Series · Markdown source

IV. Theta chains and integral periods

1. The column chain and its period
   - Integral branching at the word's level prime
   - The coefficient field of the intervening theta maps
   - The stabilizer denominator
2. The unshifted lattice and the local inputs
   - Taylor bounds at the two unramified inert classes
     - One lift in Taylor coordinates
     - The orbit exponent and the signal-prime branch
     - Successive lifts on the served inert-class lattice
     - Valuation of the signed period
3. Fitting one common row lattice
   - The integral period-divisibility theorem
4. The word height and the quantitative handoff
   - The finite covolume bound still needed for the height
   - The four divisor bounds on the selected family
   - Projected covolume and the signed height main term
     - The core slope from a deep cover
       - Uniform cusp estimates at the prime cutoff
       - The boundary divisor on a marked power-series arc
       - Good reduction and the primitive theta section
       - Stable theta spanning on the fitted cover
       - Short integral sections on the cover
     - The core slope on the definite terminal lift
5. Remaining propositions for a full prose proof

Fix an ordinary conductor-one family in the sense used by Paper III. For a real cutoff x and a selected index i at that cutoff, fix the corresponding host, translate, and decorated datum. Write E for its CM field, K for the totally complex coefficient field, $\mathcal{O} = \mathcal{O}_K$, and $\sigma : K \hookrightarrow \mathbb{C}$ for the chosen embedding. The finite word set is U. A word specifies, among other data, its own prime T_u and starting column space. The two common type primes are g_0, g_1. Write $J = \{1, \dots, m\}$ for the lifts, $\mathcal{T}$ for the top weights, and $N_j = k + j(t+1)$ for the column rank. At a finite place $w \mid p$, put $O_w = \mathcal{O}_{\mathfrak{p}_w}$, $v_w(\mathfrak{p}_w) = 1$, $e_w = v_w(p)$, and $\ell_w = \log N\mathfrak{p}_w / [K : \mathbb{Q}]$. We use $v_w(0) = +\infty$. All finite-place assertions below quantify over places of K, including distinct places above the same p. Ranks · Places and fractional shifts · Size normalization

Paper VI selects the ordinary family and its fitted host data, with the size and support hypotheses used here. Paper IV constructs the period map and establishes its arithmetic properties under the stated inputs. Paper III uses the resulting maps into one common dual row lattice and their uniform height inequality. The theorem below proves the divisibility from concrete local and global hypotheses. A **release input** is a

proposition of the pinned release whose proof is still required for a complete prose proof. **Classical** refers to an external standard theorem used here. The range of each release input is part of its statement. Ordinary comparison and cutoff data · Division interface

1. The column chain and its period

For each word u and type tuple μ, take the finite classes C_j of the rank-N_j positive unitary lattice datum, including the specified frames and levels. We suppress u, μ in the chain notation. A column at stage j has components f_c on the variable space Y_j for $c \in C_j$. Write $\Gamma_j(c)$ for its finite stabilizer and $d_j(c) = |\Gamma_j(c)|$. The level unfolds the theta step into maps $U_j(c, c')$ from the functions on Y_j to those on Y_{j+1}. For each CM-type index ρ, the nonvanishing coordinate factors are $d_{j,\rho}(y)$ and the integer exponents are $s^{\sharp}_{j,\rho}$, $s^{\mathrm{div}}_{j,\rho}$. Let $D^{\sharp}_j$ multiply by $\prod_\rho d_{j,\rho}^{s^{\sharp}_{j,\rho}}$, and let D^{-}_{j+1} multiply by $\prod_\rho d_{j+1,\rho}^{-s^{\mathrm{div}}_{j,\rho}}$. With the determinant phases ρ_j, ρ'_j of modulus one, the prescribed step is

$$(K_j f)_{c'} = \sum_{c \in C_j} \frac{\rho_j(c)\rho'_j(c')}{d_j(c)} D^{-}_{j+1} U_j(c, c') D^{\sharp}_j f_c. \tag{1}$$

The class weights belong to the source stage of a lift. The tuple changes the phases and unfolded kernels and has a specified starting identification $e_{0,u,\mu}$ from the reference coordinates. The reference tuple is denoted by μ_0. Column operations and twist data · Comparison of base and twisted columns

The terminal operation first evaluates at the framed points, then lowers the rank by one. In the release these two maps are evz and lowerMap. The source period, in its bilinear convention, is

$$\mathcal{P}_{u,\mu}(f, F) = \mathrm{period}_{u,\mu}\left(\mathrm{lowerMap}_{u,\mu}\, \mathrm{evz}_{u,\mu} K_{m-1} \cdots K_0\, e_{0,u,\mu} f,\ F\right). \tag{2}$$

Here F is a row test in the corresponding coordinates, so the row transport is included in the notation. The arithmetic period uses the actual phases, determinant factors, and branching normalization in (1) and (2). Linearity follows from the displayed composition and the bilinearity of the terminal pairing. The chain is a composition of different maps between different ranks. Terminal maps and period · Base and twisted period identities · Linearity in the column

Input 1 (framed realization and terminal branching). For every word and each tuple used at a place, as well as the reference tuple, the realization of (1) supplies compatible embeddings of rational stages $V_{u,\mu,j}$ into the complex column spaces with full $\mathcal{O}$-lattices $\Lambda_{u,\mu,j}$. For every $z \in \Lambda_{u,\mu,m}$, evaluation followed by lowering takes z to the integral point lattice $\Lambda^{\mathrm{pt}}_{u,\mu}$ in the lower-rank point space. At the base tuple the point map intertwines the coefficient-field map and (2) on the entire terminal rational space, with the norm bound in the cited statement. The twisted point map must also intertwine them on its entire terminal rational space. On a framed class the lowering is the normalized transpose of branching in the bilinear polynomial pairing. The required integral image is in the *Weyl* point lattice at the class. The assertion includes the equivariance that makes the result independent of the frame used to compute it. This input concerns the selected theta kernels and their coefficient lattices, including their placewise twists. Rational and integral point map · Base conclusion and metric · Base and twisted hypotheses · Integral transpose of branching

Input 2 (class decomposition on the selected rows). For each integral point z and integral admissible row test F, the terminal period in (2) has the exact expression

$$\sigma\big(P^{\mathrm{pt}}_{u,\mu}(z,F)\big)=\sum_{c}\frac{\sigma(y_c(z,F))}{s_c},\qquad s_c=|\operatorname{Stab}_2(c)|. \tag{3}$$

Here c runs over the classes of the lower-rank point space, $y_c\in K$, and $v_w(y_c)\ge 0$ at every finite place for the reference tuple and at every $p\notin Q(\mu)$ for a twisted tuple. The finite exceptional set $Q(\mu)$ is the release's Q_p for that tuple. The selected tuple $\mu(w)$ satisfies $p\notin Q(\mu(w))$ whenever this twisted valuation is used. The numerator is the restricted class pairing multiplied by s_c. In the concrete coefficient model it pairs a Weyl point coefficient with an integral dual Weyl row coefficient. Their actions, including the inverse character on the row, are dual actions of the whole class stabilizer. These are assertions about the actual terminal point and row. Base class numerator and sum · Twisted decomposition and integrality range · Twisted coefficients · Full stabilizer action · Selection away from the exceptional set

Integral branching at the word's level prime

Write $N=N_m$, $n=N-1$. At a point class c let M_c be its unimodular Hermitian lattice. For a CM embedding τ extend M_c to the coefficient field and denote the coordinate dual by $M_{c,\tau}^{\vee}$. The point coefficient belongs to the Weyl lattice of the lower weight on $M_{c,\tau}^{\vee}$. The row coefficient belongs to the dual Weyl lattice of that weight on $M_{c,\tau}$. These choices specify which integral duality the terminal contraction uses. Transported lattices · Class pairing

Proposition. On the framed class realization with the interlacing weights, common determinant shift and row identifications of the selected family, terminal evaluation followed by lowering preserves the indicated Weyl lattices at every finite place. For a row whose integral class coordinates agree with the selected period model, the numerator y_c in (3) is integral. Both assertions apply to the reference tuple and its prescribed twists.

Proof. Consider one embedding and write λ,ν for the upper and lower weights. Write A_α, D_α for the standard dual Weyl and Weyl lattices of weight α. The normalized bilinear Fock form pairs them perfectly. In characteristic zero let ι_0 take the lower extremal polynomial to the prescribed product of initial and last-row minors ξ. The integral multiplier in the release can be read directly from the possible interlacing weights. With indices starting at zero put

$$D=\sum_i(\nu_i-\lambda_{i+1}),\qquad I_i=[\lambda_{i+1},\min(\lambda_i,\lambda_{i+1}+D)]\cap\mathbb{Z}.$$

Every interlacer η in the same last-coordinate degree has $\sum_i(\eta_i-\lambda_{i+1})=D$ and hence $\eta_i\in I_i$. On its component the integral Capelli operator used here has eigenvalue

$$C(\eta)=\prod_i\prod_{v\in I_i\setminus\{\nu_i\}}\prod_j((v-i)-(\eta_j-j)).$$

If $\eta\neq\nu$, take $v=\eta_i$ at a differing index and $j=i$. For $\eta=\nu$ the shifted intervals I_i-i are disjoint and the only candidate $v=\nu_i$ at $j=i$ was removed. Thus $c_{\mathrm{Br}}=C(\nu)$ is nonzero. The integral bideterminant lift supplies each $a\in A_\nu$ with a representative in A_λ agreeing with $\iota_0(a)$ modulo the other components of this degree. This lift appends the fixed last row to the prescribed column minors. Lower row divided powers leave that row fixed, so comparison with the extremal vector and its divided powers detects the desired component. Applying the Capelli operator to the representative gives $c_{\mathrm{Br}}\iota_0(a)\in A_\lambda$. Multiplier and minors · Integral bideterminant lift · Projection onto the interlacing component

The source normalizes the lowering map b as the transpose of $c_{\mathrm{Br}}\iota_0$. For $a \in A_\nu$ and $v \in D_\lambda$ this gives $(a, bv)_\nu = (c_{\mathrm{Br}}\iota_0 a, v)_\lambda \in \mathbb{Z}$. Perfect duality proves $bD_\lambda \subset D_\nu$. It also explains the factor in this assertion. For $\lambda = (4,2,0)$ and $\nu = (3,0)$ the two factors in $C(\nu)$ are both -3. The lower extremal consequently maps to 9ξ, which vanishes on reduction modulo 3. The proof uses the integral image of the transpose and allows a nonprimitive image of the embedding. Extending scalars and taking the tensor product over τ gives the standard lowering map b_{std}. Integral transpose · Extension to the tensor lattice

Put $z_\tau = \tau(\overline{\pi_{T_u}})$ and $d_\tau = |\lambda_\tau| - |\nu_\tau|$, where $|\cdot|$ denotes the sum of weight entries. In the framed point coordinates the actual map is $b_z = (\prod_\tau z_\tau^{-d_\tau}) b_{\mathrm{std}}$. The upper Hermitian lattice there is $M_c \oplus \overline{\pi_{T_u}}^{-1}\mathcal{O}_E$. Its coordinate dual therefore has last line $z_\tau O_w$. If B_τ is a local basis matrix of $M_{c,\tau}^\vee$, the upper basis matrix is $H_\tau = \operatorname{diag}(B_\tau, z_\tau)$. Block equivariance and the last-coordinate degree give

$$\rho_\nu(B^{-1}) b_z(v) = b_{\mathrm{std}}\big(\rho_\lambda(H^{-1})v\big).$$

For an integral upper coefficient the right side belongs to the standard lower Weyl lattice. This proves integrality also at $w \mid T_u$. Moving the framed point by a stabilizer acts by the same finite character on both sides, and the polynomial maps commute with the coefficient embedding. This gives the descended map on the full rational terminal space. Split dual lattice · Local basis identity

For the class numerator, the integral transport of the chosen row puts its coefficient r_c in the dual Weyl lattice in the basis dual to the point basis. The row carries the inverse of the point character, so $y_c = (z_c, r_c)_\nu$ is independent of representatives and belongs to O_w. Orthogonality of different class summands in the Petersson realization identifies y_c with s_c times the restricted class pairing. This gives (3) with precisely the stabilizer denominator shown there. For a twisted tuple the constructed integral row gives the same contraction, and the rest of the division argument uses it at $p \notin Q(\mu)$. In the reference metric, summing the squared per-class branching bound with weights $1/s_c$ gives the norm bound of input 1 when the framed class map is injective with matching stabilizer orders. Those metric and row identifications are properties of the family in Paper VI. The descent and valuation estimates for the earlier theta steps are used in §2. Details of the integral branching and both row comparisons Selected integral row · Twisted contraction · Sum in the point metric

The coefficient field of the intervening theta maps

Fix a word and a prescribed tuple. In the decorated model of Paper VI all stages use one coefficient field K and one embedding σ. The ranks, kernels and the local stabilizer actions can change at each step. We give the descent calculation because algebraicity of individual values alone would allow a different field for each choice of representatives.

First consider the elliptic coefficient. On the auxiliary CM elliptic curve with period lattice Λ and invariant differential dz, take the reduced odd theta F with $F'(0) = 1$. Write $a = \pi/\operatorname{area}(\Lambda)$ in the convention that its translation is $T_x(z) = e^{-a(z\bar{x} + |x|^2/2)} F(z + x)$. Use the rigidified origin line bundle and the CM endomorphism with differential $\alpha = \sqrt{-d}$. Its odd norm $h = d$ fixes the parity of the lattice semicharacter. The pullback of the line is its h-th tensor power, with the isomorphism fixed at the origin. In an algebraic formal frame let $b(z)$ denote that pullback multiplier, using the logarithm for the chosen differential as coordinate. A normalized frame change satisfies $b(z)v(\alpha z) = v(z)^h$ and $v(0) = 1$. The coefficient of degree j of $\log v$ is the corresponding coefficient of $\log b$ divided by $h - \alpha^j$. These denominators are nonzero. Equality of absolute values would require $j = 2$, when $\alpha^2 = -h$. This proves

rationality and functoriality of the normalized frame over the field of the curve, the CM action and the differential. The derivative at zero identifies the associated divisor section with F. CM line and differential calculation · Bannai–Kobayashi, §§1–2

For a torsion translate choose $s > 0$ with $sx/2 \in \Lambda$ and set $Q = [x/(2s)]$. Over a field containing Q, choose H_Q with divisor $t_Q^*[s]^*[0] - [s]^*[0]$. This divisor is principal since $s^2 Q = 0$. The quotient $R = t_Q^*(H_Q/[-1]^* H_Q)$ is independent of the scalar used in H_Q. Set $w = x/2$. Pullback along the uniformization writes H_Q as a constant times $e^{-asz\bar{w}} F(sz+w)/F(sz)$. Its multipliers cancel since $sw \in \Lambda$. Taking the indicated quotient and using oddness gives

$$R([z]) = e^{-2a(sz+w)\bar{w}} \frac{F(sz+2w)}{F(sz)} = \frac{T_x(sz)}{F(sz)}.$$

Thus every Taylor coefficient of the prescribed translate commutes with embeddings transporting the divided point and the differential. For the even theta one divides by the nonzero constant at the specified two-torsion translate. Under a change of differential by c, its normalized coefficients in degree j scale by c^{-j}. This comparison works in characteristic zero. Its finite-place valuation is addressed in §2. Divided point, scalar and changes of differential

In a step from rank a_0 to b_0, let $h_e = \chi_F^{-1} J_e$ be the normalized derivative of the induced theta function in the multidegree required by the outgoing weight. Here J_e includes the prescribed powers of the generator of the different. The finite Heisenberg expansion expresses fourth powers of the vacuum germs in terms of the elliptic translates. In characteristic zero a nonzero leading homogeneous part determines all subsequent coefficients of a chosen fourth root recursively, and its leading coefficients are algebraic. Finite Fourier and lattice-translation sums preserve algebraicity. Finiteness of the self-dual genera, the displacement quotients, the degree bounds and the twist packet yields a finite generating set for the h_e. Arbitrary representatives are obtained by rational unitary basis changes. Their effect on derivatives is the chain rule with coefficients in the field of the embedded CM entries. The fitted field contains this generating set simultaneously for all words, stages and tuples. Finite generator and fitted-field construction

In the source and target determinant frames write $D_c, D_{c'} \in K^\times$ for their inverse character factors. The product of the two phases in (1) is $\sigma(D_c D_{c'}^{-1})\chi_F^{-1}$. Consequently the phased kernel coefficients are linear combinations over $\sigma(K)$ of the h_e with the prescribed finite character values. Sharpening multiplies the source polynomial by nonnegative powers of the frame determinants. After integral matrix substitution and contraction with these coefficients, division uses exactly the determinant shift of the target weight. The result has that weight and is equivariant for the complete target stabilizer. Summation with the rational source weights $1/d_j(c)$ yields a K-linear map on the rational stages. Its image under σ is (1), and recovery of the polynomial from its frame evaluations proves uniqueness. At the initial stage the class permutation preserves the transported lattices. Composition with terminal branching therefore defines (2) on these same K-spaces. Proof with the three terminal and nonterminal kernel formulas

The stabilizer denominator

Let $G = D_\Gamma$ be the common positive integer used for the datum. The framed point stabilizer injects into the corresponding terminal column stabilizer, whose order divides G. In the concrete true-column model G is the lcm of the orders of the unitary automorphism groups of all unimodular lattices in the indicated standard positive Hermitian spaces of ranks $N_1, \dots, N_m$ and three auxiliary ranks at most N_m. These

identifications are release inputs. For each smaller rank, orthogonal extension by standard unit lines embeds each automorphism group in one at rank N_m. Thus G is already the lcm at N_m. The standard lattice contains the signed permutation group, so $2^{N_m} N_m! \mid G$. **Classical.** Minkowski's finite-subgroup bound for rational linear groups, applied to restriction of scalars, gives $G \mid M(2fN_m)$ in the release's dimension bound, where $M(n)$ is the Minkowski integer and f is the size parameter for the embeddings. In particular a prime factor of G is at most $2fN_m + 1$. The ordinary level window and this bound give the required $T_u \nmid G$ on the selected tail. Common integer and stabilizer injection · Point stabilizer · Ranks in the lcm · Orthogonal extension · Minkowski divisibility · Level-prime coprimality

Proved here. Put $S_{u,\mu} = \operatorname{lcm}_c s_c$, with the empty lcm equal to one. At a place of numerator integrality, $S_{u,\mu} P^{\mathrm{pt}}_{u,\mu} = \sum_c (S_{u,\mu}/s_c) y_c \in O_w$. Thus the signed valuation of the terminal pairing is at least $-e_w v_p(S_{u,\mu}) \geq -e_w v_p(G)$. The bound holds on the localized point lattice since multiplying a point by a denominator outside $\mathfrak{p}_w$ preserves the valuation. For any individual nonzero value equality in the first bound holds precisely when

$$\sum_{v_p(s_c) = v_p(S_{u,\mu})} \overline{S_{u,\mu}/s_c}\,\overline{y_c} \neq 0 \qquad \text{in } O_w/\mathfrak{p}_w O_w.$$

All other terms vanish on reduction, since their integer coefficients are divisible by p. This criterion concerns the sum on the actual pair and permits cancellation among its maximal-denominator terms. The aggregate improvement from replacing the common order by the point lcm, even when $\mu = \mu(w)$ varies, is at most $2 \log G$ per word in the squared-norm height convention. Indeed $0 \leq v_p(G) - v_p(S_{u,\mu(w)}) \leq v_p(G)$ and $\sum_{w \mid p} \ell_w e_w = \log p$. Base denominator calculation · Twisted denominator calculation · Localization

At the preceding m steps a stabilizer bound is needed for (1) itself. In the ordinary good-place kernel estimate with parameter 1, the first step has no stabilizer charge and the remaining steps each cost at most $v_p(G)$ in p-normalized valuation. The terminal charge is one more copy. As a consequence of the local step inclusions, the tower and terminal loss is at most $m v_p(G)$ in these units, or $2m \log G$ after the place sum. The improved point lcm addresses the terminal copy. Archimedean unit modulus of a phase does not by itself give a finite-place unit comparison. Good-place kernel inclusions · Tower stabilizers · Per-lift correction

2. The unshifted lattice and the local inputs

Let V_u be the coefficient-field space of the starting columns on which division will be performed. At T_u the ramified part is the intersection of the initial column subspace with the kernels of the finite-unit norms $\sum_{a \in k_v^\times} \rho_v(a)$ for the selected residue fields k_v. Write M_u^{ram} for the pullback of the initial integral column lattice to V_u under the rational starting map ν_u. At a place w, the local module $S_{u,w}^{\mathrm{ram}}$ is the intersection of that stage's localized lattice with the rational ramified subspace in the stage coordinates. The lattice N_u below is the release's *unshifted* NuLat_u. Norm kernels · Pullback · Local ramified and level lattices

Write $g_{D,j}(w)$ for the assigned gain at lift j, set to zero at a prime in the drop set. Write $F_j(w)$ and $E_{u,j}(w)$ for the specified correction terms. The subscript in F_j distinguishes it from the row input F. Define

$$s_u(w) = \sum_{j=1}^{m} \big(g_{D,j}(w) + F_j(w) - E_{u,j}(w) \big), \qquad a_u(w) = \lfloor e_w s_u(w) \rfloor. \tag{4}$$

The data have a finite support $\mathcal{S}$ such that $s_u(w) = 0$ off $\mathcal{S}$. Off the word level, $E_{u,j}$ is $v_p(G)$ plus the branch correction. The correction is E_{atG} at a type prime, E_{atD} at the discriminant class ramD, E_{atC} or E_{atI} at the inert or signal primes up to the upper cutoff, and zero in the remaining branch. At T_u the term is $E_{\mathrm{atT}}(u, w)$. In particular s_u can be negative. Corrections and support · Piecewise error · Signed size and drop set · Sum over lifts

More explicitly, $F_j(w)$ is $-(D_j^+ + D_j^-)/2$ at a type prime, using the sums of the two signed row-weight lists. At the discriminant prime it is $-(k-1)/2$ times the sum of the minima of the two lift weights in each class pair. At the inert or signal primes up to X'' it is the averaged orbit expression in the lift weights, with the geometric factor $(p^2-1)^{-1}$ for a trivial inert translation. It is zero in the other cases. Finite correction formulas

Release input 3 (local construction and gain). For every word, N_u is obtained by fitting the following full local column lattices in V_u. At $p \neq T_u$ it lies in $(M_u^{\mathrm{ram}})_w$. When the served set $VV(p)$ is nonempty, its image under the comparison of rational column spaces is in the served module NP of the local chain. When $VV(p) = \varnothing$, its image is in $S_{u,w}^{\mathrm{ram}}$. At $p = T_u$ its image is exactly the rational part of

$$T_u^{m \sum_{v \in \mathcal{V}_{u,w}} b_{\mathrm{top}}(\rho(v))} \left(\prod_{v \in \mathcal{V}_{u,w}} U_v \right)^m S_{u,w}^{\mathrm{ram}}. \tag{5}$$

The product uses the selected places in their prescribed order. The set $\mathcal{V}_{u,w}$ is empty unless w sees the chosen prime element in the sense of seenP, and otherwise contains the places whose associated index is top. In (5), U_v sums over local vectors with the specified linear form equal to one, transporting each summand by its frame change. The off-level comparison and the level equality assert membership of the same rational input after these transports. Served or unserved membership · Fitted column membership · Level operator and lattice · Fitted level equality

Here $VV(p)$ uses the good orbit representatives for the Frobenius translation, provided $p < X'$ is in the residue-zero class, the orbit order is at least $x^{\varepsilon/4}$, the excluded band is avoided for $p \geq p_0(k)$, and the sum $\mathrm{Sig}(p)$ of the orbit rank costs is at most $x^{-\varepsilon}$. For a served orbit the lift gain is $\max(g_U - \mathrm{Pr}, g_0, 0)$, where g_0 is the baseline, g_U is the candidate from the divided lift, and Pr is its transport allowance. An unserved orbit uses $\max(g_0, 0)$. Their sum defines gainP subject to the cutoff and trivial-translation exclusions. The corresponding gainS also removes the exceptional, split and unused prime classes as specified in the source. The numerical comparison required below is $g_{D,j}(w) \leq \mathrm{gainS}(p, j)$. Orbit eligibility, allowances and assigned gain · Restrictions on gainS · Numerical comparison

In the served coordinates let M be the base local lattice, V^{ram} the selected rational subspace, and Y_c the operator attached to orbit c. For a subset $S \subseteq VV(p)$ let Y_S be the chosen ordered product of the Y_c^m for $c \in S$. With $S_c = \lceil \sum_j (g_U - \mathrm{Pr} - \max(g_0, 0))_+ \rceil$ on orbit c, the source module NP is

$$\left(\sum_{S \subseteq VV(p)} p^{\sum_{c \notin S,\, c \in VV(p)} S_c} Y_S M \right) \cap p^{\tau(w)} M \cap V^{\mathrm{ram}}.$$

The sum denotes the module spanned by those images under the coefficient action of the local model. Thus an omitted orbit contributes a scalar power of p, and the intersection with $p^{\tau(w)} M$ imposes the common local lattice price. Membership after the rational comparison and the length estimate are separate assertions. Products and served module

The quantitative part of input 3 asserts that off the word level

$$(N_u)_w \subseteq \mathfrak{p}_w^{e_w \tau(w)} (M_u^{\mathrm{ram}})_w, \qquad \frac{\mathrm{length}_{O_w}((M_u^{\mathrm{ram}})_w/(N_u)_w)}{e_w} \le d_u^S \, (\tau(w) + \mathrm{io}(w)). \tag{6}$$

Here d_u^S is the ramified column dimension. The concrete τ is the nonnegative integer 0 for the empty served set and otherwise

$$\tau(w) = \mathrm{ep}(p) + m \sum_{c \in \mathcal{VV}(p)} \big(\lceil t_v(p,c) \rceil_+ + D_p + 1 \big).$$

The two formulas for t_v, D_p meet at the same small-prime threshold $p_0(k)$ used for the serving construction. The parameter $\mathrm{io}(w)$ is $2\,\mathrm{Sig}(p)$ times the maximum over served orbits of the rounded sum of positive gain residuals, and is zero when the set is empty. Input 3 includes, for every inert-class or signal place with $p \le X''$, the prime-sum comparison

$$\sum_j \sum_{t \in \mathcal{T}_w} \big(\mathrm{vmClass}_{j,t}(p) - \mathrm{loG}(p) \big) - m \, \mathrm{flG}(p) \;\le\; \sum_j \mathrm{gain}_j(w) - \tau(w). \tag{7}$$

The class sum is the sum of $(a_t/p^r - \mathrm{wPart}(j,t,p,r))_+$ over odd $r < \lfloor \log_p \lfloor a_t \rfloor \rfloor + 2$. The bound also asserts $0 \le \mathrm{gain}_j(w) \le 4|\mathcal{T}|a/p$. Paper III uses (7) on the eligible primes outside the drop set. The level-prime length comparison has its own bound in terms of the selected top places and is supplied with (5). Concrete exponent and index allowance · Class sum · Local price and comparison · Level-prime length

Release input 4 (column and type-prime bounds). Here a statement that a column lies at depth r in a terminal lattice means its coordinate value is az, z in the localized integral terminal lattice and $v_w(a) \ge e_w r$. Equivalently it is in $\mathfrak{p}_w^{\lceil e_w r \rceil} \Lambda_{u,\mu,m,w}$. All three clauses refer to the actual normalized column (1), with the starting twist for $\mu(w)$.

1. For $p \notin \{T_u, g_0, g_1\}$, for every used μ with $p \notin Q(\mu)$, and every input in the served comparison module or the empty-served-set module of input 3, its terminal column lies at depth

   $$\sum_j \big(\mathrm{gainS}(p,j) + F_j(w) - E_{u,j}(w) \big) + v_p(G). \tag{8}$$

   The underlying good-place step estimate is for the localized integral coefficient lattices in both the base and twisted chain. For an empty served set with a nonzero translation or a signal prime it is augmented by the baseline gain, and the trivial residue-zero translation has the distinct E_{atC} estimate. At the discriminant class the depth on $S_{u,w}^{\mathrm{ram}}$ is $\sum_j (F_j - E_{\mathrm{atD}}(j) - v_p(G)) + v_p(G)$. In particular the discriminant estimate is required even though a good-place step excludes $p = d$. For a nonempty served set, the needed further input is the divided local operator estimate on NP, followed by transport through the twisted chain. It gives the orbit maxima in (8) on the same input, with the small-inert-prime bound on the corresponding formal group systems and the intertwining of their local and global column maps. Aggregate bound and domain · Good-place integral steps · Local-to-global column comparison and divided operators · Unserved and discriminant bounds
2. At $p = T_u$ the terminal column of every input in (5) has depth $-mE_{\mathrm{atT}}(u,w)$ in the corresponding localized terminal lattice. In the base chain this says that multiplication by any scalar with valuation at least $e_w m E_{\mathrm{atT}}$ puts the column in that lattice. Here mE_{atT} is an integer. Indeed m is even, t is odd, and the only fractional term in E_{atT}, when it is nonzero, is $|\mathcal{T}|(t+1)(m+1)/4$. Thus $T_u^{mE_{\mathrm{atT}}}$

has valuation exactly $e_w m E_{\mathrm{atT}}$, and dividing by that scalar proves the asserted base depth. The twisted chain needs its own inclusion at the same depth for every used twist outside its exceptional set. The numerical level data give $\tau(w) = g_{D,j}(w) = F_j(w) = 0$ and $E_{u,j}(w) = E_{\mathrm{atT}}(u, w)$, with $T_u \nmid G$. Sharp level assertion and size data · Base level and parity · Twisted level assertion · Denominator at the level prime

3. At $p = g_i \neq T_u$, for every used tuple and every input of M_u^{ram}, the pairing on each integral row test in that tuple has a K-value with v_w/e_w at least $-\sum_j v_{NN}(i, u, j, w) - C_1 \mathrm{errC}$. Here v_{NN} is the source weight sum of valuations of the conjugate prime generator, divided by e_w. The numerical inequality bounds $\sum_j (F_j - E_{u,j})$ above by this same lower bound, and the gains vanish. This branch requires the character projector and its integral lattice bound. The generic good-place assertion excludes these two primes. Period bound on integral tests · Conjugate generator weight · Numeric type-prime comparison · Selected type-prime source

For the complement estimate used in the row comparison one needs the bound on all of M_u^{ram} off T_u, with depth $\sum_j (F_j - E_{u,j}) + v_p(G)$ for terminal columns off the two type primes and with the direct period bound at the type primes. This is an additional clause of input 4. It yields the period bound without the gain on this larger domain by the argument below. Complement hypotheses and conclusion

Proved here (local period). For every $u, w, \phi \in N_u$ and integral admissible row test F at $\mu(w)$, the unique $y \in K$ whose image is $\mathcal{P}_{u,\mu(w)}(j_u \phi, F)$ satisfies $v_w(y) \geq e_w s_u(w)$. Off the three level and type primes, write the terminal column as az using (8). The terminal period of z is rational and has valuation at least $-e_w v_p(G)$ by inputs 1 and 2 and (3). By the composition (2), y is a times this value. The valuation bound is therefore $e_w(\sum_j (\mathrm{gainS} + F_j - E_{u,j}) + v_p(G) - v_p(G)) \geq e_w s_u(w)$. The same subtraction gives the complement estimate from its larger-domain clause. At T_u the denominator bound is zero by coprimality, so clause 2 gives the claim since $s_u = -m E_{\mathrm{atT}}$. At a type prime, for a given $\phi \in N_u$ choose $q \in \mathcal{O} \setminus \mathfrak{p}_w$ with $q\phi \in M_u^{\mathrm{ram}}$ by (6). Apply clause 3 there and divide the period by q using linearity. Its valuation is unchanged and the zero-gain formula gives (4). Uniqueness follows from injectivity of σ. This also proves $y \in \mathfrak{p}_w^{a_u(w)} O_w$. Transport of scaled columns to the period · Level and type-prime transport · Assembly of the three cases

Taylor bounds at the two unramified inert classes

The ordinary host distinguishes two unramified classes of primes inert in the auxiliary imaginary quadratic field. The formal class `signal` and the class `inertClass` are disjoint. A served prime is in the latter class because the predicate defining $VV(p)$ requires `inertClass`. Thus on the `signal` branch $VV(p) = \varnothing$ and the gain uses the baseline even though the prime contributes to the ordinary signal. A nonempty $VV(p)$ uses the additional operators below. Fix a word $u, w \mid p$, and a used tuple μ with $p \notin Q(\mu)$. Assume $5 \leq p \leq X''$ and $p \nmid T_u \, \mathrm{condN}\, g_0 g_1$, where $\mathrm{condN} = 1$ for ordinary conductor one. In the branch with $VV(p) \neq \varnothing$ use the local model selected at u, w, μ and its orbit and lattice identifications. These hypotheses specify the branches of input 4 treated in this section. Prime classes and orbit eligibility · Definition of the served set · Branch selection in the proof

One lift in Taylor coordinates

Use the p-adic valuation $v = -\log_p |\cdot|_p$ on $\mathbb{C}_p$, so $v(p) = 1$. The algebraic embedding at w identifies this valuation on K with v_w/e_w. At lift $j = i + 1$ take a pair of source and target classes. Its Taylor series Θ has variables $Z_{z,a,b}$, with z an embedding label in $\{0, 1\} \times H$ where H is the translate group, $1 \le a \le N_i$, and $1 \le b \le N_{i+1}$. The label z carries the Young diagram ν_z determined by the word and lift. Its row lengths are the $\nu_{j,\rho,r}$ from Section 4 at the embedding ρ belonging to z. For a diagram column of height h and an increasing list $I = (b_1, \dots, b_h)$ of target indices, use the symbol

$$\det\left(\sum_{a=1}^{N_i} X_{z,\ell,a} Z_{z,a,b_k}\right)_{1\le \ell,k\le h}.$$

Here $X_{z,\ell,a}$ are polynomial variables for the input frame. Multiply these determinants over the diagram columns and over z. In the resulting polynomial replace each monomial Z^α by $\alpha![Z^\alpha]\Theta$, where α is the tuple of derivative orders and $\alpha! = \prod_{z,a,b} \alpha_{z,a,b}!$. This is the signed differential operator applied at zero. The release's coefficient bound quantifies over every list choice and every coefficient in the X variables, as well as every class pair of the lift. Each product on a given orbit uses exactly one derivative per box in its diagrams. This follows by expanding each determinant of height h into products of h entries linear in Z. Diagrams and class-pair series · Determinants and derivative extraction

Release input 4, local series. On a Frobenius orbit of even length $d \ge 2$, number the positions by $s \in \mathbb{Z}/d$. The admissible Taylor series has an integral cofactor after the coordinate substitution $Z_{s,a,b} \mapsto \eta L_{s,a,b}(Z)$ with $|\eta|_p = 1$ and

$$L_{s,a,b}(Z) = \sum_{h\ge 0} p^{-h} Z_{s-2h,a,b}^{p^{2h}}.$$

In these substituted coordinates the other factor is an exponential of the signed orbit series be inserted into each matrix entry and composed with the same logarithm L without the scalar η. Its sign is -1 in the true-column construction. The hypotheses on be are coefficients in the unramified field, invariance under the roots of unity ζ with $\zeta^{p^d-1} = 1$ acting by $\zeta^{(-1)^s p^s}$ in slot s, no pure-parity terms or terms of bidegree $(1, 1)$, and integrality with constant one of the exponentiated additive defects in each parity after log substitution. The additive defect in one parity is the value at the sum of two sets of variables of that parity minus the two separate values with the other parity fixed. These hypotheses come from the normalized CM theta sections, their algebraic torsion translates and the formal group comparison at the fitted places. The true-column theorem uses, in particular, algebraicity of translates, the genus-one integral model and curve dictionary, inert integrality, and level-free comparisons. These geometric assertions on the actual kernel remain release inputs here. Logarithm · Integral cofactor and orbit shape · Substitution and admissibility · True-column construction

The analytic argument passes to the inverse of L, which exists formally by its identity linear term. Denote the inverse by g. Working after the unit rescaling of Θ leaves all valuations of the derivative coefficients unchanged, since every monomial of degree n is multiplied by η^n. In these coordinates the series has an integral-cofactor term composed with g and an exponential factor. One may also apply integral linear frame substitutions before taking a derivative. Set $\rho_p = (p^2 - 1)^{-1}$ and $\gamma_p = p/(p^2 - 1) = 1/(p - 1) - \rho_p$. For a vector of nonnegative radius parameters $\rho_p + \delta_\ell$, indexed by the variables $\ell = (s, a, b)$ on an orbit, the Gauss weight of a term cZ^β is $v(c) + \sum_\ell (\rho_p + \delta_\ell)\beta_\ell$. The inverse-logarithm estimate and the linear-

substitution estimate preserve nonnegative weight on the cofactor. For the logarithm of the exponential factor one asks for weights at least $1/(p-1)$ on the monomials with $\beta \le \alpha$ componentwise, where α is the multi-index being tested. The release's pure-cross estimate reduces the extra radius constraints to exceptional terms with two powers p^h, p^l at positions s, s' in the same orbit. They can require an extra inequality only when $s \ne s'$, $s + h = s' + l$ modulo d, and $h + l$ is odd. In that case the lower estimate for the coefficient at the base radius is $1/(p-1) - 1$, so $p^h \delta_\ell + p^l \delta_{\ell'} \ge 1$ suffices. The frame argument bounds exceptional couplings to the first rows at top positions by the prescribed list of row labels and the capacities of their partners. Inverse substitution · Pure-cross Gauss step · Which frame terms need the extra inequality · Row-label and frame conditions

Proved here (from the weighted coefficient hypotheses). The weights through α suffice for a bound on $[Z^\alpha]$ of the exponential. The logarithm has zero constant term, so only finitely many powers occur in that coefficient. Each term in its q-th power divided by $q!$ has weight at least $q/(p-1) - v_p(q!) \ge 0$, by **Classical** Legendre's formula. In the coefficient of a product every contributing pair of indices sums to α. Thus the cofactor of nonnegative weight and this exponential give the derivative bound

$$v\big(\alpha![Z^\alpha]\Theta\big) \;\ge\; \sum_\ell \alpha_\ell(\gamma_p - \delta_\ell) - dN_iN_{i+1}(\log_p A + 1). \tag{8a}$$

Here $A \ge p$ bounds the number of boxes at every position and hence each α_ℓ. In this formula Θ denotes the series in the frame in which the weighted estimate holds. Legendre gives $v_p(n!) = n/(p-1) - s_p(n)/(p-1)$, where $s_p(n)$ is the sum of the base-p digits of n. For $0 < n \le A$ there are at most $\log_p A + 1$ digits, and for $n = 0$ there is no loss. Subtracting the radius contribution from this factorial bound proves (8a). Therefore bounds on the derivatives imply bounds on their integral linear combinations in the signed polynomial. For the refined estimate the source instead evaluates the *whole* frame polynomial in residually generic variables, so estimates on this evaluation recover all its coefficients by the Gauss norm. The frame labels, generic evaluation and inverse substitution are the specific steps used to transfer the derivative estimate. Legendre bound and multiplication · Logarithm loss at the derivative · Generic frame and return to coefficients

The orbit exponent and the signal-prime branch

For the part of the derivative away from the first rows at top positions, write D_s for the number of boxes remaining at position s. A transport variable $x_{s,h}$ gives a nonnegative amount assigned to the odd step $1 \le h < d$ from s to $s + h$. The capacity at s constrains $\sum_h (p^h x_{s,h} + x_{s-h,h}) \le D_s$. If $X(D)$ is the maximum of $\sum_{s,h} x_{s,h}$, write $\mathcal{F}_p(D) = \gamma_p \sum_s D_s - X(D)$. **Classical.** The finite linear programming duality theorem gives $X(D) = \min \sum_s D_s \delta_s$ with $\delta_s \ge 0$ and $p^h \delta_s + \delta_{s+h} \ge 1$ for the same steps. The feasible primal region is nonempty and bounded for nonnegative capacities and the dual contains the constant vector 1. The form of the dual is obtained by pricing each capacity inequality. This accounts for the appearance of $\mathcal{F}_p$ when choosing radii in (8a). In particular the equal radii $\delta_s = 1/(p+1)$ alone give $\gamma_p - 1/(p+1) = \rho_p$. The refined proof uses a separate assignment at the top rows. Linear program in the size data · Radius assembly for the row degrees

Release input 4, refined jet bound. Let T be the top positions of an orbit, A_s and B_s the lengths of its first and second rows, and W_s the largest partner capacity among predecessors $s - h$ with h positive odd and p^h at most the number of boxes at the predecessor. The capacity at position y is B_y for $y \in T$ and A_y

otherwise. Take the empty maximum to be zero. Write h_s for the least positive odd step to T from s, with $h_s = \infty$ if there is no such step. Define

$$\operatorname{base}(s) = \sum_{\substack{h \geq 1 \text{ odd} \\ h \neq h_s}} \big((A_s - B_s)/p^h - \operatorname{cap}(s+h) \big)_+ .$$

The sum converges by comparison with the geometric series. Here cap is the capacity just defined. Let $w_s > 0$ be the fitted lift weight at s. With D_s as above, the hypotheses for the corollary are $D_s = (k-1)w_s + E_s$ off T and $D_s = (k-3)w_s + E_s$ on T, with $0 \leq E_s \leq 4w_s$. Also A bounds every diagram size, $p \leq A$, and $2 \log_p A + 2 \leq N_i - 1$. In the cited application $A = a^2$, where a is the host scale of (7). The bound on each signed polynomial coefficient at a single orbit is

$$\sum_{s \in T} \operatorname{base}(s) + \mathcal{F}_p\big((k-1)w\big) - \frac{11}{p} \sum_s w_s - \gamma_p \sum_{s \in T} W_s - dN_i N_{i+1}(\log_p A + 1). \tag{8b}$$

This is the split-constant version in the selected proof. The first estimate there is the frame Gauss estimate with the actual D_s and the first-row top term. The comparison of linear programs gives $\mathcal{F}_p(D) \geq \mathcal{F}_p((k-1)w) - 11 \sum_s w_s/p$. The source proves this via its odd-level gain inequality. It uses the inequality of the base term with the first-row top term to deduce (8b). The proofs of the pure-cross estimate, the required frame labels and the quantitative linear-program comparison for all diagram sizes are retained as release inputs, with the coefficient argument (8a) separated above. Top and partner terms · Frame Gauss estimate with integer diagram sizes · Comparison giving $(11, 1)$ · Selected use of the split constants

The fitted baseline can be read directly in these orbit coordinates. Denote by $\mathcal{T}_c$ the top labels on the orbit of c, by $A^*_t = a_t - 2c_G$ the reduced top weight with the host's error cushion c_G, and by $C_{j,t}(h)$ the source's adjusted fill fillG at the partner of t after h steps, plus the same c_G. The fill is the weight with the specified shift $\big(k + (2j-1)(t+1)\big)/2$ removed, where t in this rank shift is the rank parameter of N_j. For the orbit representative c, the source sets

$$g_0(p, j, c) = \sum_{t \in \mathcal{T}_c} \sum_{\substack{h \geq 1 \text{ odd} \\ h \text{ not excluded for } t}} \big(A^*_t/p^h - C_{j,t}(h) \big)_+ - \mathrm{dl}(p, c).$$

The exclusion removes the first hit on each distinct top label, if that first hit is odd. Here $\mathrm{dl}(p, c)$ is the loss $K_{\mathrm{dl1}} d N_m^2 \log_p(2a) + K_{\mathrm{dl2}} |\mathcal{T}_c| (\log_p(2a))^2$, where the two K-constants are from the fitted data. For the signal class odd steps go to Boolean coordinate one and all top labels lie on coordinate zero. Thus no odd step is excluded there. In the operated branch, $g_U(p, j, c)$ starts from $\gamma_p \sum_{t \in \mathcal{T}_c} A^*_t$ and subtracts $\mathrm{dl}(p, c)$ and the sum of $\min(A^*_t/p^h, C_{j,t}(h))$ over odd h and tops at odd orbit positions relative to c. The transport allowance $\mathrm{Pr}(p, j, c)$ then charges the different choices of the local operator at $p < p_0(k)$ and $p \geq p_0(k)$. These are the same quantities used in the generators of NP. Fitted fill and representatives · Baseline, operated value and transport allowance

For the global stage estimate let $\mathcal{R}_p$ be the orbit representatives in the full set of embedding labels. Write $g_{0,j,c} = (g_0(p, j, c))_+$ for the baseline assigned to a representative c with Boolean coordinate zero, and use zero baseline at coordinate one. In the signal class each orbit crosses between the two Boolean coordinates, and $\mathcal{R}_p$ uses representatives of coordinate zero. Write $h_p \in H$ for the Frobenius translation. The h-th step adds hh_p to the translate and flips the Boolean coordinate when h is odd. Accordingly the signal orbit

length is $2\,\mathrm{ord}(2h_p)$, whereas it is $\mathrm{ord}(h_p)$ on either side of the even-order inert-class branch. For each $r \in \mathcal{R}_p$ the numerical estimate bounds the orbit coefficients below by · Partner and orbit length $\beta_{j,r} = g_{0,j,r} + F_{\mathrm{orb}}(j,p,r) - \mathrm{errF}(j,p,r)$. Here $F_{\mathrm{orb}}(j,p,r) = \mathcal{F}_p((k-1)w)$ on that orbit. The source's errF absorbs the weight, partner-capacity and logarithmic expenses in (8b), together with the loss in comparing the fitted baseline to the first-two-row gap. On the even-order `inertClass` branch the same expression is used on the second Boolean coordinate while the first uses its own word-free estimate of $g_{0,j,r}$ alone. The passage from (8b) to the claimed $\beta_{j,r}$ requires those comparisons for the actual weights. Orbit errors and branchwise β · Comparisons from the jet corollary · First-coordinate word-free estimate

Proved here (orbit addition once the single-orbit bounds hold). The all-place logarithm is the sum of the orbit logarithms in disjoint sets of variables after the joint substitution. Expand the integral cofactor there into monomials. Only its terms of degree at most the largest degree differentiated can contribute. Indeed L and its inverse have zero constant part, so substituting either preserves the filtration by total degree. For a fixed cofactor monomial, split its exponents between the orbits. In each orbit this gives a single integral monomial cofactor in the substituted coordinates, covered by the same orbit estimate. On inverting the substitution the product still has independent variable blocks. The signed polynomial extractor factors on that product because its determinant symbols and the factorials both factor by orbit. The coefficient bounds therefore add. Taking the finite sum against the integral cofactor coefficients keeps the same lower bound by the ultrametric inequality. The resulting coefficient estimate is $\sum_{r\in\mathcal{R}_p} \beta_{j,r}$ for every class pair. The sum of F_{orb} over representatives equals $F_j(w)$, since the original sum runs over every position and divides by the orbit length. At a signal prime the errors sum to $E_{\mathrm{atI}}(j,p)$, since every orbit is in the branch with errF and the other terms of E_{atI} vanish. On an even-order inert-class prime the first-coordinate $\mathcal{F}_p$ term is absorbed by the positive part of its sum in E_{atI}. Thus the release's common lower bound for one lift before stabilizers is $\sum_c g_{0,j,c} + F_j(w) - E_{\mathrm{atI}}(j,p)$. Monomial and finite truncation argument · Average defining F_j · Error and orbit summation

To pass from this Taylor estimate to the rational stage lattice requires the integral Weyl expansion and the rationality of the class-pair map, with the unit factors in the chosen frames. With these release inputs, an entry bound B in p-normalized units gives depth B for the unweighted class-pair map: the coefficient-field valuation is at least $e_w B$, and a scalar of valuation $\lceil e_w B \rceil$ can be pulled out in the local lattice. Multiplying the i-th class sum by its stabilizer multiple clears the remaining weights there. In the reference and used twists the initial stabilizers are trivial in the selected construction. Thus for p of class `signal` the depth at lift j is at least

$$b_j = \sum_c g_{0,j,c} + F_j(w) - E_{u,j}(w) + \mathbf{1}_{j=1} v_p(G). \tag{8c}$$

The indicator restores the charge already included in $E_{u,1} = v_p(G) + E_{\mathrm{atI}}(1,p)$. For composition, the O_w-linear inclusions extend to the localized lattices, and applying the next K-linear map to a scalar times an integral vector lets one keep the scalar and apply its bound to the integral vector. Induction on the lift number gives depth $\sum_j b_j$. The reference-to-twist identification at stage zero preserves its localized lattice. This proves the signal-prime specialization of (8) from the stated entry and lattice inputs because with $VV(p) = \varnothing$ the only assigned gains are the baselines. Rational descent and common class denominator · Trivial initial stabilizer and base comparison · One-lift and successive bounds · Twisted composition and sum · Application without operators

Successive lifts on the served inert-class lattice

Fix $VV(p) \neq \varnothing$. The selected local model has the source's inert-class index condition, with nontrivial Frobenius displacement of even order. Work in the local model with valuation ring O_{loc}, uniformizer ϖ and $e_{\mathrm{loc}} = \mathrm{ord}_{\varpi}(p)$. Let M_j be the model's integral lattice at stage j, K_{j-1}^{loc} its lift to stage j, and $C_j = K_{j-1}^{\mathrm{loc}} \cdots K_0^{\mathrm{loc}}$, with C_0 the identity. The local notion of depth r is membership in $\varpi^{\lceil e_{\mathrm{loc}} r \rceil} M_j$. The source uses e_{loc} in this exponent and e_w in the coefficient-field exponent. The comparison maps identify the respective p-normalized depths. An estimate stated in the extension coefficient ring of the Taylor model can also be expressed by the condition $x^b p^{-a} \in O$ for all integers $a, b > 0$ with $a/b \leq r$, where O is the extension valuation ring. Taking rational a/b arbitrarily close to r gives exactly $v(x) \geq r$. This is the normalization of the source's `LowAt` coefficient estimate. Model depth and rational-place normalization · Taylor model coefficient test · Comparison to the global column · Local model index · Even order at a selected site

Release input 4, served operators. The chosen positive-order formal group model must supply all of the following for the same word, place and twist. Let $\mathcal{V}$ be its set of local orbits and identify the representatives used for $VV(p)$ with their images in $\mathcal{V}$. The start operators Y_c preserve M_0 and commute for distinct c. Let $Z_{j,c}$ be the divided Hecke polynomial on stage j and let $P_{j,c}$ be its commutation exponent (the integer `PrZ` from the orbit weights). The commutation assertion on the rational domain is $Z_{j,c} C_j = p^{P_{j,c}} C_j Y_c$. The polynomial uses the normalized Hecke operators of indices $r, r-1, \ldots, 0$, where r is the local operator depth, and a signed coefficient involving the complete homogeneous polynomials of the accumulated local theta parameters. Its coefficients are integral. Its zeroth-operator term has the additional valuation $\sum_s A_c(s)$ over even top positions s, where $A_c(s)$ is the initial top-row gap of the local system. The coefficient bound for that term and the Taylor bound for the other terms are imposed on the actual polynomial, with its prescribed weights and units. The model uses $r = 1$ at $p \geq p_0(k)$ and the selected small-prime depth r_G of the host at $p < p_0(k)$. Commutation and integral coefficients · Hecke polynomial and even top gaps · Commutation exponent · Model depth at each prime

More precisely, for an operator subset $B \subseteq \mathcal{V}$ the stage estimate says that $Z_{j,B} K_{j-1}^{\mathrm{loc}}$ carries M_{j-1} to depth

$$\sum_{c \in B} U_{j,c} + \sum_{c \notin B} L_{j,c} - d_{j-1} \tag{8d}$$

in M_j. Here $Z_{j,B}$ is the product in the order of the local expansion, $U_{j,c}$ is an available operated coefficient bound, $L_{j,c}$ is an unoperated coefficient bound, and d_{j-1} is the stage stabilizer loss of the local model. The source checks (8d) by expanding the product of Hecke polynomials in the span of terms with an index i_c in $0, \ldots, r$ at each operated orbit. Where $i_c = 0$ on an operated orbit, its Taylor estimate loses the even top-gap sum just described, while the scalar coefficient contributes at least the same amount. Where $i_c > 0$, the estimate is the operated Taylor bound. At orbits outside B it is the unoperated bound. The source's `form_of_place` combines these per-orbit estimates with the volume ratio and rational compatibility of the Weyl lattices. Once those estimates and the spanning expression are supplied, adding the scalar valuations cancels the even-gap loss on every summand, and the ultrametric inequality gives (8d). The formal group construction of the spanning expression and the two Taylor estimates for the actual $Z_{j,c}$ and K_{j-1}^{loc} remain release inputs. The expansion uses the exponential of the Dwork cocycle in compatible local coordinates, with an integral cofactor. In the integral model its further rescaling uses a root $\varpi_0^{p^2-1} = p$, so ϖ_0^{p+1} has p-normalized valuation $1/(p-1)$. This accounts for the integrality of $\varpi_0^{q(p+1)}/q!$ in that exponential. This

root serves in the Taylor substitution, while ϖ serves in the lattice depth. Local Taylor expansion and factorial scale · Expansion of the divided polynomial and cancellation · Proof of the simultaneous stage bound · Formal group inputs used in the proof · Adding rational lower bounds before rounding

Proved here (formal Dwork criterion). The integral expansion used for (8d) invokes the classical Dieudonné–Dwork lemma. In this special case let O^0 be an unramified p-adic coefficient ring, σ_p its Frobenius lift, and X the tuple of formal variables. For a series f of constant term one over its fraction field, the criterion is $f \in O^0[[X]]$ exactly when $\sigma_p(f)(X^p)/f(X)^p$ lies in $1 + p(X)O^0[[X]]$. Here σ_p acts on coefficients, X^p denotes componentwise p-th powers, and (X) is the ideal of series with zero constant. The forward direction follows by reduction modulo p. For the reverse direction, assume by induction that all terms of degrees below n are integral. Divide f by its truncation through degree $n-1$, which is an integral unit. The resulting series has the same ratio condition. It now starts with one plus a homogeneous term of degree n, say H_n. The degree-n term of its ratio is $-pH_n$, because substitution by X^p raises degree. Hence H_n is integral. Induction gives the criterion. If $f = \exp(P)$ with P of zero constant, its ratio is the exponential of $\sigma_p(P)(X^p) - pP(X)$. Formal exponential and logarithm identify $p(X)O^0[[X]]$ with $1 + p(X)O^0[[X]]$, since $p^j/j!$ and p^j/j belong to pO^0 for $j \geq 1$. Applying the criterion to the CM cocycle in the model still requires its Frobenius and additive-defect identities. The release's formal Dwork condition · Application of the multivariable lemma

Proved here (induction through the fitted generators). For a subset $S \subseteq VV(p)$ write $Y(S) = \prod_{c \in S} Y_c$ for the product with each operator taken once. The operator Y_S used in the definition of NP above is $Y(S)^m$. Take $L_{j,c}$ to be the nonnegative baseline $g_{0,j,c}$ on the first-coordinate representatives. Take $U_{j,c} = g_U(p,j,c)$ for the served representatives in S. On the second-coordinate orbits take for $L_{j,c}$ the unoperated Taylor bounds supplied by the model. In the source the second-coordinate bounds are $F_{\mathrm{orb}} - \mathrm{errF}$ at the chosen cutoff. Let $\mathrm{Pr}_{j,c} = \mathrm{Pr}(p,j,c)$, the same allowance used in $VV(p)$, and assume $P_{j,c} \leq \mathrm{Pr}_{j,c}$. These identifications use the numerical comparisons of the model with the global weights, including the sum of the second-coordinate bounds less the stage losses. The comparisons hold in the range $1 \leq \log_p(2a)$ in use here. Comparison of operated, unoperated and commutation terms · Comparison and sum in the local model · Second-coordinate estimates on the selected system

For $j = 1, \ldots, m$, select $B_j = \{c \in S : U_{j,c} - \mathrm{Pr}_{j,c} \geq L_{j,c}\}$. For every $x \in M_0$, the commutation identities give

$$C_j(Y(S)^j x) = p^{-\sum_{c \in B_j} P_{j,c}} Z_{j,B_j} K^{\mathrm{loc}}_{j-1} C_{j-1}\left(Y(S)^{j-1} \prod_{c \in S \setminus B_j} Y_c\, x \right). \tag{8e}$$

All products of starting operators can be ordered as needed since they commute, and the unused operators send x back into M_0. Start the induction with depth zero. Apply the induction hypothesis to that new integral input, use linearity and (8d) on the integral representative of the stage-$j-1$ column, then subtract the commutation exponent in (8e). The resulting increment is at least $\sum_{c \in S} \max(U_{j,c} - \mathrm{Pr}_{j,c}, L_{j,c}) + \sum_{c \notin S} L_{j,c} - d_{j-1}$. Depths add since the sum of the two ceilings is at least the ceiling of the sum. This proves the cumulative bound uniformly on $Y(S)^m M_0$. Induction and commutation through successive lifts

In the actual fitted generator of NP, the omitted served orbit $c \in VV(p) \setminus S$ contributes a scalar p^{S_c}, where S_c is the rounded positive residual in Section 2. Its p-normalized depth contribution is S_c, which is

at least the sum over j of $\max(U_{j,c} - \mathrm{Pr}_{j,c}, L_{j,c}) - L_{j,c}$ on that orbit. It therefore pays for all the increments that this generator foregoes. The resulting bound holds on the full span by linearity, and hence also after intersecting with $p^{\tau(w)} M_0$ and the ramified subspace in NP. The additional intersection is used in the lattice price, so τ need not be added to the exponent from this span argument. The source's remaining site comparison gives

$$\sum_{j} \sum_{\text{second-coordinate orbits } c} L_{j,c} - \sum_{j=1}^{m} d_{j-1} \;\geq\; \sum_{j} (F_j(w) - E_{u,j}(w)) + v_p(G).$$

To obtain the inequality, sum the second-coordinate Taylor bounds $F_{\mathrm{orb}} - \mathrm{errF}$. The orbit sum calculation gives at least $\sum_j (F_j - E_{\mathrm{atI}}(j,p))$. The local model supplies $d_0 = 0$ and $d_i \leq v_p(G)$ for $1 \leq i < m$, so the subtracted sum is at most $(m-1)v_p(G)$. Substituting $E_{u,j} = E_{\mathrm{atI}}(j,p) + v_p(G)$ gives the displayed right-hand side. With it, the lower bound on the entire column is $\sum_j \mathrm{gainP}(p,j) + \sum_j (F_j - E_{u,j}) + v_p(G)$ in p-normalized units. The comparison from the local model to the actual normalized chain (1) transports this bound to the terminal coefficient-field lattice. The inequality $\mathrm{gainS} \leq \mathrm{gainP}$ then gives (8). Residual bound and span · Fitted NP and depth on all generators · Quantitative site comparison · Transport and comparison of gains · Second-coordinate comparison and stabilizer sum

Valuation of the signed period

In either branch, take $\mu = \mu(w)$, the actual input $\phi \in N_u$ and an admissible integral row F. The column conclusion writes the terminal vector as $a_0 z_0$, with z_0 in the localized terminal lattice and $v_w(a_0) \geq e_w\big(s_u(w) + v_p(G)\big)$, using $g_{D,j} \leq \mathrm{gainS}(p,j)$. Inputs 1 and 2 and (3) identify its period with $a_0 \sum_c y_c(z_0, F)/s_c$ in K. If the period is nonzero, $a_0 \neq 0$ and the sum is nonzero, and its valuation is exactly

$$v_w\big(a_0 \sum_c y_c(z_0,F)/s_c\big) = v_w(a_0) - e_w v_p(S_{u,\mu}) + v_w\left(\sum_c (S_{u,\mu}/s_c) y_c(z_0,F)\right) \;\geq\; e_w s_u(w) \tag{8f}$$

The final summand on the right is nonnegative. This is the terminal cancellation of the restored $v_p(G)$ in (8) when using the common bound $v_p(S_{u,\mu}) \leq v_p(G)$. In the integral lattice the shift uses $\lfloor e_w s_u(w) \rfloor$ as in (4). A nonzero value for a chosen pair is a separate condition. For the nonzero maps needed in the height and slope argument, the rank and Gram estimates of Paper V and the common-space comparison in input 5 govern the choice of tests. Base and twisted transport to the terminal period · Cancellation in the actual division argument · Word-map interface

The remaining local work for a general word is to establish the geometric orbit expansions and coefficient-lattice transports uniformly for every admitted class pair at every used tuple, with the prescribed diagrams, unit factors, and initial stabilizer convention. The served case additionally needs the local Hecke commutation identity, the placewise Taylor and lattice estimates for (8d), the fitted site sum, and the comparison to (1). These are the hypotheses left open by the analytic and algebraic deductions here. Each comparison is required for the realization selected in Paper VI. True-column hypotheses · Admitted system and required local coefficient estimates · Quantified family interface

3. Fitting one common row lattice

Release input 5 (row quotients and rational comparison). There is a common K-space W and, for each used tuple μ, a row lattice M_μ in a rational space of row tests. Every element of M_μ, after the prescribed row-coordinate map, is in the complex admissible subspace $V_0(\mu)$ and is an $\mathcal{O}$-linear combination of generators of the integral row module $\widetilde{M}_\mu$. The integral cutting assertion is that such an element in $V_0(\mu)$ is an $\mathcal{O}$-linear combination of tests in $\widetilde{M}_\mu \cap V_0(\mu)$. This is an assertion about integral combinations in the actual row module. There is a surjective map $R_\mu = \iota_\mu P_{q,\mu}$ from the rational row space onto W, where $P_{q,\mu}$ is the quotient map and ι_μ identifies the quotient with W. Let $A_\mu = R_\mu(M_\mu)$, a full row lattice. The chosen tuple $\mu(w)$ is μ_0 away from a finite set. The metrics in the column and common row coordinates are tied to the metrics of the lattices used by Paper III. Integral row cutting and valuation transfer · Row quotient lattice · Finitely many twisted places · Coordinate and metric ties

In these coordinates the selected realization gives a K-linear $B_u : V_u \to W^\vee$, for every u, with the factorization

$$P_{u,\mu}(\phi, F) = \zeta_{u,\mu}\, B_u(\phi)(R_\mu F) \tag{9}$$

for all rational inputs, where $\sigma(P_{u,\mu})$ is the period (2) in the matched coordinates. The comparison factor is nonzero, equals one at μ_0, and satisfies $v_w(\zeta_{u,\mu(w)}) = 0$ at each finite place. At the chosen complex embedding the same B_u identifies, up to the stated isometries and one nonzero scalar per word, with the Bessel pairing of Paper III. This last identity is needed to use its Gram estimate on B_u. The proof of the identification for the realized automorphic datum is part of input 5. Period and Bessel factorization · Bessel identity at the family interface

Define the fitted row lattice by

$$RR = \{z \in W :\ z \in (A_{\mu(w)})_w \text{ for every finite } w\}, \qquad RR^\vee = \{\beta \in W^\vee :\ \beta(RR) \subseteq \mathcal{O}\} \tag{10}$$

This is the dual for the K-bilinear pairing. The metrics used to measure its slopes are Hermitian. The prime-by-prime definition is with respect to the coefficient field, so different primes above p can impose different conditions if the given tuple selection does so. Fitted row lattice · Local row identification

Proved here (gluing). If full O_w-lattices $A(w)$ agree with the localization of a fixed $\mathcal{O}$-lattice A away from finitely many w, then $A' = \bigcap_w A(w)$ is full and $(A')_w = A(w)$. Choose a nonzero $d \in \mathcal{O}$ such that $dA_w \subseteq A(w) \subseteq d^{-1}A_w$ at the finitely many exceptional places. At all other places these bounds hold as well. Since $A = \bigcap_w A_w$, it follows that $dA \subseteq A' \subseteq d^{-1}A$, which proves fullness and finite generation. For $z \in A(w)$, there are only finitely many other places v with $z \notin A(v)$. Choose exponents n_v such that $\mathfrak{p}_v^{n_v} z \subseteq A(v)$ and use the Chinese remainder theorem to find $q \in \mathcal{O}$ lying in those powers and congruent to one at $\mathfrak{p}_w$. Then $qz \in A'$, and hence $z \in (A')_w$. The reverse inclusion follows from $A' \subseteq A(w)$. The equalities $A = \bigcap_w A_w$ and $\mathcal{O} = \bigcap_w O_w$ are the usual local membership criteria over a Dedekind domain. For the first, if $z \notin A$, the proper ideal $\{a \in \mathcal{O} : az \in A\}$ is contained in a maximal ideal, and localization there still excludes z. This proof also constructs the global N_u from the local lattices of input 3 once their fullness and agreement off a finite set are supplied. Local gluing in the fitted row construction · Fitted column and level conditions

The integral period-divisibility theorem

Proved here. Under inputs 1–5 and the selected host data, for every word u define the fractional ideal and the divided lattice

$$I_u = \prod_{w \in \mathcal{S}} \mathfrak{p}_w^{-a_u(w)}, \qquad N_u' = I_u N_u.$$

Then the actual rational map used in the Bessel comparison satisfies

$$B_u(N_u') \ \subseteq \ RR^\vee. \tag{11}$$

The signs in I_u allow both division and multiplication. A bound $s_u(w) < 0$ can require the word lattice to become smaller at that place. Fractional ideal · Signed integer input · The release's dual lattice conclusion

To prove (11), fix w and put $\mu = \mu(w)$. The local period estimate from Section 2 holds for $\phi \in N_u$ and the integral tests in $\widetilde{M}_\mu \cap V_0(\mu)$. By the integral cutting assertion in input 5, it extends by bilinearity to all row tests in M_μ in the matched coordinates. It also extends to $(N_u)_w$ and $(M_\mu)_w$ since their extra scalar denominators have valuation zero. By (9) and the unit property, $B_u(\phi)(z)$ lies in $\mathfrak{p}_w^{a_u(w)} O_w$ for all $\phi \in (N_u)_w$ and $z \in (A_\mu)_w = (RR)_w$. Now $(I_u)_w = \mathfrak{p}_w^{-a_u(w)} O_w$ and B_u is linear. It follows that each value of $B_u((N_u')_w)$ on $(RR)_w$ lies in O_w. For global arguments in N_u' and RR, the same element of K lies in O_w for every w, and therefore lies in $\mathcal{O}$. The coordinate ties identify this inclusion with the release's $B_k(u)(I_u \mathrm{NuLat}_u) \subseteq RR^\vee$. Passing from the tested rows to the fitted rows · Fitted row factorization and dual · Word-level identification

The choice in (4) is the signed floor prescribed by the release. Since v_w is integral, the unrounded local bound also implies the integer ceiling bound on a nonzero entry, but (11) and its height comparison use the floor. With all metrics fixed, the normalized mean slope increases from N_u to N_u' by $2 \sum_w \ell_w a_u(w)$ in positive rank. Indeed on the rank-r determinant line the local factor is $\mathfrak{p}_w^{-r a_u(w)}$ and the squared-norm mean degree is divided by r. The difference from $2 \sum_w \ell_w e_w s_u(w)$ has absolute value at most $2 \sum_{w : s_u(w) \neq 0} \ell_w$. The bound for this last sum by CX'' is a numerical support estimate supplied by the selected data. Slope and shift formula · Global support and rounding bound

At an off-level place, (6) implies $(N_u')_w \subseteq \mathfrak{p}_w^{e_w \tau(w) - a_u(w)} (M_u^{\mathrm{ram}})_w$. For the *mean* slope the index of the inclusion $N_u \subseteq M_u^{\mathrm{ram}}$ contributes locally $2\ell_w \ \mathrm{length}_{O_w}((M_u^{\mathrm{ram}})_w / (N_u)_w) / \operatorname{rk} N_u$ when comparing equal-rank lattices. The actual quotient estimate uses the bound in (6) and the available rank comparison. At T_u it uses the level length in input 3 and $\tau(w) = 0$. Define

$$\begin{aligned} \mathrm{Ga} &= 2 \sum_w \ell_w e_w \left(\sum_j g_{D,j}(w) - \tau(w) \right), \\ \mathrm{Fo} &= 2 \sum_w \ell_w e_w \sum_j F_j(w), \\ \mathrm{Ee}(u) &= 2 \sum_w \ell_w e_w \sum_j E_{u,j}(w). \end{aligned}$$

The unrounded shift less the assigned common lattice price is exactly $\mathrm{Ga} + \mathrm{Fo} - \mathrm{Ee}(u)$, by (4). The prime weight on a term common to all places over p is $2 \log p$. In particular Ga is an assigned and priced gain.

The signed ideal of an individual word uses all the terms of (4). Gain and correction sums · Mean-slope comparison · Prime regrouping and common gain

4. The word height and the quantitative handoff

Write av_v for the uniform average over the infinite places of K. Because K is totally complex this gives the mean used for these embeddings. For a nonzero linear map between the specified Hermitian spaces let

$$\mathrm{ht}(B_u) = \mathrm{av}_v \log \|B_u\|_v^2.$$

The word height assigned to it is $h(u) = h_A(u) + \sum_j (h_{B,j}(u) - h_{X,j}) - \mathrm{av}_v \log \kappa_v$, where $\kappa_v > 0$ is the common row scale. Put $h_0(y) = y \log y - y$ on positive y, using its continuous value at zero. The release's three terms and their explicit size allowances are as follows.

- h_A bounds the terminal branching and evaluation term. Its leading term is $-\sum_{t \in \mathcal{T}} a_t \log T_u$. The allowance is $C_{\mathrm{sz}}\big(|\mathcal{T}| b' N_m + f(m\eta + S) N_m + f N_m^2\big)\mathrm{La}$, with η, S the source variation bounds, b' the thin weight scale, and $\mathrm{La} = \log(a \ \mathrm{rd}\, E)$.
- $h_{B,j}$ bounds the mean squared operator norm logarithm of lift j, or its combined contribution when several lifts are grouped into one matrix. Its leading term at an embedding ρ is $\sum_i h_0(\nu_{j,\rho,i}) - \log \Omega \sum_i \nu_{j,\rho,i}$, where $\nu_{j,\rho,i}$ is the source Harish-Chandra parameter plus the lift weight minus the displayed rank shift. Sum over ρ and add $C_{\mathrm{sz}} f N_m^2 (\mathrm{La} + m^3 \log b)$.
- $h_{X,j}$ is the lower bound for the row eigenvalue term, the sum of the logarithmic hook products at the embeddings with $C_{\mathrm{sz}} f n_V^2 \log(b \ \mathrm{rd}\, E \, g_0 g_1 n_V)$ subtracted. Here n_V is the source row rank parameter. The terms are constructed from the actual weights of the selected chain. Weight and branching formulas · Hook and height formulas · Numerical record · Height normalization

Release input 6 (reference metrics and norms). In the reference tuple, the terminal map of (2) in rational coordinates has squared norm at most $\exp(H_P)$ at every v, with $H_P \le h_A(u)$. The matrices of the rational lift chain, with the chosen stage metrics, have squared norms bounded by the exponentials of the appropriate sums of $H_j(v)$, with $\mathrm{av}_v\, H_j(v) \le h_{B,j}(u)$. The reference quotient metric makes the decomposition into parts b orthogonal both before and after ι_{μ_0}. On part b, ι_{μ_0} scales the squared norm by $\kappa_v \prod_j \lambda_j(b,v)$, with $\lambda_j(b,v) > 0$. The eigenvalue estimate gives $\mathrm{av}_v \min_b \sum_j \log \lambda_j(b,v) \ge \sum_j h_{X,j}$. All the operator and weight estimates are on the same word, stages, and reference row as in (9). Rational matrix chain · Norm, quotient and eigenvalue inputs · Numerical height and branching caps

Proved here (norm propagation). At μ_0 the scalar in (9) is one. A unit vector in $W_{\mathbb{C}}$ lifts back to the reference row quotient with squared norm at most $\big(\kappa_v \min_b \prod_j \lambda_j(b,v)\big)^{-1}$. To see this, decompose it into the orthogonal parts and add their preimage squared norms with their respective inverse scalars. Lift this quotient vector by the orthogonal section of P_{q,μ_0}, which preserves its quotient norm. Applying (9) on this lift and multiplying the squared norm bounds through (2) gives

$$\log \|B_u\|_v^2 \ \le \ H_P + \sum_j H_j(v) - \log \kappa_v - \min_b \sum_j \log \lambda_j(b,v).$$

Averaging and applying input 6 proves $\mathrm{ht}(B_u) \le h(u)$. This proves the claim for each nonzero map. The integral inclusion (11) already covers the zero map. Complex norm comparison and mean

The finite covolume bound still needed for the height

Let $\mu_C = \mu(RR^\vee)$ in the squared-norm convention. The finite part of the argument introduces a reference-row slope μ_0^{ref} after subtracting the mean log row scale. The release uses scalars $F_0(w,j)$ in a base-row presentation, with the factor-of-two tie $\mathrm{Fo} = 2\sum_{j,w\in \mathrm{Tw}} \ell_w e_w F_0(w,j)$, where Tw is the finite set of twisted places in that presentation. The covolume calculation compares the same projected good-part lattice at the reference type and the type selected at each finite place. The sign can be seen before any majorization. Base-row finite bound and passage · Signed estimate and row passage · Tie to the global Fo

Proved here (signed determinant identity). At paired places $w, \bar w$ write M for the fixed full lift lattice, A for the intermediate full lattice, θ for the invertible transport and H for the intermediate Hermitian Gram form. If $d_v(L)$ denotes the valuation of the determinant ideal of L in fixed coordinates, the relative index is $\mathrm{idx}_v(M,\theta A) = v(\det\theta) + d_v(A) - d_v(M)$. Write $j_w(H,A)$ for the negative valuation of the Gram determinant on a semilocal basis of A. Expanding that determinant gives $j_w(H,A) = -v_w(\det H) - d_w(A) - d_{\bar w}(A)$. Thus

$$\mathrm{idx}_w(M,\theta A) + \mathrm{idx}_{\bar w}(M,\theta A) + j_w(H,A) = j_w\big((\theta^{-1})^* H \theta^{-1}, M\big).$$

At a fixed place of conjugation both index terms still occur. This identity explains the signed lift contributions to the reference-row estimate. Its use in an upper bound requires control of the remaining transported forms and the reference row. Signed covolume terms · Comparison at selected and reference types

Release input 7 (quantitative covolume and cancellation of main terms). In the selected family, assume the level-three projected covolume bound, the integral spectral and quotient comparisons with the conjugate local indices, and the twisted-place length bound. In the source normalization the precise consequences used are the following. Write L_{pass} for the passage bound and E_b for the reference-row error. The input uses the same row scale and quotient as input 6. On that datum the geometric comparisons give the first two inequalities below. The passage and arithmetic budgets in the last line are estimated in the signed main-term calculation after (17).

$$\begin{aligned} \mu_C - \mathrm{av}_v \log\kappa_v &\le \mu_0^{\mathrm{ref}} + L_{\mathrm{pass}}, \\ \mu_0^{\mathrm{ref}} &\le \sum_{j,\rho} \mathrm{pq}_j\, w_{j,\rho} \log\Omega + \mathrm{Fo} + E_b, \\ L_{\mathrm{pass}} \le 1 + C_0 m\lambda_0, &\qquad E_b \le c_e m x^{\varepsilon}(B_0 + N_c). \end{aligned}$$

Here pq_j is the row dimension parameter in the numerical datum. The weights $w_{j,\rho}$ are the tied reference weights, λ_0 is the source group-log scale, and B_0, N_c are the two size budgets fixed with the host. The reference-row problem includes bounding the projected quotient Gram term at the reference type and at each chosen twist. The passage to the final row uses the good/bad rank bound, the actual quotient metrics, and the smallness condition of the source passage theorem. The numerical step, also a release input here, is that adding $h_A + \sum_j (h_{B,j} - h_{X,j})$ and cancelling the finite term gives the upper bound

$$\mathrm{MAIN}(u) + (50C_0 + 10)|\mathcal{T}|a\,\mathrm{La} + 60 c_e m x^{\varepsilon}(B_0 + N_c), \qquad \mathrm{MAIN}(u) = \sum_{t\in\mathcal{T}} \big(m(h_0(a_t) - a_t\log\Omega) - a_t\log T_u\big).$$

In detail the numerical input bounds $\mu_C + h(u) - \mathrm{Fo}$. With the first two inequalities, its remaining estimate concerns the sum $\sum_{j,\rho} \mathrm{pq}(j) w_{j,\rho}\log\Omega + h_A + \sum_j (h_{B,j} - h_{X,j}) + E_b + L_{\mathrm{pass}}$ and uses the specified weight and size comparisons. The constants C_0, c_e are uniform across the words and selected

indices at fixed family parameters. Exact hypotheses for the signed covolume inequality · Numerical inequality and sizes · Numerical main term and error · Uniform constants in the selected construction

For fixed admissible family and $\varepsilon > 0$, the quantitative handoff to Paper III takes the following form. Assume inputs 1–7 with their fixed family constants for all $x \geq x_0$ and every selected index i and word u. Assume also the family frame bound $s_{\mathrm{Frame}} + \mu_{\max}(RR) \leq x^{CH} b$ for the tied metrics and height. Then (11) holds for each word and the same B_u and RR satisfy

$$\mathrm{ht}(B_u) \leq h(u) \quad (B_u \neq 0), \qquad \mu(RR^{\vee}) + h(u) - \mathrm{Fo} \leq \mathrm{MAIN}(u) + C|\mathcal{T}|a\,\mathrm{La} + Cm(B_0 + N_c)x^{2\varepsilon} \tag{12}$$

The uniform C is obtained by taking a maximum of the numerical constants and using $x^{\varepsilon} \leq x^{2\varepsilon}$ for $x \geq 1$. Equation (12) uses the stored choice $C_{\mathrm{abs}} = 2$. The comparison at $C_{\mathrm{abs}} = 1$ is given after (17). The finite gain (7), the mean-slope correction following (11), and the comparison (12) are used together with Paper III's frame estimate on the Bessel maps identified by (9). This deduction of (12) assumes the uniform quantitative inputs just stated. Selected division exponent and handoff · Word-height estimate and assembly · Use in the global slope inequality

The four divisor bounds on the selected family

The global divisor estimate uses four assertions on the same data. These concern the absolute frame bound, the signed height inequality, the finite error of each lift, and the variation in height between two words. In the release they are clauses 2, 7, 8 and 9 of $\mathrm{LinkDIV_at}$. The word-height calculation above gives clause 7 from the projected covolume and metric comparisons, in the stored normalization $C_{\mathrm{abs}} = 2$. The other three have the following derivations. Definitions of the four bounds · Assembly on one member

Variation between words. Put $L = x^{\varepsilon}$, $b = \exp L$, $B_0 = fmbL$, and $A = \sum_{t \in TT} a_t$. The explicit height at word u is $-A \log T_u$ plus a term common to all words and a sum of $\phi(\nu) - \omega\nu$ over its lift coordinates, where $\phi(y) = y \log y - y$ and $\omega = \log \Omega$. The common term includes the negative row hook allowances and the mean logarithm of the metric scalar. Subtract $h(u') + A \log T_{u'}$ from $h(u) + A \log T_u$. The first and last entries of the original weight and all appended entries agree by the weight identifications. Only the $k - 2$ interior coordinates per lift and complex place can change. Each changes by at most the common width S, and the positive integer arguments are at most $M = 6x^6 b$ under the selected size bounds. The derivative on $[1, M]$ has absolute value at most $\log M + |\omega|$. Since $kS \leq b$, $|\omega| \leq 6 + 2L$, and $\log x \leq L/\varepsilon$, we obtain

$$\left|(h(u) + A \log T_u) - (h(u') + A \log T_{u'})\right| \leq (14 + 6/\varepsilon) B_0.$$

This is uniform over all pairs of words at the chosen member. Exact height difference · Coordinate bounds and onset

Finite error of a lift. At a type prime $p = g_i$ let $e_w = v_w(p)$ and $\ell_w = \log Nw / [K : \mathbb{Q}]$ on the full fibre in the coefficient field. The weight of a p-normalized exponent is $\mathrm{sz}_w(y) = 2e_w \ell_w y$. Each selected generator y_τ satisfies $y_\tau \bar{y}_\tau = p$. Conjugation permutes the whole fibre, so $2 \sum_{w|p} \ell_w v_w(y_\tau) = \log p$. This accounts for the factor two also when a place is fixed by conjugation. Write $P_{u,j}$ for the sum of the lift exponents over all τ and Q_j for the sum of the absolute signed Fock degrees. The type-prime error has terms $v_p(G_u) - Q_j/2 + R_g$ and $e_w^{-1} \sum_\tau (\sum_r \nu_{u,j,\tau,r}) v_w(y_\tau)$, with G_u the stabilizer denominator and R_g the local error allowance. Summing gives exactly

$$\sum_{w\mid p}\mathrm{sz}_w(E_{u,j}(w))-A\log p=(P_{u,j}-Q_j-A+2v_p(G_u)+2R_g)\log p.$$

The degree count on the two signed lists bounds $|P_{u,j}-Q_j-A|\le 166fb$ under the size hypotheses. Also $v_p(G_u)\log p\le \log G_u$ and the type-prime budget bounds $R_g\log p$. The two type fibres therefore cost at most $2C_gB_0$. Off these fibres the level, stabilizer, discriminant and remaining local contributions cost at most $(C_T+C_G+C_d+C_I)B_0$ in absolute sum, given their respective size estimates. This proves clause 8 with $C=2C_g+C_T+C_G+C_d+C_I$ before summing over j. The total allowance in the column slope comparison is CmB_0. The degree count · Type-prime sizes

Absolute frame bound. For the shifted column lattice Nu_u put $d_u=\operatorname{rank}\mathrm{Nu}_u$. Its frame slope is the weighted mean $\sum_u d_u(\mu(\mathrm{Nu}_u)-h(u))/\sum_u d_u$. Fullness, positive ranks and the signed shift comparison bound each column slope by the maximal slope on its unshifted lattice plus $\mathrm{Ga}+\mathrm{Fo}-\mathrm{Ee}(u)+C_smB_0$. Write H_κ for the mean log scale. Subtracting $h(u)$ and adding $\mu_{\max}(RR)$ gives five terms. They are the unshifted maximal slope, the signed correction, $-(h(u)+H_\kappa)$, $\mu_{\max}(RR)+H_\kappa$, and C_smB_0. The two occurrences of H_κ cancel. If the respective upper bounds are $x^{E_i}b$, their sum is at most $x^{CH}b$ for $x\ge 5$ and $CH\ge 1+\max_i E_i$. Taking the positive-rank mean proves clause 2. The unshifted column metric and the bound on the maximal slope of the fitted row are geometric inputs. Clause 7 instead uses the mean slope of the dual fitted row and its projected covolume. Rank mean and component bounds · Selected-family hypotheses · Constants for the two types and four off-type errors

Projected covolume and the signed height main term

Theorem. Fix the ordinary conductor-one family of Papers III–VI and use its reference tuple, lift weights, level and row metric at each selected index. Write $p_j=k+j-1$, $q_j=m-j$, and $E_3=B_0\log\log x$. Under the geometric and integral comparisons stated below, the level-three lift lattice M_j, in its prescribed frame $X_{W,j}$ of rank $D_0>0$, satisfies

$$\frac{\mathrm{lcv}(X_{W,j},M_j)}{D_0}\le(p_j-q_j)\sum_\rho w_{j,\rho}\log\Omega+(C_{\mathrm{low}}+C_{\mathrm{mass}}+1)E_3. \tag{13}$$

Here C_{low} bounds the error in the core slope and C_{mass} bounds the upper logarithm of the mass times the product of Weyl dimensions. The term 1 accounts for the discarded rank, uniformly on the selected tail. The signs in (13) and in $h_{B,j}$ leave exactly $m\sum_{t\in\mathcal{T}}(h_0(a_t)-a_t\log\Omega)$ as the main term after the Weyl hook comparison. The terminal term then contributes $-\sum_t a_t\log T_u$.

Proof of the projection estimate. In the selected lift let L_j be the full type lattice in the realized row space, with the realized Hermitian norms, and let e_j be the orthogonal rational projector onto its good part. Put $L_{g,j}=L_j\cap\operatorname{im}e_j$, $L_{b,j}=L_j\cap\ker e_j$, and write d_g,d_b for their ranks. Suppose $ge_jL_j\subset L_j$, where $g=|G_{\mathrm{row}}|$ is the finite row group order. The map $z\mapsto(1-e_j)z$ induces an isomorphism

$$L_j/(L_{g,j}\oplus L_{b,j})\ \simeq\ (1-e_j)L_j/L_{b,j}.$$

The right side is killed by g. At w its length is at most $d_bv_w(g)$, since $((1-e_j)L_j)_w$ is free of rank d_b. Thus the logarithm of the absolute index is at most $[K:\mathbb{Q}]d_b\log g$.

For a rank-$d>0$ lattice L over the totally complex coefficient field K, use the product Hermitian metric over one embedding per infinite place and its real Lebesgue covolume. The normalized logarithmic

covolume and squared-norm mean slope are

$$\operatorname{lcv}(L) = \frac{2}{[K:\mathbb{Q}]} \log\left(\frac{2^{[K:\mathbb{Q}]d/2}\operatorname{covol}_{\mathbb{R}}(L)}{|\operatorname{disc} K|^{d/2}} \right), \qquad \mu(L) = -\operatorname{lcv}(L)/d.$$

A complex change of frame multiplies the real covolume by the squared modulus of its determinant. Multiplication of a lattice norm by $\sqrt{R}$ at every infinite place subtracts $\log R$ from μ. At a finite place the contribution of an index of length c_w is $2\ell_w c_w$, with $\ell_w = \log N\mathfrak{p}_w/[K:\mathbb{Q}]$ and $v_w(\mathfrak{p}_w) = 1$. Orthogonality, the index formula and $\mu(L_{b,j}) \le \Lambda$ give

$$\mu(L_{g,j}) \;\ge\; \mu(L_j) - \frac{d_b}{d_g}\big(\Lambda + 2\log g - \mu(L_j)\big) \tag{14}$$

in positive good rank, with the last term taken as zero when $d_b = 0$. Here Λ can be the maximum slope of the ambient row lattice. In particular, if $d_b/(d_b+d_g) \le 3\delta$, $\delta \le 1/12$, and $\Lambda, 2\log g, -\mu(L_j) \le Q$, the loss is at most $12\delta Q$. Slope normalization · Projection and covolume comparison

The actual reference coordinates. The full realized type lattice L_j is the extension of the integral framed lattice intersected with the rational eigenspace of the selected reference level subgroup P and character. The good space is the further intersection with $\operatorname{im} e_j$. The release's rational good-part basis (b_a) spans this space with rank $D_0 = d_g > 0$. In those coordinates $\phi(c) = \sum_a c_a b_a$, and the actual lattice M_j is $\phi^{-1}(L_{g,j})$ by the definitions of the integral type, the good intersection and the coordinate preimage. Saturation follows since $az \in \operatorname{im} e_j$ for nonzero $a \in \mathcal{O}_K$ and $z \in L_j$ implies $z \in \operatorname{im} e_j$. The integral intersection spans the rational good space by clearing denominators in the full integral type lattice. Coordinate transport therefore has zero finite index. Integral type, good intersection and coordinates · Coordinate preimage · Actual reference basis and lattice · Fullness of the fitted coordinate lattice

Put $n_P = |P| \ge 1$ for the subgroup in the reference branch. At each infinite place the columns defining $X_{W,j}$ are the realized b_a scaled by $n_P^{-1/2}$, with the prescribed conjugation at that place. Here $\|\cdot\|_{R,v}$ is the realized norm on L_j before this frame scaling. The inverse-frame norm is the Gram norm of these scaled columns, so

$$\|c\|^2_{X_{W,j},v} = n_P^{-1}\|\phi(c)\|^2_{R,v}, \qquad \frac{\operatorname{lcv}(X_{W,j}, M_j)}{D_0} = -\mu(L_{g,j}) - \log n_P \le -\mu(L_{g,j}).$$

The second equality uses the normalization above for a squared-norm scalar, uniformly in the reference row and lift. Exact frame norm · Reference rank and covolume comparison · Determinant scaling

Mass and Weyl normalization on this lift. Let $L_{\mathrm{core},j}$ be the integral type lattice before transfer from its field of definition, with the probability Petersson norm, including the lift's Lie factor in the release's probability normalization. The exact metric comparison on the selected lift multiplies the *squared* norm by

$$R_j = V_j \prod_{\rho} d_{j,\rho}, \qquad d_{j,\rho} = \prod_{B\in\{+,-\}} \prod_{a<b} \frac{\lambda^B_{j,\rho,a} - \lambda^B_{j,\rho,b} + b - a}{b-a}. \tag{15}$$

Here V_j is the mass of the selected fundamental domain for the actual level in the realized adelic Haar measure, and the two λ^B are the dominant compact highest-weight lists of that lift, of lengths p_j, q_j. The CM-type index occurs once in each product and an empty product is one. Adding a determinant character shifts all entries in its block by the same integer and leaves (15) invariant. With the integral scalar-extension comparison, the normalized determinant slope therefore changes by $-\log R_j$. Indeed the mean is

invariant under field extension before rescaling, by the norm of the extended determinant ideal and by counting the extensions of each embedding. The same determinant argument gives a lower bound if the realized lattice contains the compared scalar extension. In the release the equality of lattices uses the integral span of the *same* type in the two models. No $R_j^{1/2}$ survives in the slope. Weyl dimension · Exact mass normalization · Probability metric comparison · Integral type descent

The mass and Weyl terms enter additively as $\log V_j + \sum_\rho \log d_{j,\rho}$ in the projected upper bound. For example, if the spread of both highest-weight lists is bounded by H_j, then each factor in (15) is between 1 and $1 + H_j$, so

$$0 \le \sum_\rho \log d_{j,\rho} \le \sum_\rho \left(\binom{p_j}{2} + \binom{q_j}{2} \right) \log(1 + H_j).$$

One can use the release's sharper mass estimate with its stated size condition and its exact measure identity. In that normalization, with $r = p_j + q_j$, $L = \log b$, and f the CM-type count, the coefficients of $fr^2 L$ in the logarithmic upper bounds for V_j and $\prod_\rho d_{j,\rho}$ are $38 + 81$ and $3/2$, respectively, with the additional $fr^2 \log(r + 3)$ in the mass bound. The bound used directly is $\log R_j \le 122 fr^2 L$, hence $\log R_j \le 244 E_3$ when $fr^2 L \le 2E_3$. This is a bound on the product in (15), with no extra Weyl-dimension term added to the height. For the definite lift, the upper estimate groups the full-level classes by the norm-one determinant torus, bounds the special unitary covolume in each class, and accounts for the relative level index. The opposite Tamagawa mass inequality is used in the release for a two-sided norm comparison. The upper covolume calculation here uses only $0 < R_j < \infty$ and the upper logarithmic bound. Mass and dimension budgets · Upper ratio on the selected family · Tamagawa lower comparison · Torus and special unitary mass estimates

Sign of the period on the same host. Let $K_0 \subset E$ be the imaginary quadratic host subfield of discriminant $-d < -4$ and $\chi = \chi_{-d}$. The signed scalar used in both the row and the column is

$$\omega = \log \Omega = \gamma + \log(2\pi / \operatorname{rd} E) - L'(1, \chi)/L(1, \chi). \tag{16}$$

To obtain it, for each ideal class c take $\tau_c \in \mathbb{H}$ for its ideal lattice and put $y_c = \operatorname{Im} \tau_c$. Summing the lattice Epstein series $\sum_{(a,b) \ne (0,0)} y_c^s / |a + b\tau_c|^{2s}$ over classes gives $2(\sqrt{d}/2)^s \zeta(s) L(s, \chi)$ by the ideal norm correspondence and the two units of K_0. The first Kronecker limit formula has residue π and constant $2\pi(\gamma - \log 2 - \log(\sqrt{y_c}|\eta(\tau_c)|^2))$. Comparing residues and constants gives

$$2 \operatorname{av}_c \log(\sqrt{y_c}|\eta(\tau_c)|^2) = \gamma - 2\log 2 - \log(\sqrt{d}/2) - L'(1, \chi)/L(1, \chi).$$

For the Weierstrass normalization $g_2^3 - 27 g_3^2 = 1$, the lattice area divided by π is $4\pi y_c |\eta(\tau_c)|^4$. This follows by dilating the lattice $\mathbb{Z} + \tau_c \mathbb{Z}$ by a twelfth root of its discriminant, whose modulus to the twelfth power is $(2\pi)^{12} |\eta(\tau_c)|^{24}$. Complex multiplication gives equal multiplicities of the ideal classes in the conjugate j-invariants. Thus the mean of the normalized quadratic log areas is $\gamma + \log(2\pi/\sqrt{d}) - L'/L$. In the host's period metric the integral correction is subtracted. Its value is $\log N_{E/\mathbb{Q}} \mathfrak{D}_{E/K_0} / [E : \mathbb{Q}] = \log \operatorname{rd} E - \log \sqrt{d}$ by the tower discriminant formula. This proves (16) for the normalized CM period with the prescribed different correction. All means are taken over the conjugate CM datum before applying its lift weight. The identity retains the sign of ω. Epstein normalization · CM conjugate areas · Different correction

On a deep cover for a nonterminal lift, the flat line at ρ has normalized degree $c_{1,\rho}$ satisfying $|c_{1,\rho} + (p_j - q_j)\omega/2| \le 3q_jL$. Under the integral flat twist and the rank-preserving identification with the shifted weight, a tensor power with exponent $A_{j,\rho}$ contributes $2A_{j,\rho}c_{1,\rho}$ to the squared-norm slope, up to the bad-prime correction of the twist. If

$$\left|(p_j - q_j)\sum_\rho (A_{j,\rho} - w_{j,\rho})\omega\right| \le C_A E_3, \qquad q_j L \sum_\rho |A_{j,\rho}| \le C_A E_3,$$

the flat twist contributes $-P_j$ up to $7C_AE_3$ and the correction, where $P_j = (p_j - q_j)\sum_\rho w_{j,\rho}\omega$. These are the actual anchor exponents and tied reference weights. The sign reverses when this lower slope estimate becomes an upper covolume estimate. Flat line degree · Integral twist and numerical loss · Signed reference main term

The core slope from a deep cover

Proposition (nonterminal core slope). Fix the family and its auxiliary parameters. Assume constants $C_H, C_F \ge 0$ and a threshold x_0 such that for every $x \ge x_0$ in the family range, every selected index, every fitted row datum, and every $1 \le j < m$, there is a good deep cover Y in the prescribed level and coefficient-field realization. Assume the following inputs for every such cover and every admissible quotient weight, with an admissible weight available at each lift. Each component may have its own model and section frames. All comparisons at a given lift use the same cover, and all constants in the conclusion are independent of the row, lift, cover and weight. Under these hypotheses every such lift satisfies (16j). The proof uses one cover and one weight satisfying the inputs simultaneously. Cover hypotheses and quantifiers · Family quantifiers

Write F for the totally complex coefficient field on the cover and $K_* = \mathrm{PP}\,\kappa_{\mathrm{wt}}$ for the product of the cover parameter and the weight size. Assume $1 \le b \le E_3$, $b \le K_*$ and $|P_j| \le K_*$. Let G denote the original weight and $G_\pm$ the two weights obtained by shifting it by $\pm c_0$ times the prescribed Hodge weight. Write $A_\pm$ for their full integral cover lattices and B for the cover lattice at G intersected with the prescribed type. Let U_Y be the full cover lattice at the auxiliary weight u_- with scalar weight zero. To obtain u_-, remove the anchor and scalar characters and apply the negative Hodge shift. The integral flat twist restores the removed characters and compares U_Y with A_-. All the section lattices here are finitely generated and span their indicated rational spaces, with positive definite Hermitian metrics at infinity. Assume $\operatorname{rk} A_- > 0$. With ordinary norms write $\deg_F M$ for the arithmetic degree of $\det M$ divided by $[F:\mathbb{Q}]$. Then $\mu(M) = 2\deg_F M / \operatorname{rk} M$ in the convention of (14). Weights and cover lattices · Integral twist and weight identification

Component data. Let $\mathcal{Z}$ be the finite nonempty set of components of the flag cover, $D \ge 1$ its flag dimension and $T = \lceil \mathrm{PP}^2 D^2 k/m \rceil \ge 3D + 5$. The quotient weight u_1 is obtained by taking the integral quotients of the prescribed corrected adjacent weight gaps by $T + D + 1$ and fixing the anchor coordinate. In the release $\rho_B = \rho_B(u_1)$ is the boundary error parameter. Put $0 \le \rho_B < 1$, $z = 7(D^2/T + D\rho_B)$ and assume $z \le 1/2$. At Z take δ_Z to be the degree of the projective image of the cusp linear system of weight u_1. Take h_Z to be its logarithmic Chow height, normalized by $(D+1)\delta_Z$ in the prescribed section-lattice metric. Let r_Z, d_Z be the rank and ordinary normalized arithmetic degree of the cusp lattice at u_- on Z, using the component probability metric. The degree comparison, rank shares, and component splitting with its probability masses give, for some $\delta_0 > 0$,

$$\begin{aligned} &\delta_0 \le \delta_Z \le \delta_0/(1-\rho_B), \qquad (1-z)r_Z \le Q_Z := \binom{T-1}{D}\delta_Z \le r_Z \quad (Z \in \mathcal{Z}), \\ &\operatorname{rk} U_Y = \sum_Z r_Z, \qquad \mu(U_Y) \ge 2\frac{\sum_Z d_Z}{\sum_Z r_Z}. \end{aligned} \tag{16a}$$

For the last inequality, passing from the component norm to the full norm multiplies the squared norm by the probability mass $m_{Z,\sigma} \in (0,1]$ at each infinite place. Orthogonality and the integral splitting give the component degree contribution $d_Z - \frac{1}{2} r_Z \operatorname{av}_\sigma \log m_{Z,\sigma} \ge d_Z$. The rank shares use the adjoint section system and Kawamata–Viehweg vanishing for the lower bound, and the jet estimate and the comparison of Chern numbers with volumes for the upper bound. Applying these results requires the particular geometric and integral section identifications of the selected cover. Quotient gaps · Rank shares · Degree comparison with component masses

Rank shares from jets and residues. For the upper share, suppose the top geometric degree Δ_{top} of the compactified u_1 bundle satisfies $(1-\rho_B)\Delta_{\mathrm{top}} \le \delta_Z$. This requires the fitted Chern number and volume comparison and the boundary valuation bounds. Put $n_{\mathrm{jet}} = T + D + 3$ and $a = \lfloor n_{\mathrm{jet}}/(1-\rho_B) \rfloor$. Then $n_{\mathrm{jet}}\Delta_{\mathrm{top}} < (a+1)\delta_Z$. Use the fitted jet comparison for u_- on the degree-δ_Z transverse cut. The comparison requires extension on the compactified flag component, nefness in the weights used in the jet argument, dominance of u_- and $n_{\mathrm{jet}}u_1 - u_-$, and adjacent gaps of u_1 at least $(D+3)(\binom{D+1}{2} + 4(p_j+q_j) + D + 3)$. The fitted intersection comparison bounds the mixed degree of the u_- bundle against $D-1$ copies of the generated u_1 moving bundle by $n_{\mathrm{jet}}\Delta_{\mathrm{top}}$. With this bound, the restriction count below gives $r_Z \le \delta_Z \binom{a+D}{D}$. The binomial ratio satisfies

$$\frac{\binom{T-1}{D}}{\binom{a+D}{D}} \ge \left(\frac{T-D}{a+D}\right)^D \ge 1 - D\left(\frac{3D+3}{T} + \rho_B\right) \ge 1 - z.$$

Indeed $(T-D)/(a+D) \ge (T-D)(1-\rho_B)/(T+2D+3-D\rho_B) \ge 1 - (3D+3)/T - \rho_B$. Bernoulli's inequality gives the middle bound. This proves the upper share in (16a). The residue columns constructed below give $Q_Z \le \dim_F V_T$, and the component flat transfer gives $\dim_F V_T = r_Z$. Jet estimate on the flag component · Degree comparison and rank shares

On the modification used in the jet estimate, write $\mathcal{M}$ for the pullback at u_- and $\mathcal{L}$ for the generated moving bundle at u_1. Its top degree is the cut length, denoted by δ_Z. Since $\mathcal{M}\cdot\mathcal{L}^{D-1} < (a+1)\delta_Z$, $\mathcal{M} - (a+1)\mathcal{L}$ has negative mixed degree and no nonzero section by nefness of $\mathcal{L}$. For $0 \le i \le a$ the restriction sequence to a general hyperplane in $\mathcal{L}$ bounds the contribution of $\mathcal{M} - i\mathcal{L}$ modulo sections of $\mathcal{M} - (i+1)\mathcal{L}$ by the corresponding section dimension on the hyperplane. Its mixed degree there is less than $(a+1-i)\delta_Z$. Induction on dimension therefore bounds the contribution by $\delta_Z\binom{a-i+D-1}{D-1}$, with the length bound in dimension zero. Summing gives $\delta_Z\binom{a+D}{D}$. The source uses generic hyperplanes over field extensions, which preserve the section dimensions and intersection numbers. · Mixed intersection and restriction count

The height input consists of numbers $c_s, c_c, X \ge 0$ and $s_Z \in \mathbb{R}$ such that

$$\begin{aligned} &h_Z + c_s \ge 0, \qquad (1-\rho_B)(s_Z + c_s)_+ \le h_Z + c_s + c_c \quad (Z \in \mathcal{Z}), \\ &|\mathcal{Z}|^{-1}\sum_Z s_Z \ge -X, \qquad TX,\ Tc_c,\ c_s\big(6D + 10 + T(z + 2\rho_B)\big) \le C_H E_3. \end{aligned} \tag{16b}$$

Here s_Z is a lower bound on the component volume slope at the weight used for the boundary comparison, c_s bounds the logs of the short generators, c_c includes the finite-place cost of carrying sections to weight u_1, and X bounds the loss in the mean of the volume slopes. Height comparison and budgets

Boundary volume on the selected cover. Put $r = p_j + q_j$ and let κ_{H} be the Hodge weight with entries 0 for $i < p_j$ and -1 for $i \ge p_j$, with zero-based labels. The boundary weight is $u' = u_1 - 3r\kappa_{\mathrm{H}}$, so the last q_j entries increase by $3r$. The divided-weight bounds give $g = h_0/(4T) > 0$, with h_0 the input gap parameter. They give strict decrease and $|u'_{\tau a} - u'_{\tau b}| \ge g$ whenever $a \ge p_j$ and $a \ne b$. In particular the wall gap $u'_{\tau,p_j-1} - u'_{\tau,p_j}$ is at least g. Set

$$B_{\mathrm{ht}} = 4fr^2(1 + \|u'\|_\infty/g), \qquad b_{\mathrm{ht}} = \lfloor B_{\mathrm{ht}} \rfloor, \qquad S = \{\mathfrak{p} \text{ rational prime} : \ \mathfrak{p} \le b_{\mathrm{ht}} \text{ or } \mathfrak{p} \mid 2|\operatorname{disc} E|M'\}.$$

Here M' is the actual level integer. The selected $M' = g_0 g_1\, 15^{\mathrm{expL}}$ has at most 15 distinct prime factors. The host estimate then bounds the bad set $\{\mathfrak{p} : \mathfrak{p} \mid 2|\operatorname{disc} E|M'\}$ by $4fr^2$. The selected $m^2 \le k$ gives $r \ge 3q_j + 1$, and $b_{\mathrm{ht}} \ge 4fr^2 \ge 128$. These are the tower conditions at the prime cutoff b_{ht}, with $B_{\mathrm{ht}} < b_{\mathrm{ht}} + 1$. Boundary weight · Divided-weight conventions · Selected gap and level budget · Deep-cover specialization

For the actual moduli datum, level and field of Y, the global model over $\mathcal{O}_F$ gives a root tower $\mathscr{T}$ and a bijection κ from $\mathcal{Z}$ to its components over the corresponding generic points. Write w for the real weight associated to u', $G(w) > 0$ for its full flag degree, $v_a > 0$ for the tower component volume, and J_a for the limiting arithmetic intersection PPc of $\mathrm{linRoot}(w)^{D+1}$ at b_{ht} in the full intersection system. Put $\Lambda_{\mathrm{ht}} = \log(b_{\mathrm{ht}} \operatorname{rd} E) + |L'(1,\chi)/L(1,\chi)|$, $\ell = \log \operatorname{rd} E + |L'(1,\chi)/L(1,\chi)|$, and $P(u') = \sum_\tau (\max_i u'_{\tau i} - \min_i u'_{\tau i})$. The tower comparison gives

$$a_Z = \frac{J_{\kappa Z}}{(D+1)v_{\kappa Z}G(w)}, \qquad \delta'_Z = \mathrm{topDeg}(Y, Z, u') = v_{\kappa Z}G(w).$$

The sheet degree pushforward makes all v_a equal. Hence

$$\frac{1}{|\mathcal{Z}|}\sum_Z a_Z = \frac{\sum_a J_a}{(D+1)(\sum_a v_a)G(w)} \ge -C_{\mathrm{tower}}\big(\Lambda_{\mathrm{ht}}\|u'\|_1 + q_j P(u')\big), \qquad C_{\mathrm{tower}} = 2\max(C_h + C_u/2, 0) + 17.$$

The inequality is the total intersection bound $\mathrm{T2}$. The Hodge–Tate/de Rham comparison at each prime transports the limiting Chern identity to b_{ht}. The Grassmannian test and the determinant Chern term control the two scalar coefficients there, and the Frobenius remainder has normalized absolute value at most $2q_j P(u')$. The constants C_h, C_u are the uniform cusp height and cusp volume constants in the tower proof. Tower on the same moduli model · Degree comparison · Equal component volumes · Unweighted mean · Tower estimates and constants · Derivation of the total bound

Over each such component the tower uses a reference root model and a finite extension with Hodge and flag charts at the primes in S. Its projective resolution preserves the smooth open and has simple normal crossings boundary. At fixed b_{ht}, Z regularize by $w - s\nu$ with $\nu_{\tau i} = -i$, auxiliary weight $s > 0$, and metric parameter $\exp(-s^{-2})$. The nef and semipositive comparison gives normalized degree tending to δ'_Z and normalized slope tending to a_Z. Along positive rational $s \to 0$, clear the coefficients by a multiple of eight and multiply by the effective correction to recover exact powers of u' at a log-radius cost $O(s)$ per unit weight. For each fixed perturbation take the power limit with the geometric and arithmetic intersections divided by the corresponding powers D and $D+1$ of that denominator. Then let $s \to 0$. The resulting

lower volume on the root model at radius t is $\delta'_Z(a_Z+t)_+$. Relative boundary resolution · Root tower and charts · Integral perturbation and correction · Normalized limit

The count for $V'_Z(t)$ is on the original component, with scalar-zero weight nu', pointwise homogeneous norm at most e^{nt}, and Hodge–Tate integral orders over S and reciprocal integral orders elsewhere. Its normalization is $D!\log\#/[F:\mathbb{Q}]n^{D+1}$ along $8\mid n$. After multiplying by a fixed basis in weight k_0u' with $8\mid k_0$, the weight gauge and the field trace give an injective tuple of cusp rules in weight $(n+k_0)u'$. If the generic extension has degree d_{tr}, the tuple has d_{tr} entries. This cancels the factor d_{tr} in the root-model normalization $[F:\mathbb{Q}]d_{\mathrm{tr}}$. The fixed norm factors and shift by k_0 disappear in the power limit. For a rational prime $\mathfrak{p}$ write

$$e_{\mathfrak{p}}=(\mathfrak{p}-1)^{-1}+40\,\frac{v_{\mathfrak{p}}(|\operatorname{disc}E|)}{[E:\mathbb{Q}]},\qquad A=\sum_{\mathfrak{p}\in S}(e_{\mathfrak{p}}+\mathbf{1}_{\mathfrak{p}=2})\log\mathfrak{p}.$$

The reciprocal order loss of the trace and the comparison from reciprocal to Hodge–Tate integral tests each cost at most $(e_{\mathfrak{p}}+\mathbf{1}_{\mathfrak{p}=2})\|u'\|_1$ per unit weight over $\mathfrak{p}$, in units of the order of $\mathfrak{p}$. A product of rational primes with rounded exponents clears their sum. The rounding error divided by the weight tends to zero. Thus the volume on Y satisfies $V'_Z(t)\ge\delta'_Z(s_Z+t)_+$ for $s_Z=a_Z-2A\|u'\|_1$. Since $0\le A\le 60(\log B_{\mathrm{ht}}+\log\mathrm{rd}\,E)$ and $\Lambda_{\mathrm{ht}}\le\log B_{\mathrm{ht}}+\ell$, the choice

$$X=(C_{\mathrm{tower}}+120)\big((\log B_{\mathrm{ht}}+\ell)\|u'\|_1+q_jP(u')\big)$$

gives $|\mathcal{Z}|^{-1}\sum_Z s_Z\ge -X$ with one constant on every qualifying cover. The fitted budget gives $TX\le 83400(C_{\mathrm{tower}}+120)E_3$. The perturbation and power-limit thresholds may depend on the component and weight. Mixed cusp volume and slope · Trace and denominator clearing · Prime sum bound · Mean slope and constants · Fitted numerical budget

The degree ratio $1-\rho_B$ compares the Weyl degrees of u' and u_1. The fitted comparison with the u_1 cusp image gives $\delta'_Z\ge(1-\rho_B)\delta_Z>0$. For carrying to weight u_1 use the fitted section with nonzero generic value on Z at weight $8(u_1-u')=24r\kappa_{\mathrm{H}}$, with carrying cost c_{mult}. Passing from the mixed orders to the reciprocal orders in the u_1 finite-depth count uses the opposite one-sided comparison, with cost $\|u'\|_1\sum_{\mathfrak{p}\in S}e_{\mathfrak{p}}\log\mathfrak{p}$. This sum has no added dyadic term. Thus $c_c=c_{\mathrm{mult}}+\|u'\|_1\sum_{\mathfrak{p}\in S}e_{\mathfrak{p}}\log\mathfrak{p}$, and the second term multiplied by T is at most $83400E_3$. The fitted carrying and short-generator bounds supply the rest of the height budget. Fitted degree, carrier and budget · Carrying into the finite-depth volume · Height input

Uniform cusp estimates at the prime cutoff

Use $p=p_j,q=q_j,r=p+q$ and $n_X=fpq$ for the Grassmannian base dimension. The full flag dimension in the volume normalization remains $D=f\binom{r}{2}$. The boundary-mean constant, denoted here by $C_{0,\mathrm{ht}}$, can be taken to be 1341. This follows from the witnesses $C_h\le 600$ and $C_u=4$ under the following hypotheses. The host $E/\mathbb{Q}$ is abelian CM with $[E:\mathbb{Q}]=2f$. The framed component laws are those of the selected good deep cover Y, with relative algebraic closedness on each component, trivial dominant-weight characters, good local models at every prime in the release's sense, level parameter $A=1$, and the uniformizations with the prescribed common complex normalization. Use the actual global toroidal model, its compatible level and Hecke comparisons, and the tied Hodge and root models. The required ranges are $q\ge 1,r\ge 3q+1,f\ge 2,|S_0|\le 4fr^2\le b_{\mathrm{ht}}$, and $128\le b_{\mathrm{ht}}$, with $\omega(M')\le 15$, where S_0 is the prime

support of $2|\operatorname{disc} E|M'$. The strict weight and wall-gap hypotheses on u' and $B_{\mathrm{ht}} < b_{\mathrm{ht}} + 1$ are as above. At b_{ht} the Hodge–Tate lattices must be free in the full and labelled tests over S, and the vertical twists change the de Rham lattice at $\{\mathfrak{p} \le b_{\mathrm{ht}} : \mathfrak{p} \notin S_0\}$ to the Hodge–Tate lattice. The base line already uses the Hodge–Tate lattice at S_0. Cover and weight range · Tied degree setting · Admitted origins and height clauses · Patched lattice

The height comparison uses the tied normal chart cover of the selected toroidal component, with compatible level and Hecke maps, integral lifts of the height points and their adapted semiabelian frames. The PEL marking, full generic ℓ^{R_ℓ}-torsion at each $\ell \in S$, and the smooth semiabelian family with its rigidified cubical polarization are the inputs to the chart and torsion constructions below. The projective refinement below gives the local freeness of the modified cotangent and its restricted labelled images used in height-one descent and specialization to the height traits. On the complex resolution it uses the tied log-log Hodge and bounded corrected presentations with the prescribed reference metrics, the positive auxiliary Levi margin, and the conjugation and fibre-nefness conditions. The arithmetic comparison uses characteristic-zero resolution and principalization, invariant Poincaré–Lelong, and the full intersection system with metric change and birational and relative flag projection on the prescribed finite model. Zhang's argument uses the arithmetic height provisions, including arithmetic Hilbert–Samuel and Moriwaki bigness in the release's wrapper. The polarized Lie-bundle estimate uses the large-degree line theorem and the corrected quadratic Chowla–Selberg degree. Semiabelian Hom extension, GAGA, the finite-flat quotient, equal-volume pushforward, the prime-by-prime Chern comparison and the generic nef Riemann–Roch comparison remain part of the tower and section-count construction. Geometric inputs · Boundary theorem inputs, including Zarhin and nef Riemann–Roch · Construction of the tower and flat-degree input · Grassmannian height from the intersection and volume systems · Intersection constructor · Height provisions

Point heights. Fix the degree setting on component a. Write μ for the small-block determinant and U_{cusp} for the cusp weight. Correct the Hodge potential by $U_{\mathrm{cusp}}/2$ and tensor an auxiliary ample metric $\overline{A}'$ with point heights at least $-K_A$. For integers $1 \le \alpha \le \beta$, the regularized line $\overline{L}_{\alpha,\beta,t'}$ has the twisted determinant to power β and $\overline{A}'$ to power α. If W is an arithmetic point closure, normalize its height by $\beta \deg_{\mathbb{C}} W$, where $\deg_{\mathbb{C}} W$ counts the geometric points with multiplicity. Equivalently, take the normalized number-field point height and divide by β. The two bounds are

$$\frac{\widehat{\deg}(\overline{L}_{\alpha,\beta,t'}|_W)}{\beta \deg_{\mathbb{C}} W} \ge -C_h f q\, \Lambda_{\mathrm{ht}} - \frac{\alpha}{\beta} K_A, \qquad 0 \le \frac{\int U_{\mathrm{cusp}}\, d\mu_{\mathrm{H}}}{\int d\mu_{\mathrm{H}}} \le C_u f q\, \Lambda_{\mathrm{ht}},$$

where $d\mu_{\mathrm{H}}$ is the top curvature measure of the uncorrected Hodge metric and $\int d\mu_{\mathrm{H}} > 0$. The tied regularization gives a positive t'_0 at the fixed cutoff, setting and auxiliary integers before choosing W. The point bound holds simultaneously for all W and $0 < t' \le t'_0$, including points on the boundary. To see this, lift a point to an integral point over a number field L with an adapted basis on the principally polarized abelian part and the toric part. With finite norms, the vertical twist, and the limiting archimedean Gram weights included, the product formula cancels the determinant transition in the chart and gives $-\phi\mathrm{Deg}$ for the height with the specialized weight before the auxiliary term. Replacing this weight by the regularized one can only increase that lower bound. The cusp correction majorizes the negative toric log covolume in the boundary chart. The upper semicontinuous bounded envelope then majorizes this specialized weight, and the regularized metric dominates the envelope on the setting. That last comparison supplies the threshold independently of the point or its field of definition. Height normalization · Uniform threshold · Determinant identity · Boundary metric comparison · Auxiliary term

The field of definition of the marked torsion. Let K_0 be the function field of the selected toroidal component and A_0 its generic abelian variety. At the selected odd level M', the two coprime conjugate level ideals split $A_0[M']$ geometrically as $P \oplus Q$. The adjoint identity for the action of $\mathcal{O}_E$ makes both summands isotropic. The principal polarization therefore pairs P and Q perfectly. The level labels P pointwise and the coefficient field F contains $\mu_{M'}$. For Galois g, $x \in P$ and $y \in Q$ this gives $e_{M'}(x, gy - y) = 1$. Since g preserves the two summands, $gy = y$. Thus the selected data split $A_0[M']$.

For the height calculation fix the finite set S and its depths $R_p \geq 4$, and put $D_S = \prod_{p \in S} p^{R_p}$. Over any starting function field K_1 above K_0 used in the construction, adjoin all points of $A_0[D_S]$. In characteristic zero this is a finite Galois extension, even at the factor 2^{R_2}. A symplectic pair of integral labels with $B(a, b) = 1$ has primitive Weil pairing ζ_{D_S} in this field. For $d = p^{R_p}$, multiplication of the D_S-level points by D_S/d produces their d-level labels. Their marked root is $\zeta_{D_S}^{D_S/d}$. The normal projective model in this common field supplies the generic marking simultaneously on all its charts. The root tower and the intersection and square calculations take depth four. The flag construction allows the depths to vary. Generic full-level field · Root tower · Intersection model · Square calculation · Variable-depth flag

Write $d_{\mathrm{cov}} = [\Omega : K_0]$ for the total finite function-field degree actually used over the selected component. The coefficient datum stays over F. The root model normalizes its section count by $[F : \mathbb{Q}]d_{\mathrm{cov}}$, while the trace comparison uses d_{cov} coordinates. Equivalently, pushforward of a generically finite pullback cycle of degree d_{cov} multiplies both its top intersection and its Chow degree by that degree. The subsequent projective modification of the Hodge lattices has function field Ω and degree one over the prepared model. The PEL comparison requires the selected semiabelian group and its generic polarized marking at each bad place. Projective refinement · Degree on the flag · Selected level and height normalizations

The marked Hodge chart. Fix a rational prime $\ell \in S$ and put $n = R_\ell \geq 4$ and $d_\ell = \ell^n$. Write X_1 for the input projective model on the selected component and X for its refinement below. Use a tied coordinate chart $\operatorname{Spec} A$ on either model, with normal Noetherian coordinate ring A and fraction field K. The chart ring is generated over $\mathbb{Z}_{(\ell)}$ by the coordinate ratios with the chosen nonzero coordinate product as denominator. Use the pullback G_A of the smooth separated semiabelian PEL family, with its $\mathcal{O}_E$-action and its compatible rigidified cubical polarization line. The generic pullback G_K is abelian. The comparison with the tautological PEL object must be cartesian and preserve the action and polarization on this selected component. The preceding finite extension makes every geometric point of $G_K[\ell^n]$ rational over the chart function field $K = \Omega$. The level, coefficient field, cusp data and fans here are those of the selected cover. Lan's good-reduction toroidal theorem, Theorem 6.4.1.1, or the corresponding normalized PEL model in ramified characteristics, Theorems 10.13, 11.2 and 11.4, provide the degenerating family where their hypotheses apply. The PEL inputs are the identification of that family and its level maps with the selected model, its determinant signature and polarized marking, and the compatibility of the cubical line with the marked Weil pairing. Chart and generic PEL comparison · Lan, good-reduction toroidal theorem · Lan, ramified characteristics

Choose one polarized Tate marking over $\overline{K}$ compatible with the $\mathcal{O}_E$-action, using the projective rank-r lattice Λ and its unimodular alternating form B. Full level descends its points at d_ℓ to K. A pair with $B(a, b) = 1$ exists by unimodularity and positive rank. Its Weil pairing is the primitive root ζ_n of order ℓ^n, so $\zeta_n \in K$ and $\zeta_n^{\pm 1} \in A$ by normality. Lift a basis of $\Lambda/\ell\Lambda$ over $\mathcal{O}_E/\ell$ to $u_1, \ldots, u_r \in \Lambda$. Such a basis exists by projectivity over the local factors of $\mathcal{O}_E/\ell$. These lifts generate modulo every ℓ^m by iteration, and at level n the map $(\mathcal{O}_E/\ell^n)^r \to \Lambda/\ell^n\Lambda$ is an isomorphism by equality of cardinalities. The same marking

gives the points, root and pairing on every chart. Finite marking and root · Finite lattice and pairing · Assembly on the tied charts

The scheme-theoretic kernel $G_A[d_\ell]$ is separated, quasi-finite and flat over A. For $x \in \mathcal{O}_E^r$ and $a_x = \sum_i x_i u_i$, the generic function $b \mapsto \zeta_n^{B(a_x,b)}$ is a character of its split generic fibre labelled by $\Lambda/\ell^n\Lambda$. To extend it, henselize A at each prime. Zariski's Main Theorem and the decomposition of a finite scheme over a henselian local ring give a finite open containing the whole closed fibre of the kernel. Identity, inverse and product on this open land in it because every point of a finite scheme over the local base generalizes a point of its closed fibre. The open is therefore a finite flat subgroup. Its finite flat Cartier dual extends the restricted generic character over the normal henselian local ring, since the coordinates of a point of a finite scheme are integral over the base. At a kernel stalk over the original prime, extension to a stalk after henselization is faithfully flat. For the multiplicative set $A \setminus \{0\}$, membership of a generic function in the original stalk can be checked after this flat extension by tensoring the quotient of the localization by the stalk. Thus the generic character is regular at every point of the kernel. Flatness gives schematic density of the generic fibre of the kernel and its square. It gives uniqueness of the extension f_x, identifies the inverse with the extension of the inverse character, and gives multiplicativity by comparison on the generic fibre. Quasi-finite kernel · Zariski's Main Theorem · Henselian finite algebras · Character extension

Write $\omega = \omega_{G_A}$ for the finite projective invariant cotangent and $L_A = \mathrm{Lie}(G_A)$ for its dual. The differential of multiplication by d_ℓ at the identity is $d_\ell \,\mathrm{id}$, so the kernel's cotangent at the identity is $\omega/d_\ell\omega$. The differential $d \log f_x$ is the class of $f_x - 1$ there. On the dual numbers over $A/(d_\ell)$ in direction $v \in L_A$, the function has value $1 + \varepsilon(w_x(v) \bmod d_\ell)$ for any lift $w_x \in \omega$ of this class. Such lifts differ by $d_\ell\omega$. Choose lifts on a $\mathbb{Z}$-basis of $\mathcal{O}_E^r$ and extend by additivity, using the product rule for the logarithmic differential. The adjoint identity $B(ua, b) = B(a, \bar{u}b)$ gives $f_{ux} = f_x \circ [\bar{u}]$, by checking on the generic fibre. In particular

$$w_{ux} \equiv w_x \circ d[\bar{u}] \pmod{d_\ell\omega}.$$

The action on ω itself is precomposition with $d[\bar{u}]$. On chart overlaps the regular functions agree, so the differential lifts agree modulo d_ℓ in the common cotangent. A representative on the kernel lifts locally to a function on the ambient group near the identity, since the kernel is closed. Conversely, an ambient function with the prescribed generic values restricts to f_x on the kernel where both are defined, by flatness. The linear term of the kernel ideal at the identity is $d_\ell\omega$, so the ambient gradient gives this same class. This identifies the chart lattice $\omega'_\ell = d_\ell\omega + \sum_x A w_x$ with the release's stalkwise span of the character gradients. The chart construction uses smoothness for the projectivity of ω. Differential lifts and action · Stalk gradient and dual lattice · Gluing · Marked Hodge chart construction

Simultaneous freeness on the projective refinement. On the input model X_1, the chart spans glue to coherent subsheaves $\mathcal{F}_\ell$ of the invariant cotangent, with $d_\ell\omega \subseteq \mathcal{F}_\ell \subseteq \omega$. For every label s, also take the image $\mathcal{F}_{\ell,s}$ under restriction to the dual of the generic Lie eigenspace. On each affine chart the restricted generators span a finite module in that generic dual, and localization commutes with the image. The tied coordinate ratios have fraction field K, the field Ω of the release. The label values are in K by the coefficient embedding. Since d_ℓ is invertible in K, the total span has full cotangent rank, and restriction of generic functionals is surjective, giving the labelled ranks p and q from the PEL signature. On $X_{1,\mathbb{Q}}$ the modified cotangent equals the ordinary cotangent. The trace-dual projectors split the Lie bundle there, since the label values on $\mathcal{O}_E$ are integral over $\mathbb{Z}$ and their projector denominators are cleared by a nonzero

integer. Thus the labelled images are already bundles on the whole characteristic-zero fibre. Total and labelled ranks · Tied chart fraction fields

For each sheaf $\mathcal{F}_b$ in this finite family, take a surjection $\mathcal{E}_b \twoheadrightarrow \mathcal{F}_b$ from a sum of negative powers of the coordinate line, using global generation after twisting. Close the joint generic rank quotients in the product of the quotient Grassmannians of the $\mathcal{E}_b$ over X_1. The closure is integral and has function field K because the generic quotients are defined over K. Its normalization in K gives the normal model X. The closure is projective over X_1 and its normalization is finite, since it is finite type over $\mathbb{Z}$. Hence $g : X \to X_1$ is projective and birational. Over $X_{1,\mathbb{Q}}$ the quotients already give sections, so the graph and its normalization restrict to $X_{1,\mathbb{Q}}$ itself, including the horizontal boundary. On X the tautological quotient of $g^*\mathcal{E}_b$ kills the pulled-back relations because they vanish generically in a vector bundle. The resulting surjection from $g^*\mathcal{F}_b$ is an isomorphism generically, and its kernel is exactly the torsion. Common refinement of the two families · Finite join and same-field normalization

In the pin the generic quotients are given by Plücker coordinates, the maximal minors $\Delta_{b,J}$ of the matrices of twisted generators in K-bases. Over an input affine chart of ring A_c, the joint graph on the pivot charts has ring $A_c[\Delta_{b,J}/\Delta_{b,I_b}]_{b,J}$ inside K, with each pivot minor $\Delta_{b,I_b} \neq 0$. Changing the trivialization of a twist scales its entire minor block by a common factor. Varying one factor at a time in the Segre products recovers the minor ratios of each block, so the pin's join has exactly these chart rings before normalization. Cramer's rule makes the pivot generators the tautological frame. A rank-zero image contributes the unit minor. Minor blocks and Cramer's rule · Segre ratios

Pull back the group, action and cubical line along g. The kernel commutes with base change, and the cubical pairing on the same labelled generic points gives the same Weil values. The pulled-back f_x therefore represents the actual Weil character on the new kernel, by uniqueness on its schematically dense generic fibre. Invariant differentials commute with base change, so the new modified cotangent is the image of the pulled-back generators in the smooth cotangent, with the same $d_\ell\omega$ term on the new model. Restriction to each generic labelled dual commutes with taking their span. These images are the torsion-free pullbacks just proved locally free, giving both clauses of **HTFree** for every $\ell \in S$. The refinement preserves the coefficient embedding, generic degree, level and R_ℓ, and the original coordinate blocks, whose normalized heights pull back. Properness lifts a height trait already mapped to X_1 using its unique characteristic-zero lift. The adapted toric and principally polarized abelian-part data at the height point remain inputs to the height comparison. Cubical pairing under pullback · Pullback of the modified cotangent · Freeness and retained projective data · Height-point data · Common Hodge refinement

The primitive polarization line. On the selected normal integral component let G be the smooth separated semiabelian identity family with its $\mathcal{O}_E$-action and the marked principal polarization λ on its generic abelian fibre. Let P be the Poincaré biextension there, rigidified on both axes, and write $B(x, y) = P(x, \lambda y)$. By Moret-Bailly, II, Theorem 3.6, restriction to the generic point is an equivalence for $\mathbf{G}_m$-biextensions of a pair of semistable groups over a normal integral locally noetherian base if one group has connected fibres. Apply it to (G, G). We obtain B on the whole product. Full faithfulness extends its symmetry and the Rosati identity $B(ax, y) \simeq B(x, \bar{a}y)$ with their compatibilities. Moret-Bailly, *Pinceaux de variétés abéliennes*, II, Theorem 3.6

The host contains $w = (1 + \sqrt{-d})/2$ and $w + \bar{w} = 1$. Set $L_w(x) = B(x, wx)$. For a rigidified line L put $D(L)(x, y) = L(x + y) \otimes L(x)^{-1} \otimes L(y)^{-1} \otimes L(0)$. Bilinearity, symmetry and adjunction give

$$D(L_w)(x,y) \simeq B(x,wy) \otimes B(y,wx) \simeq B(x,(w+\bar{w})y) = B(x,y).$$

Thus L_w induces λ on the abelian open. Simultaneous inversion in both variables gives its symmetry, and the biextension laws give its normalized cube trivialization. For any rigidified symmetric line L inducing λ, the difference $L \otimes L_w^{-1}$ on the abelian open belongs to $\mathrm{Pic}^0[2]$. The release requires the additional identity $L(x) \simeq D(L)(x,wx)$, which fixes the line as L_w. This argument uses the integral trace identity also at residue characteristic two. For a cartesian comparison preserving the generic marked polarization, the comparison of the two biextensions extends and pulls back along $(1,w)$. For an $\mathcal{O}_E$-linear isogeny f with $f^\vee \lambda' f = c\lambda$, the same calculation gives $f^* L'_w \simeq L_w^{\otimes c}$. Thus the multiplier matters when identifying a level map. The local torsion formulas below use the ordered commutator and the primitive root of that marking. For the selected component, the comparison of the marked abelian families and the extent of its integral continuation are as follows. Released adjoint and primitivity identities · Lan, Example 2.3 and Theorem 11.2 · Polarization, functoriality and commutator

The marked family on the characteristic-zero component. Let Z be the selected smooth component over its coefficient field $F \subset \mathbb{C}$. Use its reference $\mathcal{O}_E$-lattice Λ with signed unimodular alternating form B_0, and its prescribed signature at every embedding. In Lan's integral PEL problem take the identity family j_0 of Example 2.3 with this lattice. Use the alternating form $-(2\pi i)B_0$ valued in $\mathbb{Z}(1)$ and the principal congruence level $M = g_0 g_1\, 15^{\mathrm{expL}}$. Fix the cyclotomic coordinate by sending $2\pi i \bmod M$ to $\zeta_M = \exp(2\pi i/M) \in F$. On the component containing the selected marked period lattice, this identifies the two abelian schemes, their $\mathcal{O}_E$-actions and their principal polarizations over F. Lan, §2 and Example 2.3 · Moduli object and geometric marking

To see the role of the root, write the level ideal as $\mathfrak{b}\bar{\mathfrak{b}} = (M)$ with $\mathfrak{b} + \bar{\mathfrak{b}} = \mathcal{O}_E$. Chinese remaindering gives two summands $P_b = \bar{\mathfrak{b}}\Lambda/M\Lambda$ and $Q_b = \mathfrak{b}\Lambda/M\Lambda$. Conjugation exchanges their idempotents. The adjoint relation therefore makes each summand isotropic and makes the cross pairing perfect. The same idempotents split $A[M]$ on the abelian family. The release marks the first summand by η. There is exactly one $\mathcal{O}_E$-linear mark α of the complementary summand with values in the complementary part of $A[M]$ satisfying

$$e_M^\lambda(\eta(a),\alpha(b)) = \zeta_M^{-B_0(a,b)} \qquad (a \in P_b,\ b \in Q_b).$$

Indeed this formula prescribes the image of b under the Cartier-dual identification with the characters on the first summand. Its sign is the sign of the ordered Weil pairing on the fitted Betti lattice. A geometric integral Tate marking witnessing the release's level condition lifts the completed marking after a similitude which scales the second summand at each prime dividing M. The lift of the scalar is a unit and may be chosen to carry the old root to the specified one. Lan's integral principal level uses precisely this geometric liftability. Conversely, restriction of Lan's full marking to P_b recovers η. The inverse construction recovers the cubical line by $x \mapsto B(x,wx)$ as above. The Lie eigenspace ranks give the same determinant condition on the two moduli problems over F. The constructions commute with base change on reduced locally noetherian F-schemes and recover the universal objects. Representability gives the claimed identification on Z. If an additional full marking at N is transported, its individual labels and its cyclotomic coordinate must also be transported. Two prescribed markings in one Tate tower have to agree on their reductions at $\gcd(M,N)$. Lan, §§1.3.6 and 5.2.4 · Sign, integral lift and inverse on the moduli problems

A finite model for the change of marked level. Consider the forgetful map from the selected higher level $M_2 = M_1 s$ to the lower level M_1, after placing both cyclotomic coordinates in the same coefficient

field. In the depth enlargement $s = 15^{\mathrm{expL}}$ and the two roots satisfy $\zeta_{M_2}^{s} = \zeta_{M_1}$. The half-level idempotents agree modulo M_1, since the lower half ideal and its conjugate keep their respective factors at every prime dividing M_1. Thus multiplying the higher marked point by s recovers the lower marked point for every lattice label. The auxiliary marking at $A = 1$ is trivial. Over characteristic zero, compatible lifts of a given full marking form a finite étale cover. Indeed the torsion is finite étale, and étale-locally the liftable markings are a finite set of integral Tate orbits with their prescribed reductions. Take the selected component Z_2 over its image component Z_1. The forgetful map preserves the underlying polarized abelian scheme. Released level compatibility

Fix a normal projective integral lower model X_1 carrying the semiabelian family over an excellent discrete valuation ring. Take the normalization X_2 of X_1 in the **specified** field $F(Z_2)$. For a change of coefficient ring, first take the closure of the image component in the base change of the lower model. On an affine open with ring C the upper ring is the integral closure of C in $F(Z_2)$. It is finite over C by excellence. Localization commutes with this construction, so the rings glue, and their restriction above Z_1 gives Z_2, which is normal and finite there. The result is projective and flat over the discrete valuation ring, since the upper rings are torsion-free. Pull back the lower semiabelian family, its action and its primitive line. This gives the cartesian family comparison required in the finite level transition, with generic degree $[F(Z_2) : F(Z_1)]$. The construction works even over residue characteristics dividing M_2, where the integral finite map can ramify. Any other normal integral model already finite over this same lower model in this same field is uniquely the normalization, by integral closure on each affine. Using an independently specified toroidal model in this identification still requires finiteness of its map to X_1. Lan's Proposition 13.15 gives projective extensions for compatible cone data. Stacks Project, relative normalization · Finite transition in the release · Lan, Proposition 13.15

For the chosen base and deep levels, these normalizations can be taken inside Lan's construction at each finite place $v \mid p$. Use the same identity lattice and one auxiliary prime-to-p level containing the image of the lower group. Its neatness can be imposed at an odd level prime different from p. The deeper group maps into it as well. Proposition 4.12 gives compatible finite maps of the generic components to the auxiliary interior. Choose the same projective auxiliary cone decomposition for both levels. Since Z_2 maps onto Z_1, their images have the same closure in that auxiliary compactification. Equations (7.7)–(7.8) normalize this one closure in the two function fields. Transitivity of integral closure gives the finite transition just constructed. Proposition 6.1 and Theorem 11.2 supply the identity semiabelian families on these models. The general toroidal Hecke correspondence uses further lattice and cone choices. The proposition below gives the integral theta line on these local models. The global height theorem also requires the comparison with the prescribed normalization, cusp divisors and metrics of the selected sections over the coefficient ring. Lan, Proposition 4.12 and §7 · Local models, marked traits and the global height hypotheses

Comparison over a common modification. Work over one excellent discrete valuation ring of the coefficient field. Suppose X and X' are normal integral projective flat models containing this same Z as a dense open in their generic fibres, and each carries a smooth separated semiabelian extension of the specified abelian family. Let W be the normalization of the closure of the graph of the identification of Z in $X \times X'$, with product over the valuation ring. The normalization is finite by excellence. Both projections are projective, are isomorphisms over Z, and have generic degree one. Indeed, over Z the graph is already closed by separatedness of the other model, and Z is normal. This assertion allows the projections to modify the characteristic-zero boundary.

On W the marked generic isomorphism of abelian schemes and its inverse extend uniquely to homomorphisms of the pulled-back semiabelian schemes by Lan, Proposition 3.3.1.5. Their composites extend the identity, so the extensions are inverse. Uniqueness also identifies the $\mathcal{O}_E$-actions. Under this isomorphism the two polarized biextensions agree by full faithfulness in Moret-Bailly, II, Theorem 3.6. Pullback along $x \mapsto (x, wx)$ identifies the primitive cubical lines with their rigidifications. In particular the invariant cotangent lattices agree on W, including after restriction to a height trait through a bad fibre. A height comparison for line bundles and metrics pulled back from the same data now uses the ordinary degree-one projection formula. The degree of the earlier field extension that supplies the torsion labels is counted at that earlier step. Lan, Proposition 3.3.1.5 · Moret-Bailly, II, Theorem 3.6

The prescribed orders of vanishing at cusps require a further comparison. For example, on the regular local chart with parameters π, t, take the horizontal boundary $D = (t = 0)$ and blow up (π, t). On the chart $t = \pi u$ the total transform of D is the strict transform $(u = 0)$ plus the exceptional divisor $(\pi = 0)$. The trait $u = 1$ meets their sum with order one and the strict transform with order zero. Thus a multiplicity in the height computation requires the actual pullback of its specified divisor. Applying Lan's level morphisms in mixed characteristic also requires the labelled level transition and a compatible choice of cusp cones. Lan, Theorems 7.14 and 11.4 and Proposition 13.15

The boundary divisor on a marked power-series arc

Let V be a complete discrete valuation ring of mixed characteristic, π a uniformizer and $R = V[[u]]$. Fix a marked degeneration at the selected level M, with auxiliary level one. Its period lattice Λ carries a positive definite symmetric integral form b with $b(ay, z) = b(y, \bar{a}z)$ and $2M \mid b(y, z)$. The initial period sections $\psi_0(y)$ generate their rigidified value lines over V. Write G for the semiabelian family on the model and U for its abelian locus, in all characteristics. A marked arc is a morphism $\chi : \operatorname{Spec} R \to Z$ with the given degeneration and torsion labels. In the fourth toroidal model condition it satisfies $\chi^{-1}U = D(u)$, and $D(\pi u)$ maps into the characteristic-zero marked moduli space. Marked degeneration and arc conditions

The deformation multiplies $\psi_0(y)$ by $u^{b(y,y)/2}$. Taking the polarized second difference multiplies its biextension trivialization at (y, z) by $u^{b(y,z)}$. For a marked point of order dividing $k \mid M$ represented by a division lift of the period y, multiply its evaluation on the character z by $u^{b(y,z)/k}$ on the generic open. Taking the k-th power gives the deformed period. Replacing y by $y + kh$ changes the factor by $u^{b(h,z)}$, the deformed period translation. In the ordered Weil pairing the extra exponents cancel by symmetry of b. All these exponents are integers, including at places above two. Division lifts, pairing and deformation

Proposition. In the released construction choose the model Z together with its effective Cartier divisor B and its invariant boundary function φ. For the marked arc above there is an integer $d = \varphi(b) \geq 0$ such that

$$\mathcal{I}_B \, R = (u^d).$$

Consequently for every ring map $R \to W$ to a discrete valuation ring with nonzero image of u, the order of the pulled-back divisor is $d \, v_W(u)$, with v_W normalized at a uniformizer of W.

Proof. The construction chooses (Z, B) together. On this very model it supplies the marked arc, gives $\operatorname{Supp} B \subset Z \setminus U$, and prescribes the order $\varphi(c)$ on a complete valuation trait with marked monomial degeneration form c and generic point in the moduli locus. In particular the support exclusion applies at

the image of the prime (π) of R, where u is invertible. Selection of the model and its divisor · Boundary valuation rule

Trivialize the pulled-back Cartier line over the local ring R, and call its equation f. The support exclusion makes f a unit in $R[1/u]$. Hence $fg = u^s$ for some $g \in R$ and $s \geq 0$. The element u is prime because $R/(u) = V$ is a domain. Dividing by u successively in this identity gives $f = hu^j$ for a unit $h \in R^\times$ and $0 \leq j \leq s$. In particular this pullback is an effective Cartier divisor on the whole surface.

Evaluate u at π by the convergent coefficient-preserving map $R \to V$. The generic point of this trait lies over $D(\pi u)$. Specialization of the marked uniformization gives the period exponent $b(y, y)/2$ in an integral generator on this trait. Polarization recovers the bilinear form b on the transported lattice, and invariance of φ under lattice isomorphisms gives the prescribed order $\varphi(b)$. On the other hand, the evaluation of hu^j has order j. Thus $j = d$. Pulling back the equation hu^d gives the asserted order for every other trait through this arc. Construction of the specialized marked witness · Transport of the period form and the boundary order

For a comparison with a chosen Lan chart, suppose the two marked arcs over R have been identified and each divisor is supported over $u = 0$ on its arc. Equality of their orders on $u = \pi$ is then equivalent to equality of their ideals over R. If the Lan Cartier equation is a monomial of weight ℓ times a unit and that character evaluates to a unit times $u^{n(\ell)}$, the comparison at this arc is the numerical identity $n(\ell) = \varphi(b)$. Constructing the comparison of arcs still requires the cusp character lattice, its integral weight lines, and the divided torsion evaluations for the chosen marking. An equality of ideals determines the vanishing orders and leaves a unit in a fixed theta frame. Normalized toroidal charts, §§8 and 10 · Integral chart and marking comparison

Good reduction and the primitive theta section

Proposition (good reduction at a valuation ring). Let X be the selected proper toroidal model over an excellent discrete valuation ring R, with its semiabelian identity family G. Let V be a valuation ring over R with fraction field K. Suppose a K-point of the marked abelian locus has a good-reduction witness, an abelian scheme A/V whose generic group is identified with the given fibre. The unique extension $s : \operatorname{Spec} V \to X$ then has its entire image in the abelian locus. This proves both good-reduction implications for this model in the local level conditions. The definition also asks that the witness carry a relatively ample line with the prescribed generic identification. Its conclusion concerns the centre and its geometric fibres. Good-reduction implications

Proof. Properness gives s over the prescribed map $R \to V$. Choose an affine open of X containing the image of the closed point. It contains the entire image of s, since every point of $\operatorname{Spec} V$ generalizes the closed point. The coordinate ring of this open is of finite type over R. Write V as the filtered union of its finite type R-subalgebras B, using the image of R in V. At some stage the chart map descends. Finite-presentation approximation descends A, its group operations and their identities, and after enlargement gives a proper smooth commutative group A_B over B. Write H_B for the pullback of G to this stage. The rings $B[1/d]$, for all these stages and all $0 \neq d \in B$, have filtered union K. At one such stage the generic isomorphism and its inverse descend to inverse group maps on the dense open $D(d)$. Approximation of schemes and morphisms · Descent of smoothness · Descent of properness

Normalize the final B in its own fraction field. The normalization B' is finite by excellence and lies in V because V is integrally closed in K. The proper smooth group $A_{B'}$ has geometrically connected generic fibre, as can be checked after extension to K. The number of geometric connected components of its fibres is locally constant, hence one everywhere on the integral base. Thus $A_{B'}$ is abelian. Lan's extension theorem for semiabelian groups over normal noetherian schemes extends both group maps from $D(d)$ over B'. Their composites are identities by uniqueness. Pullback to V gives $s^*G \simeq A$. At every $v \in \operatorname{Spec} V$ the fibre of G at $s(v)$ is proper after extension to $\kappa(v)$. Properness descends along field extensions and persists under base change, as required by the definition of the abelian locus. Connected components of proper smooth fibres · Lan, Proposition 3.3.1.5 · Definition of the abelian locus

For the normed-field test, the coefficient map identifies V with the full closed unit ball of C. Each nonzero $z \in C$ or its inverse belongs to that ball, so $C = \operatorname{Frac} V$. The same argument applies to this possibly nondiscrete ring and to each complete discrete valuation ring in the other test. Unit-ball hypothesis

The same model also extends each specified prime-to-p torsion marking on its abelian locus. The torus rank is upper semicontinuous. Its zero set U is therefore open, and the semiabelian family restricted to U is an abelian scheme. For $p \nmid n$ its n-torsion is finite étale over U. A marked generic section killed by n gives a rational section of this finite scheme. On a normal integral affine open with ring D, its coordinate homomorphism takes values in $\operatorname{Frac} D$ integral over D, hence in D. These extensions agree on overlaps and with the marking on the generic abelian open by separatedness. Apply this to each of the two prescribed markings. This verifies the extension clause on the selected component with its original labels. Lan, Lemma 3.3.1.4 and Corollary 3.3.1.6 · Marked section extension

Proposition (the primitive theta on the selected integral model). Let $f : G \to X$ be the identity family on the selected Lan model and let L_G be its rigidified trace-one line. Then f_*L_G is invertible, with Zariski-local generators whose restrictions to all geometric fibres are nonzero. In the Raynaud splitting induced by the trace-one line at a boundary point, a generator can be normalized to have a generating weight-zero coefficient and all other coefficients zero on the closed fibre. This holds in residue characteristic two. The same theta line pulls back under the cartesian comparison on the common graph. On a height trait with abelian generic fibre it generates the full section module. On a height trait with positive generic torus rank it spans a saturated line in that module.

Proof. At an abelian point use cohomology and base change for an ample principal line. At a boundary point take the completed strict henselian local ring R, a normal local domain by excellence, with maximal ideal $\mathfrak{m}$ and fraction field K. Its generic point maps to the abelian locus. The actual identity family is the Mumford family on Lan's completed chart, and the Raynaud extension $G^\natural \to B$ over R has split torus and abelian quotient B. Let X_T be the character lattice and Y_{per} the period lattice. The extended Poincaré biextension on the formal group and its dual descends to the Poincaré biextension P_B on $B \times B^\vee$. Functoriality carries the polarization λ and the integral trace-one endomorphism w to λ_B and w_B. Thus the Raynaud line is the pullback of $M_B = (\mathrm{id}_B, \lambda_B w_B)^* P_B$. The chosen trivialization on the torus is pulled back from the rigidification of M_B at zero. The symmetry of the original line comes from simultaneous inversion in the two entries of the biextension. It descends by this same construction to the normalized symmetry of M_B. For $\pi : G^\natural \to B$, the fibre map on $\pi^* M_B$ is pulled back from B, so it is constant along each torus orbit. Since $[-1](tg) = t^{-1}[-1](g)$ for torus elements t, a weight χ is sent to $-\chi$ with no character shift. Lan, Theorems 10.13 and 11.2 · Lan, Propositions 3.3.3.11 and 3.4.4.1

The polarized cross terms for M_B give λ_B by the trace-one identity. The degree formula $1 = [X_T : \phi(Y_{\mathrm{per}})]^2 \deg \lambda_B$ makes the polarization on B principal and $\phi : Y_{\mathrm{per}} \to X_T$ an isomorphism. The symmetry and generic ampleness of L_G, together with the displayed descent to M_B, give the hypotheses for Lan's ample degeneration theorem. The split character decomposition gives coefficient modules $E_\chi = H^0(B, M_B \otimes \mathcal{O}_\chi)$, where $\mathcal{O}_\chi$ is the weight-χ line from the torus torsor. Each E_χ is free of rank one and commutes with base change by abelian Riemann–Roch and vanishing for ample lines, also in characteristic two. Lan's Fourier expansion over K is injective. The period recurrence connects all weights since ϕ is an isomorphism, so the zero coefficient of a nonzero generic theta is nonzero. Scale the generic section h so that its weight-zero coefficient b generates E_0. Lan, Lemma 3.4.4.3, Theorem 4.2.1.14 and §4.3.1

Let J_y be the fractional ideal of the coefficient at $\phi(y)$ in its free module, defined using any basis of that module. By the recurrence at weight zero this is the ideal of the multiplier $\psi(y)$ in its integral value line. Indeed translation of b is a generator under the integral translation isomorphism, which identifies $T^*_{c^\vee(y)} M_B$ with $M_B \otimes \mathcal{O}_{\phi(y)}$ tensored by the value line $c^\vee(y)^* M_B$. The normalized inversion acts on the generic theta by a square-one scalar and interchanges the integral coefficient modules at $\phi(y)$ and $-\phi(y)$, so $J_y = J_{-y}$. The cocycle at $(y, -y)$ gives $J_y J_{-y}$ equal to the ideal of the biextension multiplier $\tau(y, \phi(y))$. Positivity puts this ideal in R, and in $\mathfrak{m}$ for $y \neq 0$. Since J_y is principal over the local ring R, normality places its generator in R and in $\mathfrak{m}$ when $y \neq 0$. This square-ideal calculation applies at two without dividing by two. Lan, Definition 4.2.1.10 and equations (4.2.3.4) and (4.3.1.7)

To pass from coefficients to sections, choose $0 \neq d \in R$ with dh algebraic over R by localization of sections on the quasi-compact separated group. Its formal Fourier expansion has finite support modulo each power of $\mathfrak{m}$. All its coefficients lie in dE_χ. Artin–Rees gives, for each n, a single j with $dR \cap \mathfrak{m}^j \subset d\mathfrak{m}^n$. The same containment holds in every free coefficient module, so the expansion after division by d still has finite support modulo $\mathfrak{m}^n$. It defines a formal section with reduction $\pi_s^* b_s$ on the closed fibre, which is nonzero. At every point of that fibre, divisibility of dh by d in the completed local ring descends by faithful flatness. Thus h is regular there. A codimension-one pole on the normal group would lie over a height-one prime of R and be the generic point of the integral inverse image of that prime's closure. This inverse image contains the closed fibre, ruling out such a pole. Hence h is algebraic on G_R. Lan, §4.3.1 · Flat base change for sections

Every integral section over R is a scalar multiple of h over K. At a point above the closed point where h is a frame, faithful flatness of the smooth local ring over R forces the scalar to belong to R. Flat base change of sections and faithfully flat descent from R give a generator at the original base stalk. Represent it on a Zariski neighbourhood and shrink to the open image of its nonvanishing locus. The same scalar argument at each point proves generation there and throughout its subopens. This proof applies on the common graph and on any such normal model with the identified semiabelian group and trace-one line. It gives the local generation and fibrewise nonvanishing clauses used in the release. Pinned direct-image criterion

On a height trait with valuation ring V and fraction field Q, let h_V be the pullback of a model generator. If a section is generically $a h_Q$ with $a \in Q \setminus V$, multiplying by a^{-1} would force h_V to be zero on the closed fibre. Thus the sections in $Q h_Q$ form exactly $V h_V$. This is the entire section module when the generic fibre is abelian. With positive generic torus rank the generic open semiabelian fibre has infinitely many independent weights after splitting the torus, with one-dimensional abelian coefficient spaces. The pulled-

back line gives the saturated lattice in Qh_Q even in that case. A generating Fourier coefficient is a unit in a base frame of its coefficient module. A value at a prescribed point may still vanish.

Which height terms depend on a marked section. Write $s = ah$ once the given section has descended to the same rational theta line as a primitive local generator h. Each nonzero coordinate ratio t_j/t_i on a projective chart is invariant under this rescaling. At a finite place the differences $v(t_j) - \min_i v(t_i)$ are likewise invariant. At all places together the product formula cancels the scale of a single algebraic vector in a homogeneous height. The modified cotangent is generated modulo p^r by the differentials of the marked Weil characters with value 1 at the identity. It depends on the polarized group and the marking. In the intrinsic point-height proof one chooses an integral section after a finite extension and takes the complex norms of that same section. Thus these steps require the given metrics and coordinate comparisons, and allow a choice of primitive theta generator. Projective coordinate ratios · Marked cotangent model · Choice of the section in the height estimate

For an absolute value in a prescribed torsion frame e, a nonzero evaluation satisfies $v(s(x)/e) = v(h(x)/e) + v(a)$. In the symmetric Raynaud splitting above, the weight-zero coefficient of s is a times the generating coefficient of h. It generates the same integral coefficient line precisely when a is a local unit. At positive abelian rank this tests a section of the abelian theta line, whose value at an individual point can vanish. If the prescribed normalization includes a divisor, the comparison instead requires the specified height-one orders of a and the pullback of that Cartier divisor to the trait. For example, multiplication by a uniformizer preserves every projective coordinate ratio and raises the order of every nonzero evaluation in a fixed frame by one. The cusp calculation still requires the actual boundary multiplicities and its metric. Scalar in a prescribed frame · Boundary local equation

Extension of the polarization on the height traits. Let R be a DVR and G/R the specified smooth separated semiabelian group with abelian generic fibre. A rigidified line on the generic abelian variety extends to a rigidified line on G. Indeed, close a divisor representing the generic line in G. Smoothness over the regular base makes the total space regular, so the resulting Weil divisor is Cartier. Tensor the extension with the inverse of its pullback at the identity, pulled back from the base. The specified generic rigidification then identifies this normalized extension with the original line. For two such extensions the divisor of their rigidified generic isomorphism is $a[G_s]$. The special fibre G_s is geometrically integral and reduced, so it is the divisor of a uniformizer with multiplicity one. Removing that power of the uniformizer gives a regular isomorphism. Its restriction to the identity shows $a = 0$ because the original isomorphism preserves the rigidifications.

The same proof on G^k extends the normalized inversion, cube, adjoint-action and primitivity isomorphisms. Their identities hold on the dense generic fibre and hence everywhere. This applies over a DVR of residue characteristic two and commutes with passage to the completed algebraic closure used in the Hodge test by pulling back the resulting data. On a normal locally noetherian integral chart base S', assume an invertible sheaf extending the chosen line on the abelian open and the compatible semiabelian group and endomorphisms. Rigidify the sheaf along the identity as above. Each required generic isomorphism is a rational section of an invertible sheaf on a smooth fibre power over S'. At a horizontal divisor it is already regular and invertible. A vertical codimension-one point lies over a height-one point of S' by flatness and going-down. Over its DVR the preceding argument shows its order is zero. Normality of the fibre power then extends both the section and its inverse across the whole space. This proves the structural identities once the underlying line has been extended on the selected normal model. Integral extension over a trait · Cubical structure over a normal base · Adjoint and primitivity identities

The polarized torsion on a completed chart trait. For a height-one prime $\mathfrak{q} \subset A$ over ℓ, let C be the completion of an algebraic closure of the fraction field of the completed DVR $A^{\wedge}_{\mathfrak{q}}$, and $O = \mathcal{O}_C$. The map $A \hookrightarrow O$ induces $K \hookrightarrow C$, so the test has abelian generic fibre even when its centre is on the boundary. Extend $K \hookrightarrow C$ to $\overline{K}$ and transport the marking, preserving the prescribed points and root at level n. Write G/O for the pullback and $M = \omega_G = \omega \otimes_A O$. The generic labels at level m are $T_m = \Lambda/\ell^m\Lambda$, with $h = \mathrm{rk}_{\mathbb{Z}}\,\Lambda = 2\,\mathrm{rk}_O\,M$. The transported primitive roots ζ_m satisfy $\zeta_{m+1}^{\ell} = \zeta_m$, and the pairing on the generic torsion is $\beta_m(a,b) = \zeta_m^{B(a,b)}$. Marked chart trait · Compatibility of the finite marking

Let $Q_m = G[\ell^m]$. Flat multiplication on the smooth semiabelian model gives separated quasi-finite flat kernels. Over the henselian O, each Q_m has a finite open U_m containing the whole closed fibre. Indeed Zariski's Main Theorem embeds Q_m as an open in a finite scheme, whose local factors with closed point in the open are entirely in it. A map from a finite O-scheme into Q_m lands in U_m, since its inverse image contains every closed point of the source. Apply this to the identity, inversion and multiplication to obtain a finite flat subgroup. Its generic labels $J_m \subset T_m$ are exactly the points extending in Q_m. Via evaluation write $B_m = \mathcal{O}(U_m) \subset O^{J_m}$. The full integral order of the marked torsion is

$$A_m = \{f \in O^{T_m} : \ (b \mapsto f(t+b)) \in B_m \text{ for every } t \in T_m\}.$$

Choosing representatives identifies A_m with one copy of B_m for every coset of J_m. Changing representatives translates by integral sections of U_m. On a pair of cosets the group law is addition in U_m followed by translation by the sum of the representatives minus the representative of their sum. This gives the Hopf algebra with generic group T_m and rank ℓ^{mh}. The original kernel consists of U_m together with the generic-only points in the other cosets. Its ℓ-adic completion along the special fibre retains U_m. The full order supplies finite flat models of all the generic cosets while agreeing with Q_m near the identity. Zariski's Main Theorem · Finite open subgroup · Order on the cosets

The label inclusion $T_m \hookrightarrow T_{m+1}$ is multiplication by ℓ on representatives in Λ, and multiplication by ℓ down the tower reduces labels. The finite-source property gives $U_m = \ker([\ell^m] : U_{m+1} \to U_{m+1})$. Also $J_{m+1} \to J_m$ is surjective under $[\ell]$. The lift of an extending section is sought in the quasi-finite flat inverse image under $[\ell]$, with nonempty closed fibre by semiabelian surjectivity. A nonempty finite flat open there has a C-point, and properness extends that point over O. The inclusion on full orders is a closed immersion of groups because restriction surjects on the factors meeting the image, using $U_m \subset U_{m+1}$. On a coset of J_{m+1} represented by t, the kernel of $[\ell^m]$ is empty if $\ell^m t \notin J_{m+1}$. Otherwise $\ell^m t$ comes from J_1. Choose $u \in J_{m+1}$ with $\ell^m u = \ell^m t$ and use the representative $t - u$, which is killed by ℓ^m. The kernel in that coset is then precisely a translate of the scheme U_m. Thus $\mathrm{Spec}\, A_m$ is the scheme-theoretic kernel of $[\ell^m]$ at the next level. The ranks and these inclusions define an ℓ-divisible group H with $T_\ell H = \Lambda \otimes \mathbb{Z}_\ell$. Integral division · Kernel on the full order · Divisible group and Tate module

The cubical commutator gives a regular unit on $Q_m \times Q_m$ whose generic values are β_m. In particular it gives $\beta_m(b, b')$ on $U_m \times U_m$. On two cosets the pairing is obtained by multiplying this function by $\beta_m(t,s)\beta_m(b,s)\beta_m(t,b')$ for representatives t, s and $b, b' \in J_m$. The one-variable characters extend by properness of the finite flat Cartier dual of U_m. Hence the product belongs to $A_m \otimes_O A_m$ and satisfies the pairing and level identities there by flatness. It induces an alternating $\lambda : H \to H^D$. The map on Tate modules is an isomorphism by unimodularity, and the two Hodge–Tate subspaces have equal dimension $h/2$ by the cotangent calculation below. The classification over O by the Tate module and its Hodge–Tate

subspace gives an isomorphism of divisible groups. Regular cubical pairing · Pairing on the added cosets · Tate perfection · Scholze–Weinstein, Theorem 5.2.1

At the identity, on a square-zero extension the differential of $[\ell^m]$ is multiplication by ℓ^m. Because $H[\ell^m]$ and Q_m share the identity neighbourhood, the augmentation ideal $I_m \subset A_m$ has $I_m/I_m^2 \simeq M/\ell^m M$, with restriction given by reduction. Inverse limit gives $\omega_H = M$. The character of the Tate vector marked by a_x under λ is $b \mapsto \zeta_m^{B(a_x,b)}$. At level n, the base change of the chart kernel is Q_n as a scheme, and the character on the full order restricts near the identity in Q_n. There it agrees with the base change of f_x by equality on the marked generic points and flatness. The two functions have the same linear term in the common identity cotangent. Writing $\tau(a_x)$ for the marked Tate vector gives

$$\operatorname{dlog}(\lambda(\tau(a_x))) \equiv w_x \otimes 1 \pmod{\ell^n M}.$$

Both cotangent actions are precomposition with $d[\bar{u}]$, by base change and the adjoint relation. Identity cotangent and level compatibility · Kernel base change and action on the gradients

The Tate curve with parameter q_T of positive valuation exhibits the added directions. For its semiabelian identity model the generic torsion classes are $u q_T^{i/\ell^m}$ modulo $q_T^{\mathbb{Z}}$, with $u \in \mu_{\ell^m}$ and $i \in \mathbb{Z}/\ell^m\mathbb{Z}$ after choosing compatible roots of q_T. The component map is valuation modulo $v(q_T)$, so the classes extending in the identity component have $i = 0$. The finite open is μ_{ℓ^m} of rank ℓ^m while the generic kernel has order ℓ^{2m}. The full order supplies one copy of μ_{ℓ^m} for each i and gives $H = \mu_{\ell^\infty} \times \mathbb{Q}_\ell/\mathbb{Z}_\ell$ with the chosen roots. With pairing $((u,i),(v,j)) \mapsto v^i u^{-j}$, the differential in the identity neighbourhood of the second variable is $i\,dV/V$ modulo ℓ^m, where V is the multiplicative coordinate. The added direction $i = 1$ supplies a generator of the Hodge–Tate image. Silverman, Tate uniformization and reduction · Polarized torsion on the test trait

The finite lattice comparison. The marked generators span the completed Tate module over $\mathcal{O}_E \otimes \mathbb{Z}_\ell$. Normalize $v_\ell(\ell) = 1$. Fargues's integral Hodge–Tate theorem puts $\kappa_\ell M$ in the image of the polarized map when $v_\ell(\kappa_\ell) \geq 1/(\ell-1)$. Take $\kappa_\ell = \zeta_1 - 1$ at odd ℓ and $\kappa_2 = 4$, of valuations $1/(\ell-1)$ and 2, respectively. On the chart, $\zeta_1 = \zeta_n^{\ell^{n-1}}$. After extending scalars to C, exactness and the polarized adjoint relation also identify the kernel with its orthogonal, and clearing denominators gives the integral annihilator identity. The gradient congruence puts the Hodge–Tate image in $\omega'_\ell \otimes_A O$, since the reductions of the gradients for $x \in \mathcal{O}_E^r$ account for the marked Tate classes modulo ℓ^n and the lattice contains $\ell^n M$. The extension $A_{\mathfrak{q}} \to O$ is faithfully flat because it is torsion-free over the DVR and local. Apply faithful flatness to $(\omega/\omega'_\ell)_{\mathfrak{q}}$ to descend the vanishing of the class of $\kappa_\ell w$ for each $w \in \omega$. At height one away from ℓ, ℓ^n is a unit and the lattices coincide. Normality of A and projectivity of ω'_ℓ give its equality with the intersection of its height-one localizations in the generic cotangent. Hence $\kappa_\ell \omega \subset \omega'_\ell \subset \omega$ on the chart. The projective refinement gives this projectivity on the tied charts of X. Fargues, §5.3.2, Theorem 2 · Height-one descent

Specialization to the height trait. Let $R = \mathcal{O}_{L,v}$ for a height point over $\mathcal{O}_L$, with $v \mid \ell$. The map to X factors through a chart containing its centre, and its generic point can lie on a horizontal boundary. The chart map to R can have a kernel. The regular function f_x pulls back to the multiplication kernel and its differential remains the reduction of the pulled-back w_x. For a semiabelian variety of dimension d and toric rank $t > 0$ in characteristic zero the torsion at level n has order $\ell^{n(2d-t)}$, so the height trait uses the evaluated lattice inclusion from the chart. The height-one test above uses the full abelian generic marking.

Let Λ_L be the integral Lie lattice of the whole semiabelian fibre and let V_s be its fraction-field eigenspace for $\sigma_s : E \hookrightarrow L$. Put $A_s = \Lambda_L \cap V_s$ and $A_{s,v} = A_s \otimes_{\mathcal{O}_L} R$. The embedding values on $\mathcal{O}_E$ are integral in the chart stalk by normality and evaluate to the given label. The trace-dual basis gives the generic character projector

$$P_s = \sum_i \sigma_s(a_i^\vee)\rho(a_i), \qquad \sum_i \sigma_s(a_i^\vee)\sigma_u(a_i) = \delta_{su}.$$

Here a_i is an integral basis of $\mathcal{O}_E$ and $a_i^\vee$ is its trace dual. A nonzero rational integer in the different clears all denominators. After inverting that integer the projector is a direct summand of the smooth Lie module of constant rank, and the evaluated projector has image V_s over L. At the centre take a local basis of the restricted image of ω'_ℓ in the generic labelled cotangent and lift it to ordinary cotangent sections. Its dual frame on the generic Lie eigenspace lies in $\kappa_\ell^{-1} L_A$ at the stalk by the chart sandwich. Write k_ℓ for the evaluated value of κ_ℓ. The cyclotomic equation at odd ℓ and the value 4 at 2 give $k_\ell \neq 0$ over L. The dual frame thus evaluates by multiplying its vectors first by κ_ℓ, specializing in the smooth Lie model, and dividing by k_ℓ over L. It spans V_s by the projector rank and the specialized duality identities. The lifts of the restricted cotangent basis pair integrally with $A_{s,v}$, and the evaluated dual vectors lie in $k_\ell^{-1} A_{s,v}$. These vectors span the evaluated dual lattice, since every dual-lattice germ at the centre has regular coefficients in the dual frame. At small labels this is the patched fibre $F_{s,v}$. At large labels the patch allows the prescribed enlargement controlled by the different. Outside S the discriminant is a unit, so the projectors split integrally and the two lattices agree. In particular, for $v \mid \ell$ with $\ell \in S$,

$$A_{s,v} \subset F_{s,v}, \qquad k_\ell F_{s,v} \subset A_{s,v} \quad (s = \mathrm{inr}(\tau)), \qquad k_\ell \sigma_s(\mathfrak{D}_{E/\mathbb{Q}}) R F_{s,v} \subset A_{s,v} \quad (s = \mathrm{inl}(\tau)).$$

The cancellation in the generic expansion is why the different clears denominators in the proof and contributes no multiplier at the small labels. Specialization of the dual frames · Cleared projectors and evaluation · Small and large patches

The values k_ℓ are fixed by evaluation at the generic point of the lift. For odd ℓ the cyclotomic equation gives a primitive ℓ-th root $k_\ell + 1$. All $1 - \zeta_\ell^a$, $1 \le a < \ell$, are associates by geometric sums in both directions, and their product is ℓ. Hence $(k_\ell)^{\ell-1} = (\ell)$ and $(k_2) = (4)$. For the common ideal $\mathfrak{a}_s = \sigma_s(\mathfrak{D}_{E/\mathbb{Q}})\mathcal{O}_L(\prod_{\ell \in S} k_\ell)$, its exponent at $v \mid \ell$ is $h_{s,v} = e_{v/w} d_w + \mathbf{1}_{\ell \in S} e_{v/\ell} c_\ell$, where $w = v|_{\sigma_s(E)}$, d_w is the local different exponent, e denotes ramification indices and $c_\ell = 1/(\ell - 1)$ at odd ℓ, $c_2 = 2$. Thus the exact normalized cost of the common ideal is

$$\frac{1}{[L:\mathbb{Q}]} \sum_{v \mid \ell} h_{s,v} \log |\mathcal{O}_L / v| = \left(\frac{\mathrm{ord}_\ell \, |\operatorname{disc} E|}{[E:\mathbb{Q}]} + \mathbf{1}_{\ell \in S} c_\ell \right) \log \ell, \qquad X_{\mathrm{exact}} = \log \mathrm{rd}\, E + \sum_{\substack{\ell \in S \\ \ell \neq 2}} \frac{\log \ell}{\ell - 1} + 2 \log 2.$$

Here the first equality is the contribution of one rational prime. The identity uses the norm of the different and $\sum_{v \mid \ell} e_{v/\ell} f_{v/\ell} = [L:\mathbb{Q}]$, where $f_{v/\ell}$ is the residue degree. The determinant index at each place is the sum of the elementary-divisor exponents, each bounded by the exponent of the multiplying ideal. The sharper small-label bound therefore costs at most $fq(X_{\mathrm{exact}} - \log \mathrm{rd}\, E)$ across the f labels of rank q. The degree convention divides by $[L:\mathbb{Q}]$, so further extensions preserve the normalized cost. In the product-formula identity for the point the replacement of F_s by A_s affects only these labelled indices, with the toric and archimedean Gram terms computed in the same adapted frame. Normalized cyclotomic norms and different · Determinant comparison in the point height

For the stated constants retain the bound $X_{\mathrm{lat}} = 2\sum_{\ell\in S}\log \ell/(\ell-1) + 2\log \mathrm{rd}\, E \le 18\Lambda_{\mathrm{ht}}$ from the common sandwich. The primes at most b_{ht} cost at most $\log b_{\mathrm{ht}} + 2.2$. The set of extra primes has at most b_{ht} elements, each costing at most $3\log b_{\mathrm{ht}}/b_{\mathrm{ht}}$. The common coefficient 8 for the prime sum gives the released determinant allowance fqX_{lat}. This point-height loss is distinct from the finite order costs in the section count. Uniform prime sum · Finite Hodge–Tate sandwich

For the remaining degree let x_s, x_b be the sums of normalized degrees of the small and conjugate large saturated eigenspaces in the Hermitian Lie bundle $\mathcal{B}$. Put $h = -\deg_{\mathrm{norm}} \mathcal{B}$. Orthogonality gives $x_s + x_b \le -h$. The conjugate degree differences are bounded by their pair ranks times $|\gamma + \log(2\pi) - \log \mathrm{rd}\, E - L'(1,\chi)/L(1,\chi)|$. This is the flat CM degree after the finite different correction. Use $C_f = 1 + \gamma + |\log(2\pi)|$ and let C_s be the coefficient in the maximal-slope bound $C_s \log(2 + \mathrm{rk}\,\mathcal{B} + \max(h,0))$. Put $D_0 = C_s + 10C_f$. Then $x_s - x_b \le D_0 fr\Lambda_{\mathrm{ht}}$. If $x_s > 0$ this forces $h \le D_0 fr\Lambda_{\mathrm{ht}}$. The small ranks are at most q per label, $fr \le e^{\Lambda_{\mathrm{ht}}}$, and $\Lambda_{\mathrm{ht}} \ge 1$. Applying the slope bound gives $x_s \le D_0(2 + \log(3 + D_0))fq\Lambda_{\mathrm{ht}}$, a bound that also holds when $x_s \le 0$. In $\phi\mathrm{Deg}$ the toric index cancels using $[E:\mathbb{Q}] = 2f$. The remaining toric terms for toric rank $r_0 \le q$ are $-fr_0 \log \mathrm{rd}\, E + (fr_0/2)\log(4\pi) \le fq\Lambda_{\mathrm{ht}}$. Hence the source permits $C_h = 19 + D_0(2 + \log(3 + D_0))$. Flat CM degree · Conjugate degree and period bound Period constant · Maximal slope · Degree dichotomy · Combined degree estimate

In the normalized theta estimate put $c_\theta = 1 + \log(72e^4)/(2\log 2)$, from the Gaussian-sum bound, and $C_M = 64(33 + c_\theta)$ from the torsion-mean comparison. The sup log-mean coefficient is $C_\theta = \frac{7}{2} + 2c_\theta + \frac{1}{2}\log(1 + 3C_M)$. The height of the polarized line gives $C_s = 12 + C_\theta(1 + \log 2)$ after short-vector extraction. These formulas give $C_s < 50$, $C_f < 3.5$, and $D_0 < 85$. The expression for C_h is about 556.6, so we may use 600 uniformly. Gaussian constant · Theta normalization · Torsion mean · Sup log-mean constant · Key-degree coefficient · Line estimate Parameter choice

Cusp mass. On a matrix-ball sheet, let λ_1 be the first minimum in the polarized period lattice. Minkowski's theorem and the isotropic rank bound give

$$0 \le U_{\mathrm{cusp}} \le fq\big(2\log^+(1/\lambda_1) + \log(2fq\ \mathrm{rd}\, E)\big).$$

Use $r_* = (2f\ \mathrm{rd}\, E)^{-1/2}$. The $2f$ rationally independent integral elements with conjugates bounded by $2f\ \mathrm{rd}\, E$ and integrality of the trace pairing force the vectors shorter than r_* into a common isotropic span. Packing bounds the count below $t \le r_*$ by $(3t/\lambda_1)^a$ on $\{\lambda_1 < t\}$, where $a = 2fq$. In the arithmetic quotient use the isotropic count with bounded line index. It contains a shortest vector when $\lambda_1 < t \le r_*$. Its average $M(t)$ in the quotient probability measure scales by $M(\theta t) = \theta^\kappa M(t)$, $\kappa = 2f(r-1) \ge 3a$, by unfolding. Decomposing into shells at $r_* 12^{-s}$ gives $M(r_*) \le 36^a + M(r_*)/2$. The first shell uses the packing bound and the later shells use the indicator of $\lambda_1 < r_* 12^{-s}$ bounded by the count there. The resulting coefficients sum to at most $\sum_{s\ge1}(36\,12^{-2s})^a \le 1/2$. Consequently

$$\Pr(\lambda_1 < r_* e^{-x}) \le \min(1, 2\cdot 36^a e^{-\kappa x}), \qquad \mathbb{E}\log^+(r_*/\lambda_1) \le \frac{\log(2\cdot 36^a)+1}{\kappa} \le 3.$$

The mass bound on a sheet is therefore $fq[6 + \log(2f\ \mathrm{rd}\, E) + \log(2fq\ \mathrm{rd}\, E)]$. The Hodge measure and the cusp integral use the same sheet degrees and mass normalization, so the bound holds for their ratio. Both logarithms are at most $\log(b_{\mathrm{ht}}\ \mathrm{rd}\, E)$ and $6 \le 2\log(b_{\mathrm{ht}}\ \mathrm{rd}\, E)$. This gives $C_u = 4$. The analytic estimate uses the root discriminant and the cutoff. The logarithmic L-discrepancy enters through the point-

height estimate. First minimum · Integral elements · Bounded-index count and its scaling · Shell estimate Tail integral · Sheet measure Normalized mass bound

Why the count scales on each sheet. Set $d = 2f$ and $v = d \operatorname{rd} E$. If y realizes the first minimum, the d short independent integral multipliers give vectors in $\mathcal{O}_E y$ of length at most $v\lambda_1$. Hadamard bounds their covolume by $(v\lambda_1)^d$. A ball of radius $\lambda_1/2$ injects into the torus of the line lattice $(Ey) \cap \Lambda$ and contains an open cube of volume $(\lambda_1/\sqrt{d})^d$. Thus $[(Ey) \cap \Lambda : \mathcal{O}_E y] \le (v\sqrt{d})^d$. There are finitely many arithmetic classes of isotropic lines. On each line there are finitely many sublattices of this bounded index. Generators of a principal sublattice differ by a unit. A congruence subgroup of the units acts integrally by scaling the two vectors of a rational hyperbolic pair by u and $\bar{u}^{-1}$, since the two rational projectors have a common denominator. The counted vectors therefore comprise finitely many orbits, also at finite-index level. Index for a shortest vector · Arithmetic vector orbits

Unfold one orbit with representative y to the quotient by its vector stabilizer H. A rational Witt partner identifies H with a semidirect product of its Heisenberg unipotent radical and $U(W)$ for $W = y^{\perp}/Ey$. The action of $g \in U(W)$ on the radical has determinant $|\det_{\mathbb{C}} g|^2 = 1$ at each real place, since it acts on the underlying real space of W and fixes the central direction. In particular H is unimodular. A rational character is trivial on the unipotent radical, and every character of the unitary quotient over the totally real field is trivial because its determinant takes values in the anisotropic norm-one torus. This also gives trivial rational characters after restriction of scalars. The arithmetic lattice theorem gives finite volume to $(\Gamma \cap H)\backslash H$ for the selected sheet group Γ. With the compatible quotient measure, integration along this fibre multiplies the orbital integral by that volume. This is an identity for the sheet average at every radius. On one cone $Q(z) = h(z,z) = 0$ in $\mathbb{C}^r$ the invariant measure is proportional to the coarea density $dS/\|\nabla Q\|$, because the unitary group preserves both Q and Lebesgue volume. Dilation by θ multiplies this measure by θ^{2r-2}. The link on the unit sphere is compact and smooth, so the measure of a bounded ball is finite. Over the f places the degree is $\kappa = 2f(r-1)$. The finite orbit sum gives $M(t) = Ct^{\kappa}$ with $C < \infty$ for the sheet, also after the fixed rescaling of its majorant norm. This proves the finiteness used when absorbing the shell sum above. Arithmetic lattice theorem · Unfolded integral · Orbit sum

On the selected ranks the previous estimate even gives $\operatorname{mean}_{\mu} U_{\mathrm{cusp}} \le 2fq \log(b_{\mathrm{ht}} \operatorname{rd} E)$. Indeed $b_{\mathrm{ht}} \ge 4fr^2$, $r \ge 3q+1$ and $r \ge 4$ imply $b_{\mathrm{ht}}^2/(4f^2q) \ge 12r^3 > e^6$. This absorbs the additive 6 in the sheet bound. We use the released choice $C_u = 4$ in the following intersection calculation.

Stokes on the resolved component. Put $a_0 = -C_h f q \Lambda_{\mathrm{ht}}$, $\epsilon = \alpha/\beta$, and let M_X be the common smooth compact complex resolution of the degree setting of a. Write π for the reading from multihomogeneous coordinates z to M_X. Use smooth reference norms $h_{\mathrm{ref}}, h_{\mathrm{corr}}$ of the Hodge multidegree with scalar perturbations $\varphi_t, \psi_{t'}$ on M_X. On the open chart locus their limits satisfy $\log h_{\mathrm{ref}} + \varphi \circ \pi = \phi_{\mu}$ and $\log h_{\mathrm{corr}} + \psi \circ \pi = \phi_{\mu} + U_{\mathrm{cusp}}/2$, with ψ bounded after subtracting the smooth reference. If V_b is the product of the vertical factors for $S \setminus S_0$ at $b = b_{\mathrm{ht}}$, the norms on the integral tensor power are

$$h_N = (h_{\mathrm{ref}} e^{\varphi_t \circ \pi})^{\beta} V_b^{\beta} h_{A'}^{\alpha}, \qquad h_C = (h_{\mathrm{corr}} e^{\psi_{t'} \circ \pi})^{\beta} V_b^{\beta} h_{A'}^{\alpha}.$$

Both lines have the same finite data and multidegree, so the normalized log ratio is

$$w = \beta^{-1}(\log h_C - \log h_N) = \log(h_{\mathrm{corr}}/h_{\mathrm{ref}}) + (\psi_{t'} - \varphi_t) \circ \pi.$$

The common reference multidegree makes the logarithmic ratio descend to M_X. For curvature on the chart locus, a vertical factor in a local holomorphic section is the absolute value of a nonvanishing holomorphic function obtained from the rational vertical frame and the block chart monomial. Its log is locally the real part of a holomorphic logarithm. A change of holomorphic section contributes the same type of pluriharmonic terms from the block coordinate ratios. These terms have zero Levi form. Thus the normalized reference curvature is $\Omega_N = \omega_t + \epsilon\eta$, with $\omega_t = dd^c(\theta + \varphi_t)$ for the smooth reference potential θ and $\eta = c_1(A')$. The arithmetic intersections retain the vertical factors, and the chart integration formula identifies the curvature integrals off the analytic null sets. Metrics and normalization · Vertical log factors and chart potentials · Chart integration formula

For fixed α, β, take $t > 0$ small enough that the positive curvature of A' absorbs the reference Levi error. For sufficiently small $t' > 0$ the corrected metric is smooth, conjugation invariant and semipositive, with fibrewise nef underlying line. Let T_N, T_C be the two top arithmetic intersections divided by β^{n_X+1}, and let d_ϵ be the common geometric degree divided by β^{n_X}. Equality of generic and geometric degree makes d_ϵ independent of the smooth metrics at the fixed ratio. Zhang's lower successive-minima inequality gives $T_C \geq (n_X + 1)(a_0 - \epsilon K_A)d_\epsilon$. Arithmetic metric change in one factor at a time gives

$$T_C - T_N = \sum_{j=0}^{n_X} \int_{M_X} w\, \Omega_C^j \wedge \Omega_N^{n_X - j}.$$

Here Ω_C is the corrected normalized curvature and $dd^c w = \Omega_C - \Omega_N$ in the log-norm convention. Replacing one reference factor by a corrected factor changes its integral by $\int w\, dd^c w \wedge F = -\int dw \wedge d^c w \wedge F \leq 0$, since the remaining mixed product F is closed and positive. Hence $T_C - T_N \leq (n_X + 1) \int w\, \Omega_N^{n_X}$. Reference Levi bound · Corrected positivity and Zhang normalization · Zhang, Theorem 5.2 · Gillet–Soulé, arithmetic metric change · Integration by parts · Stokes telescope

First let $t' \downarrow 0$ at fixed t and ϵ. The smooth descended reference ratio and the fixed smooth φ_t are bounded by compactness. The bounded regularizations $\psi_{t'}$ have a uniform bound for small t'. This gives a constant majorant against the fixed smooth finite measure $\Omega_N^{n_X}$. Away from the boundary, the limit of w is $U_{\mathrm{cusp}}/2 + \varphi - \varphi_t$. The tied log-log regularization gives $\varphi_t \geq \varphi$ there. Dominated convergence and positivity give

$$T_N(t, \epsilon) \geq (n_X + 1) \left[(a_0 - \epsilon K_A) d_\epsilon - \int_{M_X} \frac{U_{\mathrm{cusp}}}{2} \, (\omega_t + \epsilon\eta)^{n_X} \right].$$

The point-height threshold is common to all W once the setting and auxiliary integers are fixed. The comparison requires convergence of the ratio integral at fixed t, with a uniform lower bound for T_C. Bounded convergence and upper bound on the difference · Fixed-reference majorant

Next let $t \downarrow 0$ at fixed ϵ. Write M_X° for the divisor complement in the tied charts. A positive boundary parameter $\rho \geq 1$ there is a bounded smooth term plus $e + \sum_j -\log|z_j|^2$ in normal-crossings coordinates. Its sublevel sets exhaust M_X° by compact subsets. The tied truncation, convolution and gluing give $-C \log \rho \leq \varphi \leq \varphi_t \leq C$ and $\varphi_t \to \varphi$ in C^2 on compact subsets. On the same open set, $U_{\mathrm{cusp}}/2 = \psi + \log(h_{\mathrm{corr}}/h_{\mathrm{ref}}) - \varphi$ is measurable and has absolute value at most $C'(1 + \log \rho)$. This is locally integrable against any fixed smooth chart density. The varying densities require a uniform tail bound. Normal-crossings growth data · Convergence of the potentials and derivatives · Measurable logarithmic majorant

For the needed mixed curvature products, add a fixed positive smooth form to each factor as needed to obtain semipositivity. Multilinearity expresses the original density as a finite signed sum of positive products $d\nu_t$. On those products one has, at fixed ϵ and for all sufficiently small t,

$$\nu_t(\rho \geq R) \ll \frac{(\log R)^{n_X}}{R}, \qquad \int_{\rho \geq R} \log \rho \, d\nu_t \ll \frac{(\log R)^{n_X+1}}{R} \qquad (R \geq e),$$

with the tails taken inside M_X°. To prove the first bound, dominate the smooth backgrounds by a Kähler form Ξ and shift the perturbations down by their uniform upper bound. After changing to the common background Ξ by positivity, truncate those perturbations by smooth maxima at depth $O(\log R)$. They agree up to constants on $\{\rho < R\}$, and compact Stokes preserves the total mass of their product, so the tail mass can be bounded with the truncated potentials. Scaling these to $[-1, 0]$ costs $O((\log R)^{n_X})$ by positivity and expansion. The boundary parameter agrees up to a bounded smooth term with $\log \|s_\partial\|^{-2}$ for a metrized section cutting out the divisor. For a fixed small $c > 0$, take a smooth maximum of $-c\rho$ and $-cR$, with the bounded smoothing adjustment and constant extension at the divisor. Off the divisor $dd^c\rho$ is smooth and bounded above by a multiple of Ξ. Hence the resulting function w_R has $dd^c w_R \geq -\Xi$ and $-w_R \gg R$ on the tail. The integral of $-w_R$ against Ξ^{n_X} is bounded uniformly because the log norm of the section is integrable. For a normalized perturbation $-1 \leq v \leq 0$, integrating $(-w_R)dd^c v$ against the other positive factors gives $(-v)dd^c w_R$ against them. Its integral is bounded above by the integral of Ξ against those factors, because $dd^c w_R + \Xi$ is positive and the integral of $dd^c w_R$ against the closed product is zero. Replacing the normalized factors one at a time bounds the weighted integral by the smooth-background integral plus $n_X \int \Xi^{n_X}$. Dividing by R proves the mass bound. Integration in layers with dR/R, including the term $\log R \, \nu_t(\rho \geq R)$, gives the second bound. Expansion into positive products · Truncating the potentials and transferring the tail · Smooth tail test and energy bound · One-factor integration by parts · Mass and logarithmic tail bounds

On compact sublevels the densities converge uniformly and the weight is bounded. The logarithmic tail bound controls the missing integral uniformly, and passes to the limits of the positive products by Fatou. Summing over the finite chart partition proves convergence with the cusp weight to $\int_{M_X^\circ} (U_{\mathrm{cusp}}/2)(\omega + \epsilon\eta)^{n_X}$, where $\omega = dd^c(\theta + \varphi)$ on M_X°. The same argument applies to the geometric degree using the weight 1. For the arithmetic intersection, expand in powers of ϵ with coefficients given by mixed products of the twisted uncorrected line at t and A'. To take a coefficient limit, replace the varying factors one at a time by a fixed smooth twisted reference line at t_0. Each metric change has weight $\varphi_{t_0} - \varphi_t$ in the reference-minus-varying orientation. Its absolute value is $O(1 + \log \rho)$ and it converges uniformly on compact sublevels. The mixed curvature densities are formed with $\omega_t, \omega_{t_0}, \eta$, so the same exhaustion proves convergence of those integrals. The release also obtains the pure Hodge limit, $\mathrm{arLim}(b_{\mathrm{ht}})$, by comparison with a fixed smooth Fubini–Study metric of the same twisted multidegree. Chart exhaustion and integrable tails · Cusp integral on the same model · Mixed arithmetic coefficients · Pure Hodge limit

Finally take $\alpha = 1, \beta = s + 1$, and send $s \to \infty$ in the inequality obtained after the first two limits for each s. The limiting arithmetic top intersection, cusp integral and degree are polynomials in ϵ, of degrees at most $n_X + 1, n_X, n_X$, respectively. Equivalently the source extracts their coefficients at finitely many distinct positive $1/\beta$ by Vandermonde inversion. The arithmetic constant coefficient is $\mathrm{arLim}(b_{\mathrm{ht}})$ and the degree constant coefficient is $d_\mu = \int d\mu_{\mathrm{H}}$. Indeed $d\mu_{\mathrm{H}} = \omega^{n_X}$ on the divisor complement, by the uncorrected chart potential and the vertical log-unit calculation. The cusp constant coefficient is therefore

$\frac{1}{2}\int U_{\mathrm{cusp}}\,d\mu_{\mathrm{H}}$. The factor $1/2$ is the coefficient of U_{cusp} in the limiting corrected log norm. The auxiliary penalty ϵK_A disappears for the fixed degree setting, giving

$$\frac{\operatorname{arLim}(b_{\mathrm{ht}})}{(n_X+1)d_\mu} \ge a_0 - \frac{\int U_{\mathrm{cusp}}\,d\mu_{\mathrm{H}}}{2d_\mu}.$$

The smoothing thresholds and tail constants can depend on the setting and the fixed auxiliary ratio. The witnesses C_h, C_u remain uniform. Uncorrected mass and cusp numerator · Polynomial coefficients and constant terms · Order of the limits and thresholds · Regularized Stokes limit

The two polynomial estimates. Set $P_{\mathrm{test}} = \mu^{n_X+1}\mathrm{fibClass}(p)$ with $\mathrm{fibClass}(p)$ the product of the within-block root differences, $J_{\mathrm{test}} = \deg_{\mathrm{Fl}}(\mu^{n_X}\mathrm{fibClass}(p)) > 0$, and $K_{\mathrm{test}} = (n_X+1)J_{\mathrm{test}}$.

The projection formula on the selected relative flag presentation gives $\delta\,\mathscr{T}.\mathrm{PPc}(b_{\mathrm{ht}}, a, P_{\mathrm{test}}) = F_{\mathrm{rel}}\operatorname{arLim}(b_{\mathrm{ht}})$ and $\delta\,v_a J_{\mathrm{test}} = F_{\mathrm{rel}}\,d_\mu$. Both identities use the generic point of a, the same positive normalization δ and the same fibre intersection F_{rel}. Since $v_a J_{\mathrm{test}} > 0$ and $d_\mu \ge 0$, both F_{rel} and d_μ are positive. Dividing cancels the common fibre factor. Applying the cusp-mean bound gives the first estimate below. Arithmetic and geometric projection on the selected component

$$\frac{\mathscr{T}.\mathrm{PPc}(b_{\mathrm{ht}}, a, P_{\mathrm{test}})}{v_a K_{\mathrm{test}}} \ge -(C_h + C_u/2) f q \Lambda_{\mathrm{ht}}, \qquad \frac{h h_{\mathrm{Fl}}(P_{\mathrm{test}})}{K_{\mathrm{test}}} \le \frac{fq}{2}\log(r+1).$$

For the upper bound on the fixed flag functional, push the within-block factors down to the product of split Grassmannians. On a dense set use graph matrices with roots of unity for entries. The identity minor yields a primitive integral Plücker vector at finite places. Cauchy–Binet and Hadamard give the squared norm bound $\det(I_q + ZZ^*) \le (1+p)^q$ at each complex embedding. The upper inequality for the essential minimum gives the height bound with the stated normalization. Test polynomial and degree · Grassmannian projection · Integral graph points and Zhang upper bound · Flag height conversion Hadamard estimate

The boundary-mean constant. Put $C_* = C_h + C_u/2$. After division by total volume the exact cutoff identity on P_{test} is $\sum_a \mathscr{T}.\mathrm{PPc}(b_{\mathrm{ht}}, a, P_{\mathrm{test}})/\sum_a v_a = h h_{\mathrm{Fl}}(P_{\mathrm{test}}) - fqK_{\mathrm{test}}\beta_{\mathrm{ht}} + R_{\mathrm{ht}}$. The shift test vanishes, the total test equals qK_{test} per label, and the good-prime remainder above b_{ht} satisfies $|R_{\mathrm{ht}}| \le 1.6K_{\mathrm{test}}$. Thus $\beta_{\mathrm{ht}} \le C_*\Lambda_{\mathrm{ht}} + \log(r+1)/2 + 1.6/(fq)$. The determinant test bounds $|pA_{\mathrm{ht}} + r\beta_{\mathrm{ht}}| \le 4r\Lambda_{\mathrm{ht}}$ for the shift coefficient A_{ht}. Since $r/p \le 2$ and $\log b_{\mathrm{ht}} \ge \max(1.6, \log(r+1))$, this gives $A_{\mathrm{ht}} \ge -[2(C_*+4)+3]\Lambda_{\mathrm{ht}}$. Exact cutoff comparison · Shift and total Frobenius remainder · Determinant test Coefficient bound

For each label τ let $s_\tau = \sum_{i<p} w_{\tau i}$, $t_\tau = \sum_i w_{\tau i}$, and $a_\tau = \sum_i |w_{\tau i}|$. Dominance gives $0 \le s_\tau - pt_\tau/r \le a_\tau$ and $|t_\tau| \le a_\tau$. The coefficient contribution after division by $(D+1)G(w)$ and the total volume is the sum of

$$A_{\mathrm{ht}}s_\tau + \beta_{\mathrm{ht}}t_\tau = A_{\mathrm{ht}}(s_\tau - pt_\tau/r) + (pA_{\mathrm{ht}} + r\beta_{\mathrm{ht}})t_\tau/r.$$

The two coefficient bounds give a lower bound $-[2(C_*+4)+7]\Lambda_{\mathrm{ht}}\|u'\|_1$ for this sum. The flag height term is nonnegative for w, and the normalized remainder for $\mathrm{linRoot}(w)^{D+1}$ has absolute value at most $2qP(u')$. Together they give the T2 coefficient $C_{\mathrm{tower}} = 2C_* + 17 \le 1221$. The cost $2A\|u'\|_1$ from trace and denominator clearing in the section count adds at most 120 to the coefficient of the boundary error. Since $\Lambda_{\mathrm{ht}} \le \log B_{\mathrm{ht}} + \ell$, the constant $C_{0,\mathrm{ht}} = C_{\mathrm{tower}} + 120 = 2C_* + 137 \le 1341$ suffices for X in (16b), with $TX \le 83400 \cdot 1341\,E_3$. All numerical coefficients are independent of the host, ranks, level

and auxiliary field. · Labelwise coefficient estimate · Total intersection · Trace-cost budget · Portable proof with the full constant calculation

The counted sections in the height input. Fix a good deep cover, an admissible quotient weight u_1, a component Z, and its coefficient field F. Let W_Z be the degree-one integral cusp lattice at scalar-zero weight u_1 with its component probability Petersson metric. The generic framed evaluation ev of cusp rules is faithful on the component, and the ratios of the degree-one evaluations generate the function field $F(Z)$. Use the fitted valuation tests, including integral frame extension over the testing DVRs, and the complex cusp realization on this same component. For $n \geq 1$, let $C_Z(n,t)$ be the cardinality of the set of rational cusp sections s of weight nu_1 satisfying

$$\nu(s) \geq n \inf_{w \in W_Z} \nu(w) \qquad \text{for every recorded finite-place framed test } \nu, \qquad \|s(y)\|^2_{\mathrm{hom},n,\sigma} \leq e^{2nt} B_{\sigma,u_1}(y)^n \qquad \text{at every prescribed complex frame } y.$$

Here ν is the additive order in the tested frame, $\|\cdot\|_{\mathrm{hom},n,\sigma}$ is the pointwise homogeneous norm in weight nu_1, and B_{σ,u_1} is the squared degree-one Bergman evaluation norm in the fitted complex realization at a good embedding σ. These are the conditions of the release's `smallL`. Its `volL` uses this count along the powers $8a$. Counted set and volume · Volume normalization

Arithmetic image model and injection. Choose integral generators w_k of W_Z and take the normalization $\mathscr{X}$ of the closure over $\mathcal{O}_F$ of the projective image given by their generic values. It is integral, projective and flat over $\mathcal{O}_F$, with function field $F(Z)$ and relative dimension D. Flatness follows from torsion-freeness over the Dedekind base, and the normalization is finite. The pullback $\mathscr{L}$ of $\mathcal{O}(1)$ is generated by the sections $\widetilde{w}_k$ given by the coordinates. It has the Fubini–Study metric defined by the component Petersson metric on FW_Z. In any local frame the squared metric divides the squared coefficient by the sum of the squared coefficients of the sections from an orthonormal basis of $FW_Z \otimes_{F,\sigma} \mathbb{C}$. Those sections have no common zero because they span the same complex system as the $\widetilde{w}_k$. For t at least the logarithm of the common generator norm bound the norm shift e^{-t} makes every integral generator small by Cauchy–Schwarz. The shifted bundle $\overline{\mathscr{L}}(t)$ is nef by semipositivity of the Fubini–Study metric, generation on the special fibres, and the small generators on horizontal curves. Construction of the arithmetic image model · Nef metric and model comparison

For any s counted by $C_Z(n,t)$ and w_0 with nonzero generic value, set

$$\tilde{s} = \frac{\mathrm{ev}(s)}{\mathrm{ev}(w_0)^n}\, \widetilde{w}_0^{\,n}.$$

The coordinate ratios make this rational section independent of w_0. To prove extension, take a valuation of $F(Z)$ above a finite prime and choose w_h of least order among the generators. If the generic value of a nonzero s had order below n times the order of $\mathrm{ev}(w_h)$, then all ratios $\mathrm{ev}(w_k)/\,\mathrm{ev}(w_h)$ would be in the valuation ring and $\mathrm{ev}(w_h)^n/\,\mathrm{ev}(s)$ in its maximal ideal. The ratios and this last quotient generate a finite-type algebra over $\mathcal{O}_F$. The specialization lemma preserves these weak and strict inequalities on a complete DVR over the same prime and keeps the relevant denominators and values nonzero. It uses a horizontal curve through a closed specialization of the centre whose generic point avoids the denominators, followed by normalization and completion. The fitted integral frame extension after ramification gives a recorded test. The change of frame multiplies all weight-u_1 values by one common factor and the weight-nu_1 value by its n-th power. The specialized generator w_h thus realizes the minimum at that test, since integral linear combinations cannot lower it. The strict inequality contradicts the finite condition of $C_Z(n,t)$. It follows that the displayed ratio with w_h in the denominator is integral at the chosen valuation. Applying this on a

vertical codimension-one chart proves regularity there. For a horizontal codimension-one valuation, compose it with an extension to its residue field of a finite valuation of F. The composite is above a finite prime and its ring is contained in the original valuation ring, so the same ratio argument gives regularity. Normality then extends $\tilde{s}$ everywhere by the intersection of the height-one local rings on each affine chart. Specialization to a complete DVR · Fitted frame comparison · Valuation refinement and extension

For the norm comparison, the fitted complex span identifies $B_{\sigma,u_1}(y)$ with the squared evaluation norm on the degree-one cusp space on Z. In orthonormal coordinates this is the sum of the squared basis evaluations with their common homogeneous metric factor. At a generic framed complex reading the degree-n homogeneous factor is the n-th power of the degree-one factor. It cancels in the second inequality defining $C_Z(n,t)$, yielding $\|\tilde{s}\|_{\mathscr{L}^n} \le e^{nt}$ in the Fubini–Study metric. Generic readings suffice for the supremum over the complex fibre. In each affine chart the complex points avoiding extra algebraic relations over the countable field of definition are dense, since those relations define countably many proper analytic subsets on the relevant component. The norm estimate extends by continuity, using the regular extension and the positive Fubini–Study denominator. Faithfulness of generic framed evaluation makes $s \mapsto \tilde{s}$ injective. In particular, for every integer $a \ge 1$ the model has at least $C_Z(8a,t)$ integral sections of $\overline{\mathscr{L}}(t)^{8a}$ of sup norm at most one. The proof uses the same radius t on every admitted cover. Framed evaluation and Bergman bound · Complex cusp span on the cover · Passage from generic readings to complex points · Injection

Let $I(t) = \hat{c}_1(\overline{\mathscr{L}}(t))^{D+1}$ before division by $[F:\mathbb{Q}]$. The arithmetic Chow comparison uses the actual section lattice and its Fubini–Study frame and gives $I(t) \le [F:\mathbb{Q}](D+1)\delta_Z(h_Z+t)$. The constant metric shift contributes $[F:\mathbb{Q}](D+1)\delta_Z t$. For the fixed nef model and radius the small-section form of arithmetic Hilbert–Samuel gives

$$\frac{(D+1)!}{n^{D+1}} \log \#\widehat{H}^0(\mathscr{X}, \overline{\mathscr{L}}(t)^n) \longrightarrow I(t) \tag{16c}$$

Together with the injection and the intersection estimate this gives

$$V_Z(t) := \limsup_{a\to\infty} \frac{D! \log C_Z(8a,t)}{[F:\mathbb{Q}](8a)^{D+1}} \le \delta_Z(h_Z+t) \tag{16d}$$

The power limit is taken separately for each component, model, weight and radius, with onset allowed to depend on these choices. The family comparison uses the resulting volume inequality on each admitted cover, and the determinant bound at the selected T uses its finite-degree calculation. Chow and arithmetic projection formula · Nef arithmetic Hilbert–Samuel · Subsequence and comparison · Yuan, *Big line bundles over arithmetic varieties*

For the other direction assume the boundary comparison gives a volume $V'_Z(t) \ge \delta'_Z(s_Z+t)_+$ with $\delta'_Z \ge (1-\rho_B)\delta_Z > 0$. If this is obtained from rational boundary perturbations, fix the perturbation on the selected component, clear its denominator and take the power limit first. The resulting lower bound has the form $(\alpha_r/(D+1) + g_r(t-c-Cr))_+$ for small positive rational r, with $g_r \to W \ge \delta'_Z > 0$ and $\alpha_r/((D+1)g_r) \to s_0$. Here g_r, α_r are the geometric degree and field-normalized arithmetic intersection after division by the respective powers $D, D+1$ of the denominator. The bound at the chosen radius requires its own section comparison. For each fixed t, taking $r \downarrow 0$ gives the asserted bound with $s_Z = s_0 - c$. Now assume the carrying-section comparison into the pointwise and finite-order count yields $V'_Z(c_s) \le V_Z(c_s + c_c + \eta)$ for every $\eta > 0$, with both radii in the range of the upper bound. Apply (16d)

with that fixed radius, divide by δ_Z, then let $\eta \downarrow 0$. The first inequality of (16b) follows from $V_Z(c_s) \geq 0$. The mean lower bound and error budgets in (16b) require the intersection and carrying-section estimates of the fitted cover. Boundary perturbation limit · Comparison at the shifted radius

Finite-degree determinant bound. Fix Z and write $\delta = \delta_Z$ and $Q = Q_Z$. The fitted geometry gives a smooth projective geometrically integral modification $\widetilde{X}/F$ of the flag component, with its open part U_{off}. Write κ_{can} for the flag canonical weight and $u_2 = u_1 + \mathrm{rem}$, where rem is the anchored remainder of the same corrected gaps used for u_1. The section weight in degree i is $w(i) = \kappa_{\mathrm{can}} + u_2 + (D+i)u_1$. Let L, L_2 be the moving bundles for the u_1, u_2 cusp systems, obtained by principalizing and removing their base divisors. Let A be the pullback of the weight-$u_2 + Du_1$ bundle. The fitted comparison gives a section ϵ of $A \otimes (L_2 \otimes L^D)^{-1}$ invertible on U_{off}. Assume A, L_2 are nef and big and the chosen system in L is base-point-free and birational onto its image. In particular, L is nef. Weight identities · Bundles on the modification · Positivity and birationality

Put $B_i = K_{\widetilde{X}} \otimes A \otimes L^i$ and let G_i be the restriction of $H^0(\widetilde{X}, B_i)$ to U_{off}. For $0 \leq i \leq T$, $V_i \subset H^0(U_{\mathrm{off}}, B_i)$ is the image of the scalar-zero cusp space at $w(i)$ under pullback and the canonical identification on U_{off}. The fitted identifications give $G_i \subset V_i$ and $V_0 = G_0$. They transport the full cusp lattice M_i and the component probability Petersson norm $\| \cdot \|_{i,\sigma}$. Write $d = \dim_F V_T$ and $g = \deg_F M_T$. The image degree and height of the weight-u_1 system are δ and h_Z in the frame described below. Adjoint identification · Exact lattices and metrics · Component degree and height comparison

Use N coordinates on the linear system in L, and write $L(u)$ for the section with coefficient vector $u \in F^N$. Let W be the N-by-n matrix of the integral spanning columns of the u_1 cusp lattice, and A_σ the invertible complex frame matrix at σ. The frame norm is $\|u\|_{\mathrm{fr},\sigma} = \|A_\sigma^{-1} u\|_2$. Each column of W has norm at most Λ in every frame. Assume that multiplication by each column carries M_i into M_{i+1} for $i < T$ and that

$$\|L(u)v\|_{i+1,\sigma} \leq \sqrt{\beta}\, \|u\|_{\mathrm{fr},\sigma} \|v\|_{i,\sigma}$$

for arbitrary complex u and $v \in V_i \otimes_{F,\sigma} \mathbb{C}$. In degrees $0, D, T$, assume F-spanning sets in M_i with respective norm bounds $\Lambda_0, \Lambda_D, \Lambda_T$ at all embeddings simultaneously. Take all four norm bounds and β at least one. Integral and complex multiplication hypotheses · Fitted short families

Transverse cuts and leading symbols. For $s = (s_1, \dots, s_D)$ in a nonempty open of the space of cut coefficients, the sections $c_i = L(s_i)$ cut out a degree-δ finite étale subscheme C_s contained in U_{off}. Indeed the proper birational image map is an isomorphism over a dense open of its image. Shrink this open into the smooth locus and away from the image of the boundary. General linear sections of codimension D avoid the complement by dimension and meet the smooth locus transversely by characteristic-zero Bertini. The degree is $c_1(L)^D$ by the projection formula and generic degree one. The good locus contains a principal open defined over F. The irreducible Chow form $f_Z(s, \ell)$ has degree δ in each of the $D+1$ blocks. Over a splitting field F_s of a good cut, Chow hyperplane specialization gives $f_Z(s,\ell) = \eta_s \prod_{p \in C_s(F_s)} L(\ell)(p)$ for $\eta_s \neq 0$, with the evaluations in chosen frames of $L|_p$. All intersection multiplicities are one. Bertini · Good-cut open and degree · Chow hyperplane specialization · Chow forms and heights

At a split cut point the D local equations are regular parameters. For the ideal I_s of C_s, the identity $\mathrm{gr}_{I_s} \mathcal{O}_{\widetilde{X}} = \mathrm{Sym}_{\mathcal{O}_{C_s}}(I_s/I_s^2)$ makes the monomials c^α of a fixed degree a independent in the corresponding symbol space. Trivialize $B_0|_p$ and $L|_p$ and use these frames in all degrees. The symbol of $c^\alpha L(\ell)^{T-a} v$ for

$v \in V_0 \otimes F_s$ has only the α block, with entries $L(\ell)(p)^{T-a} v(p)$. For any subspace of sections filtered by I_s^a, the kernel of the symbol is the next step. Field extension preserves the dimensions of the section spans and commutes with global sections and multiplication. Conormal algebra · Field extension · Symbol maps on the fitted spaces

The restriction $V_0 \otimes F_s \to F_s^{\delta}$ is onto. Tensor the Koszul resolution of $\mathcal{O}_{C_s}$ by $K_{\widetilde{X}} \otimes L_2 \otimes L^D$. Its terms are sums of $K_{\widetilde{X}} \otimes L_2 \otimes L^{D-i}$ for $0 \le i \le D$. Kawamata–Viehweg vanishing for the nef and big twists of the canonical bundle makes these terms acyclic in positive degrees. Taking global sections yields a surjective restriction to C_s. Multiplication by ϵ gives sections of B_0 and an isomorphism on the restricted bundles along C_s. Now use $V_0 = G_0$. Koszul resolution · Kawamata–Viehweg vanishing · Degree-zero restriction in the release

Write $t_0 = T - D$ within this determinant argument. For $r \ge D$, tensor the Koszul resolution of I_s by B_r. The terms are sums of B_{r-i} for $1 \le i \le D$ and are acyclic in positive degree by Kawamata–Viehweg for $A \otimes L^{r-i}$. Thus every section in G_r vanishing on C_s is a sum $\sum_i c_i g_i$, with $g_i \in G_{r-1}$. By induction, a section of G_{j+a} vanishing to order a, with $j \ge D$, is a sum $\sum_{|\alpha|=a} c^{\alpha} g_{\alpha}$ with $g_{\alpha} \in G_j$. For the induction step, write the section using order $a-1$ and coefficients in G_{j+1}. The order-$a-1$ symbol is zero, so each coefficient vanishes on C_s by conormal independence. Apply the order-one division to each coefficient and collect monomials. Global division by the cut ideal · Induction on the order

The release filters $S_s = G_T + P_s \subset V_T$ by vanishing along C_s, where $P_s = \sum_{|\alpha|=t_0} c^{\alpha} V_D$. Multiplication places P_s in the degree-T space and gives vanishing to order at least t_0. For a section of S_s vanishing to that order, subtract its P_s summand and apply the global division just proved to its G_T summand, with $j = D$ and $a = t_0$. Thus the deepest step is exactly P_s. For any $a < t_0$ and any good cut, take a last hyperplane avoiding the cut points. The symbol formula and surjectivity in degree zero then span the whole grade using $c^{\alpha} L(\ell)^{T-a} V_0 \subset G_T$. These grades and P_s account for S_s, with the possible complement taken in V_T. Ambient filtered space · Deepest step

Polynomial determinant and Chow multiplicity. Choose a good F-rational cut s^0 maximizing $y = \dim_F P_s$ over the good rational cuts. Ranks are bounded integers and the good open has rational points over the infinite field F. Choose ℓ^0 avoiding the cut points, and choose $v_1, \dots, v_{\delta}$ among the short degree-zero spanning vectors whose residues at s^0 form a basis over its splitting field. In any fixed F-basis of V_T, form a polynomial determinant $\Delta(s, \ell)$ with the following columns.

- For each α with $|\alpha| < t_0$, take $L(s)^{\alpha} L(\ell)^{T-|\alpha|} v_i$ for $1 \le i \le \delta$. There are Q such columns. At (s^0, ℓ^0) they are independent modulo P_{s^0} by their independent leading symbols.
- Choose y products $L(s)^{\alpha} v$ with $|\alpha| = t_0$ and v from the short degree-D spanning set, giving a basis of P_{s^0} at s^0.
- Complete at (s^0, ℓ^0) using $d - Q - y$ short degree-T vectors as constant columns. Thus $\Delta(s^0, \ell^0) \neq 0$.

For any good F-rational cut s, dependence of the y deepest columns makes $\Delta(s, \ell) = 0$ identically in ℓ. When independent, they span P_s by maximality of y. Over the splitting field, a basis adapted to the filtration of S_s and extended to V_T gives a block triangular determinant. The deepest columns and constant completion contribute factors independent of ℓ. In the grade for each α, the symbol block is the residue

matrix of the v_i with row p multiplied by $L(\ell)(p)^{T-|\alpha|}$. This calculation also gives zero when a residue block is singular. Hence $\Delta(s,\ell)$ is a scalar times $\prod_{p\in C_s(F_s)} L(\ell)(p)^e$, with exponent

$$e = \sum_{|\alpha|<T-D} (T-|\alpha|) = \binom{T-1}{D}\left(D+\frac{T}{D+1}\right). \tag{16e}$$

Indeed appending a slack exponent gives $D+1$ exponents of sum $T-D-1$. Their count is $\binom{T-1}{D}$ and the mean slack is $(T-D-1)/(D+1)$. Each summand is $D+1$ plus the slack. This also gives the count Q of residue columns. Independent columns and specialized determinant

Choose $\ell_* \in F^N$ with $f_*(s) = f_Z(s,\ell_*)$ nonzero as a polynomial. The two specialized product formulas give $\Delta(s,\ell) f_*(s)^e = f_Z(s,\ell)^e \Delta(s,\ell_*)$ for every good rational s. Density gives the identity in the polynomial ring over F. The irreducible polynomial f_Z has positive ℓ-degree while f_* is independent of ℓ. Unique factorization therefore gives $\Delta = f_Z^e R$. The degree of Δ in the last block is δe by column counting, so the nonzero cofactor R is independent of ℓ and the Chow order is exactly e. The determinant and cofactor are multihomogeneous. In total degree the columns give

$$e' = \deg R = QT + y(T-D) - (D+1)\delta e = y(T-D) - QD(D+1), \qquad 0 \le e' \le T(d-Q), \tag{16e.1}$$

Here $y \le d-Q$ by independence. Global divisibility · Degree accounting

Integral coefficients and archimedean bounds. In every block substitute $u_b = W z_b$. The substitution is injective on polynomials because W has an F-linear right inverse. It can also give a witness with integral cuts and last hyperplane. If the good principal open is defined by $q(s) \neq 0$, apply the substitution to $q(s)\Delta(s,\ell)$ and evaluate at integer coordinates between zero and its total degree. Some value is nonzero by iterating the univariate root bound. The arithmetic estimate uses the coefficients of the substituted determinant, so it incurs no norm cost for this witness. Spanning substitution

After substitution, each coefficient of every column belongs to M_T. Expand the product in its hyperplane factors. An ordered selection of one frame column per factor is integral by repeated lattice multiplication, and sums over selections with the same monomial remain integral. Let $\mathbf{b}$ be the fixed rational basis of V_T, J_T the fractional ideal of $\bigwedge^d M_T$ in the coordinates $\bigwedge^d \mathbf{b}$, and $G_{\sigma,\mathbf{b}}$ the Gram matrix of $\mathbf{b}$. Multilinearity puts each coefficient of $\Delta(Wz)$ in J_T. Normalize $|c|_{\mathfrak{p}} = (N\mathfrak{p})^{-v_{\mathfrak{p}}(c)}$ and use the maximum coefficient norm $\|\cdot\|_{\mathfrak{p},\mathrm{G}}$. The coefficient bound and the determinant-line formula are

$$\begin{aligned} \log \|\Delta(Wz)\|_{\mathfrak{p},\mathrm{G}} &\le -v_{\mathfrak{p}}(J_T)\log N\mathfrak{p}, \\ [F:\mathbb{Q}]g &= -\log N(J_T) - \tfrac{1}{2}\sum_{\sigma} \log\det G_{\sigma,\mathbf{b}}. \end{aligned} \tag{16e.2}$$

The first inequality compares the minimum coefficient valuation with the valuation of the determinant ideal. The degree uses ordinary norms and counts both complex embeddings of each infinite place. Coefficient integrality · Finite determinant comparison

For a nonzero multihomogeneous polynomial P, write $m_\sigma(P)$ for the mean of $\log|\sigma(P)(A_\sigma z)|$ over the product of complex unit spheres, with A_σ acting in each block. Set

$$H(P) = \frac{\sum_{\mathfrak{p}} \log\|P(Wz)\|_{\mathfrak{p},\mathrm{G}} + \sum_\sigma m_\sigma(P)}{[F:\mathbb{Q}]}.$$

The Gauss norms multiply by reduction after rescaling over each discrete valuation ring. The sphere logarithms are integrable by local integrability in projective charts, and their means add under multiplication. Thus $H(\Delta) = eH(f_Z) + H(R)$. The product formula gives invariance under multiplication by a nonzero scalar from F. With $H_{N-1} = \sum_{i=1}^{N-1} 1/i$, the Fubini–Study normalization in the prescribed frame gives $h' = H(f_Z)/((D+1)\delta) = h_Z - H_{N-1}/2$. Since $N \le n$, one has $H_{N-1} \le 1 + \log n$. Polynomial height · Harmonic correction · Fitted Chow height

At every unit-sphere input, repeated complex multiplication bounds the column norms by $\beta^{T/2}\Lambda_0$, $\beta^{(T-D)/2}\Lambda_D$, and Λ_T in the three groups, respectively. Hadamard's inequality bounds $|\sigma(\Delta)(A_\sigma z)|\sqrt{\det G_{\sigma,\mathbf{b}}}$ by the product of those norms. Average its logarithm at each embedding and combine with (16e.2) to obtain

$$\begin{aligned} g \ge\ & eH(f_Z) + H(R) - Q\log(\beta^{T/2}\Lambda_0) \\ & - y\log(\beta^{(T-D)/2}\Lambda_D) - (d-Q-y)\log\Lambda_T. \end{aligned} \tag{16e.3}$$

Complex Hadamard bound and arithmetic degree

For any nonzero P of total multidegree a, we also have $H(P) \ge -ac_*$ for $c_* = \log(n\Lambda) + (\log n + 1)/2$. To see this, first work in the metric coordinates $A_\sigma z$. In one block of degree a_b, rotate a chosen unit vector to the first coordinate direction. At fixed other coordinates, the coefficient of the first coordinate to power a_b is the evaluation at that direction. Jensen's formula bounds the phase mean from below by the logarithm of the absolute value of that evaluation plus $a_b \log|z_1|$. The mean of $\log|z_1|$ is $-H_{N-1}/2$, since $|z_1|^2$ has the beta $(1, N-1)$ law for $N > 1$. For $N = 1$ the mean is zero. Applying this one block at a time bounds the logarithm of the maximum on the product of frame-unit spheres by $m_\sigma(P) + aH_{N-1}/2$. On the polytorus $|z_{b,i}| = 1$ we have $\|\sigma(W)z_b\|_{\mathrm{fr},\sigma} \le n\Lambda$. Homogeneity and Cauchy's coefficient formula bound every coefficient of $\sigma(P(Wz))$ in absolute value by $\exp(m_\sigma(P) + aH_{N-1}/2)(n\Lambda)^a$. Apply the product formula to a single nonzero coefficient, using these archimedean bounds and the finite Gauss norms. This yields $H(P) \ge -a(\log(n\Lambda) + H_{N-1}/2) \ge -ac_*$. Sphere mean inequality · Coefficient lower bound for the height

The prescribed degree and the flat transfer. Apply the coefficient estimate to f_Z and R to get $h' \ge -c_*$ and $H(R) \ge -e'c_*$. Use the exact cofactor degree from (16e.1) to obtain

$$eH(f_Z) + H(R) \ge QTh' + QD(D+1)(h' + c_*) - y(T-D)c_*.$$

The term with $h' + c_*$ is nonnegative. Writing $\gamma = H_{N-1}/2$ and using $h' = h_Z - \gamma$ in (16e.3) gives

$$\begin{aligned} g &\ge Q(Th_Z - a_0) - ya_D - (d-Q-y)a_T, \\ a_0 &= T\gamma + \log\Lambda_0 + \tfrac{T}{2}\log\beta, \\ a_D &= (T-D)(c_* + \tfrac{1}{2}\log\beta) + \log\Lambda_D, \qquad a_T = \log\Lambda_T. \end{aligned} \tag{16f}$$

The $d - Q$ completion columns therefore cost at most $(d-Q)\max(a_D, a_T)$. Finite numerical inequality in the release

Put $L_b = \log b$ and $\Xi = f(p_j + q_j)L_b$, where f is the CM-type size. Use $L_b \ge 1$, $\Xi^{12} \le \mathrm{PP}$, $T\mathrm{PP} \le \sqrt{b}$ and $zK_* \le b$. Require the following bounds in the fitted metrics at every Z.

$$\log n, \log \beta \le \mathrm{PP}, \qquad \log \Lambda \le K_*/T + \mathrm{PP}, \qquad \log \Lambda_D, \log \Lambda_T \le K_*,$$
$$\log \Lambda_0 \le D(K_*/T + \mathrm{PP}).$$

Stable theta spanning on the fitted cover

Proposition (stable theta spanning on the fitted cover). Fix a good nonterminal cover Y of odd principal level $M \ge 3$ and auxiliary parameter $A = 1$, over its prescribed coefficient field F. Write $r = p_j + q_j$, $[E : \mathbb{Q}] = 2f$, and use an even theta exponent $n_\theta \ge 8fr + 4$, dominant nonnegative exponents $\mathbf{a}$, and positive last exponent in the p_j-block at each active embedding. Put $B_Y = Q_S q_{\mathrm{host}} d_{\mathrm{host}} M$ and $M_* = 2B_Y M$. Assume the fitted marked uniformizations, the cubical polarization and normalized fibre generators, and the compatible charts, connection, weight and evaluation laws specified below. Then the full F-space of cusp rules at $\mathrm{wt}\mathrm{Z}(\mathbf{a}, n_\theta)$ is spanned by a finite family of algebraic theta rules. The rules use Hermitian self-dual lattices $\Lambda \subset (E^{n_\theta}, 1)$ standard at every prime above M_*, odd N divisible by M with $N\mathcal{O}_E^{n_\theta} \subseteq \Lambda \cap \mathcal{O}_E^{n_\theta}$, and integral harmonic polynomials in the letters of Λ. Their coefficient is the level Fourier transform of the indicator of a tuple generating the dual-half marking module.

The marked lattice family. At one good embedding $\sigma_0 : F \hookrightarrow \mathbb{C}$, use the uniformizations of all components, with the same Hermitian matrix, CM labels, signature and imaginary period scalar. They give a common Hermitian space $W = E^r$, alternating trace polarization e and frame domain $\mathcal{D}$. The uniformization gives the prescribed principal congruence subgroup, trivial stabilizer character for the weight, and the dictionary from frame characters and tangent directions to the algebraic weight and letters. Fix a self-dual reference period lattice L_0 with symplectic basis and a basis of its reduction modulo the marked half-level ideal $\mathfrak{m}$, where $\mathfrak{m}\bar{\mathfrak{m}} = M\mathcal{O}_E$ and the halves are coprime. Let $\mathscr{L}$ consist of e-self-dual lattices agreeing with L_0 at primes over M. Let Δ be the rational unitary group preserving the localized marking modulo $\mathfrak{m}$. The auxiliary comparison uses the identity cover of Y, so the allowed reduction matrix is the identity and the auxiliary theta multiplier is one.

The fitted lattice moves send the period lattices to $L_i \in \mathscr{L}$ with this common marking. The first r entries of a common tuple x_0 give the marked basis and the lower entries are zero. The perfect pairing with the conjugate half recovers the original marking modulo M. Thus a move in Δ identifying L_i and $L_{i'}$ identifies the corresponding marked uniformizations, which forces $i = i'$. At L_i the stabilizer acts by full principal congruence. The vacuum transformation uses the host's adapted symplectic basis, in which the relevant off-diagonal blocks are zero modulo 8 and the lower diagonal block has determinant 1 modulo 8. The square theta law in these coordinates gives the period determinant, which cancels with the frame determinant in even exponent. Hence the vacuum multiplier is one on the stabilizers. These calculations use the same moved markings in the theta reading and in the lattice action. Classical level · Simultaneous markings and lattice moves · Trivial stabilizer multiplier · Adapted period transformation

The finite Gram matrix. For an auxiliary lattice Λ as in the statement, the tensor lattice $A_{\Lambda,L} = \Lambda \otimes_{\mathcal{O}_E} (\mathfrak{m}/\bar{\mathfrak{m}})L$ is self-dual for the tensor alternating trace pairing. The factor $\mathfrak{m}/\bar{\mathfrak{m}}$ cancels with its conjugate under duality, and testing against a local Hermitian dual basis of Λ reduces the assertion to the alternating duality of L. The host element $w + \bar{w} = 1$ gives the sign $\exp(\pi i\, e(y, wy))$ on the tensor. Its cross term is $2e(y, wy') - e(y, y')$, so the signed Heisenberg operators give a genuine lattice action. In a symplectic splitting one half of the operators determines the support by the additive characters, and the

other determines the translates from one value. This gives a one-dimensional vacuum line, with the same sign convention also at primes over 2.

For two self-dual signed lattices A, A' in these coordinates, the signed translates of a unit A'-vacuum by $A/(A \cap A')$ are orthogonal. For any nonzero coset, commutation with A' detects a nontrivial additive character by self-duality. Averaging the translates projects to the A-vacuum line, so the squared overlap is $[A : A \cap A']^{-1}$. The nonzero projection of the standard unit vacuum onto the Λ-vacuum, divided by its norm, fixes the phase. Let $\Phi_{L,\Lambda}$ be its Heisenberg translate with the prescribed first r rows modulo $\mathfrak{m}$ and zero lower rows. The translate is taken in $\Lambda \otimes 2\bar{\mathfrak{m}}L$ modulo $\Lambda \otimes 2\bar{\mathfrak{m}}\mathfrak{m}L$. The Chinese remainder theorem gives a common representative for L, L' in their intersection, because they have the same localizations at the marking. The factor 2 preserves the phase on changing a representative. The identity cover makes the free-lower-row test identical, with lower residues zero modulo the same $\mathfrak{m}$.

The space of auxiliary classes consists of these definite lattices modulo rational unitaries congruent to 1 modulo M_* in the standard localizations. This marked class set is finite, with positive mass weights $w_\Lambda = |\operatorname{Aut}_{M_*} \Lambda|^{-1}$. Choose one representative of each class and sum $\langle \Phi_{L,\Lambda}, \Phi_{L',\Lambda} \rangle$ with those weights to form the matrix $C(L, L')$, using inner products linear in the first entry. Its diagonal is $a_0 = \sum_\Lambda w_\Lambda > 0$. For a fixed Λ a local basis gives $[A_{\Lambda,L} : A_{\Lambda,L} \cap A_{\Lambda,L'}] = [L : L \cap L']^{n_\theta}$. The overlap bound therefore gives $|C(L, L')| \le a_0 [L : L \cap L']^{-n_\theta/2}$. Put $d = [E : \mathbb{Q}]r$. For fixed L, there are at most a^{2d} lattices with $[L : L \cap L'] = a$. There are at most a^d choices for the intersection I by counting chains of sublattices of prime index. Each prime step ℓ has at most ℓ^d choices, as the sublattice is the kernel of a map to $\mathbb{F}_\ell$. For each I, duality gives $L' \subseteq I^\vee$ of index a, yielding another factor a^d. Index one forces equality. Consequently

$$a_0^{-1} \sum_{L' \ne L} |C(L, L')| \le \sum_{a \ge 2} a^{2d - n_\theta/2} \le \tfrac{1}{4} + \tfrac{1}{9} + \int_3^\infty t^{-2}\, dt = \tfrac{25}{36}.$$

The last inequality uses $n_\theta \ge 4d + 4$. Thus $C = a_0(1 + T_0)$ has a Neumann inverse on uniformly bounded families in the homogeneous norm. The sum converges locally uniformly in frame coordinates. If $\mu(\gamma, L)$ is the unit modulus multiplier of the finite test, the covariance of C under simultaneous transport preserves the law $G(\gamma L, \gamma x) = \mu(\gamma, L)^{-1} G(L, x)$. Signed vacuum and orthogonal projection · Overlap · Tensor index · Marked finite Gram · Lattice count and tail · Mass weights and covariance · Finiteness of the marked classes and stabilizers

The doubled kernel. Use the pinned convergent definite Siegel–Weil class-sum identity in Siegel degree r on the split doubled space $W \oplus (-W)$. Its convergence condition is $n_\theta > 4r$, or $s_0 = (n_\theta - 2r)/2 > r$ in the normalized Eisenstein parameter. This identity uses the same class weights, finite additive character and even-weight period splitting as the Gram calculation. For the two marked tests use the product of one with the conjugate of the other, transported to the diagonal polarization by a linear change and partial Fourier transform. Its value at zero is a fixed nonzero scalar times their finite Hermitian pairing. For the unshifted vacua its support is $\Lambda \otimes (L^s + L'^s)$, where $L^s = (\mathfrak{m}/\bar{\mathfrak{m}})L$. The momentum kernel is dual to this position lattice, and the lift in the doubled lattice gives a quadratic trace phase on the support. The scalar is common across the auxiliary classes. Its modulus follows from the overlap, and its phase follows by taking the inner product with the doubled standard vacua, where the common quadratic phases cancel and both normalized projection pairings are positive. Reinsert the translates using a common residue lift for the two lattices. This gives translations and modulations of the same indicator. Clearing denominators in the support and phase gives the support and period bounds by a common integer. Equality of local lattices at its

primes and at the marking gives locality, and unitary transport preserves the quadratic trace and the residues. These are the test-family conditions of the convergent class-sum identity.

In the resulting Eisenstein sum a maximal isotropic subspace with injective projections is a unitary graph. Otherwise it contains an isotropic line in one summand. On the tensor of that line the semicharacter is trivial, and the translated vacuum has a nontrivial commutator character. Locally at the half-level, the first r rows of the mark are a basis. At the conjugate prime the shift enlarges the line lattice, so the perfect trace pairing detects a nontrivial character against those rows. The translated term vanishes on averaging over the isotropic unipotent translations. For a graph with nonzero pairing, the displacement belongs to the sum of the two vacuum lattices, by commutation with their intersection. In the marked half both correction arrays are zero modulo $\mathfrak{m}$. Their first r rows are still bases, so the graph map is integral and invertible there and respects the marking. This leaves the graph sum over Δ.

Take a Fock-orthonormal basis P_j of the row-highest harmonic space. This space contains products of highest minors in disjoint sets of auxiliary coordinates for the two signature blocks, since $n_\theta \geq r$. The mixed contractions then vanish. Its tensor sum $\sum_j P_j \otimes \overline{P_j}$ is invariant under the compact auxiliary unitary group. The zero-Jacobi class-sum identity suffices for these derivatives. The theta-minus-Eisenstein difference satisfies the heat equations $\sum_k \partial_{z_{ka}} \partial_{z'_{kb}} D = 2\pi i\, \partial_{Z_{ab}} D$ in the doubled Jacobi variables. After applying the product rule to the common frame Gaussian, the invariant operator in the remaining Jacobi derivatives is a polynomial in the pairings of vectors and covectors. Each pairing acts by the heat equation and all the resulting derivatives vanish at zero Jacobi argument. Absolute convergence and the compact-set bound on the Jacobi quadratic factors justify differentiation. Thus the pinned identity for the theta derivatives $\theta_{\Lambda,j}$ and the corresponding frame Gaussian kernel K_∞ reads

$$\sum_{\Lambda,j} w_\Lambda\, \theta_{\Lambda,j}(L,x)\overline{\theta_{\Lambda,j}(L',y)} = c_*\sum_{\gamma\in\Delta} C(L,\gamma L')\,\overline{\mu(\gamma,L')}\,K_\infty(x,\gamma y), \qquad c_*\neq 0.$$

The constant comes from the Fourier change of polarization and is common to the components and the auxiliary classes. Siegel–Weil by degree · Doubled vacuum and common phase · Admissibility of the doubled tests · Isotropic line detection · Graph support · Graph expansion and differential kernel · Normalized kernel identity · Nonzero invariant harmonic basis

Inversion and reproduction. First work with bounded holomorphic target functions Q_i on the moved components, of the indicated right character. Extend to their lattice orbits by the inverse multiplier law and put zero on the other orbits. The stabilizers act trivially and the orbits are distinct, so this gives a uniformly bounded family $\mathcal{F}$. Set $G = C^{-1}\mathcal{F}$. Take a small unitary-invariant slab S_δ of frames whose Gram matrices lie near the orthonormal signature matrices. For fixed x, the kernel $K_\infty(x,\cdot)$ is bounded and integrable on this slab for $n_\theta \geq 2r$. The leading-minor estimate bounds it by a constant times the product of the positive-block determinant moduli to power $-n_\theta$. On unitary frames these moduli are products of $\cosh t_\ell$ in Cartan coordinates. The Haar density grows at most as $\exp(2\sum_\ell (r-2\ell+1)t_\ell)$, and the bounded Gram coordinates of the slab preserve the estimate.

The integral of $K_\infty(x,y)G(L',y)$ over the slab is $\rho_\delta G(L',x)$ with $\rho_\delta > 0$. At an orthonormal x_*, average the integrand for the left diagonal unitary torus in the right weight character. The resulting holomorphic function is determined by its value at x_*. On the dense triangular cell this follows by conjugating lower triangular coordinates by $\operatorname{diag}(1, z, \dots, z^{r-1})$ and using holomorphy on the contracting disk. The positive leading span stays positive under the contraction, since the negative rows contract faster. The value at the diagonal therefore fixes the function. Applied also to $K_\infty(\cdot, x_*)$, this gives

$\rho_\delta = \int_{S_\delta} |K_\infty(x_*, y)|^2 \, dy / K_\infty(x_*, x_*)$. The denominator is positive by the diagonal Gaussian and the nonzero harmonic basis. Left unitary invariance and right equivariance give the formula on the full domain.

Multiply the doubled kernel by $G(L', y)$ and integrate over measurable slices for the simultaneous lattice and frame action. The conjugate multiplier combines with G to give $G(\gamma L', \gamma y)$. Unfolding to $\mathscr{L} \times S_\delta$ is absolutely convergent, bounded by $a_0(1 + 25/36) \sup |G| \, \|K_\infty(x, \cdot)\|_{L^1(S_\delta)}$, with the supremum on the slab. The result after removing the common factor c_* is $\rho_\delta(CG)(L, x) = \rho_\delta \mathcal{F}(L, x)$. On the product-theta side the result belongs to the finite theta span. To justify this last integration, choose a basis of the first factors and finitely many frame evaluations giving an invertible matrix on it. The coefficient functions in that basis are linear combinations of those finitely many kernel evaluations, each with the absolute integral just proved. This yields the simultaneous complex span on the moved lattice family. Slab integrability · Lowest-weight uniqueness · Reproduction on the slab · Neumann inverse and unfolding

The algebraic section. Use the cover selected for the fitted row by the call on the row core with depth $e_{\mathrm{cov}} = \operatorname{expL}(L)$ in the release's notation. The construction uses a half-ideal over $15^{e_{mcov}}$ and the compositum F of the core coefficient field and the specified ray field, which contains the primitive root ζ_M for the half-level pairing. On this cover use the identity auxiliary comparison. Its parameter is 1 and its comparison section has value 1. Fix the common dual-half tuple x_0 and its lift. Put $\Lambda_0 = \Lambda \cap \mathcal{O}_E^{n_\theta}$ and choose $N = Ms$ with $s > 0$ prime to M_* and $N\mathcal{O}_E^{n_\theta} \subseteq \Lambda_0$. The standard localizations at M_* make this possible and ensure N is odd. The row-highest and harmonic equations are linear over F in the given multidegree, so extension of scalars preserves their kernel. Clearing coefficient denominators in the lattice letters gives a spanning family of integral harmonic polynomials. Selected cover call · Level and coefficient field · Identity comparison and indicator transform

Let $\mathcal{Q}$ be the n_θ-fold product of the marked half-level quotient, $h_{\mathcal{Q}} = |\mathcal{Q}|$, and a_α a representative of $\alpha \in \mathcal{Q}$. Use the unnormalized Fourier coefficient $w(\alpha) = \zeta_M^{\operatorname{pairL}(x_0, a_\alpha)}$, the transform of the indicator of the dual-half class x_0. These roots of unity come from F on every test. For a framed point v over an F-algebra S, choose over a faithfully flat S-algebra an enumeration of the fitted isotropic torsion kernel of size $[\Lambda : N\mathcal{O}_E^{n_\theta}]^r$ and N-division points of the marked level tuple, killed by a power of the marked half-ideal $\bar{\mathfrak{m}}$. Let $s_{\alpha,\nu,k}$ be the section in column k obtained from the framed fibre generator by the rigidified pullback under $[N]$ and the odd-order symmetric theta-group lifts of the kernel point and the marked division point, in that order. For a homogeneous harmonic polynomial P of degree d_P, the raw rule has value

$$y_{\Lambda,N,P,x_0}(v) = \frac{1}{[\Lambda_0 : N\mathcal{O}_E^{n_\theta}]^r h_{\mathcal{Q}} N^{d_P}} \sum_{\alpha \in \mathcal{Q}} w(\alpha) \sum_{\nu} \sum_{\mu \in \operatorname{supp} P} P_\mu \prod_{k=1}^{n_\theta} \nabla_{v,N^2}(s_{\alpha,\nu,k}; \mu_k),$$

where μ_k is the column part of the monomial and ∇_{v,N^2} uses the fitted Lie frame and connection on the N^2-th power of the polarization line. Changing the division points reindexes the isotropic kernel, by standardness of Λ at M. The product of column phase changes has square one by the isotropic commutator law and N^2-th power one by the lift law. It is one since N is odd. After pullback to the overlap ring, chart independence and base change of the differential identify the two values. The faithfully flat equalizer gives a unique value in S and the same comparison gives naturality under change of test. The triangular frame law then gives one section over F in the fitted weight. The positive last exponent and the fitted weight dictionary give cusp membership at every good embedding. Torsion witnesses · Independence of choices · Faithfully flat descent · Frame weight law · Identity cover and natural rules · Cusp membership

Comparison on the marked components. On a unit frame for L_i, the rigidified polarization line has Hermitian form H_x with $\operatorname{Im} H_x = e$ on the period lattice. Let χ be its Appell–Humbert semicharacter. The fitted polarization isomorphism identifies the line with the alternating product of its pulls by $0, 1, w, 1+w$, with exponents $+1, -1, -1, +1$, where $w + \bar{w} = 1$ is the host endomorphism. The Hermitian form of this product is H_x, by the Rosati adjoint identity. Its semicharacter is $\chi(v+wv)/(\chi(v)\chi(wv)) = \exp(\pi i e(v, wv))$. Appell–Humbert uniqueness therefore identifies the characteristic with the sign of the vacuum. The space of fibre theta sections of the principal polarization has dimension one. The product of the framed generators in the n_θ columns is consequently the Fock theta of the standard signed vacuum times a common factor $\kappa_i(x)$. Its modulus on unit frames is fixed by the mean-one Gaussian normalization and the determinant of H_x in the prescribed Lie frame. Uniformized fibre and unit generator · Cubical characteristic · Birkenhake–Lange, Appell–Humbert · Product of fibre theta sections · Fixed modulus

In the complex reading chart independence and the letter-to-tangent dictionary identify the harmonic differential with the Jacobi derivative. The rigidified pullback dilates the Jacobi variables by N, so homogeneity cancels N^{d_P}. The even representatives of the odd torsion kernel project the standard vacuum onto the $\Lambda \otimes L_i$ line. The finite Fourier average of the marked N-division points projects further to the line for $\Lambda \otimes (\mathfrak{m}/\bar{\mathfrak{m}})L_i$ and applies the translate for x_0. Here the free-row test has one residue with coefficient 1, since the auxiliary comparison is the identity. Evaluating its Fourier coefficient at $\alpha = 0$ fixes the unit in the Fourier comparison to 1. The factor h_Q from the finite sum cancels the outer average. The shift of L_i is supported at M and the change of the auxiliary lattice is trivial there. The coprime projection identity therefore allows the shift of the standard vacuum first. Its overlap with the shifted standard unit vacuum gives a nonzero factor t_i depending on the component normalization alone. The subsequent projection onto the Λ-vacuum has norm $J^{-r/2}$ for $J = [\Lambda : \Lambda_0] = [\mathcal{O}_E^{n_\theta} : \Lambda_0]$, using self-duality and the tensor index. The kernel-count ratio gives J^r. In the normalized short rule there is also a denominator $\beta_H(\Lambda) \in F^\times$. Multiplying by this one field element gives the raw rule above. Thus the raw lattice factor is $u_{\Lambda,N} = J^{r/2}$, the same on every component and independent of the clearing degree N. Values, derivatives and kernel counts · Finite Fourier formula · Finite projection and shift · Overlap of shifted vacua · Separation of scalar factors

On each unit frame the remaining factor $\kappa_i(x)t_i$ is common to all (Λ, N, P) and has fixed positive modulus. Both readings have the same holomorphic right weight law. Triangular change extends the comparison to all frames, and the quotient for any nonvanishing summand is holomorphic of constant modulus on a ball. It is constant there, and the identity theorem on the connected frame domain gives the same constant for every summand. If all summands vanish, any constant of that modulus works. With the fixed moved marking and frame relabelling g_i, this gives

$$R_i(y_{\Lambda,N,P,x_0})(x) = C_i\, J^{r/2}\, \theta^{\mathrm{free}}_{\Lambda,\sigma_0 P}(L_i, g_i x).$$

The $C_i \in \mathbb{C}^\times$ include the phases of the chosen component trivializations. They are fixed for the tuple x_0 and weight as the summands vary. The formula over F uses the single set of coefficients $w(\alpha)$ and P_μ on the entire cover. Simultaneous complex reading · Constancy on the frame domain

Faithful evaluation. At σ_0 use every component. Each complex point of a prescribed reduced affine test ring S gives a framed value equal to a component reading times a multiplier common to all sections of the weight. A complex relation among the simultaneous readings therefore gives zero at every complex point of $S \otimes_{F,\sigma_0} \mathbb{C}$. This ring is reduced and of finite type, so the Nullstellensatz detects zero. Applying arbitrary F-

linear functionals $\mathbb{C} \to F$ to the scalar coefficients gives relations of rules over F. This proves that F-independent rules have independent simultaneous readings over $\mathbb{C}$. The construction and this descent use framed evaluation. The Fourier coefficients used above are on a finite level module.

To span an F-rational cusp rule, first apply the analytic span to its component readings divided by C_i and absorb the inverse factors $J^{-r/2}$ in the coefficients. Take a maximal F-independent subfamily of the finitely many algebraic rules used. If the target lay outside its span, adjoining it would contradict faithful evaluation. The union of such subfamilies for an F-basis spans the cusp space. The short harmonic family may be substituted before this step since it spans the same harmonic space after extension of scalars. Test value as a multiple of the reading · Faithful reading after extension · Span over the coefficient field · Integral harmonic polynomials · Free lower rows and simultaneous span · Passage to algebraic readings

Fitted-cover realization and full torsion calculation · Stable theta spanning proof and source details

Short integral sections on the cover

Write $p = p_j, q = q_j, r = p + q$, and use the fitted good cover with odd principal level $M \geq 3$ and auxiliary parameter $A = 1$. Use the same marked cover and connection. Identify its integral cusp lattice by the prescribed framed complete-DVR tests, with the local chart and torsion comparisons below. Use $f > 2$, $L_b = \log b \geq 1, \Xi = frL_b, \log M \leq 30L_b, \operatorname{fieldSz}(E) \leq 3L_b, \log(fr) \leq L_b$, and $\log[F : \mathbb{Q}], \log \operatorname{rd} F \leq 391\Xi^2$. We use the ordinary homogeneous norm, with the Lie factor at the prescribed good embedding representatives. Use the fitted identification of this pointwise norm with the integrand of the component probability Petersson metric. Its supremum then bounds that metric on each component.

Use the algebraic spanning family just constructed at $\operatorname{wtZ}(\mathbf{a}, n_\theta)$. Put $\Lambda_0 = \Lambda \cap \mathcal{O}_E^{n_\theta}$, write $d_P = \sum_{\tau,i} a_{\tau i}$ for the total degree in the active blocks, and put $K = d_P + frn_\theta/2$. For the short estimate require $\log K \leq 20L_b$. The fitted short harmonic family bounds $\operatorname{polyNorm}(P)$ by $\exp(20000\Xi^2 K)$. This uses the per-embedding coefficient ℓ^1 bound for products of minors in short integral Hermitian-orthogonal columns and the bound on $\log[F : \mathbb{Q}]$ to sum the squares over embeddings. Raw theta family · Identity cover and differential reading · Coefficient bound and descent

Put $J = [\Lambda : \Lambda_0]$ and $\mathfrak{a} = \det \Lambda_0 (\det \Lambda)^{-1}$. Local elementary divisors give $\mathrm{N}_{E/\mathbb{Q}}\mathfrak{a} = J$. If $\iota_\tau : E \to F$ are the labelled coefficient maps and $m_\tau = p$ on the CM type and q on its conjugate, the ideal

$$\mathfrak{B} = \prod_\tau (\iota_\tau(\mathfrak{a})\mathcal{O}_F)^{m_\tau} \quad \text{has norm} \quad \mathrm{N}_{F/\mathbb{Q}}\mathfrak{B} = J^{r[F:\mathbb{Q}]/2}.$$

For each raw rule choose $0 \neq s \in \mathfrak{B}^{-1}$ with $|\sigma(s)| J^{r/2} \leq \sqrt{\operatorname{rd} F}$ for every σ. In the complex Minkowski embedding the inverse ideal has covolume $2^{-[F:\mathbb{Q}]/2}\sqrt{|\operatorname{disc} F|}\, J^{-r[F:\mathbb{Q}]/2}$. The product of disks of radius $\sqrt{\operatorname{rd} F}\, J^{-r/2}$ has volume greater than $2^{[F:\mathbb{Q}]}$ times that covolume, since $\pi > 2$.

The half-level marking uses $\mathfrak{m}\bar{\mathfrak{m}} = M\mathcal{O}_E$. Its rank-$r$ projective module reduced modulo $\mathfrak{m}$ has cardinality M^{fr}, by local freeness over the factors of $\mathcal{O}_E/\mathfrak{m}$. Thus the raw reading averages over $Q_M = M^{frn_\theta}$ tuples with root-of-unity coefficients. Each inner sum uses the invariant derivatives of the chosen connection on the prescribed cubical translations and divides by $[\Lambda_0 : N\mathcal{O}_E^{n_\theta}]^r N^{d_P}$. The kernel has cardinality $[\Lambda : N\mathcal{O}_E^{n_\theta}]^r$. In the local short-rule formula the scalar β_H is removed by using $g = \beta_H^{-1} y$, so the one-copy

bound on $\beta_H Z_{\mathrm{den}} g$ is a bound on $Z_{\mathrm{den}} y$, where Z_{den} is defined below. Reflex ideal · Half-level quotient · Framed short-value formula · Raw and normalized rules

On each integral test, use the fitted chart and torsion lifts, allowing a finite extension of the complete DVR and dividing its valuation by ramification. The chart loss is at most

$$\rho_v = \frac{d_P}{2} \sum_{\ell' \mid \mathrm{cond}_{\mathrm{host}}} v(\ell') \, \frac{v_{\ell'}(|\operatorname{disc} E|)}{[E:\mathbb{Q}]} \le v\big(|\operatorname{disc} E|^{d_P}\big),$$

where the sum is over prime factors of the host conductor. Off the level the lifted point has order prime to the residue characteristic ℓ by the torsion conductor condition. In the adapted kernel formula, the determinant of the matrix P_ℓ has local ideal $\mathfrak{a}$. Indeed the relations $P_\ell Q_\ell = kNU_\ell$ and $G_\ell Q_\ell = k_1 NI$, with k, k_1 units, give this ideal since $U_\ell \Lambda$ is standard and G_ℓ gives a basis of Λ_0. In this basis the primary denominator has the valuation of the norm of $\det Q_\ell$ to power r, by elementary divisors. The factor $s \prod_\tau \iota_\tau(\det P_\ell)^{m_\tau}$ is integral. The coprime primary-kernel estimate gives the lower valuation $-\rho_v$ on the corrected inner term. At $\ell \mid M$ the lattice is standard and s is integral. The level correction is ℓ^{d_P} when $v_\ell(M) = 1$ and $\ell^{d_P v_\ell(M)}$ in the prime-power argument. The latter uses $N = \ell^e N'$, $M = \ell^k M'$ with $e \ge k$, $M' \mid N'$, and the residual similitude ℓ^{e-k} of conjugate norm $\ell^{2(e-k)}$. This supplies the torsion reach in the primary-kernel estimate. Thus in each level branch M^{d_P} suffices for the lower bound $-\rho_v$. The discriminant absorbs ρ_v and Q_M clears the average. The nonarchimedean inequality then gives an integral rule on the entire test set,

$$Z_{\mathrm{den}} s \, y, \qquad Z_{\mathrm{den}} = |\operatorname{disc} E|^{d_P} M^{f r n_\theta + d_P}.$$

The three branches give this one-copy bound directly. Adapted determinant · Off-level test · Prime-level test · Primary factorization and one-copy bound

For the archimedean reading, the fitted first-minimum and Gaussian theta estimate in the homogeneous norm give $\|y\|_{\infty,\sigma} \le J^{r/2} \operatorname{polyNorm}(P) \exp(9000 \Xi^2 K)$. The kernel count ratio is J^r and the projected vacuum supplies $J^{-r/2}$. The uniformization scalar is the common positive value $\xi_R = |\tau(\xi_0)|$ in the prescribed complex charts. The normalization used here multiplies the frame value by $(2\pi\xi_R)^{-d_P/2}$. The same choice of s bounds the remaining factor $|\sigma(s)| J^{r/2}$ by $\sqrt{\operatorname{rd} F}$ at every embedding. The nonzero multiples still span over F and their log norms are bounded by

$$B_\theta = 29000 \Xi^2 K + \log Z_{\mathrm{den}} + \tfrac{1}{2} \log \operatorname{rd} F \le 29066 \Xi^2 K + \tfrac{1}{2} \log \operatorname{rd} F.$$

Here $f r n_\theta + d_P \le 2K$ and $\log |\operatorname{disc} E| \le 6 f L_b$. Vacuum and kernel counts · Fitted homogeneous norm

For a scalar-zero weight λ put $S(\lambda) = \sum_{\tau,i} |\lambda_{\tau i}|$. Use dominance, $\log S(\lambda) \le 10 L_b$ and a gap of at least $8fr + 10$ between the two signature blocks. Put $m_* = fr + 1$, $v_* = 4m_*$, $n_\theta = 2v_*$, and $c_\tau = 1 + v_* - \lambda_{\tau,p-1}$, with indices starting at zero. In the shifted list $\lambda_\tau + c_\tau$ the last positive-block entry is $v_* + 1$ and the first negative-block entry is at most $-v_* - 1$. Subtract v_* from the positive entries and from the negatives of the reversed negative entries to obtain the dominant nonnegative even theta exponents. With $C_\lambda = \sum_\tau |c_\tau|$ we have

$$C_\lambda \le S(\lambda) + f(v_* + 1), \qquad K \le (1+r) S(\lambda) + fr(v_* + 1) \le 2\Xi S(\lambda) + 9\Xi^2.$$

The sharper middle bound with $\log(fr) \le L_b$ gives $\log K \le 20 L_b$.

Project integrally to Z and multiply by the fitted key power of scalar weight $-v_*$ and the flat section of parallel weight $-c_\tau$. These sections are supported on Z and multiplication identifies the rational component spaces. For the key use relative algebraic closedness in the component function fields, good level at every prime and the prescribed uniformizations with common absolute uniformization scalar. Its finite orders are nonnegative and its ordinary log norm cost is at most $2m_* fr \log(2\pi)$. The flat comparison requires the fitted points and their unramified-point property outside the designated bad primes. Let $\rho_\mathfrak{p}$ and $r_\sigma > 0$ be the lower valuations and the constant ordinary norm factors of the rational flat section. Multiplication by an element of $\prod_\mathfrak{p} \mathfrak{p}^{\lceil -\rho_\mathfrak{p} \rceil}$ makes those lower valuations nonnegative. Apply the disk argument in this ideal with radii proportional to r_σ^{-1}. It bounds the ordinary log norm cost by $-\deg_{\mathrm{lo}} + \Delta_\rho + \frac{1}{2}\log \mathrm{rd}\, F$, where

$$\deg_{\mathrm{lo}} = \frac{\sum_\mathfrak{p} \rho_\mathfrak{p} \log \mathrm{N}\mathfrak{p} - \sum_\sigma \log r_\sigma}{[F:\mathbb{Q}]}, \qquad \Delta_\rho = \frac{\sum_\mathfrak{p} (\rho_\mathfrak{p} + \lceil -\rho_\mathfrak{p} \rceil) \log \mathrm{N}\mathfrak{p}}{[F:\mathbb{Q}]}.$$

The sum over σ counts all embeddings, with conjugate-compatible norm factors. The flat comparison gives $-\deg_{\mathrm{lo}} \le (|c_1| + q \log \mathrm{rd}\, E) C_\lambda$, where $|c_1| \le 9 r L_b$. Orders are integral outside the bad-prime set, whose log sum bounds $\Delta_\rho \le 37 L_b$. Multiplication in the fitted homogeneous norm and the component projection therefore give an integral F-basis at weight λ with both log norms at most

$$B_\lambda \le B_\theta + 2m_* fr \log(2\pi) + (|c_1| + q \log \mathrm{rd}\, E) C_\lambda + \Delta_\rho + \tfrac{1}{2} \log \mathrm{rd}\, F.$$

Even weight and twists · Key power and norm · Flat section and degree · Bad-prime sum

For a full integral frame when $R = \mathrm{rk}\, M_Z(\lambda) > 0$, take $l = [F:\mathbb{Q}]$. Euclidean transference in the Minkowski embedding, with complex coordinates weighted by $\sqrt{2}$, gives l independent integers with every conjugate bounded by $\sqrt{l}\ \mathrm{rd}\, F$. Indeed the dual is the conjugate codifferent, whose nonzero vectors have Euclidean length at least $\sqrt{l}/\mathrm{rd}\, F$ by the ideal norm and arithmetic-geometric mean. Products with the short F-basis are lR independent lattice vectors over $\mathbb{Q}$. Reduce a $\mathbb{Z}$-basis into their half-open parallelepiped. The products and the reduced vectors give at most $2lR$ integral generators, of log norm at most $B_\lambda + \frac{3}{2}\log l + \log R + \log \mathrm{rd}\, F$. In rank zero the family is empty. The fitted scalar-zero Bergman bound gives $\log R \le 4000\Xi^3$ by integration of the evaluation kernel on the component. For u_1, its squared evaluation bound $\beta = e^{4000\Xi^3}$ yields the multiplication bound $\sqrt{\beta}$ by integrating the squared pointwise product in the same probability metric. The integral multiplication and the identification of that metric with the section and Chow frames use the fitted bundle and lattice comparisons.

The theta term contributes at most $58132\Xi^3 S(\lambda) + 261594\Xi^4$ apart from the field term. The flat term contributes at most $12\Xi^3 S(\lambda) + 108\Xi^2$. The key, rounding and two field terms contribute at most $(8 + 37 + 391)\Xi^2$. Adding the generator overhead gives log norms at most $326000(\Xi^3 S(\lambda) + \Xi^5)$. On the fitted quotient and remainder path use

$$\begin{aligned} S(u_1), S(u_2) &\le \Xi^2(\kappa_{\mathrm{wt}}/T + 1), \\ S(w(0)) &\le 3D\Xi^2(\kappa_{\mathrm{wt}}/T + 1), \\ S(w(i)) &\le \Xi^2 \kappa_{\mathrm{wt}} \qquad (0 \le i \le T). \end{aligned}$$

The path gaps give the signature-break hypothesis, using $D + 1 \ge 8fr + 10$ in the fitted range and the lower bound $(D+1)\,\mathrm{gap}(u_1) - 2$ at degree zero. The path size bounds and $\log \kappa_{\mathrm{wt}}, \log \Xi \le 2L_b$ give the required logarithmic weight range. Eventually $4 \cdot 326000\Xi^5 \le \mathrm{PP}$ and $1 + 391\Xi^2 + 4000\Xi^3 \le \mathrm{PP}$

under the fitted growth and $\mathrm{PP} \ge \Xi^{12}$. The full frames at u_1, u_2 then have log norm at most $K_*/T + \mathrm{PP}$ and log cardinality at most PP. The F-spans at $0, D, T$ have log norm bounds $D(K_*/T + \mathrm{PP}), K_*, K_*$, respectively. This supplies the degree-zero improvement and the bounds for the finite determinant. Banaszczyk, Euclidean transference · Bergman and rank bound · Fitted degrees · Quotient weight frame Short-section calculation · Finite-degree determinant

The adjacent-gap identity gives $u_- - w(T)$ as a parallel character with entries bounded in absolute value by $3(p_j + q_j)$. Require the corresponding flat section to induce an isomorphism on the rational component spaces and to supply finite lower valuations and archimedean norm factors in the fitted metrics. The fractional ideal with exponents the negatives of the floored lower valuations times the image of M_T is contained in the component lattice at u_-. The floor costs at most the log sum of bad primes per rank. With the fitted flat constant $|c_1| \le 9(p_j + q_j)L_b$, $\log \mathrm{rd}\, E \le 3L_b$, and the release's bad-prime support and level bounds, the total degree loss per rank is at most $120\Xi^2$. Thus $d = r_Z$ and $d_Z \ge g - 120d\Xi^2$. Fractional ideal comparison · Flat degree loss · Component rank comparison

By (16a), $(d - Q)/Q \le z/(1 - z)$ and $d/Q \le 1/(1 - z)$. In (16f), $\gamma \le \mathrm{PP}$ and $c_* \le K_*/T + 3\mathrm{PP}$, so $\max(a_D, a_T) \le 2K_* + \frac{7}{2}T\mathrm{PP}$. Also $D^2/T \le 1/14$ and $D^2K_*/T \le b/7$. Hence $DK_*/T \le b/7$, $D\mathrm{PP} \le b/14$, and $a_0 \le 12b/7$. The completion costs at most $4zK_* + 7zT\mathrm{PP} \le 15b/2$ per Q. Since $1 \le \Xi$ and $\Xi^{12} \le \mathrm{PP} \le \sqrt{b}$, we have $\Xi^2 \le b$. Combining the costs gives

$$\begin{aligned} &d_Z \ge Q_Z(Th_Z - \mathcal{E}_Z), \\ &\mathcal{E}_Z = a_0 + \frac{ya_D + (d - Q - y)a_T + 120d\Xi^2}{Q} \le 10b + 240\Xi^2 \le 250E_3. \end{aligned} \tag{16g}$$

The arithmetic Hilbert–Samuel limit in (16d) enters the height average below. The determinant bounds use only the spaces at the prescribed finite degrees.

Averaging and descent. Put $y_Z = h_Z + c_s \ge 0$, $S = \sum_Z r_Z$, $W = \sum_Z Q_Z$. By (16b) the unweighted mean of y_Z is at least $((1 - \rho_B)c_s - X - c_c)_+$. By (16a), $Q_Z/W \ge (1 - \rho_B)/|\mathcal{Z}|$ and $1 - z \le W/S \le 1$. Consequently

$$S^{-1}\sum_Z Q_Z y_Z \ge (1 - z)(1 - \rho_B)^2 c_s - X - c_c.$$

Subtract $c_s W/S \le c_s$ to get a lower bound for the height average with the same weights. Since $1 - (1 - z)(1 - \rho_B)^2 \le z + 2\rho_B$, summing (16g) and using (16b) and (16a) proves

$$\mu(U_Y) \ge -C_U E_3, \qquad C_U = 2(3C_H + 250). \tag{16h}$$

The shift by c_s makes the terms nonnegative when the unweighted mean is converted to the weighted mean. Release averaging with the weaker component estimate

Assume the flat twisting section gives the equal-rank determinant comparison $\mu(A_-) \ge \mu(U_Y) - P_j - C_F E_3$. Under the anchor bounds above, the source's comparison gives $C_F = 11C_A + 83$ when the logarithmic root discriminant is at most $2L$, $10 \le L \le E_3$, the even scalar numerator $2v$ satisfies $0 \le 2v \le n_V$, $f(p_j + q_j)n_V L \le b$, and the logarithmic bad-prime sum is at most $40L$. Indeed the signed anchor contribution costs at most $7C_A E_3$ beyond $-P_j$. The different term on the anchor costs at most $4C_A E_3$, the scalar twist costs at most $3b$, and the bad primes at most $80L$. Integral rank and slope comparison · Numerical losses

Assume there are n_ϕ integral F-linear maps from the span of A_- into the span of B, each sending A_- into B, whose common kernel is zero. Assume that for every embedding and every a in the source,

$$\sum_h \|\phi_h a\|_\sigma^2 \le e^{2b}\|a\|_\sigma^2, \qquad n_\phi \operatorname{rk} B \le \operatorname{rk} A_+ \le (1 + b/K_*)\operatorname{rk} A_-. \tag{16i}$$

Require the integral ring comparison $\mu(L_{\mathrm{core},j}) \ge \mu(B) - 2b$ and a spanning set of integral vectors in B of ordinary norm at most e^{K_*} at every embedding. For the latter, the release identifies B with the integral eigensublattice for a finite-order F-valued character χ of a finite level group. If the group acts integrally and isometrically, the sum $\sum_h \chi(h)^{-1} h$ carries short integral generators of the full rational space to B and acts as the group order on the eigenspace. Roots of unity are integral units of absolute value one at every embedding. The norm loss is at most the group order. The log norm bound $C_G \Xi^2 \kappa_{\mathrm{wt}}$ for the full even theta family and the group-order bound $45\Xi^2$ on the log scale give the needed bound K_* when $\kappa_{\mathrm{wt}} \ge 1$ and $\Xi \ge C_G + 45$. Integral character sum · Type and short-span comparison

Let $r_- = \operatorname{rk} A_-$ and $N = n_\phi \operatorname{rk} B$. The jointly injective map into the orthogonal sum $B^{\oplus n_\phi}$ has operator norm at most e^b. Its top exterior power has norm at most e^{br_-}. The arithmetic degree increases under saturation of its image, giving twice the degree of the saturation at least $r_-(\mu(A_-) - 2b)$. The torsion-free quotient with quotient norms is spanned over F by images of the short coordinate vectors. A maximal independent set of those images gives a full-rank free sublattice whose determinant norm is at most $e^{(N-r_-)K_*}$ at every embedding by Hadamard's inequality. Thus twice the degree of the quotient is at least $-2(N - r_-)K_*$, including the zero-rank case. Additivity of the degree gives $N\mu(B) \ge r_-(\mu(A_-) - 2b) - 2(N - r_-)K_*$. With $\epsilon_* = b/K_*$ we have $(N - r_-)/N \le \epsilon_*$ by (16i), so $\mu(B) \ge t - \epsilon_*(|t| + 2K_*)$ for any $t \le \mu(A_-) - 2b$. Descent through the orthogonal sum

Take $C_1 = C_F + C_U$ and $t = -P_j - C_1 E_3 - 2b$. Since $b, |P_j| \le K_*$, the rank loss is at most $C_1 E_3 + 5b$. The ring comparison and $b \le E_3$ give the claimed lower slope

$$\mu(L_{\mathrm{core},j}) \ge -P_j - \big[2(C_F + C_U) + 9\big] E_3. \tag{16j}$$

This holds uniformly on the family tail where all the stated inputs hold.

The core slope on the definite terminal lift

Proposition (terminal core slope). Use the fitted terminal row of the ordinary family in the normalization of (14)–(16). Write $p = p_m = k + m - 1$, $q = q_m = 0$, F for the core coefficient field, Ψ for the chosen CM type, $f = |\Psi|$, $\omega = \log \Omega$, and $\ell_E = \log \mathrm{rd}\, E$. Here F is totally complex. In the notation of the release, the terminal weight is $\mathrm{wtZ}(a, n)$ with $a = \mathrm{wtA}(\kappa_{m-1})$, where the source lift index starts at zero. The references for this argument are pinned to commit `e7c8972c701d3396947808ab20cc16dc14623d33`. Let $a_\tau = (a(\tau, 0), \dots, a(\tau, p-1))$, an antitone list of nonnegative integers, d_τ its Weyl dimension, and $n \ge 0$ the even theta exponent. Set

$$\lambda_{\tau,i} = a(\tau, i) + n/2 \quad (0 \le i < p), \qquad g_A = fp, \qquad c_* = -p\omega/2.$$

Here g_A is the abelian dimension of the terminal object, and ℓ_E is the logarithm of the root discriminant of the CM **host** E. Put $\Delta = \log N_{E/\mathbb{Q}} \mathfrak{D}_{E/k_0} / [E : \mathbb{Q}]$ for the relative different of the host over the inducing quadratic subfield. Under the fitted PEL, integral test, frame and level-action hypotheses below, the core type lattice has positive rank and satisfies

$$\begin{aligned}
&\mu(L_{\mathrm{core},m}) \geq 2B_{\mathrm{def}},\\
&\qquad B_{\mathrm{def}} = \sum_{\tau\in\Psi}\left[\left(\frac{|a_\tau|}{p}+\frac{n}{2}\right)c_* - \frac{\log d_\tau}{2}\right] + \frac{ng_A}{4}\big(\log(2\pi)-\Delta\big),\\
&\qquad 2B_{\mathrm{def}} = -\omega\sum_{\tau\in\Psi}\sum_{i=0}^{p-1}\lambda_{\tau,i} - \sum_{\tau\in\Psi}\log d_\tau + ng_A\left(\frac{\log(2\pi)}{2}-\Delta/2\right).
\end{aligned} \tag{16k}$$

In particular, under the selected size and anchor estimates below,

$$\mu(L_{\mathrm{core},m}) \geq -P_m - 135E_3, \qquad P_m = p_m\sum_{\rho} w_{m,\rho}\omega, \quad E_3 = B_0\log\log x. \tag{16l}$$

The first bound is pointwise. For the second, fix the family and its auxiliary parameters, and take the maximum of the thresholds for the stated fitted and size inputs. Thereafter the bound holds for every x in the selected range, every selected index and every fitted row datum to which the inputs apply. The components and frame matrices can vary with the row. The estimates must apply to the same terminal realization as the core lattice and its probability Petersson metric. Terminal weight and definite hypotheses · Core slope · Core lattice and norm · Family quantifiers

The framed-point lattice. On the fitted base model X/F, use signature $(p,0)$, auxiliary level parameter $A=1$, good local models at every rational prime in the source's sense, and the prescribed complex uniformizations and frame domains. Assume F is algebraically closed in every component function field and the components form a finite nonempty set $\mathcal{Z}$. In signature $(p,0)$ the components have dimension zero. On each $Z\in\mathcal{Z}$ there is a framed F-point y_Z whose underlying point gives the unique point of that component after base change, in the stated framed moduli problem. At each finite w, choose $b_{Z,w}\in F^\times$ and invertible matrices $B_{Z,w,\tau}$ such that the changes from y_Z to the integral frames over any of the testing DVRs above w have theta scalar $b_{Z,w}u$ and Lie matrices $(B_{Z,w,\tau}U_\tau)^{\mathsf{T}}$, with u a unit and U_τ integral invertible. The scalar and matrices may be taken equal to 1 and the identity outside a finite set. This follows by taking an integral frame over a completion, approximating the change of frame over F modulo integral units, and using that integral frames over the same point differ by integral units. The rational frame spreads out over a localization of $\mathcal{O}_F$. At $\sigma:F\hookrightarrow\mathbb{C}$, use invertible $X_{Z,\sigma,\tau}$ which put the Lie frame into its metric normalization by the changes $X_{Z,\sigma,\tau}^{\mathsf{T}}$, and let $N_{Z,\sigma}>0$ be the resulting squared theta frame norm parameter. A change from there to unit metric frames uses unitary Lie changes and a theta scalar of modulus $N_{Z,\sigma}^{-1/2}$. The data at conjugate embeddings give equal $N_{Z,\sigma}$ and equal moduli of the Lie frame determinants. Relative algebraic closedness · Finite and infinite frame conditions · Existence of the local and rational frames · Almost everywhere triviality and assembly

Use the integral standard bideterminant basis of the dual Weyl polynomials of shape a_τ. Let $\mathcal{I}$ be the product of their basis index sets and $d_{\mathrm{fib}}=\prod_{\tau\in\Psi}d_\tau$ its cardinality. Embeddings outside Ψ have Lie multiplicity zero and contribute a one-dimensional scalar factor. In each factor let $\rho_\tau(M)$ be the matrix in the integral basis of the polynomial action by M^{T} as in the source's rhoMat. Let C_τ be the change from the integral basis to orthonormal coordinates for the compact dual Weyl squared norm $d_\tau\int_{\mathrm{U}(p)}|P(u)|^2\,du$ on polynomials P, with probability Haar. In coordinates $x\in F^{\mathcal{I}}$ the adelic lattice and comparison norm on Z are

$$A_{Z,w} = b_{Z,w}^{n} \bigotimes_{\tau\in\Psi} \rho_\tau(B_{Z,w,\tau}), \qquad L_Z = \{x \in F^{\mathcal{I}} : A_{Z,w}x \in \mathcal{O}_{F,w}^{\mathcal{I}} \text{ for every finite } w\},$$
$$Y_{Z,\sigma} = N_{Z,\sigma}^{-n/2} \bigotimes_{\tau\in\Psi} \big(C_\tau \rho_\tau(X_{Z,\sigma,\tau})\big), \quad \|x\|_{Z,\sigma} = \|Y_{Z,\sigma}\sigma(x)\|_2.$$

These are the matrices matA and matY. The local matrices give the localized lattices $A_{Z,w}^{-1}\mathcal{O}_{F,w}^{\mathcal{I}}$. They are standard almost everywhere, so L_Z is a full lattice of rank $d_{\mathrm{fib}} > 0$. In particular the inverse matrices belong in the description of the local *lattice*, and $Y_{Z,\sigma}$ acts directly in the norm. Actual coordinate definitions · Tensor matrices · Adelic lattice construction

Evaluation in the rational frame identifies $F^{\mathcal{I}}$ with the space of cusp sections supported on Z, via a map ϕ_Z. On a new frame of change M_τ^{T} and theta scalar s, the value of the section is s^n times the tensor polynomial with coefficients x evaluated on the M_τ. The frame laws and the rational point make this expression a natural section on the component, extended by zero using its idempotent. The triangular weight law recovers these same dual Weyl coordinates from a section. The compact definite cusp condition gives the entire component space. On any tested integral frame, transport first replaces the coordinate vector by $A_{Z,w}x$, and the remaining evaluation uses integral polynomial vectors on integral matrices and multiplication by u^n. It follows that $\phi_Z(L_Z) \subset \mathrm{latW}$, where latW means nonnegative valuation on every test of the algebraic section at every finite place. At infinity the remaining Lie changes are unitary, and compact matrix-coefficient evaluation at every unitary frame has norm at most one in the tensor comparison norm. This is the Schur-orthogonality bound with the factor d_τ in the squared polynomial norms. The pointwise normalized value homN, including the source's Lie factor on its uniformization charts, is consequently bounded by $\|x\|_{Z,\sigma}$. This comparison uses the good representative of each conjugate pair in the definition of normW. Frame reading and surjectivity · Finite and infinite comparison · Passage to the uniformization metric · Section lattice

Determinants and finite-place losses. Write $\ell_w = \log N\mathfrak{p}_w/[F:\mathbb{Q}]$, $v_w(\mathfrak{p}_w) = 1$, and av_σ for the average over all embeddings. Define

$$D_{\tau,Z} = \sum_w \ell_w v_w(\det B_{Z,w,\tau}) - \mathrm{av}_\sigma \log|\det X_{Z,\sigma,\tau}|,$$
$$D_{\theta,Z} = \sum_w \ell_w v_w(b_{Z,w}) + \tfrac{1}{2}\,\mathrm{av}_\sigma \log N_{Z,\sigma}.$$

The determinant ideal of L_Z is specified locally by $(\det A_{Z,w})^{-1}$, whence the arithmetic degree per rank, equal to $\mu(L_Z)/2$, is $\big(\sum_w \ell_w v_w(\det A_{Z,w}) - \mathrm{av}_\sigma \log|\det Y_{Z,\sigma}|\big)/d_{\mathrm{fib}}$. The character $\det\rho_\tau$ is $(\det)^{d_\tau|a_\tau|/p}$. Indeed the central scalar tI acts as $t^{|a_\tau|}$ on each coordinate and all algebraic determinant characters of GL_p are integral powers of $\det$. The tensor determinant identity then gives

$$\frac{\mu(L_Z)}{2} = nD_{\theta,Z} + \sum_{\tau\in\Psi}\left(\frac{|a_\tau|}{p}D_{\tau,Z} - \frac{\log|\det C_\tau|}{d_\tau}\right) \;\geq\; nD_{\theta,Z} + \sum_{\tau\in\Psi}\left(\frac{|a_\tau|}{p}D_{\tau,Z} - \tfrac{1}{2}\log d_\tau\right).$$

Every basis bideterminant is a product of minors of modulus at most one on unitary matrices. Its norm squared is at most d_τ, and Hadamard gives $|\det C_\tau|^2 \leq d_\tau^{d_\tau}$. The local matrices can have denominators, all accounted for by the displayed signed valuations. The determinant character introduces no prime denominator p in the lattice. The apparent fraction $|a_\tau|/p$ results from division by the rank. Determinants and slope · Tensor determinant calculation · Hadamard bound in the actual metric

Integral Lie model and determinant. Write k_0 for the imaginary quadratic subfield from which Ψ is induced, $\mathfrak{d} = \mathfrak{D}_{E/k_0}$ for its relative different, and $\Delta = \log N_{E/\mathbb{Q}}\mathfrak{d}/[E:\mathbb{Q}]$. Choose a principally polarized CM elliptic curve C of the specified k_0-type, with projective period module J after removing a common period. If Γ is the projective $\mathcal{O}_E$-period module of the terminal point, set $M = J^{-1} \otimes_{\mathcal{O}_{k_0}} \Gamma$. The point is the Serre tensor $C \otimes_{\mathcal{O}_{k_0}} M$ on the generic fiber after finite extension. Indeed the complex uniformizations have the same period module by multiplication and the same complex structure by the induced type. The resulting isomorphism of complex abelian varieties is algebraic and descends to a finite extension. Transfer the prescribed polarization along this isomorphism.

Choose a finite extension carrying the generic isomorphism and good reduction of C everywhere, which exists by semistable reduction and potential good reduction of CM elliptic curves. The projective module M realizes the tensor by an idempotent on a power of the good elliptic scheme, so its Lie module over a testing DVR R is $M \otimes_{\mathcal{O}_{k_0}} \operatorname{Lie}(C/R)$. The comparison below identifies this module with the base-changed full Lie lattice of the test. This uses the actual Cartesian generic identification of the group in the PEL test, with its action, line and rigidification.

Let L be this full Lie lattice and $L_\tau = L \cap V_\tau$ its intersection with the generic τ-eigenspace. The L_τ commute with flat extensions of DVRs by their kernel description using an integral basis of $\mathcal{O}_E$. The order $S_R = \mathcal{O}_E \otimes_{\mathcal{O}_{k_0}} R$ embeds by evaluation in R^f, and its trace dual is the image of $\mathfrak{d}^{-1} \otimes_{\mathcal{O}_{k_0}} R$. Projection of this dual to coordinate τ is $(\tau\mathfrak{d})^{-1}R$. Dualizing on that axis gives $S_R \cap \operatorname{Frac}(R)e_\tau = (\tau\mathfrak{d})Re_\tau$. The squared determinant ideal of the evaluation map is the extended relative discriminant, since the trace Gram matrix is the transpose of the evaluation matrix times that matrix. Thus the projection sum contains S_R with colength $\frac{1}{2}\sum_\tau v_R(\tau\mathfrak{d})$. As $M \otimes_{\mathcal{O}_{k_0}} R$ is free of rank p over the semilocal order, the two calculations give

$$L_\tau = (\tau\mathfrak{d})R\,(M \otimes_{\mathcal{O}_E,\tau} R) \otimes_R \operatorname{Lie}(C/R), \qquad c_R = \operatorname{length}_R(L/\bigoplus_\tau L_\tau) = \frac{p}{2}\sum_\tau v_R(\tau\mathfrak{d}).$$

Here the first equality is inside the generic eigenspace. For the second, subtract the colength of L in the projection sum from the colength of $\bigoplus_\tau L_\tau$ there. Both lengths and valuations use a uniformizer of R. Put $c_w = c_R/e(R/O_{F,w})$ when testing above w of the core field F. Under finite extension each length scales by ramification, and the weights $\ell_w = \log Nw/[F:\mathbb{Q}]$ account for the residue degrees via $\sum_{w'|w} e(w'/w)f(w'/w) = [F':F]$. Thus $\sum_w \ell_w c_w = g_A\Delta/2$, since the weighted valuation sum of $\tau\mathfrak{d}$ equals Δ for each fixed label. This argument also applies at 2.

For the polarization, write the elliptic form as $\operatorname{Tr}_{k_0/\mathbb{Q}}(t_0 x\bar{y})$ with $(t_0)J\bar{J} = \mathfrak{D}_{k_0/\mathbb{Q}}^{-1}$. The point form is $\operatorname{Tr}_{E/\mathbb{Q}}(t_0 h(x,y))$ for a positive Hermitian form h on $M \otimes_{\mathcal{O}_E} E$. By the tower trace pairing, unimodularity on JM says that the h-dual of M is $\mathfrak{d}M$. If $\mathfrak{a} = \det M$ in an E-basis and $\gamma_h = \det h$, then $(\gamma_h)\mathfrak{a}\bar{\mathfrak{a}} = \mathfrak{d}^{-p}$. The squared determinant norm in the induced Appell–Humbert metric is the p-th power of the squared elliptic tangent norm times the conjugate of γ_h. The finite determinant ideal has the elliptic factor to power p times $(\tau\mathfrak{d})^p\tau\mathfrak{a}$. Thus, with arithmetic degrees divided by $[F:\mathbb{Q}]$,

$$-D_{\tau,Z} = \widehat{\deg}\det L_\tau = p\,\widehat{\deg}\operatorname{Lie} C - p\Delta - \frac{\log N\mathfrak{a} + \frac{1}{2}\log|N_{E/\mathbb{Q}}\gamma_h|}{[E:\mathbb{Q}]} = p\big(\widehat{\deg}\operatorname{Lie} C - \tfrac{1}{2}\Delta\big).$$

The first sign follows because $B_{Z,w,\tau}$ gives the integral determinant ideal and the rational frame has determinant norm $|\det X_{Z,\sigma,\tau}|^{-1}$. With the discriminant-one elliptic differential the squared tangent norm is π/area, and the elliptic Kronecker limit calculation of (16) gives $2\widehat{\deg}\operatorname{Lie} C = \omega + \Delta$, since

$\Delta = \log \operatorname{rd} E - \frac{1}{2}\log|\operatorname{disc} k_0|$. Therefore $D_{\tau,Z} = -p\omega/2$. The relative different contributes $-2p\Delta$ to the doubled Lie degree from the eigenlattice and $+p\Delta$ from the polarization determinant. Serre tensors and polarizations · Pure-type tensor supplier · Lie degree on the fitted point · Reader note

The global theta degree. Take a finite extension F'/F over which the elliptic curve has good reduction everywhere and the generic tensor isomorphism is defined. Write $\mathcal{A}/\mathcal{O}_{F'}$ for the resulting good abelian scheme. To compare at $w' \mid w$, pass from the complete testing DVR and the completed local ring at w' to a common finite extension. The actual integral test uses a smooth, separated, finite-type group with geometrically connected semiabelian fibers. Its two Cartesian generic maps preserve the $\mathcal{O}_E$-action and identify the rigidified cubical lines and the theta section.

Over the common DVR, the Néron mapping property extends the prescribed generic isomorphism to a homomorphism from the test group to $\mathcal{A}$. For an auxiliary prime ℓ invertible in the DVR, a geometric ℓ-torsion point of the special-fiber kernel lifts over the strict henselization because multiplication by ℓ is étale on the test group. Its image in the good model's finite étale ℓ-torsion is zero, so the generic isomorphism and separatedness force the lift to be zero. A positive-dimensional kernel would have a positive-dimensional reduced identity component in the geometric semiabelian fiber and hence nontrivial ℓ-torsion. The map is therefore quasi-finite. The source is integral by smoothness, flatness and integrality of the generic fiber. Zariski's main theorem makes the separated birational quasi-finite map to the normal good model an open immersion. Its image is the full generic fiber and, on the geometric special fiber, an open subgroup of a connected group. Thus the map is an isomorphism. This argument also works above 2, using an odd auxiliary ℓ.

Extend the specified generic line by closing a divisor on the regular abelian scheme and rigidifying at zero. Over a DVR the divisor of the generic identification of two extensions is a multiple of the single integral special fiber. Its pullback to the rigidifying section has order zero, forcing that multiple to vanish. This identifies the test line with the global extension $\mathcal{L}$. Symmetry and the cube structure extend by the same uniqueness and the theorem of the cube. The polarization homomorphism and its generic inverse extend to the dual abelian schemes, so the translation kernel is trivial on each fiber. Upper semicontinuity gives a nonzero section on each fiber, and an effective line with finite translation kernel on an abelian variety is ample. Thus $\mathcal{L}$ is relatively ample and principal. Riemann–Roch, vanishing for ample lines and cohomology and base change give an invertible $P = \pi_*\mathcal{L}$ with formation commuting with base change. Pazuki's Definition 3.1 uses the original cubist line and a finite flat model for the subgroup $K(L^{\otimes 2})$. Here that model is $\mathcal{A}[2]$, including above 2. Choose the auxiliary generic isomorphism of Proposition 2.5 after finite extension, with the marked set empty.

The test explicitly requires θ_R to generate $H^0(G, L_G)$, the module of global sections of the tested line. Flat base change of H^0 on the quasi-compact separated test group follows from a finite affine cover and its equalizer for sections. The generator survives comparison to $\mathcal{A}$. The frame change gives $\theta_R = b_{Z,w} u s$ for an integral unit u and for s the section in the rational frame. Thus in the rational basis s the local ideal of P has valuation $v_w(b_{Z,w})$ in the normalization of F. The comparison over the common DVR gives equality over the completion at w' by dividing by ramification. Summing over $w' \mid w$ preserves the normalized finite term, since $\sum_{w'|w} e(w'/w) f(w'/w) = [F' : F]$.

At infinity the release represents the same section s by a function ϑ in an Appell–Humbert trivialization that sends the rigidifying vector at zero to 1. Its pointwise squared norm $|\vartheta(z)|^2 e^{-\pi H(z,z)}$ descends by the multiplier. It has translation-invariant curvature and makes the rigidifying vector unit, which characterize

the cubist metric. Dividing the integral over the fundamental domain by its volume gives Haar probability, so $N_{Z,\sigma} = \|s\|_{P,\sigma}^2$. Using the ordinary norm in the degree gives

$$-\widehat{\deg}_{\mathrm{nor}} P = \sum_{w} \ell_w v_w(b_{Z,w}) + \tfrac{1}{2}\operatorname{av}_\sigma \log N_{Z,\sigma} = D_{\theta,Z}.$$

For the Hodge metric, take symplectic periods (I, T) and $Y = \operatorname{Im} T$. The polarization form has Hermitian matrix Y^{-1}, so the fitted Lie matrix is πY^{-1}. Its dual determinant has squared norm $\det Y / \pi^{g_A}$ on $dz_1 \wedge \cdots \wedge dz_{g_A}$. This is Deligne's squared norm $i^{g_A^2}(2\pi)^{-g_A}\int \alpha \wedge \overline{\alpha}$, by integrating on the period parallelepiped of volume $\det Y$. The full Lie determinant consequently has normalized degree $-h_F(A)$, where $h_F(A)$ denotes the stable Faltings height. The slope theorem of Moret-Bailly and Bost in Pazuki's Theorem 3.4(iii)–(v) gives

$$\frac{\widehat{\deg}_{\mathrm{nor}}(\pi_* \mathcal{L})}{\rho} = -\tfrac{1}{2} h_F(A) + \tfrac{1}{4}\log\!\left(\frac{\rho}{(2\pi)^{g_A}}\right), \qquad \rho = \frac{c_1(L)^{g_A}}{g_A!}.$$

For the principal line $\rho = 1$, so $D_{\theta,Z} = \tfrac{1}{2} h_F(A) + \tfrac{g_A}{4}\log(2\pi)$.

Distinct Lie eigenspaces are orthogonal at infinity for the polarization metric, since the $\mathcal{O}_E$-action has conjugate adjoint. The saturated eigenspace sum has finite colength c_w in the full Lie lattice. Its normalized determinant degree is the full Lie degree minus $\sum_w \ell_w c_w$, so $\sum_\tau D_{\tau,Z} = h_F(A) + \sum_w \ell_w c_w$. Inserting the relative different calculation $\sum_w \ell_w c_w = g_A \Delta/2$ gives

$$D_{\theta,Z} = \tfrac{1}{2}\sum_{\tau} D_{\tau,Z} + \tfrac{g_A}{4}\log(2\pi) - \tfrac{g_A}{\Delta}/4.$$

Pazuki, Definition 3.1 and Theorem 3.4 · Integral test · Fitted complex metric · Reader note

Substituting $D_{\tau,Z} = -p\omega/2$ and the exact theta degree into the fiber determinant bound gives

$$\mu(L_Z) \geq -\omega \sum_{\tau,i} \lambda_{\tau,i} - \sum_{\tau} \log d_\tau + n g_A\big(\tfrac{1}{2}\log(2\pi) - \tfrac{1}{2}\Delta\big) = 2B_{\mathrm{def}}.$$

The theta term contributes $-n g_A \omega/2$ to the signed weight and $-n g_A \Delta/2$ from the finite colength. Since $0 \leq \Delta \leq \ell_E$, this also implies the original absolute-different bound and preserves its size allowance in (16l).

The fourth-power identity in full frames. Suppress the component index Z in the frame data in this paragraph. Evaluation of P gives a symmetric relative effective Cartier divisor Θ. Its section is nonzero on every fiber by cohomology and base change, and inversion preserves its line and hence its zero scheme. Evaluation and rigidification give $\mathcal{O}(\Theta) \simeq \mathcal{L} \otimes \pi^* P^{-1}$ and $e^*\mathcal{O}(\Theta) \simeq P^{-1}$. Moret-Bailly's Theorem 0.2 (1990), citing Chai–Faltings, gives a generating section of $(P^{\otimes 2} \otimes \lambda)^{\otimes 4}$ on the stack of principally polarized abelian schemes with a symmetric theta divisor over $\mathbb{Z}$, with λ the Hodge determinant. This applies above 2, where the closures of the parity components meet (Proposition 1.4(i)). Write $\kappa(s, \mathbf{v})$ for the coefficient of this section in $s^{\otimes 8} \otimes \eta_{\mathbf{v}}^{\otimes 4}$, with $\eta_{\mathbf{v}}$ the differential dual to a full Lie frame $\mathbf{v}$. The coefficient is a unit in the full integral frame with the section generator. Under $(s, \mathbf{v}) \mapsto (as, Q\mathbf{v})$ it is multiplied by $a^{-8}(\det Q)^4$.

Moret-Bailly's Theorem 0.5 fixes the even classical trivialization up to sign. Theorem 3.3 gives ordinary norm $(2\pi)^{4g_A}$ on both complex components, using the product of an odd divisor with itself for the odd case. In the metrics of that theorem the canonical divisor section $s \otimes s^\vee$ has squared Haar norm $2^{-g_A/2}$.

Thus $s^\vee$ has ordinary norm $2^{-g_A/4}N_\sigma^{-1/2}$ in $e^*\mathcal{O}(\Theta)$. The ordinary Hodge norm is $\pi^{g_A/2}n_\eta^{1/2}$, with n_η the squared Deligne norm of $\eta_{\mathbf{v}}$. The tensor frame then has ordinary norm $(2\pi)^{2g_A}N_\sigma^4 n_\eta^2$. In the rational frame, $n_\eta = \prod_\tau |\det X_{\sigma,\tau}|^2$, because the metric Lie frame is orthonormal. Write κ_0 for the coefficient in that rational frame. The finite change from a full integral Lie frame to the integral eigenframe has normalized determinant valuation c_w. The two identities in the fitted rational frame are therefore

$$v_w(\kappa_0) - 8v_w(b_w) + 4\sum_\tau v_w(\det B_{w,\tau}) = 4c_w,$$
$$|\kappa_0|_\sigma N_\sigma^4 \prod_\tau |\det X_{\sigma,\tau}|^4 = (2\pi)^{2g_A}.$$

These give the unit and absolute-value assertions for the integral trivialization. The degree argument uses the slope theorem directly. Moret-Bailly, Theorems 0.2, 0.5 and 3.3 · Reader note

Free level orbits and probability mass. The group here is the fitted reference level group $G = (\mathrm{lv2.lvl(false)}).P$, acting by the actual level moves $\rho(g) = \mathrm{actLev}(g)$ on sections. The character χ is the descent to $F^\times$ of the fitted row character $\chi_{\mathrm{FW}}\psi^{-1}$. The comparison of the level moves under the row realization identifies $L_{\mathrm{core},m}$ with $\mathrm{latW} \cap V[\chi]$ inside the definite cusp space V. The moves preserve latW, permute component idempotents, and preserve the pointwise normalized norms under transport. Their action on components is free because the components are rational points and the fitted level has the free-point property. The release deduces this property from the base level $M = g_0 g_1$ with distinct primes, unipotence of the level moves modulo both prime ideals used in the level condition, and faithfulness of the fitted action modulo $\mathfrak{M}_b$. More explicitly, a fixed point supplies a finite-order automorphism of the polarized fiber and a compatible action on its period lattice. On the sublattice of periods lifting the tested level, the finite-order operator is unipotent modulo each of the two distinct primes. The two-prime lemma makes it the identity on that sublattice, and faithfulness on the level then forces $g = 1$. Actual group, character and lattice equality · Fitted two-prime conditions · Free action and the period lattice · Integral level transport

Choose one representative o per component orbit. The map to the eigenspace is

$$\Phi((x_o)_o) = \sum_o \sum_{g\in G} \chi(g)^{-1}\rho(g)\phi_o(x_o).$$

Each summand lies on go, and each component occurs once, so restriction to o recovers $\phi_o(x_o)$. Conversely, a χ-eigenvector has component go equal to $\chi(g)^{-1}\rho(g)$ applied to its component at o. Thus Φ is a rational isomorphism onto $V[\chi]$, giving the exact core rank $|\mathcal{Z}/G|d_{\mathrm{fib}} > 0$, with $|\mathcal{Z}/G| = |\mathcal{Z}|/|G|$. Each $\chi(g)^{-1}$ is an algebraic unit of absolute value one at every embedding. The unnormalized orbit sum therefore sends $\bigoplus_o L_o$ into $L_{\mathrm{core},m}$ with full rank, without a factor $1/|G|$ or a discarded rank. Any finite index of its image in $L_{\mathrm{core},m}$ contributes $+2\log[L_{\mathrm{core},m} : \Phi(\bigoplus_o L_o)]/([F:\mathbb{Q}]\,\mathrm{rk}_F\, V[\chi])$ to the squared-norm slope when evaluated with the transported core metric. The inclusion suffices because this index term is nonnegative. Integral unnormalized orbit sum and rank comparison

The core norm squared is $\int \mathrm{homN}^2\, d\mu_{Q,1}$ over the actual automorphic quotient, where $\mu_{Q,1}$ is the normalized probability measure. In particular homN includes the factor $\exp(-\frac{1}{2}\sum_{\tau,i} a(\tau,i)\log(2\pi\xi_R))$ from $\mathrm{lieFacZ}$ on this weight. If $m_{\sigma,Z}$ is the mass of the test frames on Z under that probability measure at σ, then $\sum_Z m_{\sigma,Z} = 1$. Pointwise transport gives

$$\|\Phi(x)\|_{\mathrm{core},\sigma}^2 \le \sum_o \left(\sum_{g\in G} m_{\sigma,go} \right) \|x_o\|_{o,\sigma}^2 \le \sum_o \|x_o\|_{o,\sigma}^2.$$

The source uses the uniform pointwise bound and the total probability mass of one. The displayed inequality makes Φ a contraction from the orthogonal sum, so its top exterior power has norm at most one. The index inequality and additivity of degree give $\mu(L_{\mathrm{core},m}) \ge \sum_o d_{\mathrm{fib}} \mu(L_o) / \operatorname{rk}_F V[\chi] \ge 2B_{\mathrm{def}}$, proving (16k). Quotient probability normalization · Lie factor and probability Petersson norm · Petersson pointwise bound · Orthogonal-sum determinant comparison

Comparison with the tied reference weights. Let A_τ be the entry $\lambda_{\tau,k-2}$, with zero-based indexing. This is the source anchor with argument $k-1$, since the anchor definition uses coordinate $(k-1)-1$. The ordinary parameters ensure this entry exists. Put $S_\tau = \sum_{i=0}^{p-1} |\lambda_{\tau,i} - A_\tau|$ and $P_A = p \sum_\tau A_\tau \omega$. The triangle inequality, $\log(2\pi) > 0$, and $n \ge 0$ give

$$2B_{\mathrm{def}} \;\ge\; -P_A - |\omega| \sum_\tau S_\tau - \sum_\tau \log d_\tau - n g_A |\ell_E|.$$

The main term in the release uses the tied reference weights $w_{m,\rho}$, so the anchor comparison gives $|P_m - P_A| \le 100 E_3$. The last three errors together are at most $35E_3$ in the selected normalization. For details, write $L = \log b = x^\varepsilon$. The weight comparison bounds the first $k-1$ entries relative to the anchor by $\mathrm{Sc} + 2p$ each, where Sc is the weight spread budget of the family, and bounds the remaining m entries by $8b$ each. Since $k\,\mathrm{Sc} \le b$, $p^2 \le b$ and $m \ge 2$, the spread is at most $(k-1)(\mathrm{Sc}+2p) + 8mb \le 3b + 8mb \le 10mb$. The Weyl formula and determinant-shift invariance give

$$1 \le d_\tau = \prod_{0\le i<h<p} \frac{\lambda_{\tau,i} - \lambda_{\tau,h} + h - i}{h-i} \le (1+S_\tau)^{\binom{p}{2}}.$$

Using $L \ge 1$, $\log(11m) \le 1 + L$, and $p^2 L \le b$ yields $\log d_\tau \le 3b$. With $|\omega| \le 3L$, $n = n_V$, $|\ell_E| \le 2L$, and $f p n_V L \le b$, the three errors are bounded by $30 f m b L$, $3fb$, and $2b$, respectively. Their sum is at most $35E_3$ because $E_3 = fmbL \log\log x$ on the selected size range, where $\log\log x \ge 1$, $f \ge 1$, and $m \ge 2$. For the anchor error the source identifies the CM indices with the tied reference coordinates, bounds the per-index anchor/reference difference by $\mathrm{Sc} + n_V + m + t + 2$, and uses $p(\mathrm{Sc} + n_V + m + t + 2) \le 20b$, with $t = n_V - p$. Summing after multiplication by $|\omega|$ gives at most $60 fbL \le 100 E_3$. This comparison preserves the sign of ω, and passing from lower slope to upper covolume reverses the sign of the main term. Anchor and degree conventions · Signed terminal arithmetic · Weight and dimension bounds · Sum of terminal errors · Anchor comparison · Numerical assembly

The fitted row pins and coefficient-field comparison give relative algebraic closedness in the component function fields, which is the branch in the source with zero error in (16k). The framed-point lattice and integral orbit sum are constructed on that base model, so (16k)–(16l) require no deep cover. Combining (16j) and (16l) gives $C_{\mathrm{low}} = \max\{135, 2(C_F + C_U) + 9\}$ on the common family tail. Relative algebraic closedness from the fitted pins · Zero error in the core comparison · Elimination of the other case · Disjoint cases at the final slope assembly

With (16j) and (16l), suppose $\log R_j \le C_{\mathrm{mass}} E_3$. Then $\mu(L_j) \ge -P_j - (C_{\mathrm{low}} + C_{\mathrm{mass}}) E_3$ by (15). The selected bounds $1 \le b \le E_3$, $|P_j| \le x^B b$, $\Lambda \le x^A b$, $2\log g \le b$, and $d_b/(d_b + d_g) \le 3\exp(-x^{e_z/7})$ allow $Q = (x^A + x^B + C_{\mathrm{low}} + C_{\mathrm{mass}} + 1)E_3$ in (14). For fixed $A, B, C_{\mathrm{low}}, C_{\mathrm{mass}}$ and $e_z > 0$,

eventually $12\exp(-x^{e_z/7})Q \le E_3$, proving (13) at every selected index and lift. This proof also gives the sharper pointwise term in (14) when the actual bad rank and index are available. Quantitative projection on the selected tail

Finite place and word-height comparisons. Write $\mathscr{E} = LG + \sum_j (E_{1,j} + E_{2,j}) + E_3^{\mathrm{fin}}$ for the reference-row comparison errors, with E_3^{fin} distinct from the logarithmic scale E_3. In the notation of Paper IV §4, the spectral and quotient comparisons on the actual fitted data are required to give

$$\mu_0^{\mathrm{ref}} \le \sum_j \frac{\mathrm{lcv}(X_{W,j}, M_j)}{D_0} + 2 \sum_{j,w \in \mathrm{Tw}} \ell_w e_w F_0(w,j) + \mathscr{E}.$$

The comparison pairs the index at w with that at $\bar{w}$ and the local Hermitian Gram valuation. If H is the intermediate form with lattice A and transport θ, the Gram determinant on A has valuation $v_w(\det H) + d_w(A) + d_{\bar{w}}(A)$. Subtracting this valuation and adding both signed indices $d_v(\theta A) - d_v(M_j)$ cancels the two $d_v(A)$. The resulting negative valuation is that of the transported form on M_j. At $w = \bar{w}$ both index terms still count. In particular the displayed finite main term is Fo of Paper IV, twice the release's undoubled place sum. The remaining fitted-row passage requires $\mu(RR^\vee) - \mathrm{av}_v \log \kappa_v \le \mu_0^{\mathrm{ref}} + L_{\mathrm{pass}}$, with the same κ_v used for the word height. Finite signed Gram comparison · Reference bound and passage · Factor of two in the fitted data

Use $\nu_{j,\rho,i}$ for the nonnegative shifted column weights of §4 and $\lambda_{j,\rho}$ for the Young diagram of the sorted nonnegative integral positive row weights. The hook term $H_{j,\rho} = \log \prod_{\square \in \lambda_{j,\rho}} h(\square)$ in $h_{X,j}$ comes from the Weyl norm and differs from the dimension factors of (15). For a partition with ℓ rows, zero padding allowed, its hook product is $\prod_i (\lambda_i + \ell - i)! / \prod_{i<j} (\lambda_i - \lambda_j + j - i)$. The denominator is the numerator product in the Weyl dimension formula. This follows by removing from row i of the full factorial $(\lambda_i + \ell - i)!$ the factors $\lambda_i - \lambda_j + j - i$ for $j > i$. Set $a^*_\rho = a_t$ for the selected top coordinate of t and zero at the other coordinates. The weight comparison used by the release supplies

$$\sum_i h_0(\nu_{j,\rho,i}) - H_{j,\rho} \le h_0(a^*_\rho) + e_{j,\rho},$$
$$\sum_i \nu_{j,\rho,i} = a^*_\rho + k w_{j,\rho} + \Delta_{j,\rho}.$$

For instance, the first bound uses the hook factorial lower bound and the interlacing of the fitted original k columns against the row coordinates, with the first top entry and the appended blocks treated separately. The numerical caps on the two defects used here are

$$e_{j,\rho} \le C_0\big(w_{j,\rho} + S + N_m^2 + m(t+1)\eta\big)\mathrm{La},$$
$$|\Delta_{j,\rho}| \le C_0\big(w_{j,\rho} + k(N_m + S) + m(t+1)\eta + \mathbf{1}_{\rho\,\mathrm{top}} k b'\big).$$

The coefficient of ω in the combined row and column expression is exactly

$$(p_j - q_j) w_{j,\rho} - \sum_i \nu_{j,\rho,i} = -a^*_\rho + (2j - m - 1) w_{j,\rho} - \Delta_{j,\rho}.$$

Therefore the signed contribution is bounded above by $h_0(a^*_\rho) - a^*_\rho \omega + e_{j,\rho} + |\omega|(|2j - m - 1|\,|w_{j,\rho}| + |\Delta_{j,\rho}|)$. Summing this bound, using Fo once and inserting the bound for the terminal branching and evaluation term h_A yields the quantitative conclusion

$$\mu(RR^{\vee})+h(u)-\mathrm{Fo}\le \mathrm{MAIN}(u)+L_{\mathrm{pass}}+\mathscr{E}+m(C_{\mathrm{low}}+C_{\mathrm{mass}}+1)E_3 \\ +E_A(u)+\sum_j (E_{B,j}(u)+E_{X,j})+\sum_{j,\rho}\bigl(e_{j,\rho}+|\omega|(|2j-m-1|\,|w_{j,\rho}|+|\Delta_{j,\rho}|)\bigr). \tag{17}$$

Here $E_A, E_{B,j}, E_{X,j}$ are the three nonnegative allowances in Paper IV §4, with $E_{X,j}$ subtracted in $h_{X,j}$ and therefore added in (17). $\mathrm{MAIN}(u)=\sum_t [m(h_0(a_t)-a_t\omega)-a_t\log T_u]$. The occurrence of $\mathrm{av}_v \log\kappa_v$ in the row bound cancels its subtraction in $h(u)$. The signed formula applies even when $\omega<0$. Fitted weight and hook bounds · Height allowances · Signed per-lift cancellation

For the uniform estimate in (17), put $U=mx^{\varepsilon}(B_0+N_c)$ and $C_{\mathrm{pr}}=C_{\mathrm{low}}+C_{\mathrm{mass}}+1$. Use $B_0=fmbL$, $L=\log b=x^{\varepsilon}\ge 1$, $N_c\ge 0$, $0\le \mathrm{La}\le 3L$, $|\omega|\le \mathrm{La}$, and at most f coordinates. The selected size bounds are $N_m^2, m^3N_m^2L, m\eta N_mL, kN_m, N_mw''\le b, kS, m(t+1)\eta\le w'', N_m, k\ge 1$, and $b'N_m, mkb'\le a$.

The selected numerical record takes $E_{B,j}=C_hfN_m^2(\mathrm{La}+m^3L)$ and $E_{X,j}=C_hfn_V^2L_p$, where $L_p=\log(b\ \mathrm{rd}\,E\,g_0g_1n_V)$. The fitted norm and level-six eigenvalue comparisons use these allowances on the same word and row. The record gives $0\le C_h\le C_0$, $n_V\le N_m$ and $C_hL_p\le C_0\mathrm{La}$. Since $N_m^2\le b$, $m^3N_m^2L\le b$ and $0\le \mathrm{La}\le 3L$, the per-lift allowances are at most $C_0f(3bL+b)\le 4C_0fbL$ and $3C_0fbL$, respectively. Therefore $\sum_j(E_{B,j}+E_{X,j})\le 7C_0B_0\le 7C_0U$.

Put $V=x^{\varepsilon}(B_0+N_c)=U/m$. Use the finite Gram, quotient and index comparisons on the same projected row as the signed finite main term. At the selected auxiliary level $s=1$ they allow

$$\mathscr{E}=LG+\sum_j\bigl(C_{\mathrm{row}}V+bd(j)+\log d_\gamma\bigr)+1.$$

Here $C_{\mathrm{row}}\ge 0$ is the coefficient in the fixed row scale, $LG=\log|G|\le 2U$ is the group term, and the final 1 caps the weighted finite index-pair remainder. The twisted-place length bound gives the per-lift cap $bd(j)=C_eV$ after identifying the twisted places with a subset of the record support and identifying the positive base rank with the record rank D_0. For the common denominator d_γ, the selected family bound gives $\log d_\gamma\le c_DV$ with $c_D=2+9\max(C_d,0)$ from the family denominator constant C_d. The selected bound $\log\log x\le L$ gives $E_3=B_0\log\log x\le V$. Since $1\le U$, these estimates give

$$\mathscr{E}+mC_{\mathrm{pr}}E_3\le\bigl(C_{\mathrm{pr}}+2+C_{\mathrm{row}}+C_e+c_D+1\bigr)U\le c_eU.$$

The source level-three comparison has allowance $C_{\mathrm{row}}E_3$ per lift and chooses $c_e=\max(C_0,\ C_{\mathrm{row}}+2+C_{\mathrm{row}}+C_e+c_D+1)$. Use a common family constant c_e dominating both the stored choice and the coefficient with C_{pr} from the geometric comparison. The norm, projected covolume, finite comparison and row passage hypotheses all enter on the fitted lattices and metrics.

For even $m=2r$, each of the odd shifts in absolute value occurs twice, so

$$\sum_{j=1}^{m}|2j-m-1|=2\sum_{\ell=1}^{r}(2\ell-1)=\frac{m^2}{2}.$$

With $0\le w_{j,\rho}\le 2b$ and $|\omega|\le \mathrm{La}\le 3L$, the rank-shift error is at most $3fm^2bL=3mB_0\le 3U$.

The two weight-error sums cost $15C_0U$ each, with an additional $C_0|\mathcal{T}|a\,\mathrm{La}$ in the Δ sum by $mkb'\le a$. The terminal allowance is at most $C_0\bigl(|\mathcal{T}|b'N_m+f(m\eta+S)N_m+fN_m^2\bigr)\mathrm{La}$. Here $b'N_m\le a$, $m\eta N_m\mathrm{La}\le 3b$, $SN_m\mathrm{La}\le 3bL$ and $N_m^2\mathrm{La}\le 3bL$, giving $9C_0U+C_0|\mathcal{T}|a\,\mathrm{La}$. The fitted quotient

passage with $3\,de\,Q \le 1$ and $C_0 \ge 28$ gives $L_{\mathrm{pass}} \le 1 + C_0 m \log |G|$. The size bound $m \log |G| \le b$ then gives $L_{\mathrm{pass}} \le (1 + C_0)U$. Together these estimates bound the error in (17) by $2C_0|\mathcal{T}|a\,\mathrm{La} + (47C_0 + 4 + c_e)U$.

The release's recurrence yields $\mathrm{MAIN} + (50C_0 + 10)|\mathcal{T}|a\,\mathrm{La} + 60c_eU$ on the tied record. The record comparison takes a common coefficient dominating $50C_0 + 10$ and $60c_e$ and uses $x^{\varepsilon} \le x^{C_{\mathrm{abs}}\varepsilon}$ for $C_{\mathrm{abs}} \ge 1$. The powers agree at $C_{\mathrm{abs}} = 1$, which the parameterized division wire accepts with the same row and arithmetic data. The stored selected wire and top assembly take $C_{\mathrm{abs}} = 2$, with the matching global divisor exponent $e_D = 2$. Derivation of the two error caps and pinned source ranges

Residual assumptions and scope. These statements must hold on one selected realization, eventually at every selected index, with the same weight lists, measure, level, and lattice coordinates as the word-period maps.

- For (16a)–(16h), supply the good cover and admissible weight, the top degree and boundary valuation comparisons and the fitted jet-rank bound, the common component degree comparison, and the integral orthogonal splitting with its component masses. The volume-height comparison uses the fitted finite tests and complex cusp realization for the model injection, the arithmetic Chow comparison, and the boundary and carrying-section bounds. At finite degree the fitted construction must supply the smooth flag modification, the two moving systems and their positivity, the canonical and adjoint identifications, and the inclusion of globally extending sections with equality in degree zero. The proof derives the good cuts by Bertini and obtains the residues and deepest filtration step from the two Koszul sequences with Kawamata–Viehweg vanishing. Cover assumptions · Fitted construction · Geometric and metric inputs
- For the theta realization on the selected cover, require the marked uniformizations of the selected model, the rigidified cubical polarization and unit fibre generators, the compatible harmonic connection and charts, and the triangular and cusp identifications with the actual weight spaces. Require base change of the algebraic frames and coverage of the complex test points by the component readings. With these data the torsion sum constructs the rational theta rules and proves their simultaneous comparison and faithful reading. The analytic span uses the marked Gram and doubled kernel in their prescribed normalizations. Fitted-cover realization
- For short integral sections require the identification of the integral cusp lattice by all prescribed framed DVR tests, including the chart and primary torsion estimates off the level and at both types of level primes. The norm bound uses the short harmonic coefficients, the first-minimum and Gaussian comparisons in the fitted Lie normalization, and the identification of the component Petersson metrics with their probability realizations. The full frame at u_1 and the spans at $0, D, T$ also require the integral component projections, the key and flat sections with their valuation and metric comparisons, the scalar-zero Bergman bound, and integral multiplication and the Chow frame comparison. The final flat transfer from $w(T)$ to u_- uses its own rational rank identification, finite lower valuations and archimedean factors. Short-section calculation · Finite-degree estimate · Stable theta spanning proof and source details
- For (16j), supply the rank-preserving integral twist with its signed degree and index bounds, the joint maps and both rank comparisons in (16i), the short type span from the integral isometric level action and even theta family, and the integral ring comparison, all on the same cover. Slope descent and last-lift branch

- The mass comparison requires the exact Petersson realization and the upper mass bound in the chosen Haar measure on the selected lift, together with integral rational descent of the type lattice. Formula (15) fixes the required scalar. The projection requires the positive good rank, orthogonality, denominator and rank bounds, and the ambient maximum slope cap. The calculation of the actual reference covolume uses the fitted integral type and its complex Gram frame. Release arithmetic route
- Equation (17) additionally requires the local twisted- and reference-type Gram, quotient and index comparisons including the finite error sum, the row passage with its smallness bound, and the fitted word-weight and norm allowances used in §4. In particular the reference-row error budget includes both the summed level-three errors and the finite comparison errors. None of these bounds follows from the signed determinant identity alone. The theorem supplies the projected determinant estimate from the stated geometric inputs and the ensuing signed cancellation. Local comparison hypotheses and errors · Passage smallness

5. Remaining propositions for a full prose proof

The uniform existence statement needed here is that for every ordinary eligible family, after one family-dependent lower cutoff, every selected index admits **one** realization carrying the column, point, and row identifications of inputs 1, 2, and 5. On that realization the lattices, maps, levels, gain and losses must satisfy inputs 3 and 4 at every finite place and inputs 6 and 7 at every word with the indicated uniform constants. The release assembles eventual point and class-period assertions, the level and type-prime assertions, the served construction, and the selected division package on the same family. Point and local input assembly · Selected package and its hypotheses · Assembly on the theorem route

The remaining propositions on that realization are the following.

- Identify the selected automorphic realization with the framed class model of §1, including its integral rows and reference metric. For the coefficient-field argument in §1, identify the finite vacuum and induced theta expansions with the prescribed operators and determinant frames. Establish the Taylor and fitted-site inputs of §2 and the discriminant and level-prime bounds on those models. The local comparisons must give the specific valuation in (8), the image of (5), and the unit in (9).
- For the terminal row, the degree calculation uses the fitted rational PEL points of induced type, the actual integral and complex test frames, and their prescribed line and metrics. The passage to (16k)–(16l) also requires the fiber evaluation inclusion, free reference level action, probability Petersson comparison and anchor bounds. For the nonterminal lifts, supply the uniform cover and adjoint construction, component height and rank comparisons, integral and metric bounds on the fitted spaces, and flat transfer and descent on the same cover. Fitted base-model assertions · Lan, Arithmetic Compactifications of PEL-Type Shimura Varieties · Nonterminal construction
- Identify the core and realized type lattices with the Petersson scaling (15), and supply the projector, rank bounds and fitted reference frame used with (14). Establish the finite Gram, quotient and index comparisons, the passage to the final row and the word-weight and norm bounds of §4 with their uniform budgets.