\xi \in \widehat{B}$, the two orthogonality relations are

$$\frac{1}{|B|} \sum_{\xi' \in \widehat{B}} \xi'(z) = \mathbf{1}_{\{z=0\}}, \qquad \frac{1}{|B|} \sum_{z' \in B} \xi(z') = \mathbf{1}_{\{\xi=1\}}.$$

On each cyclic factor the first sum is a geometric series. For the second, translating the sum by an element with character value different from 1 forces it to vanish for a nontrivial character. Consequently, for every $F : \Gamma \to \mathbf{C}$ and $g \in \Gamma$, Fourier inversion in the convention used here is

$$F(g) = \frac{1}{2f} \sum_{\xi \in \widehat{\Gamma}} \overline{\xi(g)} \sum_{z \in \Gamma} F(z)\xi(z).$$

[74, Theorem 4.4, (4.12)–(4.13)] Apply this with $F(z) = |I_p|^{-1}\mathbf{1}_{\{z-hF_p\in I_p\}}$. The inner sum is $v_\xi(p^h)$, giving

$$\frac{1}{f}\sum_{\xi\in\widehat{\Gamma}}\overline{\xi(g)}v_\xi(p^h) = \frac{2}{|I_p|}\mathbf{1}_{\{g-hF_p\in I_p\}}.$$

Multiplying by $\Lambda(p^h) = \log p$ on $\Gamma\setminus P$ recovers the coefficient in (95). When $p \mid Q_S q$, both the Frobenius lift and inertia lie in P since p splits in $\mathbf{Q}(\sqrt{-d})$. For the classical host these subgroups are explicit in (4). At $p = d$, the two elements of $hF_d + I_d$ have opposite c-values, so there is one element outside P for each $h \ge 1$, of weight $\log d$. For $p \nmid Q_S q d$, the subgroup I_p is trivial. The unique term of the coset has weight $2\log p$ if $c(F_p) = -1$ and h is odd. These are the support conditions used in Section 4.

For $r = p_t(g)$, factor the positive part as $(a/p^h - w_r)_+ = w_r((a/w_r)/p^h - 1)_+$. At $a = a_t$ and $w = w_{u,j}$, the common cutoff N in (96) satisfies $a/w_r \le N$. Thus the sum through N equals the unrestricted sum in (95) used to define $D(u)$. The weighted uncentered sum is $w_r G_\xi(a/w_r) + a(\log a - \log w_r - 1 - \gamma)$ for $\xi = 1$, and $w_r G_\xi(a/w_r) - a\ell_\xi$ for $\xi \ne 1$. On $\Gamma\setminus P$, character orthogonality on P makes the mean of $\overline{\xi(g)}$ vanish unless ξ is trivial there. The two remaining characters are $1, c$, with means $1, -1$, respectively. The bijection p_t identifies the sum of $\log w_{p_t(g)}$ over $\Gamma\setminus P$ with the sum of $\log w_r$ over P. The uncentered sum therefore equals the sum of weighted G_ξ's plus

$$a(\log a - 1 - \gamma) - \frac{a}{f}\sum_{r\in P}\log w_r - a(-1)\ell_c = M(a, w).$$

Here the last term uses the same primitive η_d as the CM normalization in (164). We get

$$A(a, t, w) - M(a, w) = \frac{1}{f}\sum_{g\notin P}\sum_{\xi\in\widehat{\Gamma}}\overline{\xi(g)}\, w_{p_t(g)}G_\xi\big(a/w_{p_t(g)}\big). \tag{169}$$

Taking real parts and summing over (t, j) with the weights at $i(u)$ gives $D(u)$. Formal identity.

5.3 Reduction to the large weights

Let $H_0 \le H$ be the distinguished cyclic subgroup from §1, with $o = [H : H_0]$ equal to the order of the target character ψ. Embed H_0 into Γ by $h \mapsto (0, 0, h)$. Put $\delta_2 = N_2/n$. For each $u \in H_0$ and stage j, the positions with $s \in T_2 \subset H_0$ in the first row of the weight table are distinct. Section 5.1 places their weights in $[b, b + \eta_{\mathrm{wt}} + 1]$ and all other weights in $[1, 2b']$.

For every real $z > 0$, the Chebyshev function satisfies $\mathcal{V}(z) := \sum_{n\le z}\Lambda(n) < 1.03883\, z < 3z$. [72, Theorem 12, (3.35)] Write $z = a_t/w$. In the nonprincipal case $wG_\xi(a_t/w) = \sum_{n\ge1}\Lambda(n)\chi_\xi(n)(a_t/n - w)_+ + a_t\ell_\xi$. The last term is constant in w. Between ramp breakpoints the derivative is $-\sum_{n<z}\Lambda(n)\chi_\xi(n)$ for $\xi \ne 1$, and $z - \sum_{n<z}\Lambda(n)$ for $\xi = 1$. In both cases its modulus is at most $z + \mathcal{V}(z) < 3z$. On the replacement interval $w \ge b$ we have $z \le x + 1$. Continuity at the breakpoints bounds the change by $3(x + 1)(\eta_{\mathrm{wt}} + 1)$.

For a small weight w, $b'X' = xb$ and $x \ge 3$ imply $a_t/w \ge (x - 1)b/(2b') \ge X'/3$. The base host satisfies

$$|G_\xi(y)| \le \vartheta y \qquad (\xi \in \widehat{\Gamma},\ y \ge X'/3).$$

Thus each omitted term has modulus at most $(x + 1)b\vartheta$. The estimate covers every ξ, because the small weights are omitted before taking the average over H_0. The sum in (169) has f positions, $2f$ characters and prefactor $1/f$. Multiplying each of the two error bounds by $2f$ and applying (168) gives a total error at most $2b/x$ for each (t, j, u). This uses the tail (11) at i_0.

5.4 The translate average

Write $e(h) = (1, 0, h)$. The preimage of $(0, 1, s+u)$ under p_t is $e(s+u-t)$. Its conjugate character value is

$$\overline{\xi(e(s+u-t))} = \xi(e(t))\,\overline{\xi(0,0,s)}\,\overline{\xi(0,0,u)}.$$

Indeed $e(s+u-t) = -e(t) + (0,0,s) + (0,0,u)$ since the first coordinate has order two. By character orthogonality on H_0, the mean over u vanishes unless $\xi \in H_0^{\perp} = \{\xi : \xi|_{H_0} = 1\}$. For such a character, $s \in H_0$ also gives a factor of 1, so all N_2 large weights contribute equally. Set $Q = N_T m b$. The mean of $D(u)$ is at most E_x. The replacement estimate bounds the real part of the reduced sum at a fixed u by $D(u) + 2Q/x$, since there are mN_T pairs (t, j). Each replaced term uses $b\,G_\xi(y_t)$, which is independent of the stage and the translate. It follows that

$$\frac{Q\delta_2}{2}\,\mathrm{Re}\,\frac{1}{N_T}\sum_{t\in T}\sum_{\xi\in H_0^{\perp}}\xi(e(t))G_\xi(y_t) \le E_x + \frac{2Q}{x}.$$

The factor $Q\delta_2/2 = mN_T b N_2/f$ comes from the m stages, the N_2 positions and the $1/f$ in (169). It is positive since $N_T, m, b > 0$ and $\delta_2 \ge x^{-\varepsilon}$.

5.5 Selecting the target characters

The characters of H/H_0 pull back along h_{nat} to the powers ψ^k, $0 \le k < o$. View each such character ν on H. As ν and $i, j \in \{0, 1\}$ vary, the factor $\nu(h_-)^{e_q}$ permutes the sign choices since $2h_- = 0$. Thus all $4o$ characters in $H_0^{\perp}$ have the form

$$\xi(e_d, e_q, h) = (-1)^{ie_q + je_d}\,\nu(h)\nu(h_-)^{e_q}.$$

In (3) the two factors $\nu(h_-)^{e_q(a)}$ cancel, so on units modulo $Q_S qd$ the pullback is $\psi^k \eta_q^i \eta_d^j$ as in the Dirichlet twist decomposition. Its primitive associate is $\chi_k \eta_q^i \eta_d^j$, where χ_k induces ψ^k modulo M and is primitive, with $\chi_0 = 1$. Omit a quadratic factor with exponent zero. The conductor is $\mathrm{cond}(\chi_k) q^i d^j$ since each selected factor retains its conductor in the coprime product. Thus every term with $(i, j) \ne (0, 0)$ satisfies the nontrivial quadratic twist hypotheses of Section 1. At $e(t) = (1, 0, t)$ the phase is $(-1)^j \nu(t)$. Let $\mathcal{X}$ consist of the characters with $i = j = 0$. Their inducing primitive characters form the fixed set $\mathcal{K}$ attached to the subgroup generated by ψ. Different powers give different primitive associates since an associate determines its inducing values on units modulo M. Thus $\mathcal{K}$ has o elements, contains 1, ψ and is closed under conjugation. The phase $\xi(e(t))$ is $\nu(t)$ for $\xi \in \mathcal{X}$, and ν is nontrivial when $\xi \ne 1$.

For $v \in H/H_0$, the proportion of T in v differs from $(3 + \mathbf{1}_{\{v=0\}})/(3o+1)$ by at most $x^{-1+\varepsilon}$. For a nontrivial character ν of the quotient, orthogonality gives

$$\frac{3\sum_{v\in H/H_0}\nu(v) + \nu(0)}{3o+1} = \frac{1}{3o+1}.$$

Thus $N_T^{-1}\sum_{t\in T}\xi(e(t))$ differs from $(3o+1)^{-1}$ by at most $o x^{-1+\varepsilon}$ for a nontrivial retained character. For $\xi = 1$ the coefficient is exactly 1. Put $a_1 = 1$ and $a_\chi = (3o+1)^{-1}$ for $\chi \in \mathcal{K} \setminus \{1\}$, and set

$$G_{\mathcal{K}}(x) = \sum_{\chi\in\mathcal{K}} a_\chi \Re G_\chi(x).$$

For every real $z \ge 2$, Mertens' estimate gives $\sum_{n\le z}\Lambda(n)/n = \log z + O(1)$, with an absolute implied constant. [74, Theorem 2.7(a)] Indeed, for $v = \lfloor z \rfloor$, summing $\log n = \sum_{d\mid n}\Lambda(d)$ up to v gives $\log(v!) = \sum_{d\le v}\Lambda(d)\lfloor v/d \rfloor$. Replacing each floor by v/d costs at most $\mathcal{V}(v) \ll v$. The integral comparison for $\log(v!)$ gives $\log(v!)/v = \log z + O(1)$. Between breakpoints, the derivative of the uncentered ramp sum at y is $\sum_{n<y}\Lambda(n)\chi_\xi(n)/n$, of modulus at most $\log y + O(1)$ since $|\chi_\xi(n)| \le 1$. For $\xi = 1$ the centering adds

$-\log y + \gamma$ to the derivative. For $\xi \neq 1$ it adds ℓ_ξ, which is fixed for $\mathcal{K} \setminus \{1\}$ and is $O(\log x)$ for the auxiliary twists by §1. Hence these continuous centered sums satisfy $|G_\xi(y_t) - G_\xi(x)| = O(\log x)$ on $[x-1, x+1]$. The same Mertens estimate gives an uncentered sum of modulus at most $x(\log x + O(1))$, so $|G_\xi(x)| = O(x \log x)$. Replacing y_t by x in the normalized sum costs $O(o \log x)$. Replacing the mean phase of a retained nontrivial character by its limit costs at most $o x^{-1+\varepsilon}|G_\xi(x)|$. The combined error is $O(x^\varepsilon \log x)$ for fixed target and precision, hence at most $4x^{1/2+\varepsilon}$ for large x. Write $\Sigma(x) = \sum_{\xi \in H_0^\perp \setminus \mathcal{X}} |G_\xi(x)|$. Each term with $\xi \notin \mathcal{X}$ has real part at least $-|G_\xi(x)|$, since its multiplier is the mean of unit complex numbers. The lower estimate for the average then gives

$$\frac{Q\delta_2}{2}\big(G_{\mathcal{K}}(x) - \Sigma(x) - 4x^{1/2+\varepsilon}\big) \le E_x + \frac{2Q}{x}. \tag{170}$$

The bound $|G_\xi(x)| \le x^{1/2+\varepsilon}$ for these twists gives $\Sigma(x) \le 4o\, x^{1/2+\varepsilon}$.

For $\psi = 1$, $o = 1$ and $H_0 = H$. The expression in the translate average at each t has the four terms

$$G_1(y_t) + G_{\eta_q}(y_t) - G_{\eta_d}(y_t) - G_{\eta_q\eta_d}(y_t).$$

The signs are the evaluations at $(1, 0, 0)$. The three auxiliary characters have conductors q, d, qd and enter $\Sigma(x)$ through the quadratic twist bounds. In particular the classical average is at least $G_1(x) - 3x^{1/2+\varepsilon} - C \log x$, with a constant uniform in the selected q, d, since $|\ell_\xi| \le 4 \log x$ for those twists. This calculation uses the size of T and the degree interval. The retained sum is $G_{\mathcal{K}} = G_1$.

5.6 The square-root scale

Write $C_{\mathrm{err}} = e_D + e_L + 3$. Dividing the budget of (167) by $Qx^{C_{\mathrm{err}}\varepsilon}$ gives

$$\frac{a_E}{mb} + \frac{a_E m}{bk} + 2\frac{n}{N_T} m x^\varepsilon + \frac{N_2}{N_T}\frac{k}{n^2} x^{2\varepsilon}.$$

In addition to the common bounds (70), we have $a_E = xb$, $N_T \ge n x^{-\varepsilon/8}/2$, and

$$x^{1/2-\varepsilon} \le m, \qquad k \le n^2 x^{1/2-\varepsilon}.$$

For $\psi = 1$, the sample counts are (10) and the bounds for m, k follow from the prescribed ranks in Section 1. The same size inequalities hold for the selected Dirichlet data. Also $N_2 \le n$ because $T_2 \subseteq H$. The first and third terms have bounds $2x/m$ and $4mx^{9\varepsilon/8}$, which balance at the square-root scale up to powers $x^{O(\varepsilon)}$. The four bounds in order are $2x^{1/2+\varepsilon}$, 2, $4x^{1/2+13\varepsilon/8}$, and $2x^{1/2+9\varepsilon/8}$. Their sum is at most $10x^{1/2+13\varepsilon/8}$. At every scale for which (170) holds, divide by $Q\delta_2/2 > 0$ and use $\delta_2 \ge x^{-\varepsilon}$. The factors from the main budget have exponent loss $C_{\mathrm{err}} + 13/8 + 1 = C_{\mathrm{err}} + 21/8$, giving

$$G_{\mathcal{K}}(x) - \Sigma(x) \le 20x^{1/2+(C_{\mathrm{err}}+21/8)\varepsilon} + 4x^{-1+\varepsilon} + 4x^{1/2+\varepsilon} \le 26x^{1/2+(C_{\mathrm{err}}+21/8)\varepsilon}.$$

For sufficiently large x, $26 \le x^{3\varepsilon/8}$ and $1 + 4o \le x^\varepsilon$. The first inequality bounds $G_{\mathcal{K}}(x) - \Sigma(x)$ by $x^{1/2+(C_{\mathrm{err}}+3)\varepsilon}$. The bound on $\Sigma(x)$ brings the coefficient to at most $1 + 4o$. Absorbing it gives

$$G_{\mathcal{K}}(x) \le x^{1/2+(e_D+e_L+7)\varepsilon}. \tag{171}$$

This estimate applies at each sufficiently large scale satisfying the stated hypotheses. The threshold depends on the target, the fixed precision and the uniform comparison constants. The set $\mathcal{K}$ and its weights depend only on the target. For the conductor-one target with the loss bounds $e_D = 2$ and $e_L = 0$, it gives $G_1(x) \le x^{1/2+9\varepsilon}$. Section 6.1 uses the same family and the inequality (167) at every scale of the chosen intervals. Formal pointwise estimate.

6 Bounds on intervals and the critical line

6.1 One family for the lattice and Fourier estimates

Take the conductor-one target $\psi = 1$. Fix the loss exponents $e_D = 2$, $e_L = 0$ and the balance constant $C_0 = 3$ used in the window construction.

The uniform comparisons of Sections 1–5 have finitely many precision cutoffs and size requirements. A lower bound may be a function of earlier constants. At fixed precision take the pointwise maximum K_* of the lower bounds in each position. In the position for C_H include the bound $1 + \max\{4P_t + 1, 9P_t + 2, 9P_t + 1, C_{F,0}\}$ from (158). Here P_t comes from the twist cutoff and $C_{F,0} = c_F + 1$ comes from the row height after the finite level change. Both are functions of the earlier data. Thus the same C_H is available for the height and the trace accuracy before choosing a scale or translate. A family meeting K_* therefore meets all of the size requirements. Each comparison also requires its hypotheses on the local models and maps. Staged lower bounds

Simultaneous family selection (input). Assume the availability of the arithmetic hosts on intervals from Section 1. There is an existence cutoff $\varepsilon_{\mathrm{ex}} > 0$. For each $0 < \varepsilon < \min(1/52, \varepsilon_{\mathrm{ex}})$ and each system K of staged lower bounds there exist one family $\mathfrak{F}$ and one exponent $C_H \geq 0$ meeting K. The family has the following data and conditions.

1. The set $\mathrm{Win}(\mathfrak{F})$ of real scales contains intervals $[Y, Y^{3/2}]$ with arbitrarily large Y. Each scale in this set has a nonempty index set I_x. Beyond a family threshold the indices use the prescribed arithmetic hosts at x.

2. At each such scale an index specifies E/F, a coefficient field with its places, the ranks and weights of the towers, the levels, the words and their local twists, the period spaces with their row and column coordinates, and the corresponding test and integral lattice data. A word, tuple, stage and place specify its local Hecke data and intended test in these coordinates. At a fixed x, the indices have specified identifications of the ranks, word labels and their letters, q, d, Q_S, the Artin coordinates and T, T_2. The rank, level-prime and sample cardinality conditions are those in Section 1. The stage weights and the range for b_t are the raw intervals in Section 5.1.

   With $b = \exp(x^{\epsilon})$, the level of each word satisfies $b^{m/2} \leq T_w$ and $\log T_w \leq C_T m \log b$. Write $f = [F : \mathbf{Q}]$ and $R_m = k + m(t+1) = \operatorname{rank} X_m$. The positive integer D_Γ at this index is a common multiple of the class stabilizer orders in the column chains. Its prime divisors are at most $2fR_m + 1$ and $\log D_\Gamma \leq C_\Delta f R_m \log(fR_m)$. Here C_T, C_Δ are uniform family constants. The column dimensions satisfy (55) and

   $$C_0 \Gamma^2 x^{C_H} \leq |W| \leq C_0 \Gamma^2 x^{C_H} + 1$$

   in the trace notation.

3. The record includes admissible integers a_t, b_t on T. In particular $a_E \leq a_t \leq a_E + a_E/(kx^{13})$, where $a_E = xb$. For any seed and any $u \in H$ there is an index whose translate is u and whose two arrays agree with the seed's under the common labeling of T. The weights at u are read from the stage data of this translated index.

This statement uses the selection part of Alpöge's design existence theorem and the common data and the translate property. The character estimates and the Selmer input enter through the arithmetic hosts of Section 1. The local, trace and height comparisons enter through separate hypotheses in the subsequent assembly in the release.

The rows and field on which the two section estimates hold. Section 4 constructs the two auxiliary section systems on the base level $g_0 g_1$ with their respective additional primes. Its numerical estimates use a common geometric coefficient field K_0. This field is chosen before the larger field containing the moving character values. They also use the Harish–Chandra rows with specified spacings. Consecutive entries in the initial block differ by 4Δ. In each subsequent block the nodes are $B_\tau + lH$ and $B_\tau + (l + \frac{1}{2})H$, with the same prescribed Δ, H for all words. These choices give the gaps in the flag weights and the bound on the degree and root discriminant of K_0 used in the section heights.

The remaining identification is simultaneous. For the same family $\mathfrak{F}$ furnishing the selected period and relative trace data, fix a threshold before choosing a real scale. At each scale in its windows above that threshold, one needs a seed i_0 and, for every $u \in H$, an index $i(u)$ with that seed's sample arrays. Each $i(u)$ must retain the specified rows and its initial embedded K_0 at both auxiliary levels. The periods and test columns must be those of $i(u)$. One such simultaneous choice suffices. The fixed-level construction makes the indicated choices before constructing the stages. Passage to the family preserves properties supplied at each index. Item 2 of the selection statement gives intervals for the rows and an ambient coefficient field. An interval can contain differently spaced entries, and an ambient field can have arbitrarily large degree. This text proves the two estimates for the specified rows and field. Their simultaneous identification at the selected translates remains open in the reconstruction. The lattice and window estimates below are conditional on that identification. Field and stages over fixed level and row choices.

The integral lattice in the ramified columns. The other outstanding identification concerns each word w at its split level prime T_w. In the complex column space, Section 3 computes the subspace on which the averages in all f residue-unit directions vanish. Its dimension is

$$d_r = d_M \left(\frac{T_w - 2}{T_w - 1} \right)^f, \qquad d_M = \operatorname{rk} M_w.$$

The slope estimate uses the integral sum-zero lattice H_w in the definite class lattice M_w, with its specified metrics at every complex embedding. The needed identification realizes the actual undivided source N_w isometrically as a full-rank sublattice of H_w, with the zero-stage norm used by the column operators. At a coefficient place $v \nmid T_w$ of absolute ramification index e_v, it must give

$$(N_w)_v \subseteq \mathfrak{p}_v^{e_v \tau_v} (H_w)_v, \qquad \operatorname{length}_{\mathcal{O}_v}((H_w)_v/(N_w)_v) \le e_v d_M(\tau_v + \iota_v).$$

At $v \mid T_w$ it must give the separate weighted level-index bound in the calculation of the divided-source slopes. On the sum-zero model the Gram determinant gives the global correction $\mu(H_w) - \mu(M_w) = -f \log(T_w - 1)/(T_w - 2)$. Applying this formula to the actual columns requires the integral and metric identifications just stated. The reconstruction proves the slope and norm deductions with this identification as a hypothesis. The Frobenius divisibility and loss bound applies in its stated domain.

Which character is used at a final index? The order of choices matters for the simultaneous specialization in Section 4. Use the family produced by the selection above. Its preliminary index has the host, words, reference levels and placed class data before the moving characters are chosen. There is a finite set of marked places, defined by the row indices, local shifts and content conditions. It is the inverse image of a conjugation-stable set in a fixed coefficient subfield K_0. The Galois signatures, toric distributions and projected base theta columns at this index determine the tests for specialization. For each conjugacy orbit of marked places of K_0 take one of the ray characters furnished by simultaneous specialization. Carry the previously chosen rational conductor primes into the next avoidance set. There are finitely many orbits, so the supports for distinct orbits can be disjoint. Keep the support, order and reciprocity reading of each of these particular characters. At the other places use the trivial character. The three marking conditions · The preliminary arithmetic data · The choice of ray data.

The second pass of the family construction takes this finite list as input. It adjoins its character values and the algebraic constants in the test models to the large coefficient field. It keeps K_0 with its given embedding. The new marks are the inverse images of the old ones. The base finite-level equations, class representatives and real frames are carried to the same base data. The construction realizes the moving characters with the same ray groups, toric distributions and finite Hecke correspondences. The final index retains the choice of side, support and order. The geometric input here is the existence of this matched adelic realization for the chosen preliminary data, at every final index above one family threshold and before any Fourier translate is picked. Realization of the chosen characters · The compatible geometric enlargement · Passage to the final member.

The normalized theta series and its integral lattices under enlargement. Fix a test with the same stage, side, ray support Q, order and character in the two realizations. Write $K \subset K'$ for the large coefficient fields. The smaller fields and their signed embeddings and local rings O_i, $O_{i,k}$ are identified. Let ω be the Weil action and let $\mathcal{F}(v, P)$ be the Fock vector of the finite section v and harmonic polynomial P in the prescribed infinite vacuum. Its theta series at a finite source representative h and target point g has value

$$T_v(h, P; g) = \big(\omega(h, g)\mathcal{F}(v, P)\big)(0).$$

The matched lattices and split halves identify the base local sections, the determinant test functions at Q and the translations placing their product in the Weil model. The splitting characters and harmonic frames agree, so the displayed value agrees. For the universal series let v_a be the section at ray coordinate a, and let $d_Q(g)$ be the ray coordinate of the target determinant. Multiplication by its group-ring basis element translates the coefficient at a to the value $T_{v_{a-d_Q(g)}}(h, P; g)$. Thus the coefficients already agree before evaluation at any character. The specialized determinant character on g also agrees. For a class form F the full theta value is the sum of $T_{v_\psi}(h, F(h); g)$ divided by the order of the stabilizer of h, over all source classes. At corresponding levels the representatives and stabilizers agree term by term. The determinant shift · The class sum.

Use the same scalar ν for the real normalization at this stage and side. Indeed, transport the extremal vector and a compatible polynomial frame. Their scale relation is unchanged on a K-basis and hence on its K'-span. Changing between two admissible frames of the irreducible lowest type multiplies the map by a scalar by Schur's lemma. Its value on a nonzero rational coordinate places it in the common CM subfield. The scalar preserves a nonzero finitely generated integral lattice and is therefore integral. Preservation of the Hermitian pairing gives $a\bar{a} = 1$. All its conjugates have absolute value one, so a is a root of unity. The scale condition is agreement of positive powers. Taking a multiple of the order of a preserves it in both directions. On a tensor $z \otimes f$ the normalized evaluation is $\nu^{-1}\iota(z)\,\mathrm{ev}_{v_0}(f)(g)$, where ι is the signed embedding and v_0 the extremal vector. All three factors agree. Pure tensors span, so the rational spaces of values and their membership conditions agree. The integral phase · Admissibility of the same scalar · The normalized evaluation.

For the integral comparison, let J be the finite source level with character λ_J, and M_h the Hermitian lattice obtained by translating the local source lattices by h. At h the coefficient functions of an integral harmonic polynomial lie in the $O_{i,k}$-span of evaluations of the dual Weyl tensors of M_h. In matched local bases the standard integral weight modules, their duals and the determinant twists have the same matrices. Substitution in the polynomials commutes with the identification of the small fields, hence these integral harmonic lattices coincide. A class form must take values in them and satisfy

$$F(\gamma h u) = \lambda_J(u)\rho(\gamma)F(h) \qquad (\gamma \in G(\mathbf{Q}),\ u \in J),$$

where ρ is the polynomial action and G the source group over $\mathbf{Q}$. The stabilizer equations and this law identify the integral forms in both directions. To pass back to the coefficient coordinates, let $m_j(g)$ be the finite matrix of a geometric argument g at stage j. The harmonic value of a coefficient vector f at $m_j(g)$ is

the image of the realized value $f(g)$ under the fixed injective coefficient map. The matched adelic matrices agree entry by entry, and the finite representation and small-field action identify $f(g)$. The inverse images of the integral class forms under this map are therefore the same at the two tests. Harmonic evaluation on the finite representation · The finite matrix at this argument. In particular all admissible base translates have the same support, coefficients and theta values. Applying the same rational Hecke projection and taking the span of their columns commutes with any further extension of the coefficient DVR, even in residue characteristic dividing the order of the level group. The universal coefficient identity gives the corresponding assertion for the specialized columns. The index comparing the base and the genuinely twisted lattice at the final level is estimated in Section 4. Integral harmonic values and the level law · Integral class coordinates in the two realizations.

Assume the three estimates of Section 4 on the preliminary signatures, toric distributions and theta tests. To transfer them, take a marked place w of the final coefficient field above ℓ, and let α be the preliminary character at its restriction. The two readings of the same complex roots of order ℓ^k differ by a power $u_w \in (\mathbf{Z}/\ell^k\mathbf{Z})^\times$. The normalized Frobenius values away from the exceptional places agree after this change of roots. Chebotarev on the finite ray quotient therefore identifies the entire final character with α^{u_w}. A primitive power on either prescribed side is now a power $\alpha^{\pm a u_w}$ with $\ell \nmid a u_w$. It has the same conductor and the same trivial restrictions at the places over ℓ. The specialization in Section 4 works for every such power on one set of tests. When a coefficient DVR is extended, an integral determinant valuation and the valuation of ℓ are both multiplied by its ramification degree. This preserves each normalized signed lattice index and each toric valuation. Extension is finite free on the Galois lattices. In an integral basis it commutes with cohomology, the inverse image of the rational finite condition and the kernels of the localization maps. The resulting normalized Selmer lengths are also the same. Thus the three estimates apply to the same character at this index, with their original tolerances. The structural selection identifies their arguments; their bounds come from the corresponding local and global calculations. Identification by Frobenius values · The same primitive powers in all the tests.

Fix ε below $1/52$, the existence cutoff and all the comparison cutoffs. Apply selection once with K_*. Use the same $C_H \geq 1$ for the trace accuracy x^{-C_H} and the height test $H_* \leq x^{C_H} b$. The following application uses the comparisons in Sections 1–5 on this family, including the choice of rows and field at each translated index and the integral identification of the ramified columns.

- At a fixed sufficiently large $x \in \mathrm{Win}(\mathfrak{F})$ choose $i_0 \in I_x$ and use its arrays as the seed. For each $u \in H$ choose $i(u) \in I_x$ by the translate property. Its arrays agree with the seed also for $u = 0$, whose index may differ from i_0. At i_0, Section 1 gives the three auxiliary quadratic twist estimates at x and the tail estimate (11) for every primitive associate of a host character. At each $i(u)$, the bound (6) applies to the selected field at its actual twist cutoff $t \leq \lfloor x^{400} \rfloor$.

- Use the matched period maps at $i(u)$. The full base image and small-level projections enter the rank comparison. For the actual tuple at a finite place, (51) gives the required prime avoidance, and (52) identifies the localized row lattice with the span of the projected integral forms. The theta–Bessel normalization unit enters the lattice comparison, and the Bessel maps into the corresponding Petersson space enter the Gram comparison. At a nonbase twisted place the local Hecke ranks sum to the dimension of the full definite source $M^L_{\mu,j}$ in (175). At the base twist the rank of the good target M_j is $d_j = D_0$ by (173). It is this rank that occurs in (91) for $1 \leq j \leq m$. The lower bound $m \geq x^{1/2-\varepsilon}$ and evenness give $m \geq 4$ for large x, so the bottom comparison has $m/2 - 1 \geq 1$.

- The period maps on the divided sources and into the one row dual at $i(u)$ carry the metrics of Sections 3 and 4. The gain and lattice estimates use the Selmer defect bound, the uniform Gram and reference-row bounds, the height and slope comparisons in those metrics, and the centered level count. On these hypotheses the lattice implication gives $D(u) \leq E_x$ by (167) for each u, with uniform constants.

At the same $i(u)$ each h_w bounds the period map on its divided source. The lower level bound and $S_a = \sum_t a_t > 0$ give $h_w \le h^*(u)$ for every word. Replacing each h_w by this common upper bound lowers the adjusted mean to $s \le s_{\mathrm{raw}}$. Since the target is the row dual with its squared Petersson norm, $\mu_{\min}(C) = -\mu_{\max}(R)$. The height cap (158) therefore gives

$$\max(0, s - \mu_{\min}(C)) \le x^{C_H} b.$$

The common shift cancels in the difference of the donor slopes. This proves the height condition for the frame estimate with the exponent C_H already used in (62).

Two local rank comparisons. At a base-tuple stage, let L_j be the small-level isotypic space and $V_j \subseteq L_j$ its selected irreducible-type part. The inclusion follows because the central type idempotent commutes with the level idempotent, and V_j is the image of their product. Let $Z_j^{\mathrm{coord}} \subset V_j$ be the coordinate row subspace used for the lifted tests. The actual block projection from L_j to Z_j^{coord} first projects to V_j, then projects for its positive Hermitian form onto Z_j^{coord}. Its restriction to the included Z_j^{coord} is the identity, so it is onto. This is the right-inverse identity for the two small-level maps in the projected coordinates.

Put $i = j-1$. The map entering this rank is the auxiliary lift $\Theta_{P_0,\boldsymbol{\mu},j}$ from the definite group of rank $r+t$ to the group of rank r. The descent through the spaces Y_k of ranks $r + kt$ supplies the bottom image used in the period comparison. For the auxiliary lift use the source class-coordinate inclusion $I_{i,\boldsymbol{\mu}}$, the fitted projection $P_Z V^{\flat} E_J$, and the target basis map $j_{W,i}$ of (40). The one-step norm of Section 2.6.5 is positive on each occurring target summand. Applying the commuting good and level projectors to a preimage gives surjectivity on the cut. Thus the restriction to the rational kernel complement has invertible matrix $U_{i,\boldsymbol{\mu}}$. For a standard target coordinate e_a, set

$$x_a = E_J^{\flat} V Z\, U_{i,\boldsymbol{\mu}}^{-1} e_a.$$

Applying the right-inverse identity and then the column formula gives

$$(P_Z V^{\flat} E_J) x_a = U_{i,\boldsymbol{\mu}}^{-1} e_a, \qquad \Theta_{P_0,\boldsymbol{\mu},j}(I_{i,\boldsymbol{\mu}} x_a) = j_{W,i}(e_a).$$

The vectors on the right span the stage multiplicity space M_j. The kernel of the adjoint $\Theta_{T,j}$ is orthogonal to their span, so the adjoint is injective.

The base record and its finite-type blocks. Use the common small level (P_0, ψ_0) and base tuple $\boldsymbol{\mu}_0$. The bottom row space is the selected space on $U(Y_0)$, with the real model used for the bottom period. For $j = i + 1$, the space M_j in the record is read on the auxiliary rank-r group $U(W_i)$. Its fitted theta map comes from the rank-$(r + t)$ auxiliary group $U(V_i)$, as in (40). The descent from the definite top of the Y-tower enters through the bottom period and its spherical eigenvalue transfer.

First compare the finite group actions. Let Γ be the product of the rank-r residue groups at the type places. The principal level is the kernel of the rank-r reduction within the subgroup acting in each selected frame. The suffix character is included in the lift of this action. In the adelic realization one takes right translation by a lift of $g \in \Gamma$, divided by its prescribed suffix character. The covariance under the principal level shows that the result is independent of the lift. For a definite class value, write the translated representative as $\gamma g_c u$, with γ rational and u in the level subgroup. The resulting value is the same coefficient action of γ on the value at g_c, with the level factor from u. This is the class-coordinate matrix $\rho(g)$ of the action. For a character η of P_0, a central element $z = \sum_g z_g g$ acts on its level coordinates by

$$E_{P_0}^{\flat}\Big(\sum_g z_g \rho(g)\Big) E_{P_0}, \qquad e_{P_0,\eta} = |P_0|^{-1} \sum_{h \in P_0} \eta(h)^{-1} \rho(h).$$

Here E_{P_0} and $E_{P_0}^{\flat}$ are the prescribed inclusion and orthogonal coordinate projection, whose composite in the other order is the level projector displayed on the right. For both normalized actions in the fitted

comparison the residue character is $\eta = \psi_0^{-1}$. On the auxiliary source this comes from the ratio of the enlarged level character to the source flag character. On the row it is the level character of the read model. Thus the coefficient of $\rho(h)$ in the level average is $\psi_0(h)/|P_0|$ on both sides. Centrality gives commutation with that projector. Taking z to be the sum of the central idempotents of the allowed Green types gives the prescribed good cut. The same averaging acts on the actual automorphic vectors. Complex conjugation between the two bottom realizations inverts the finite level character and conjugates the central projector. The allowed list is stable under duality, so the central projector is unchanged by conjugation. The scalar $\sqrt{n_{P_0}}$ in the level embedding of M_j cancels when transporting the group actions. The inverse character in the auxiliary projector · Level coordinates and central matrices · The same claim for the prescribed bottom level · The level embeddings and their normalizations.

Now let σ be a finite adelic isomorphism class occurring in either rank-r good space. It contains a tensor of allowed finite types in its principal-level fixed space. Indeed, the good projector is a central sum of isotypic projectors, and a nonzero vector of its image has a nonzero component for one summand. At the base level P_0, the factors of ψ_0 are nondegenerate Whittaker characters of the chosen maximal unipotent groups. Every allowed tensor contains such a line. In each local parabolic induction the open parabolic–unipotent cell gives the tensor product of the cuspidal Whittaker lines [44, 45]. Conjugation by the alternating diagonal inverts the character. A type occurring at either of the two bottom levels therefore occurs at the base level in the same finite adelic representation. The lifts of the residue action preserve the common principal level and all level conditions away from the type places. Level away from the type groups · The two signs on the Whittaker line and The resulting base vector.

By the Green parametrization [44] each allowed local factor is induced from the distinct cuspidal factors associated with the irreducible factors of its squarefree polynomial. If an $(a, r-a)$ unipotent radical fixed a nonzero vector, the finite Mackey formula would admit a double coset for which each cuspidal block lies entirely on one side of the cut. A block meeting both sides would give a proper unipotent radical of that block and its cuspidality gives zero. Thus a must be a subsum of the factor degrees. The first type-prime condition excludes a common subsum of degree a at all the selected places over g_0 for $1 \le a \le r/2$. Taking complements gives the exclusion for $r/2 < a < r$. For each a, therefore, some actual type factor in σ has zero invariants under that radical. The criterion also holds for the dual types. The Green test for a surviving type · Insertion into the occurring finite representation.

The same test excludes a global proper block of size a. At the chosen split place, the group is GL_r and the principal compact subgroup is $K(1)$. By Iwasawa decomposition, the $K(1)$-invariants of normalized induction from the $(a, r-a)$ parabolic are the finite parabolic induction of the inducing Levi's principal invariants. Taking these invariants is exact by compact averaging, and finite Frobenius reciprocity shows that each irreducible finite type in a subquotient has nonzero invariants under the finite unipotent radical. The same conclusion applies to a refinement of the parabolic by induction in stages. This contradicts the type factor just found. In particular, the discrete automorphic occurrences with this finite cut are cuspidal, since a residual constituent is a subquotient of induction from a proper global Levi and its block sizes specialize at each split place.

We compare multiplicities for this particular σ. The two rank-r groups are Hermitian pure inner forms of the same quasi-split unitary group. Their matched finite lattice frames identify the local groups by inner isomorphisms, carrying the level subgroup, its character and the residue homomorphisms. At all real places the conjugate bottom model and the model on W_i have the same list of Harish–Chandra parameters with signs forgotten. The selected signatures and assignments of the entries to the two signs determine a discrete-series member on each group. The lists are integral at this odd rank, with adjacent gaps greater than 1. The real weights and their gaps · Matched frames on the same finite isotypic space.

For either actual occurrence, Labesse's weak base change identifies local base change at every split place as well as at the unramified places. The hypotheses on the CM extension and real field hold for the chosen host, including $[F : \mathbf{Q}] \ge 2$. The archimedean sign condition in that theorem follows from the gaps. In fact Salamanca–Riba's classification, applied to $\mathrm{SU}(p, q)$ with the central character fixed on

$U(p, q)$, makes a unitary representation of this infinitesimal character an $A_{\mathfrak{q}}(\lambda)$ induced from a character of a θ-stable Levi. In any Levi block of size greater than one, both λ and $\rho(\mathfrak{u})$ are constant on the block and $\rho(\mathfrak{l})$ has two entries at distance one. The infinitesimal character would contain two entries at distance one. Hence the Levi is toral, giving a discrete-series member of the same real packet. These are precisely the representations allowed by Labesse's criterion for the pseudo-coefficient at infinity [104, 105].

Write the base change as an isobaric sum of discrete representations of general linear groups. With several summands, the rank a of one summand yields the proper parabolic test above at each split place. With a single noncuspidal discrete summand, the Mœglin–Waldspurger description is Speh from a fixed cuspidal rank $a < r$. Its local components are subquotients of induction of twists of that same rank, so the same test applies. Therefore the stable base change is a single cuspidal Π of $\mathrm{GL}_r(\mathbb{A}_E)$. The comparisons here use Labesse's local statement at the ramified split type places as well [104, 106].

In the Kaletha–Mínguez–Shin–White classification take $\kappa = +1$, $\chi_+ = 1$, which gives this same stable base change. Their Corollary 3.3.2 attaches the Satake records of the actual discrete occurrence to a global parameter. At almost every place of E its general-linear record equals that of Π, by local base change at split and inert places. Theorem 4.4 of Jacquet and Shalika identifies the cuspidal inducing factors, up to permutation, from this equality of unramified components outside a finite set [107]. There is just one factor, namely Π, of rank r and with Arthur SL_2 of dimension one. For a simple parameter the centralizer modulo the invariant dual center is trivial. Apply the proved generic, pure-inner case of Theorem* 1.7.1. The transported finite factors stay in the same local packets by the inner isomorphisms. At infinity the common unsigned list gives the same local discrete parameter and the chosen signs give the specified member on the other group. With trivial global component group this tensor therefore occurs there with multiplicity one. Starting on either group gives the same argument [108].

Consequently the contribution of σ to either level is the dimension of the same finite space $e_{\mathrm{good}}\sigma^{J_0,\psi_0^{-1}}$, where J_0 denotes the transported level. Use its conjugate realization on the bottom. The record's D_0 is the size of the chosen basis (b_k) of that bottom space with type labels $\mathrm{ft}_0(k)$. The stage labels are read by the actual finite isotypic spaces on W_i. In particular, if $M_j[\sigma]$ denotes that stage isotype, then at each labelled σ the identity is

$$\dim_{\mathbb{C}} M_j[\sigma] = \#\{k : \mathrm{ft}_0(k) = \sigma\}. \tag{172}$$

The matched groups also give the reverse inclusion on occurring labels, so these spaces exhaust M_j. The actual basis, finite-type labels and auxiliary stage · Reading the stage isotype in its original frame.

Finally match the spherical labels. Off the type and exceptional places, each finite representation σ has a spherical line. Its scalar for the double coset of weight d on W_i, computed in the matched lattice frame, is denoted $c_\sigma(v, d)$. Inner finite transport makes this independent of i. For the conjugate bottom representation the scalar is $\overline{c_\sigma(v, d)} = c_\sigma(v, d^\vee)$, with $d^\vee = (-d_r, \dots, -d_1)$. The equality follows by taking the adjoint, which is the inverse double coset. This accounts for the dual convention in the bottom eigenspace. At inert places the same argument uses the relative Cartan double cosets.

The top base record p consists of the prescribed split Hecke polynomials and the inert double-coset scalars of the definite top space. The bottom space for c_σ is carried by p when its long-chain adjoint lies in the top block p. On each chosen split and inert spherical operator, the adjoint carries the eigencharacter by the local spherical theta transfer. At each arrow the unitary Satake list has the old entries with the chosen common character twist and the prescribed new string, with the dual convention just described if using the conjugate model. The spherical calculation applies to the hyperspecial lattice vectors and unramified characters at these places. For an inert place orient the correspondence from the smaller group to the larger one. In Harris's notation this is $m \geq n$. The lattice vector gives the spherical line, and the map of spherical parameters on the domain is a morphism. Its agreement with the string formula on a dense open set consequently holds for each parameter that occurs here. The additive and splitting characters are unramified and the lattices self-dual. The residue characteristic is odd since the places over 2 split in the chosen host. The finite level and central eigencuts commute with these operators, which act at

disjoint places. The long-chain adjoint is injective on the selected bottom space by the base row descent, so its nonzero image belongs to the unique top block with these joint scalars. The rational top blocks are the joint kernels of the same operators with those eigenvalues. Their base changes are the corresponding complex kernels by the class-coordinate comparison. Thus the assigned label $p(\sigma)$ is an actual index of the prescribed top record. Conversely this index gives the full split Hecke polynomials and the scalars of the inert Cartan generators. At an inert place the relative Satake class is determined by its reciprocal pairs, with middle entry 1 in odd rank. The relative Weyl group permutes the pairs and interchanges their two entries. At each lift remove the prescribed string with its multiplicities and undo the common character twist. This recovers the spherical character of the bottom representation on every operator in the record. Several finite representations may have the same label $p(\sigma)$, and the spherical character on that block is well defined [109]. The dual bottom convention and the carried label · The prescribed split and inert index · Reading the label on the actual vectors and Recovering the spherical summand.

The bottom coordinate summands E_p are the preimages of the corresponding analytic cuts under the base theta columns. The top block projectors are defined over K. The long-chain adjoint maps each bottom eigenspace into a single top block, so the projectors preserve the image of the adjoint, as well as the rational kernel complement whose scalar extension is this image. Taking adjoints shows that each rational summand maps into its indicated analytic cut. These summands span the complement, the map is injective there, and its complexification has image the whole selected bottom space by the base row descent. Orthogonality of distinct spherical parts gives surjectivity block by block. Similarly the auxiliary coefficient summand $W_{i,p}$ is the preimage of its analytic cut under the class-basis embedding of W_i. The base-tuple complementary stage coordinates split by the same transported top labels. Their images under the fitted matrix belong to the corresponding analytic cuts by (40) and the auxiliary spherical transfer. These coordinate summands exhaust the complementary coordinates, and the base matrix on the complement is invertible. Thus their images span the full good class-basis space over K and split it by the given labels. The prescribed good basis and its scaled level reading give the complexification of each such summand. Equivalently, these are the direct sums of the actual isotypes with $p(\sigma) = p$. The finite isotypic summands are orthogonal because they are distinct irreducible representations of the finite adelic group in the unitary discrete spectrum. Definitions of the rational preimages · Passage to the same individual parts · The good basis on the auxiliary part.

It follows from (172) that

$$\begin{gathered}\dim_K E_p = \dim_K W_{i,p} = \sum_{\sigma\,:\,p(\sigma)=p} \#\{k \,:\, \mathrm{ft}_0(k) = \sigma\},\\ \sum_p \dim_K E_p = D_0 = \dim_{\mathbf{C}} M_j.\end{gathered} \tag{173}$$

This is the rank of the indicated base row space on the same level. With the standard-vector equality and the level coordinate right inverse, the fitted map onto this M_j is surjective. Its Hilbert adjoint therefore has zero kernel and rank D_0. The block partition and its ranks · The base record rank and its use for the fitted map.

On a nonzero carried block, the auxiliary Petersson scalar in (43) is the scalar used when constructing the record at the distinguished embedding. Indeed, the scalar is chosen there by the equation $\Theta\Theta^* y = \lambda_{j,\mu,p} y$ for the same phase-one auxiliary theta map and the same carried space. On this block its value is determined by the ratio $\|\Theta^* y\|^2/\|y\|^2$ for any nonzero y. The level-index scaling of the two Petersson forms is the same and therefore cancels in the ratio. The normalized comparison (40) gives the same ratio for the fixed record map and its adjoint. Consequently, for each used tuple and each nonzero $M_j[\sigma]$, the record scalar $\lambda_{j,\mu,\sigma}$ equals $\lambda_{j,\mu,p(\sigma)}$ at this embedding. Selection of the actual eigenvalue · Auxiliary theta and its Petersson adjoint · The scalar on a nonzero finite isotype.

The residual Hecke summands at a nonbase twist. Fix the stage j and coefficient place w of a nonbase twist $\mu = \mu_{\mathrm{tw}}(w)$. Let $M^L_{\mu,j}$ be the full rational space of definite forms of rank $r+t$ at the prescribed level character. It uses the class representatives and coefficient representation of (39). The good-type cut will be taken inside the resulting local lattices. Write O for the local coefficient DVR, L for its fraction field, $M = L \otimes_K M^L_{\mu,j}$, and $A \subseteq \mathrm{End}_L(M)$ for the O-algebra generated by the prescribed normalized Hecke operators. To specify them, put $n = r + t$ and $\ell = \mathrm{char}\, k(w)$. Each generator is indexed by a split place v of E away from ℓ, the two type primes, the designated scalar-character prime and the conductor set of the twist. The last set includes both halves above each place of F in its support. Use the split realization at $\bar{v}$, put $q = |\mathcal{O}_E/v|$, and write J for the hyperspecial group in those coordinates. For $1 \le k \le n$ its action on the definite class forms is

$$h_{v,k} f(g) = q^{-k(n-1)} \sum_{a \in J\, \mathrm{diag}(\varpi I_k, I_{n-k}) J / J} f(ga). \tag{174}$$

Here ϖ is a uniformizer at the underlying place of F. Using the conjugate half fixes which eigenvalues occur in this formula. In the normalized Satake convention with parameters $\alpha_1, \dots, \alpha_n$ at that half, the coset sum has eigenvalue $q^{k(n-k)/2} e_k(\alpha_1, \dots, \alpha_n)$. Indeed the right cosets correspond to sublattices between $\varpi\mathcal{O}^n$ and $\mathcal{O}^n$ of residual codimension k. In upper triangular echelon form fix the k positions I with diagonal ϖ. There is one free residue for each $a \in I$, $b \notin I$ with $a < b$, hence $\sum_{a \in I}(n-a) - k(k-1)/2$ free residues. On this cell the normalized spherical vector takes the value $q^{-\sum_{a \in I}(n+1-2a)/2} \prod_{a \in I} \alpha_a$. The cell size times the power of q is $q^{k(n-k)/2}$, independently of I. Summing proves the asserted eigenvalue. Consequently the monic polynomial with coefficient $(-1)^k q^{k(k-1)/2} h_{v,k}$ in degree $n-k$ has roots $q^{(1-n)/2}\alpha_a$ on this eigenspace. At the indexed half v the Satake list is the inverse list, by the inverse-transpose identification of the split unitary group. Arithmetic Frobenius in the convention of the record inverts that list once more before applying the norm shift. The half-place and Frobenius conventions.

Each representative and its inverse have denominators supported at the rational prime below v. Thus the two value lattices at g and ga have the same localization at w, including for the coefficient representation. The finite sum preserves the localized class lattice, and q is a unit there. Its value at an adelic argument is reconstructed from the rational coefficient action and the level character on its class representative. This puts the coset operator in the prescribed rational class model. The resulting operators have integral characteristic polynomials. They commute by the spherical Hecke algebra at each place and independence across places. The L used here already contains their eigenvalues: the commuting endomorphisms of the finite-dimensional class space generate a finite-dimensional algebra, so only finitely many of them are needed to generate it over the fraction field. One may split those generators in a finite extension. The local coefficient choice at (j, w) takes the compositum of that extension with the prescribed coefficient and twist fields before forming the residual ideals. The split commuting generators can be triangularized simultaneously, so their algebra splits over this one L. Its generator eigenvalues belong to the valuation ring O. Choice of the split half · Localized coefficient lattices under the coset sum · The common splitting field at the coefficient place.

Successive primary decompositions give joint generalized eigenspaces. Indeed each generator preserves the spaces already obtained, and each proper refinement increases the number of nonzero summands, which is at most $\dim_L M$. On a joint summand the commuting nilpotent parts generate a nilpotent ideal in the finite-dimensional commutative image algebra. Taking the scalar part therefore gives a character $A \to O$. Its reduction modulo the uniformizer is onto the residue field since A contains O.

For a maximal ideal $\mathfrak{m}$ of A, let $M_{\mathfrak{m}}$ be the sum of the joint spaces with that residual kernel and $p_{\mathfrak{m}}$ the projection along the other summands. The local indices over (j, w) run once through the maximal ideals of this algebra. For such an index the lattice is the O-span of $p_{\mathfrak{m}} b$, where b runs through the realized generators of the specified twisted class lattice. These generators span M over L. Applying $p_{\mathfrak{m}}$ shows that

the local lattice spans exactly $M_{\mathfrak{m}}$ over L. It is finite free, being a finitely generated torsion-free module over a DVR. Consequently, writing Λ_a for the lattice with local index a, we have

$$\sum_{a \in \mathrm{Loc}(j,w)} \operatorname{rank}_O \Lambda_a = \sum_{\mathfrak{m}} \dim_L M_{\mathfrak{m}} = \dim_L M = \dim_K M^L_{\mu,j}. \tag{175}$$

If $M = 0$, its lattice is zero and the index set is empty. In every case the local projections jointly detect each vector of M, since their sum is the identity. This rank computation takes place over the fraction field. The colength in Section 4 also depends on the position of the projected integral lattice relative to the intersected lattice. The local indices and the projected generator lattice · Exhaustion of the residual ideals and the normalized Hecke comparison.

The two estimates use the same $D(u)$. Indeed T, T_2, a_t agree by the translate property, the weights $w_{u,j}$ in both sums come from $i(u)$, and the inertia classes and sign character are read at i_0, with the ramified inertia groups of (4). By (96) every positive-part summand with $p^h > N = \lceil 2xb \rceil$ vanishes. Thus the cutoff and the untruncated Fourier ramp give the same $A(a_t, t, w_{u,j})$. Both comparisons subtract the real part of $M(a_t, w_{u,j})$ in (164), with γ and the sign-character log derivative at 1 in the same normalization. Hence the Fourier estimate uses exactly (166).

With the family, precision and constants fixed, take x_* to be the maximum of the family, arithmetic, period, trace, lattice and Fourier thresholds, increasing it to 3 if necessary. Each threshold in this maximum works for every required index at the scale. At every $x \geq x_*$ in $\mathrm{Win}(\mathfrak{F})$ the preceding applications give the lattice bound for all u and the base-host character estimates. The sampling and weight comparisons of Section 5 refer to these same records. Thus (171) applies with $e_D + e_L + 7 = 9$. The retained set for the conductor-one target is $\{1\}$ with weight 1, so the function is G_1 at every scale.

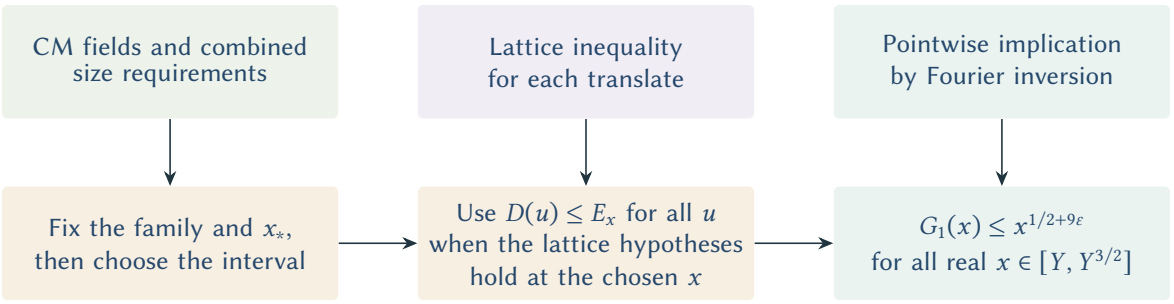

Figure 3. Fix the target, the loss bounds $e_D = 2$ and $e_L = 0$, and a sufficiently small $\varepsilon > 0$. Choose one family $(\mathfrak{F}, C_H)$ meeting the joint size requirements, including $C_H > 0$. Fix its comparison constants and common threshold x_*. For each $Y_0 \geq 2$, choose $Y \geq \max(Y_0, x_*)$ with $[Y, Y^{3/2}] \subseteq \mathrm{Win}(\mathfrak{F})$. At each real scale in this interval, obtain the lattice bound after supplying its period, rank and height hypotheses at the tied indices $i(u)$. The Artin and character data come from i_0. The shared degrees, cutoff and centering term identify the sum used in the pointwise implication.

Given $Y_0 \geq 2$, the window property gives $Y \geq \max(Y_0, x_*)$ with $[Y, Y^{3/2}] \subseteq \mathrm{Win}(\mathfrak{F})$. Every real x in the interval has $x \geq x_*$, so the pointwise implication applies at every scale in Figure 3.

$$G_1(x) \leq x^{1/2+9\varepsilon} \qquad \text{for every real } x \in [Y, Y^{3/2}]. \tag{176}$$

The same family and x_* work as Y_0 varies. The precision can be any fixed positive number below the minimum of the cutoffs.

6.2 The centered zeta transform

Write

$$\Phi_1(s) = \int_1^\infty G_1(x) x^{-s-1}\, dx.$$

The finite ramp sum is continuous in x. From $\Lambda(n) \le \log n$, its absolute value is at most $x\sum_{n\le x}(\log n)/n$. Including the centering terms gives $G_1(x) = O(x(1+\log^2 x))$, so the integral converges absolutely on $\Re s > 1$. The Euler product converges there and gives a convergent logarithm. Differentiating locally uniformly yields

$$-\frac{\zeta'(s)}{\zeta(s)} = \sum_{n\ge 1}\frac{\Lambda(n)}{n^s}.$$

[73, (25.2.11)] For $\Re s > 1$ the transform of a ramp is

$$\int_n^\infty (x/n-1)x^{-s-1}\,dx = n^{-s}\left(\frac{1}{s-1}-\frac{1}{s}\right).$$

The sum of the absolute integrals at $\sigma = \Re s > 1$ is bounded by $(\sigma(\sigma-1))^{-1}\sum_n(\log n)n^{-\sigma}$. Interchanging sum and integral and transforming the two centering terms of (1) gives

$$\Phi_1(s) = \frac{-\zeta'(s)/\zeta(s)}{s(s-1)} - \frac{1}{(s-1)^2} + \frac{1+\gamma}{s-1}.$$

The continuation of ζ is meromorphic with unique pole at 1 and $\zeta(1+z) = z^{-1}+\gamma+O(z)$. [73, §25.2(i) and (25.2.4)–(25.2.5)] Hence $-\zeta'(1+z)/\zeta(1+z) = z^{-1}-\gamma+O(z)$. Dividing by $z(1+z)$ gives $z^{-2}-(1+\gamma)z^{-1}+O(1)$, which the centering cancels.

The Euler product gives nonvanishing for $\Re s > 1$. On the line at $1+it$, $t \ne 0$, use the nonnegative sum

$$-\Re\left(3\frac{\zeta'(\sigma)}{\zeta(\sigma)} + 4\frac{\zeta'(\sigma+it)}{\zeta(\sigma+it)} + \frac{\zeta'(\sigma+2it)}{\zeta(\sigma+2it)}\right) = 2\sum_{n\ge 1}\frac{\Lambda(n)}{n^\sigma}\big(1+\cos(t\log n)\big)^2 \quad (\sigma > 1).$$

If $1+it$ were a zero of multiplicity $m \ge 1$, the coefficient of $(\sigma-1)^{-1}$ as $\sigma \downarrow 1$ would be $3-4m-m_2 < 0$, where $m_2 \ge 0$ is the zero multiplicity at $1+2it$ (zero if the value is nonzero). This contradicts the displayed inequality. [73, §25.10(i)] On the real segment $0 < \sigma < 1$, pairing consecutive terms in the alternating zeta series gives

$$(1-2^{1-\sigma})\zeta(\sigma) = \sum_{j\ge 1}\big((2j-1)^{-\sigma} - (2j)^{-\sigma}\big) > 0.$$

The prefactor is negative, so $\zeta(\sigma) < 0$. [73, (25.2.3)]

It follows that Φ_1 is meromorphic on $\Re s > 0$ with one simple pole at each zero ρ in the open strip and with residue

$$\operatorname{Res}_{s=\rho}\Phi_1(s) = -\frac{m_\rho}{\rho(\rho-1)} = \frac{m_\rho}{\rho(1-\rho)}. \tag{177}$$

Indeed, the factor $(s-\rho)^{m_\rho}$ in $\zeta(s)$ gives residue $-m_\rho$ in its negative logarithmic derivative, and $s(s-1)$ is nonzero at ρ. These poles are nonreal, come in conjugate pairs and are locally finite in $\Re s > 0$.

6.3 One-sided windows and the critical line

Use the following simple-pole case of Kaczorowski and Pintz. Let h be real and locally integrable on $(0,\infty)$, with Mellin integral $\int_0^\infty h(x)x^{-s-1}\,dx$ absolutely convergent on some closed right half-plane. Suppose this transform has a meromorphic continuation to a neighborhood of $\Re s \ge \Theta - c_0$, for $\Theta > 0$, $c_0 > 0$, holomorphic on $\Re s > \Theta$. Suppose its poles in the closed half-plane are simple, nonreal, locally finite and paired by conjugation, with supremum of real parts Θ. For every $0 < \delta < 1$, all sufficiently distant intervals $[X^{1-\delta}, X]$ contain both a strictly positive and a strictly negative value of h.

To apply Part II, cut the half-plane leftward from each of these poles. The transform restricts holomorphically to the cut domain, and near each pole its expansion in (1.5) has exponent -1, logarithmic degree 0, and nonzero coefficient equal to the residue. The pole supremum is the boundary of the half-plane of holomorphy whether or not a pole attains it. Conclusion (1.7) gives the two signs. The same case in Part I uses zero logarithmic term. [76, (1.5) and (1.7)] [75, Thm. 1 and (2.2)]

Suppose $\zeta(\rho) = 0$ with $1/2 < \beta = \Re\rho < 1$. Choose $1/2 < \theta < \beta$, then choose $0 < \varepsilon$ below all the preceding cutoffs and with $9\varepsilon \le \theta - 1/2$. The fixed-precision bound (176) gives arbitrarily distant intervals $[Y, Y^{3/2}]$ on which $G_1(x) \le x^\theta$.

Set $h(x) = G_1(x) - x^\theta$ for $x \ge 1$ and $h(x) = 0$ for $0 < x < 1$. It is real and locally integrable, and its Mellin integral converges absolutely on $\Re s \ge 2$. The transform there is $\Phi_1(s) - 1/(s-\theta)$. Let Θ be the supremum of the real parts of the poles of Φ_1. By (177), $\Theta \ge \beta > \theta$ and $\Theta \le 1$. Choose $c_0 = (\Theta - \theta)/2 > 0$. On the closed half-plane $\Re s \ge \Theta - c_0$ the correction term is holomorphic, and the poles are exactly those of Φ_1 there. They satisfy all the sign-change hypotheses and still have supremum Θ. Apply the theorem with $\delta = 1/3$ at $X = Y^{3/2}$, taking Y large enough to exceed its threshold for this fixed h and δ. An interval with $X^{1-\delta} = Y$ must contain $h > 0$, contradicting $h \le 0$ on the entire interval.

We have excluded all zeros to the right of the critical line. For $\Re s < 1/2$, $s \ne 0$, use the functional equation

$$\zeta(s) = 2^s \pi^{s-1} \sin(\pi s/2) \Gamma(1-s) \zeta(1-s).$$

[73, (25.4.2)] The argument $1-s$ lies on the zero-free side and differs from 1. The gamma factor is holomorphic and nonzero there. [73, §5.2(i)] The sine factor vanishes precisely when s is a negative even integer in this region. Also $\zeta(0) = -1/2$. [73, (25.6.1)] Thus (176) implies that every zero other than the negative even integers lies on $\Re s = 1/2$, the classical Riemann hypothesis.

6.4 Extension to Dirichlet L-functions

For a fixed primitive target ψ, assume that the arithmetic and period comparisons hold on its selected family. The window argument then gives

$$G_{\mathcal{K}}(x) \le x^{1/2+9\varepsilon} \qquad (x \in [Y, Y^{3/2}]), \tag{178}$$

for arbitrarily large Y at each sufficiently small fixed $\varepsilon > 0$. The retained $\mathcal{K}$ consists of the primitive associates of the powers of ψ, with fixed positive weights, and contains ψ. The fixed-target window comparison is applied to the primitive character class in the Dirichlet construction.

The Mellin poles. For any finite set $\mathcal{F}$ of primitive Dirichlet characters with weights $a_\chi > 0$, put

$$G_{\mathcal{F}}(x) = \sum_{\chi \in \mathcal{F}} a_\chi \Re G_\chi(x), \qquad \Phi(s) = \int_1^\infty G_{\mathcal{F}}(x) x^{-s-1}\, dx.$$

The same absolute-convergence calculation as for Φ_1 works, since $|\chi(n)| \le 1$. For primitive $\chi \ne 1$ its Euler product, differentiated on $\Re s > 1$, gives $-L'(s,\chi)/L(s,\chi) = \sum_{n\ge1} \Lambda(n)\chi(n)n^{-s}$. [73, (25.15.1)–(25.15.2)] Using the ramp transform and (2) yields

$$\Phi_\chi(s) = \frac{-L'(s,\chi)/L(s,\chi)}{s(s-1)} + \frac{L'(1,\chi)/L(1,\chi)}{s-1}.$$

Here $L(s,\chi)$ is entire and $L(1,\chi) \ne 0$. [73, §25.15(i), (25.15.9)] The centering cancels the residue of the quotient at $s = 1$. The Euler product gives nonvanishing for $\Re s > 1$, and Montgomery and Vaughan's zero-free region gives $L(1+it,\chi) \ne 0$ for all real t for nonprincipal χ. The exceptional zero in that theorem is real and strictly less than 1. [74, Thm. 11.3] Thus the poles of the transforms in $\Re s > 0$ lie in the open strip.

Conjugating the Dirichlet series and continuing gives $L(s,\bar\chi) = \overline{L(\bar s,\chi)}$, so $G_{\bar\chi}(x) = \overline{G_\chi(x)}$. The transform of $\Re G_\chi$ is $(\Phi_\chi + \Phi_{\bar\chi})/2$. Write $m_\chi(\rho)$ for the zero multiplicity at ρ, with value zero at nonzeros. The calculation of (177) gives

$$\operatorname{Res}_{s=\rho} \Phi(s) = \frac{\sum_{\chi\in\mathcal{F}} a_\chi \big(m_\chi(\rho) + m_{\bar\chi}(\rho)\big)}{2\rho(1-\rho)} \qquad (0 < \Re\rho < 1).$$

Thus a zero of any retained character gives a simple pole even when shared by several characters. At a real pole the residue is positive, since $0 < \rho < 1$.

Positive majorants and real poles. The release uses a one-sided criterion for a real measurable g on $[0,\infty)$ with $|g(u)| \le Me^{2u}$. Assume its Laplace transform Φ_g extends meromorphically from $\Re s > 2$ to $\Re s > 0$, with simple poles and positive residues at real poles. Fix $0 < \theta < 1$ and $\lambda > 1$. Assume that for every $\eta > 0$ there are arbitrarily large U with

$$\int_U^{\lambda U} g_+(u)e^{-\theta u}\,du \le e^{\eta U}, \qquad g_+(u) = \max(g(u), 0). \tag{179}$$

Then Φ_g has no pole in $\Re s > \theta$. The positive part is the least nonnegative majorant of g. Alpöge proves the criterion by estimating derivatives at a real center and a vertical translate, then comparing their growth at the nearest pole. The derivative proof in the release

Here is the deduction from the published sign-change theorem. If Φ_g has a pole ρ with $\Re\rho > \theta$, fix $\eta > 0$ with $\alpha = \theta + \eta < \Re\rho$. Extend g by zero to $u < 0$. The average $\widetilde{g}(v) = \int_{v-1}^{v} g(u)\,du$ has Laplace transform

$$\widetilde{\Phi}(s) = \Phi_g(s)\frac{1 - e^{-s}}{s}$$

by Fubini on $\Re s > 2$. The multiplier is nonzero on $\Re s > 0$ and positive for real $s > 0$. Thus the average preserves both the poles there and the signs of real residues. On each supplied window, for $U+1 \le v \le \lambda U$,

$$\widetilde{g}(v) \le e^{\theta v}\int_{v-1}^{v} g_+(u)e^{-\theta u}\,du \le e^{\theta v+\eta U} \le e^{\alpha v}.$$

Choose $1 < \lambda' < \lambda$. For all sufficiently large U, $[V, \lambda' V] \subseteq [U+1, \lambda U]$ at $V = U+1$.

The real poles B of $\widetilde{\Phi}$ above α are finite by meromorphy around the compact segment $[\alpha, 2]$ and absolute convergence on $\Re s > 2$. Write $b_\beta > 0$ for their residues. Choose $t \in [1,2]$ such that none of the $\beta \pm it$, $\beta \in B$, is a pole of $\widetilde{\Phi}$. Only finitely many values of t are excluded. Put

$$h(x) = \widetilde{g}(\log x) - x^\alpha - \sum_{\beta\in B} b_\beta x^\beta(1 + \cos(t\log x)) \qquad (x \ge 1)$$

and $h(x) = 0$ below 1. This function is real and locally integrable and its Mellin integral converges absolutely on $\Re s \ge 3$. The transform is

$$\widetilde{\Phi}(s) - \frac{1}{s-\alpha} - \sum_{\beta\in B} b_\beta\left(\frac{1}{s-\beta} + \frac{1}{2(s-\beta-it)} + \frac{1}{2(s-\beta+it)}\right).$$

Each pole at $\beta \in B$ is canceled and two simple conjugate poles at $\beta \pm it$ take its place, with residues $-b_\beta/2$. Thus the supremum $\Theta > \alpha$ of the pole real parts stays the same. For $0 < c_0 < \Theta - \alpha$, the poles to the right of $\Theta - c_0$, including that line, are all nonreal and simple. The transform satisfies the other sign-change hypotheses by meromorphy on $\Re s > 0$ and conjugation. The correction is nonnegative before subtraction, so $h(x) \le 0$ throughout $[e^V, e^{\lambda' V}]$ for arbitrarily large V. With $X = e^{\lambda' V}$, $\delta = 1 - 1/\lambda'$, this contradicts the sign-change theorem.

Primitive and imprimitive targets. For a hypothetical zero ρ of $L(s,\psi)$ with $1/2 < \Re\rho < 1$, choose $1/2 < \theta < \Re\rho$ and a fixed sufficiently small $\varepsilon > 0$ with $9\varepsilon \le \theta - 1/2$. Use $\mathcal{F} = \mathcal{K}$ and put $g(u) = G_{\mathcal{K}}(e^u)$. The growth bound $O(e^u(1+u^2))$ gives the required Me^{2u}. By (178), the integral in (179) is at most $(\lambda-1)U$ on arbitrarily distant windows at $\lambda = 3/2$. For each $\eta > 0$ this is at most $e^{\eta U}$ once U is large. The Mellin residue computation and $\psi \in \mathcal{K}$ give a pole at ρ, contradicting the criterion.

Let χ be a primitive nonprincipal character of conductor d and parity $a \in \{0,1\}$, defined by $\chi(-1) = (-1)^a$. Its completed function is entire and satisfies

$$\begin{aligned}\Lambda_\chi(s) &= (d/\pi)^{(s+a)/2}\Gamma\big((s+a)/2\big)L(s,\chi),\\ \Lambda_\chi(s) &= \epsilon_\chi\Lambda_{\bar\chi}(1-s) \qquad (s \in \mathbf{C}),\end{aligned} \tag{180}$$

where $\epsilon_\chi = i^{-a} d^{-1/2} \sum_{r \bmod d} \chi(r) e^{2\pi i r/d}$ has modulus 1. [74, (10.17) and Cor. 10.8] The gamma and conductor factors are finite and nonzero in the open strip, so a zero there at s gives one at $1 - s$ for $\bar{\chi}$. Thus the window bound (178) for the two fixed targets χ and $\bar{\chi}$ puts every zero in the open strip on $\Re s = 1/2$.

For any Dirichlet character χ modulo $N \geq 1$, let χ^* be its inducing primitive character of conductor $d \mid N$. The change-of-conductor formula is a meromorphic identity

$$L(s, \chi) = L(s, \chi^*) \prod_{\substack{p \mid N \\ p \nmid d}} (1 - \chi^*(p) p^{-s}).$$

[73, (25.15.4)] Each displayed Euler factor is holomorphic and nonzero in $\Re s > 0$, since $|\chi^*(p) p^{-s}| < 1$ there. For a principal character $\chi^* = 1$, so $L(s, \chi)$ has a simple pole at $s = 1$ of residue $\prod_{p \mid N} (1 - p^{-1}) = \varphi(N)/N$. Consequently, given the window bound for each primitive target, every $L(s, \chi)$ is nonzero in $\Re s > 1/2$ away from the principal pole. Every zero in $0 < \Re s < 1$ then has real part $1/2$. Alpöge uses the same change of conductor for the right-hand nonvanishing in the formal Dirichlet theorem.

Acknowledgment. We thank Sahibzada Allahyar for making his reconstruction available. It helped us follow the passage from the theta lifts to the prime-sum estimate [3].

References

[1] Levent Alpöge. Public announcement on the Riemann hypothesis, 20 November 2026.

[2] Levent Alpöge. Formal proof accompanying the announcement, 20 November 2026. Lean 4.33.1. Public source.

[3] Sahibzada Allahyar. *From theta lifts to zeta zeros: A reconstruction of Levent Alpöge's announced argument*. Expository manuscript, revision 2, 20 November 2026.

[4] Jesse Thorner and Asif Zaman. Refinements to the prime number theorem for arithmetic progressions. *Mathematische Zeitschrift* 306 (2024), no. 3, article 54. arXiv v2, Theorem 1.1, (1.2) and (1.7), Corollary 1.4 and Remark 1.5.

[5] Jesse Thorner and Asif Zaman. An explicit version of Bombieri's log-free density estimate and Sárkőzy's theorem for shifted primes. *Forum Mathematicum* 36 (2024), 1059–1080. arXiv v4, Theorem 1.2 and (1.1)–(1.2).

[6] D. R. Heath-Brown. A mean value estimate for real character sums. *Acta Arithmetica* 72 (1995), 235–275. Corollary 1, p. 237, and Theorem 3, p. 240.

[7] Yasutaka Ihara, V. Kumar Murty and Mahoro Shimura. On the logarithmic derivatives of Dirichlet L-functions at $s = 1$. *Acta Arithmetica* 137 (2009), 253–276. §2.1, p. 256, and Theorem 1, (2.2.1), p. 257.

[8] Michael A. Bennett, Greg Martin, Kevin O'Bryant and Andrew Rechnitzer. Counting zeros of Dirichlet L-functions. *Mathematics of Computation* 90 (2021), 1455–1482. Theorem 1.1 and Corollary 1.2, p. 1456.

[9] Hugh L. Montgomery and Robert C. Vaughan. Mean values of character sums. *Canadian Journal of Mathematics* 31 (1979), 476–487. Lemma 9, (7), p. 480.

[10] A. Page. On the number of primes in an arithmetic progression. *Proceedings of the London Mathematical Society* (2) 39 (1935), 116–141. For the Landau–Page theorem in the form used here, see Ford, Green, Konyagin, Maynard and Tao, *Long gaps between primes*, arXiv v3, Lemma 7.1.

[11] Youness Lamzouri and Alessandro Languasco. Small values of $|L'/L(1, \chi)|$. *Experimental Mathematics* 32 (2023), 362–377. arXiv v1, Proposition 3.2, p. 8.

[12] J. S. Milne. *Class Field Theory*, version 4.03, 2020. Ch. V, Examples 3.1–3.2, pp. 156–157, and Example 3.11, p. 159.

[13] J. S. Milne. *Arithmetic Duality Theorems*, 2nd ed., BookSurge, 2006. Theorem I.4.10(a), p. 57, Corollary I.4.15, p. 62, and Proposition II.2.9, pp. 170–171.

[14] Charles A. Weibel. *The K-book: An Introduction to Algebraic K-theory*. Graduate Studies in Mathematics 145, American Mathematical Society, 2013. Theorem VI.8.2 and Theorem VI.8.4.

[15] Meng Fai Lim. On the codescent of étale wild kernels in p-adic Lie extensions. *Kyoto Journal of Mathematics* 65 (2025), 315–347. arXiv v2, Proposition 5.1, p. 20.

[16] Eva Bayer-Fluckiger, Vincent Emery and Julien Houriet. Hermitian lattices and bounds in K-theory of algebraic integers. *Documenta Mathematica*, Extra Volume Merkurjev (2015), 71–83. Theorem 1.1, p. 72.

[17] David Burns and Matthias Flach. On the equivariant Tamagawa number conjecture for Tate motives, Part II. *Documenta Mathematica*, Extra Volume Coates (2006), 133–163. Corollary 1.2, p. 134.

[18] Christophe Soulé. A bound for the torsion in the K-theory of algebraic integers. *Documenta Mathematica*, Extra Volume Kato (2003), 761–788. §5.4, (26), p. 786.

[19] Wee Teck Gan, Yannan Qiu and Shuichiro Takeda. The regularized Siegel–Weil formula (the second term identity) and the Rallis inner product formula. *Inventiones Mathematicae* 198 (2014), 739–831. arXiv v3, §3.1, §7.1 and Theorem 7.1, Theorem 7.3(i), the residue for a complementary pair, §§11.3–11.6, and Propositions 11.1(ii) and 11.2.

[20] Ellen Eischen, Michael Harris, Jianshu Li and Christopher Skinner. p-adic L-functions for unitary groups. *Forum of Mathematics, Pi* 8 (2020), e9. Hypotheses 1.4.4(1), p. 13, and §4.2.1, (53), p. 49.

[21] Stephen S. Kudla. Seesaw dual reductive pairs. In *Automorphic Forms of Several Variables*, edited by I. Satake and Y. Morita, Progress in Mathematics 46, Birkhäuser, 1984, pp. 244–268. Section 1, Definition 1.3 and (1.14), p. 249.

[22] Wee Teck Gan, Benedict H. Gross and Dipendra Prasad. Symplectic local root numbers, central critical L-values, and restriction problems in the representation theory of classical groups. *Astérisque* 346 (2012), 1–109. Sections 12 and 23 treat Bessel and Fourier–Jacobi coefficients.

[23] Raphaël Beuzart-Plessis, Pierre-Henri Chaudouard and Michał Zydor. The global Gan–Gross–Prasad conjecture for unitary groups: the endoscopic case. *Publications Mathématiques de l'IHÉS* 135 (2022), 183–336. Theorems 1.1.5.1 and 1.1.5.3.

[24] Hongyu He. On the Gan–Gross–Prasad conjecture for $U(p, q)$. *Inventiones Mathematicae* 209 (2017), 837–884. Theorem 0.1 and the signed interlacing rule in the author's version.

[25] Michael Harris. Testing rationality of coherent cohomology of Shimura varieties. In J. Cogdell et al., eds., *Automorphic Forms and Related Geometry*, Contemporary Mathematics 614 (2014), 81–95. Proposition 2.6 in the author's November 2012 version.

[26] Raphaël Beuzart-Plessis and Pierre-Henri Chaudouard. The global Gan–Gross–Prasad conjecture for unitary groups. II. From Eisenstein series to Bessel periods. *Forum of Mathematics, Pi* 13 (2025), e16, 1–98. §§1.3.5–1.3.6, Theorem 1.3.6.1 and (1.3.6.1), pp. 8–9.

[27] Michael Harris. L-functions and periods of adjoint motives. *Algebra & Number Theory* 7 (2013), 117–155. Lemma 2.5.4, pp. 137–138.

[28] Raphaël Beuzart-Plessis. A local trace formula for the Gan–Gross–Prasad conjecture for unitary groups: the Archimedean case. *Astérisque* 418 (2020). arXiv v3, §6.2, Lemma 6.5.1(ii), §7.1 and Theorem 7.2.1.

[29] Neelima Borade, Jonas Franzel, Johannes Girsch, Wei Yao, Qiyao Yu and Elad Zelingher. On tori periods of Weil representations of unitary groups. *Selecta Mathematica* (N.S.) 31 (2025), article 49. Sections 3.1.1 and 4.1.1.

[30] Dihua Jiang, Binyong Sun and Chen-Bo Zhu. Uniqueness of Bessel models: the Archimedean case. *Geometric and Functional Analysis* 20 (2010), 690–709. Theorem A, pp. 690–691.

[31] Kenichi Bannai and Shinichi Kobayashi. Algebraic theta functions and the p-adic interpolation of Eisenstein–Kronecker numbers. *Duke Mathematical Journal* 153 (2010), 229–295. Example 1.9, Corollary 2.2 and Theorem 2.9 in the authors' version.

[32] John Tate. Fourier analysis in number fields and Hecke's zeta-functions. In J. W. S. Cassels and A. Fröhlich, eds., *Algebraic Number Theory*, Thompson Book Company, 1967, Chapter XV, pp. 305–347, especially Section 4 on the global zeta integrals.

[33] Don Blasius. On the critical values of Hecke L-series. *Annals of Mathematics* (2) 124 (1986), 23–63.

[34] Jie Lin. Period relations for automorphic induction and applications, I. arXiv:1511.03517v1 (2015). Section 1.2, Proposition 1.1, equation (1.1), and Theorem 1.1.

[35] Don Blasius. Period relations and critical values of L-functions. *Pacific Journal of Mathematics* 181 (1997), no. 3, 53–83, Sections L.6–7 and H.6.

[36] Hang Xue. Bessel models for unitary groups and Schwartz homology. *American Journal of Mathematics* 148 (2026), no. 4, 1103–1150. Accepted manuscript, §§1.1–1.2 and Theorem 1.1. For the tempered case, see the same author's *Bessel models for real unitary groups: The tempered case*, Duke Mathematical Journal 172 (2023), 995–1031.

[37] Hiraku Atobe. On the non-vanishing of theta liftings of tempered representations of $U(p, q)$. *Advances in Mathematics* 363 (2020), article 106984. §2.2 and Theorem 2.1(4), (6), (7), pp. 13–15 of the accepted manuscript, Appendix A.3, p. 42, Appendix A.4, p. 45, and Theorem A.4, p. 48.

[38] Wee Teck Gan, Benedict H. Gross and Dipendra Prasad. Restrictions of representations of classical groups: examples. *Astérisque* 346 (2012), 111–170. Proposition 2.1, p. 114.

[39] Atsushi Ichino. Theta lifting for tempered representations of real unitary groups. *Advances in Mathematics* 398 (2022), article 108188. arXiv v2, §4.1, p. 10, and Proposition 7.1, Remark 7.2 and §7.2, pp. 25–26.

[40] Jayce R. Getz and Heekyoung Hahn. A general simple relative trace formula. *Pacific Journal of Mathematics* 277 (2015), 99–118. §4C, (4C.8) and the following unfolding, p. 114.

[41] Emmanuel Kowalski, Abhishek Saha and Jacob Tsimerman. Local spectral equidistribution for Siegel modular forms and applications. *Compositio Mathematica* 148 (2012), 335–384. arXiv v4, Proposition 3.3, p. 18 (published p. 355).

[42] Sug Woo Shin and Nicolas Templier. Sato–Tate theorem for families and low-lying zeros of automorphic L-functions. *Inventiones Mathematicae* 203 (2016), 1–177. Theorem 9.16, p. 111, Theorem 9.19, p. 112, and Corollary 9.25, p. 118.

[43] Zhining Wei. Linear relations of Siegel Poincaré series and non-vanishing of the central values of spinor L-functions. *Journal de Théorie des Nombres de Bordeaux* 34 (2022), 755–766. Lemma 2.3, p. 760, and §3.1, p. 762.

[44] J. A. Green. The characters of the finite general linear groups. *Transactions of the American Mathematical Society* 80 (1955), 402–447. Lemma 7.7 and Theorems 12–14, pp. 439–444.

[45] S. I. Gel'fand. Representations of the full linear group over a finite field. *Mathematics of the USSR-Sbornik* 12 (1970), 13–39. English translation, Theorem 0.5 and Proposition 1.2, pp. 17–18.

[46] Yifeng Liu. Arithmetic theta lifting and L-derivatives for unitary groups, I. *Algebra & Number Theory* 5 (2011), 849–921. Theorem 2.1(1), p. 859.

[47] Shunsuke Yamana. The Siegel–Weil formula for unitary groups. *Pacific Journal of Mathematics* 264 (2013), 235–257. Theorem 2, p. 237, and Lemma 3.1, p. 246.

[48] Yumiko Hironaka. Spherical functions and local densities on hermitian forms. *Journal of the Mathematical Society of Japan* 51 (1999), 553–581. Formula (3.1), p. 572, and Corollary 3.3(i), p. 576.

[49] Shayne Waldron. Generalized Welch bound equality sequences are tight frames. *IEEE Transactions on Information Theory* 49 (2003), 2307–2309. Author's version, Theorem 1.3 and §2.

[50] Éric Gaudron. Pentes des fibrès vectoriels adéliques sur un corps global. *Rendiconti del Seminario Matematico della Università di Padova* 119 (2008), 21–95. arXiv v2, Proposition 4.22, Proposition 5.13 and Definition 6.1.

[51] Huayi Chen. Convergence des polygones de Harder–Narasimhan. *Mémoires de la Société Mathématique de France* (N.S.) 120 (2010). Equation (20), p. 34, and Proposition 2.2.4, p. 42.

[52] Jean-Benoît Bost and Huayi Chen. Concerning the semistability of tensor products in Arakelov geometry. *Journal de Mathématiques Pures et Appliquées* 99 (2013), 436–488. §§1.2–1.3, the Harder–Narasimhan polygon, and §§2.2–2.4, especially (36), p. 457, Definition 2.3, p. 458, and §2.6, p. 461.

[53] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen* 296 (1993), 625–635, Theorem (2.2).

[54] Jean-Benoît Bost. Algebraic leaves of algebraic foliations over number fields. *Publications Mathématiques de l'IHÉS* 93 (2001), 161–221. Proposition 3.9(2), p. 192, and Propositions 4.5–4.6, pp. 199–201.

[55] Christian Krattenthaler and Thomas W. Müller. Truncated versions of Dwork's lemma for exponentials of power series and p-divisibility of arithmetic functions. *Advances in Mathematics* 283 (2015), 489–529. arXiv v2, Theorem 2, (2.4)–(2.6), pp. 4–5.

[56] Sinnou David and Patrice Philippon. Minorations des hauteurs normalisées des sous-variétés des tores. *Annali della Scuola Normale Superiore di Pisa, Classe di Scienze* (4) 28 (1999), 489–543. Proposition 4.10, p. 527.

[57] Huayi Chen. Explicit uniform estimation of rational points I. Estimation of heights. *Journal für die reine und angewandte Mathematik* 668 (2012), 59–88. Ahead-of-print version, Theorem 4.8, p. 22.

[58] Sergio Caracciolo, Alan D. Sokal and Andrea Sportiello. Algebraic/combinatorial proofs of Cayley-type identities for derivatives of determinants and pfaffians. *Advances in Applied Mathematics* 50 (2013), 474–594. Theorem 2.1, (2.1), p. 477.

[59] Hélène Esnault and Eckart Viehweg. *Lectures on Vanishing Theorems*. DMV Seminar 20, Birkhäuser, 1992. Publisher's edition. Lemma 5.8(a), p. 47, and Corollary 5.12(c), p. 49.

[60] Jean-Pierre Demailly. Extension of holomorphic functions defined on non reduced analytic subvarieties. In *The Legacy of Bernhard Riemann after One Hundred and Fifty Years*, vol. I, Advanced Lectures in Mathematics 35, International Press, 2016, pp. 191–222. arXiv v1, (2.4)–(2.5), (2.7)–(2.8) and Remark (2.9)(a), pp. 3–5.

[61] Harish-Chandra. Discrete series for semisimple Lie groups. II. Explicit determination of the characters. *Acta Mathematica* 116 (1966), 1–111. § 41, Theorem 16, p. 96.

[62] Michael F. Atiyah and Wilfried Schmid. A geometric construction of the discrete series for semisimple Lie groups. *Inventiones Mathematicae* 42 (1977), 1–62. Sections 3 and 8, especially (3.10), (3.12) and Theorem (8.1).

[63] Wee Teck Gan, Jonathan P. Hanke and Jiu-Kang Yu. On an exact mass formula of Shimura. *Duke Mathematical Journal* 107 (2001), 101–131. Proposition 2.13 and Section 4, especially (4.2)–(4.3) and Propositions 4.4–4.5.

[64] Spencer Bloch and Kazuya Kato. L-functions and Tamagawa numbers of motives. In *The Grothendieck Festschrift*, vol. I, Progress in Mathematics 86, Birkhäuser, 1990, 333–400. Lemma 3.8.1, Corollary 3.8.4 and Example 3.9.

[65] Laurent Berger. Limites de représentations cristallines. *Compositio Mathematica* 140 (2004), 1473–1498. Sections II.1–II.3, III.2–III.4 and IV.1.

[66] Seunghwan Chang and Fred Diamond. Extensions of rank one (φ, Γ)-modules and crystalline representations. *Compositio Mathematica* 147 (2011), 375–427. Example 7.1 and Lemma 7.2.

[67] Hui June Zhu. Crystalline representations of $G_{\mathbb{Q}_{p^a}}$ with coefficients. arXiv:0807.1078v3, Theorems B and 4.15.

[68] Pavel Čoupek, Evangelia Gazaki and Adriano Marmora. Exactness property of Breuil–Kisin functors and Bloch–Kato Selmer groups. Proposition 3.4 and arXiv:2603.26035v3 (21 October 2026), Theorem 3.12(2) and Corollary 3.16.

[69] Adrian Barquero-Sanchez, Lindsay Cadwallader, Olivia Cannon, Tyler Genao and Riad Masri. Faltings heights of CM elliptic curves and special Gamma values. *Research in Number Theory* 3 (2017), article 13. Theorem 1.3, (5.13), and (6.2).

[70] Carl Ludwig Siegel. *On Advanced Analytic Number Theory*. Tata Institute of Fundamental Research, Lectures on Mathematics 23, original notes 1961, online reissue. Chapter I, §1, Theorem 1, p. 13.

[71] Ehud de Shalit. Artin L Functions. In *An Introduction to the Langlands Program*, edited by Joseph Bernstein and Stephen Gelbart, Birkhäuser Boston, 2004, pp. 73–87. §§3.2, 3.4 and 4.2.

[72] J. Barkley Rosser and Lowell Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois Journal of Mathematics* 6 (1962), 64–94. Theorems 9 and 12, (3.32) and (3.35), p. 71.

[73] Tom M. Apostol. Zeta and related functions. In *NIST Digital Library of Mathematical Functions*, Chapter 25. For zeta see the continuation, alternating series, Laurent expansion, Euler product, reflection formula, value at zero and zero distribution. For Dirichlet L-functions see the series and continuation, Euler product, change of conductor, functional equation and nonvanishing at one. For the gamma function see R. A. Askey and R. Roy, Chapter 5, §5.2(i).

[74] Hugh L. Montgomery and Robert C. Vaughan. *Multiplicative Number Theory I. Classical Theory*. Cambridge Studies in Advanced Mathematics 97. Cambridge University Press, 2007. Theorem 2.7(a), p. 50, Theorem 4.4, p. 116, (10.17) and Corollary 10.8, pp. 332–333, and Theorem 11.3, p. 360.

[75] J. Kaczorowski and J. Pintz. Oscillatory properties of arithmetical functions. I. *Acta Mathematica Hungarica* 48 (1986), 173–185. Theorem 1 and (2.2), p. 175.

[76] J. Kaczorowski and J. Pintz. Oscillatory properties of arithmetical functions. II. *Acta Mathematica Hungarica* 49 (1987), 441–453. Theorem, (1.5) and (1.7), pp. 441–442.

[77] Anthony W. Knapp and David A. Vogan, Jr. *Cohomological Induction and Unitary Representations*. Princeton Mathematical Series 45, Princeton University Press, 1995. Book page.

[78] Nicolas Prudhon. Translation of Dolbeault representations. *Journal of Lie Theory* 29 (2019), 1103–1118. Proposition 3.6 and (14).

[79] Joseph Bernstein and Bernhard Krötz. Smooth Fréchet globalizations of Harish-Chandra modules. *Israel Journal of Mathematics* 199 (2014), 45–111. Theorems 1.1 and 10.6.

[80] Alexander Molev, Maxim Nazarov and Grigori Olshanski. Yangians and classical Lie algebras. *Russian Mathematical Surveys* 51 (1996), 205–282. Proposition 2.7(4) and Remark 2.11.

[81] Binyong Sun. Matrix coefficients of cohomologically induced representations. *Compositio Mathematica* 143 (2007), 201–221. See Theorem 6.3 and (25)–(26).

[82] Joseph N. Bernstein. P-invariant distributions on GL(N) and the classification of unitary representations of GL(N) (non-Archimedean case). In *Lie Group Representations II*, Lecture Notes in Mathematics 1041, Springer, 1984, 50–102. Author's copy.

[83] Andrei V. Zelevinsky. Induced representations of reductive p-adic groups. II. On irreducible representations of GL(n). *Annales Scientifiques de l'École Normale Supérieure* 13 (1980), 165–210. Published article.

[84] Vladimir Platonov and Andrei Rapinchuk. *Algebraic Groups and Number Theory*. Academic Press, 1994. Author's copy.

[85] Chenyan Wu. Irreducibility of theta lifting for unitary groups. *Journal of Number Theory* 133 (2013), 3296–3318. Author's version.

[86] Shunsuke Yamana. On the Siegel–Weil formula: the case of singular forms. *Compositio Mathematica* 147 (2011), 1003–1021. Published article.

[87] Joseph Bernstein and Erez Lapid. On the meromorphic continuation of Eisenstein series. *Journal of the American Mathematical Society* 37 (2024), 187–234. Author's copy.

[88] Stephen D. Miller and Wilfried Schmid. On the rapid decay of cuspidal automorphic forms. *Advances in Mathematics* 231 (2012), 940–964. Theorem A, Theorems 2.23–2.24.

[89] Philippe Graftieaux. Formal groups and the isogeny theorem. *Duke Mathematical Journal* 106 (2001), 81–121. See Proposition 2.8, Theorem 2.10 and Propositions 2.11 and 2.13 for filtered slopes, cubist models and jets along formal subgroups.

[90] David Mumford. On the equations defining abelian varieties. I. *Inventiones Mathematicae* 1 (1966), 287–354. Section 1, Theorem 4, pp. 302–303.

[91] Barry Mazur and Karl Rubin, with an appendix by Michael Larsen. *Diophantine stability*. Proposition 9.17, Theorem 9.21 and Remark 10.4.

[92] Alexander Smith. *The distribution of ℓ^∞-Selmer groups in degree ℓ twist families II*. Theorem 2.14.

[93] Annegret Paul and Peter E. Trapa. One-dimensional representations of $U(p, q)$ and the Howe correspondence. *Journal of Functional Analysis* 195 (2002), 129–166. Theorem 4.1.

[94] Yixin Bao and Binyong Sun. *Coincidence of algebraic and smooth theta correspondences*. Proposition 5.1, unitary dual pairs.

[95] Annegret Paul. First occurrence for the dual pairs $(U(p, q), U(r, s))$. *Canadian Journal of Mathematics* 51 (1999), 636–657. See Theorem 2.7 and Proposition 3.4.

[96] Dragan Miličić. Asymptotic behaviour of matrix coefficients of the discrete series. *Duke Mathematical Journal* 44 (1977), 59–88. Part III, Theorem 1 and its corollary.

[97] Roger Howe. Transcending classical invariant theory. *Journal of the American Mathematical Society* 2 (1989), 535–552. Theorems 1 and 2.1.

[98] Wee Teck Gan and Shuichiro Takeda. A proof of the Howe duality conjecture. *Journal of the American Mathematical Society* 29 (2016), 473–493. Theorem 1.2.

[99] W. Casselman. The unramified principal series of p-adic groups. I. The spherical function. *Compositio Mathematica* 40 (1980), 387–406. Proposition 4.1 and Theorem 4.2 (Macdonald).

[100] Fritz Hörmann. Generalized automorphic sheaves and the proportionality principle of Hirzebruch–Mumford. Theorem 4.2.1.

[101] Jan Hendrik Bruinier and Benjamin Howard. Arithmetic volumes of unitary Shimura varieties. Theorem 8.6.1.

[102] Ziqi Guo. Modular heights of unitary Shimura varieties III: Proof of the main theorem. Preprint (2026). Section 1.1 and Theorem 1.1.

[103] Vincent Maillot and Damian Rössler. Conjectures on the logarithmic derivatives of Artin L-functions II. Theorem 7.26.

[104] Jean-Pierre Labesse. Changement de base CM et séries discrètes. Condition $(*)$ and Corollary 5.3.

[105] Susana Salamanca–Riba. On the unitary dual of some classical Lie groups. *Compositio Mathematica* 68 (1988), 251–303. Theorem 1.2 and Proposition 4.3.

[106] Colette Mœglin and Jean-Loup Waldspurger. Le spectre résiduel de GL(n). *Annales scientifiques de l'École Normale Supérieure* 22 (1989), 605–674. Theorem of the introduction.

[107] Hervé Jacquet and Joseph A. Shalika. On Euler products and the classification of automorphic forms II. *American Journal of Mathematics* 103 (1981), 777–815. Theorem 4.4, p. 809.

[108] Tasho Kaletha, Alberto Mínguez, Sug Woo Shin and Paul-James White. Endoscopic classification of representations: inner forms of unitary groups. Corollary 3.3.2 and Theorem* 1.7.1 with its proved generic, pure-inner scope.

[109] Michael Harris. Cohomological automorphic forms on unitary groups, II: period relations and values of L-functions. Corollary 2.2.7 and Appendix 2.A, especially Corollaries 2.A.3 and 2.A.11.

[110] Henning Haahr Andersen. Filtrations of cohomology modules for Chevalley groups. *Annales scientifiques de l'École Normale Supérieure* 16 (1983), 495–528. Propositions 3.2 and 4.1, pp. 505–506 and 509–510.

[111] Rudolf Tange. A bideterminant basis for a reductive monoid. *Journal of Pure and Applied Algebra* 216 (2012), 1207–1221. Theorem 1.1 and its corollary.

[112] Zongzhu Lin. Representations of Chevalley groups arising from admissible lattices. Sections 1.1 and 2.2–2.3.

[113] Bhargav Bhatt, Toby Gee and Mark Kisin. Reduction modulo p of crystalline Galois representations via μ_p-equivariance. Preprint, arXiv:2607.08660v2 (2026). Theorem 5.6.2 for unramified base fields.

[114] Ryosuke Nakahama. Norm computation and analytic continuation of vector valued holomorphic discrete series representations. *Journal of Lie Theory* 26 (2016), 927–990. Theorem 1.2 (Theorems 4.3 and 5.1), the ratio to the Fischer norm.

[115] Chi-Fang Chen, Jorge Garza-Vargas and Ramon van Handel. A new approach to strong convergence II. The classical ensembles. Proof of Lemma 7.4, the multiplicity of a fixed content in a partition (2024 preprint).

[116] Nicholas M. Katz. Formal groups and p-adic interpolation. *Astérisque* 41–42 (1977), 55–65. Section V, height two.

[117] Peter Schneider and Jeremy Teitelbaum. p-adic Fourier theory. Appendix, Theorem (c), the valuation of the fundamental period.

[118] János Kollár. Singularities of pairs. Theorem 6.4, an isolated non-klt point of a divisor in a prescribed numerical class.

[119] Robert Lazarsfeld. A short course on multiplier ideals. Theorem 2.4 and its proof, Nadel vanishing for a nef and big difference.

[120] Siegfried Bosch. *Lectures on Formal and Rigid Geometry*. Section 2.2, Corollaries 12–13, residue fields and geometric points of the Tate algebra.

[121] Thomas Barnet-Lamb, Toby Gee, David Geraghty and Richard Taylor. Potential automorphy and change of weight. *Annals of Mathematics* 179 (2014), 501–609. Theorem 2.1.1(2)–(4), purity away from ℓ, labelled Hodge–Tate weights and the Iwahori case at ℓ.

[122] Thomas Barnet-Lamb, Toby Gee, David Geraghty and Richard Taylor. Local-global compatibility for $l = p$, II. *Annales scientifiques de l'École Normale Supérieure* 47 (2014), 165–179. Theorem 1.1, comparison of the monodromy partitions above ℓ.

[123] Jean-Benoît Bost. Intrinsic heights of stable varieties and abelian varieties. *Duke Mathematical Journal* 82 (1996), 21–70, Theorem 4.10(v).

[124] Henri Skoda. Application des techniques L^2 à la théorie des idéaux d'une algèbre de fonctions holomorphes avec poids. *Annales scientifiques de l'École Normale Supérieure* 5 (1972), 545–579. Theorem 1, pp. 555–556, equations (3.1)–(3.2).

[125] Jean-Pierre Demailly. *Analytic methods in algebraic geometry*, Grenoble lectures, December 2009. §5.A, Theorem 5.1, the complete Kähler estimate with the inverse curvature operator on bundle-valued forms.

[126] Hiraku Atobe, Masataka Chida, Tomoyoshi Ibukiyama, Hidenori Katsurada and Takuya Yamauchi. Harder's conjecture II. *Advances in Mathematics* 503 (2026), article 111233. Extended version, Lemma 7.6, on subquotient extensions over the residue field of a short crystalline lattice.

[127] Kai-Wen Lan. *Arithmetic Compactifications of PEL-Type Shimura Varieties*. Princeton University Press, 2013. Equation (4.3.1.7), the Fourier coefficient recurrence, Theorem 4.4.16, Mumford's equivalence with the polarized degeneration data, and §4.5.3, torsion over the abelian locus.

[128] Hector Pasten. *Shimura curves and the abc conjecture*. §§13–14, Theorem 14.1 and Proposition 14.2, on Heegner point heights and norms of integral differentials.

[129] Olav K. Richter and Martin Westerholt-Raum. Sturm bounds for Siegel modular forms. *Research in Number Theory* 1 (2015), article 5. Proof of Proposition 2.4, torsion specialization of Jacobi forms.

[130] Robin de Jong and Farbod Shokrieh. Faltings height and Néron–Tate height of a theta divisor. *Compositio Mathematica* 158 (2022), 1–32. Equations (12.1)–(12.2) give the Faltings metric and the probability L^2-metric. Theorem 13.1 and Remark 13.1 give Bost's key formula for a principally polarized Moret-Bailly model in these conventions.

[131] Thomas Zink. The display of a formal p-divisible group. *Astérisque* 278 (2002), 127–248. §4.3, equations (216)–(218), the biextension associated with a pairing of displays.