Series · Markdown source

Paper V. Relative trace, archimedean norms and the rank of period columns

1. Period maps and arithmetic height
   The ordered chain and the common Bessel pairing
2. Relative trace identity
3. Regular real orbits and their energy
   Removing the real cutoff on the chosen periods
   Local Bessel endpoints and the uniform minor estimate
      Formal-degree calibration of the shifted waves
   The rotated sum over integral block points
4. Boundary terms in the rational sum and the orbit count
   The rational zero-Gram contribution for distinct words
   The rational nonzero-Gram contribution for distinct words
      Weighted same-row expansion
      Loop counts
      Nongeneric corank-one pairs
      The selected finite densities
      Generic scalar corank-one pairs
         The generic length-one Fourier residual
         The scalar recurrence
      Combining the three estimates
   The equal-word diagonal
5. Finite-dimensional implications
   Word weights and the lower spectral tail
6. Interface with Paper III and outstanding inputs

Paper VI furnishes the selected family at each scale. Paper IV constructs the common integral period maps from its stated inputs. This paper compares those maps to a single complex spectral model, proves the rank and norm implications of the Gram estimates, and derives the archimedean operator-norm inequality from the prescribed metrics. For the relative trace it proves the primitive finite-place count on the selected frame lattice, including at the comparison primes and for Grams with singular reduction. It computes the normalized Euler-product error and derives the signed row-kernel bound from the marked-corner and orbital mass identities. Paper III uses the resulting complex rank and height bounds together with the finite-place estimates. **Proved here** marks the prose proofs. **Release input** marks a pinned theorem used with its indicated hypotheses. **Classical** identifies standard results used in the proofs. Global inputs · Lattice and matrix interface · Finite density details

1. Period maps and arithmetic height

Fix a selected family point (x, i). Let $\mathcal{U}$ be the finite set of words, $P_u : C_u \to H$ the original complex period maps, $d_u = \dim C_u$, and $D_{\mathrm{all}} = \dim H$. Use Hermitian inner products conjugate-linear in the first variable and write $a \otimes b^*$ for $z \mapsto a\langle b, z\rangle$. With π_u the compact stabilizer type, put

$$c_u = \frac{\mathrm{volN}(u)\,\mathrm{volH}(u)}{\dim \pi_u} > 0, \qquad A_u = c_u^{-1/2} P_u, \qquad K_{uv} = A_u^* A_v : C_v \to C_u.$$

Thus $\|K_{uv}\|_{\mathrm{HS}}^2 = \sum_{j,\ell} |\langle A_u e_j, A_v e_\ell\rangle|^2$ in orthonormal bases. Original normalization · Gram coordinates

Integral period and metric inputs. Paper IV supplies the common integral period from its local and global hypotheses. Over the coefficient field L the common row lattice is RR. The L-linear maps $B_u = \mathrm{Bk}(u)$ send the shifted source lattices $\mathrm{Nu}(u)$ into $\mathrm{RR}^\vee$, and these lattices span their respective spaces. At every infinite place w there is a row quotient $p : V_0 \twoheadrightarrow Q$ with its quotient metric, an isomorphism $\iota : Q \to E$ to the common row space, and finitely many subspaces Q_α spanning Q. For all $z \in Q_\alpha, y \in Q$, the metric comparison is $\langle \iota z, \iota y\rangle_E = s_\alpha \langle z, y\rangle_Q$ with $s_\alpha = \kappa(w) \prod_j \lambda_{j,\alpha}(w) > 0$. The base period factors as $B_u(\phi)(\iota p F) = P_u^{\mathrm{base}}(J_m \cdots J_1 \phi)(F)$ for all inputs and rows, where J_ℓ are the prescribed chain maps. The base comparison has scalar one. Common pairing and base normalization · Metric and factorization hypotheses · Lattice inclusion

Proposition 1 (proved here). For bounds $p_u(w) \ge \|P_u^{\mathrm{base}}\|_{\mathrm{op,w}}$ and $k_\ell(w) \ge \|J_\ell\|_{\mathrm{op,w}}$ in the prescribed metrics,

$$\|B_u\|_{\mathrm{op,w}}^2 \le \frac{p_u(w)^2 \prod_\ell k_\ell(w)^2}{\kappa(w) \min_\alpha \prod_j \lambda_{j,\alpha}(w)}. \tag{1}$$

Indeed $\iota^*\iota$ acts by s_α on Q_α by the mixed pairing identity. The subspaces span after scalar extension, so the least eigenvalue is at least $\min_\alpha s_\alpha$. Every quotient vector y has a lift in $(\ker p)^\perp$ of norm $\|y\|_Q$. Evaluate the base factorization on this lift, bound the product of the source maps, and precompose by ι^{-1}, whose norm is at most $(\min s_\alpha)^{-1/2}$. This gives (1) on the whole source space. Operator-norm theorem

In frames X_w, X'_w sending Euclidean into algebraic coordinates, the norm is that of $(X'_w)^{-1} B_u^{(w)} X_w$ with the **dual** frame on the target. For $B_u \neq 0$ the source height is the mean over the infinite places of the totally complex L of $\log \|B_u\|_{\mathrm{op,w}}^2$. The supplied bounds are $p_u(w)^2 \le e^{h_P}$, $k_\ell(w)^2 \le \exp(\sum_{j:\,\ell(j)=\ell} h_j(w))$, and $\mathrm{mean}_w \min_\alpha \sum_j \log \lambda_{j,\alpha}(w) \ge \sum_j l_6(j)$. Taking means in (1) bounds the height by $h_P + \sum_j \mathrm{mean}_w\, h_j(w) - \sum_j l_6(j) - \mathrm{mean}_w \log \kappa(w)$. The comparison to $h(u) = R.h_u - \mathrm{mean}_w \log \kappa(w)$ and the upper bound for $\mu(\mathrm{RR}^\vee) + h(u) - \mathrm{Fo}$ use the geometric chain and arithmetic budgets used by Paper III and supplied in $\mathrm{LinkDIV}$. The zero map satisfies every positive norm envelope. Height definition · Height chain · Height transport · Arithmetic budget

The ordered chain and the common Bessel pairing

Finite-level maps. Let M_u be the complex realization of the ramified starting space, and $e_{\mu,u} : M_u \to \prod_{c \in C_0} \mathbb{C}^{Y_0}$ the starting map at type tuple μ. At μ_0 it is $\mathrm{eMu}_u : M_u \simeq M_u^{\mathrm{ram}}$ followed by inclusion. At $\mu \neq \mu_0$ it is this map followed by $\mathrm{e0Tw}_{u,\mu}$ into the starting function space of the twisted chain. For $j = 0, \dots, m-1$, write c, c' for a source and target lattice class, $D_j^\sharp, D_{j+1}^-$ for multiplication

by the determinant powers $\prod_\tau d_{j,\tau}^{s^\sharp_{j,\tau}}$, $\prod_\tau d_{j+1,\tau}^{-s^{\mathrm{div}}_{j,\tau}}$, and $U_{j,\mu}(c,c')$ for the unfolded theta map on the indicated function spaces. The source defines

$$(K_{j,\mu}f)_{c'} = \sum_c \frac{\rho_{j,\mu}(c)\rho'_{j,\mu}(c')}{|\operatorname{Stab}_j(c)|} D^-_{j+1} U_{j,\mu}(c,c') D^\sharp_j f_c, \qquad f_m = K_{m-1,\mu}\cdots K_{0,\mu} e_{\mu,u}\phi.$$

The terminal pair is $\widetilde{P}_{u,\mu}(\phi,F) = \mathrm{period}_{u,\mu}(\mathrm{lowerMap}_{u,\mu}\,\mathrm{evz}_{u,\mu}\,f_m, F)$ in matched row coordinates. The row argument in the terminal pairing is $\mathrm{eMt}_u^{-1}F$ at base and $\mathrm{eMtTw}_{u,\mu}^{-1}F$ at a twisted tuple. With these row arguments the release identifies $\widetilde{P}_{u,\mu}$ with $\mathrm{perCol}(u,\mu,\cdot,\cdot)$ at both levels. On the rational column face, $\mathrm{Pt}_{u,\mu}(x,y)$ is the corresponding bilinear map obtained by taking the ramified basis on input and the specified degree-zero evaluation map on rows. Its embedding agrees with the ordered period on those inputs. In particular at base the comparison is on every rational basis pair, with no scalar in this coordinate identification. Ordered kernels · Base maps · Twisted maps · Terminal maps · Rational column face · Base coordinate equality

Release input (seesaw). The comparison uses compatible families of definite lattice kernels and indefinite Bessel and Fourier–Jacobi models, with common Weil data, the chosen quotient measures, and the local level and support laws. One substantive equality is the terminal two-step identity. Write t for a tuple in the family, $j = m-2$, S_j for the specified map on ramified columns in that identity, and $\phi_\alpha = \mathrm{rd}_{j,\alpha}(\phi,y)$ for the read at frame variable y in class α. For each allowed pair of polynomials P_1, P_2 there is **one** real two-step datum $\mathcal{C}$ for all t. With the two prescribed upward-in-index row lifts written as $L_j^{P_1}$ and $L_{j+1}^{P_2}$, the identity is

$$\mathrm{perF}_{t,j}((\phi_\alpha)_\alpha, L_j^{P_1}L_{j+1}^{P_2}f) = c_V(t,j)\,\mathrm{perT}_{t,m}\big(\mathrm{two}_{t,j}((\mathrm{rd}_{j,\alpha}(S_j\phi,y))_\alpha, \mathcal{C}), f\big). \tag{1a}$$

Here f is a terminal test from the prescribed multiplicity space evaluated on a smooth tensor. On the left perF is the sum over α of Bessel integrals against finite weighted right translates of its row argument. On the right two sums over the source and intermediate classes α,γ for each final class β, and perT sums the terminal integrals $\int_{\mathcal{F}_T}(\mathrm{trc}(e_{T,\beta})f)(h)\,f_\beta(\mathrm{emb}_T(h,1))\,d\mu_T(h)$. The definition of $\mathrm{trc}(e)$ is $\sum_g e(g)R(g)$, with $R(g)f(h) = f(hg)$. The proof expands the two lifts inside the Bessel integral, uses the disjoint supports of the places serving the finite modifications, and applies the local collapse and the common-real-datum contraction. In the terminal integration the common datum is obtained by applying the real contraction operator to the two vacuum inputs. Independence of the auxiliary flat representative is used to compare this datum between tuples. Two-step and terminal definitions · Terminal seesaw statement · Unfolded integrals · Common datum and assembly

Adelic dictionaries. In the spectral model the column of a seed T is the finite-level map

$$\mathrm{colDict}_u(T)(c,y) = J_u^{\mathrm{val}}(\xi_{T,c})(e_Y y), \qquad \langle \xi_{T,c}, \xi\rangle = T(\xi)(1,h_c) \quad (\xi\in\pi_u).$$

Here h_c represents the finite stabilizer class and J_u^{val} is the conjugate-linear compact evaluation map. The two conjugate-linear operations make the dictionary linear. The matched level identifies K_u with the finite column level after conjugation by the specified level change h_0, and matches their characters. The class decomposition is $h = \gamma h_c k$, $k\in K_u$, with a unique class c. In the tied model this dictionary is bijective onto the full column space, preserves its metric with the chosen column norm, and conjugates the local diagonal right translations into the actions used to impose the ramified conditions. The isometry that transports the original P_u to the concrete Poincaré map has this same column dictionary as its inverse

coordinate identification. These are **release inputs**. Evaluation and Riesz dictionary · Level and character matching · Isometry and diagonal action · Pinned inverse

On rows, write ι_X for the group identification and t_S for the chosen translation. A fixed tensor of local real vectors v_0 gives the translated evaluation of a row multiplicity vector R, with automorphic function $f_R = R(\bigotimes_w v_{0,w})$. Its L^2 representative is built from $g \mapsto f_R(\iota_X^{-1}(g)t_S)$. The row isometry $j : E_{\mathbb{C}} \to EEC$ identifies the common row space with its spectral realization, and the ambient isometry identifies $s_E F$ with $\mathrm{conjHH}(jF)$, where $s_E : E_{\mathbb{C}} \to H$ is the original conjugate-linear isometry into H. A reference normalization in the join compares the two row evaluations by a single nonzero scalar, expressed as $\kappa_S \kappa_0$ in the construction. This comparison holds on all common rows. The construction restricts to the chosen good-level row space. Translated evaluation · Common row identification · Nonzero evaluation factor

For explicit measures, define the adelic Bessel functional in the matched decomposition by

$$\mathscr{B}_u(\varphi, f) = \int_{\mathcal{F}_{H_u}} \int_{\mathcal{F}_{N_u}} \varphi(h) f(n\,\mathrm{emb}_u(h))\, \overline{\psi_{\mathrm{Bes}}(n)}\, d\nu_{N_u}(n)\, d\mu_{H_u}(h).$$

The pinned character is the one with parameter 1 and the chosen distinguished line in BesLev. With $\mathrm{colForm}_u(T, \xi)(h) = T(\xi)(\mathrm{toH}_u(h))$, the unfolding for the L^2 representative of a bounded continuous row function f with matching rational invariance and inverse level character gives

$$\langle \overline{f}, P_{u,T} \rangle = c_u^{\mathrm{Haar}} \int_{G_\infty} \mathscr{B}_u(\mathrm{colForm}_u(T, W_u(y)),\ g \mapsto f(\iota_X(g)(y, x_u)))\ d\nu_{Q,u}(y). \tag{1b}$$

Here $\nu_{Q,u}$ is the release real quotient measure represented on G_∞, and x_u is the fixed finite translate of the seed. The release first unfolds to the rational stabilizer fundamental domain, then disintegrates over the Bessel quotient and y. The factor c_u^{Haar} is positive by this Haar comparison. The compact quotient and the bound by a constant times $\|W_u(y)\|$ justify the integrations. The complex scalar relating the ordered chain to the spectral pairing below also includes the normalization of the row evaluation and the finite comparison. Bessel functional · Parameter · Level and column form · Inner product unfolding · Positive Haar factor · Measurable disintegration and L^1 bound

Nonvanishing and the common comparison. The terminal identity and the earlier Fourier–Jacobi steps feed a joint statement on the selected package. It chooses one finite central translation z for all u. Its word data give a source map of norm $r_u^{\mathrm{sc}}\|\phi\|$ with $r_u^{\mathrm{sc}} > 0$, and factors $c(u, \mu) \neq 0$ such that, when $c(u, \mu(w))/c(u, \mu_0)$ comes from the coefficient field, its preimage has valuation zero at w. The matching transports two-term linear relations among the normalized finite chain periods to the bottom Bessel periods, and thence to the spectral pairings at zt_S. The bottom group alignment uses the same z for the wordwise finite transport corrections and allows a right factor in the common row level. These are **release inputs** on the bottom, family and served-place hypotheses. Joint condition and bottom transport · Preservation of two-term relations · Selected assembly

The constants in the repeated comparison can be read from the ordered recursion. Write $n = m/2$. In the tuple notation, the forward column recursion in the two-step proof applies $K_{2i+1}K_{2i}S_{2i}$ at step i. The backward row recursion composes the chosen maps io_j. The total factors in this comparison are

$$c_{\mathrm{top}}(t) = c_P(t) \prod_{i=0}^{n-1} c_V(t, 2i) c_K(t, 2i) c_K(t, 2i+1), \qquad c_{\mathrm{bot}}(t) = \prod_{j=0}^{m-1} c_I(t, j).$$

The comparison of terminal coordinates supplies a nonzero word factor s_u and a column similitude S_u. The source uses $c(u,t) = s_u c_{\mathrm{top}}(t)/c_{\mathrm{bot}}(t)$ and the column change S_u^{-1}. The c_I are nonzero by the row lift laws, $c_{\mathrm{top}} \neq 0$ is part of the tied factor comparison, and the ratio to the base cancels s_u. The unit-ratio theorem multiplies the two-step ratios $c_V c_K c_K/(c_I c_I)$ and the terminal factor ratio. The served-place condition is used to deduce valuation zero for an algebraic preimage of that ratio at the coefficient place in question. Ordered recursions and factors · Factor ratios · Nonzero factors and linear relations

For the reference test, (9c) gives $\phi \in C_u^{\mathrm{ram}}$ with $P_u\phi \neq 0$ in the tied full spectral row space. Under the bottom matching, real gap and fixed-vector hypotheses, using discrete-series classification and cusp rapid decay, the pinned conjugate row evaluation for the fixed real vector at the common translated level zt_S is onto that space. Choose a row test whose conjugate evaluation is $P_u\phi$, using the same column and row dictionaries as in the comparison. Its spectral reference value is $b_0 = \langle P_u\phi, P_u\phi\rangle > 0$. This test precedes the arithmetic row projection. Kernel argument on the selected family · Row evaluation and reference test

Proved here (comparison from linear relations). Fix the nonzero reference value b_0 for a word u. Let p_0 be its reference scalar on the chain side of the relation law, and let $L = c(u,\mu)\widetilde{P}_{u,\mu}(\phi, F)$ be a tested chain value. The law states that $\alpha p_0 + \beta L = 0$ implies $\alpha b_0 + \beta b = 0$ for the matched spectral value b, after the stated word change of columns. Taking the relation with zero second coefficient gives $p_0 \neq 0$, since otherwise $b_0 = 0$. Taking $(\alpha,\beta) = (L, -p_0)$ gives $Lb_0 = p_0 b$. Thus the comparison scalar is independent of the test. In transported coordinates, with a common choice of embedding or conjugate embedding σ_0, the source packages these identities as

$$\zeta^{\mathbb{C}}_{u,\mu}\,\sigma_0(\mathrm{Pt}_{u,\mu}(x,y)) = \alpha_u\left\langle s_\infty(V_\infty \Theta_\mu j_{R,\mu} e_{V,\mu} y),\, P_u^{\mathrm{ind}}(U_{\infty,u}\sigma_0(x))\right\rangle. \tag{1c}$$

Here $e_{V,\mu}$ gives the algebraic rows, $j_{R,\mu}$ their complex realization, Θ_μ their map into the common good-level row space, and s_∞ is conjugate-linear. The single V_∞ is the common translation on rows, and the $U_{\infty,u}$ are the isometries after division by the positive source norm scales. At the normalization step the source takes $\zeta^{\mathbb{C}}_{u,\mu} = c(u,\mu)/c(u,\mu_0)$, $\alpha_u = p_0/(b_0 c(u,\mu_0))$, where p_0, b_0, c denote the factors in the transported coordinates of the anchored statement. All denominators are nonzero, and $\zeta^{\mathbb{C}}_{u,\mu_0} = 1$. Linear-relation proof · Transport of the anchors · Normalization and nonvanishing · Packaged identity

The scalar convention in (1c) puts $\zeta^{\mathbb{C}}$ on the *chain* side. With $R_\mu = \mathrm{rowMat}_\mu P_{q,\mu}$ the matched common row map, the source proves $\zeta^{\mathbb{C}}_{u,\mu}\sigma_0(\mathrm{Pt}_{u,\mu}(x,y)) = \sigma_0(B_u^{\mathrm{def}}(x)(R_\mu y))$. The multiplier of the common period in Paper IV's convention is the reciprocal of the chain-side scalar. Its valuation is also zero whenever the chain-side scalar is a unit. If some period value at the tuple is nonzero, the displayed equality realizes $\zeta^{\mathbb{C}}_{u,\mu}$ as the embedding of a quotient in the coefficient field. If the scalar has no preimage, all the period values at that tuple vanish, and the rational comparison can use 1. The base chain kills the kernel of P_{q,μ_0} because the right side vanishes on it. This explains the section formula $B_u^{\mathbb{C}}(\phi, F) = \mathrm{perCol}(u, \mu_0, \phi, \mathrm{sec}_{\mu_0} F)$ in the record. Here sec_{μ_0} inverts the row map on the orthogonal complement in the degree-zero space of its kernel there. Finite tuple comparison and kernel · Units and absent preimage · Degree-zero section

Release input (Rallis and doubling). These supply the preliminary lift and norm theorems used to construct the comparison. For source A and target B with $a = \mathrm{rank}\, A < b = \mathrm{rank}\, B \leq 2a$, the convergent source statement assumes B anisotropic or $a + \mathrm{Witt}(B) < b$. It further assumes the compatible doubled Weil data, the covolume-one polarization measure, the vanishing of the lower lifts on the chosen cuspidal representation, and the specified compact-finiteness of the doubled Schwartz vector. In

the source convention set $s_0 = (b-a)/2$. The test section at s_0 is $\Phi_{\mathrm{SW}}(g) = \int_{\mathcal{F}_{\mathrm{pol}}} (\omega^{\mathrm{dbl}}(g,1)\delta(\Phi_1,\Phi_2))(\ell w)\,dw$, with $\delta(\Phi_1,\Phi_2)(x) = \Phi_1(\mathrm{res}_L\, x)\overline{\Phi_2(-\mathrm{res}_R\, x)}$ and ℓ the polarization map. The convergent Rallis identity says the doubled zeta value paired with f_1, f_2 at s_0, for the flat section through Φ_{SW} and inverse character χ_B^{-1}, equals

$$\frac{\langle \theta_{\Phi_1} f_1, \theta_{\Phi_2} f_2 \rangle_{\mathrm{Pet,B}}}{\mathrm{vol}([U(B)])}$$

in the release Petersson convention. The theorem includes convergence and L^2 of the two lifts. The source obtains the lift inner products in both orientations via the unitary Weil–Siegel formula and Eisenstein continuation. This inner-product theorem supplies norms and surjectivity inputs for the comparison. The equalities of linear functionals in (1a)–(1c) use the seesaw and its transported linear relations. Doubled vector, section and point · Convergent Rallis statement · Both orientations

The specific surjectivity input in the family is $\mathrm{OntoAll}$. Under the node gap and the chosen step conditions, every reached composition of maps on the selected multiplicity spaces, from stage $k \le m$ to stage zero, is onto. The source derives it from the stated Howe, Weil representation, convergent Siegel–Weil and doubling inputs, then invokes it in constructing the family comparison. For the local zeta nonvanishing step, on $b\kappa$ with b real and κ in the finite level, the conjugate level characters cancel. The integrand becomes the real Weil coefficient times the conjugate real matrix coefficient times $\overline{\mathrm{ip}(q,q)}$, where $q \neq 0$ is the multiplicity vector. Outside the level it vanishes on the prescribed integration domain. Haar disintegration then carries the nonzero real integral to the adelic one, since $\mathrm{ip}(q,q) > 0$. The actual real test and its integral, the unramified doubling factors, the normalizations at modified finite places, and the global doubling continuation, Weil–Siegel and lift-surjectivity theorems remain inputs here. In particular the release's proposition named $\mathrm{DoublingMethodLocalZetaIntegrals}$ includes the global flat-section basic identity in the right half-plane and continuation. The symbols $\zeta_{u,\mu}$, α_u in (1c) come from the matched period comparison. Their values in the rank argument are fixed by the period identities and nonvanishing above. Surjectivity statement · Surjectivity from theta and doubling · Family use of doubling · Local zeta nonvanishing · Doubling input name · Scope of doubling input

Proved here (one projection for all pairs). In the join, let $\jmath_R$ be the common map denoted by e_R in the source, taken with its one fixed conjugation choice ϵ. The row comparison intertwines $\jmath_R$ and V_∞ through a common similitude. Pullback gives one complex-linear isometry $V_R = \jmath_R^{-1} V_\infty \jmath_R$ on the common rows because the positive similitude factor cancels in its norm. If $c_{\phi,u} : M_u \to \mathbb{C}^{r_u}$ is the column coordinate map, the pulled-back column map is $c_{\phi,u}^{-1}\epsilon U_{\infty,u}\epsilon c_{\phi,u}$, with ϵ acting entrywise. The map $\epsilon c_{\phi,u}$ scales norms by one positive constant in the prescribed column frame. That constant cancels after pullback and the two occurrences of ϵ make the result complex-linear. This gives the isometry U_u of M_u, using the pinned inverse dictionary to identify its target. In the source equation $\Theta_{\mu_0} \jmath_R e_V(y) = \lambda\, \jmath_R(\iota(c_F y))$ the row scalar $\lambda \neq 0$ is common. Because the row realization is conjugate-linear and the inner product conjugates the first slot, λ appears as λ in the bilinear finite expression. Substituting the tied Bessel pairing with comparison factor $\kappa_u \neq 0$ gives $a_u = \epsilon(\alpha_u \lambda)\kappa_u \neq 0$. Equality on the rational coordinate basis pairs suffices to extend to all complex inputs by bilinearity. The base row map is onto the common rows by the join theorem, so for any row F one can choose its degree-zero lift in the section formula. With the pinned column dictionary, this yields for all u, ϕ, F

$$B_u^{\mathbb{C}}(\phi, F) = a_u \langle s_E(V_R F), P_u(i_u U_u \phi) \rangle, \tag{1d}$$

The isometric inclusion is $i_u = (s_{C_u})^{-1} \circ \mathrm{incl}_{M_u^{\mathrm{ram}} \subseteq M_u^0} \circ \mathrm{eMu}_u$, with s_{C_u} the original coordinate equivalence onto the full initial column space M_u^0. At rational inputs, $B_u^{\mathbb{C}}(j_{N,u}x, j_E y) = \sigma(B_u(x)(y))$; this is the separate rationality tie used for Bk. Row, column and pairing ties · Transport and common scalar · Section and original Bessel map · Rationality tie · Isometric inclusion · Surjectivity of the base row map

Set $H_{\mathrm{rows}} = s_E(V_R E_{\mathbb{C}})$ and let Π be its orthogonal projection. The same orthonormal basis of H_{rows} works in (1d) for every u. With $V_u = i_u U_u$ in orthonormal source coordinates and the common row basis, the matrix in (1d) is $a_u[\Pi P_u V_u]$. The source implements this as the row-basis transpose, the embedded rational matrix, and a wordwise source-basis change multiplied by a_u^{-1}. Passing from the chosen embedding to the representative of its infinite place can conjugate all matrices at once. In the chosen embedding, the Gram blocks of these projected columns have factors $\overline{a_u} a_v \sqrt{c_u c_v}$ times $V_u^* A_u^* \Pi A_v V_v$. These factors come from one scalar per word. The unprojected spectral bound applies to the synthesis map before projection, and the rank estimate in §5 charges the row codimension after projection. The same map gives $\|B_u\|_{\mathrm{op}} \le |a_u| \sqrt{c_u}\, \|A_u\|_{\mathrm{op}}$ in the tied metrics at this embedding. Equation (1) supplies the separate all-place height bound. The family selector uses the same join, multiplicity package and orthonormal row realization at each (x, i) for every word. These choices can depend on that point. Orthonormal row extension · Exact rational block and scalars · Simultaneous selection · Insertion in the division package

2. Relative trace identity

Let $\Gamma = G(F)$ be the rational unitary group, where F is the totally real base field. Let Γ_u, Γ_v be the rational stabilizers of the two seeds. A concrete input T gives a seed $\Phi_{u,T}$ supported on the cell $g_f = nhx_u k$, $k \in K_2$, with value $\psi_u(n)\chi_2(k) T(W_u(g_\infty))(1, h)$. The concrete period map sends T to the L^2 Poincaré series $P_{u,T}(g) = \sum_{\alpha \in \Gamma_u \backslash \Gamma} \Phi_{u,T}(\alpha g)$. In this model the positive constant is

$$c_u^X = \mathrm{vol}(K_2)\, \mu_{N_u}(F_{N_u})\, \mu_{H_u}(H_u) / \dim \pi_u.$$

The isometric model comparison preserves the corresponding normalized constants. Seed and series · Concrete pairing and volume · Model comparison

Proposition 2 (unfolding, proved here given the release inputs). Under absolute convergence, with the quotient measures of the model,

$$\begin{aligned} \frac{\langle P_{u,T}, P_{v,T'} \rangle_{L^2}}{\sqrt{c_u^X c_v^X}} &= \frac{1}{\sqrt{c_u^X c_v^X}} \sum_{\gamma \in \Gamma_u \backslash \Gamma / \Gamma_v} \int_{(\Gamma_u \cap \gamma \Gamma_v \gamma^{-1}) \backslash G(\mathbb{A}_F)} \overline{\Phi_{u,T}(g)}\, \Phi_{v,T'}(\gamma^{-1} g)\, dg \\ &= \Big\langle \Big(\sum_t a_t \Big) J_v T', J_u T \Big\rangle, \end{aligned} \tag{2}$$

where $J_u = \mathrm{sharp}$ is the conjugate-linear Riesz map to the direct sum of compact type spaces over the finite column classes. Expand the first series and translate each summand to unfold the integral to $\Gamma_u \backslash G(\mathbb{A}_F)$. In the remaining series the right Γ_u-orbit of $\Gamma_v \gamma^{-1}$ is indexed by $(\Gamma_u \cap \gamma \Gamma_v \gamma^{-1}) \backslash \Gamma_u$. Unfold that orbit by the same translation. Absolute convergence justifies these exchanges. The second equality uses the **release input** of typed finite reduction and compatible fundamental domains for the finite slabs. Unfolding hypotheses · Exact kernel expansion

In the reduction, $t = (\Lambda, \Lambda', [\gamma])$ ranges over the two sets of selected lattices and double classes of rational isometries carrying Λ' to Λ and satisfying the lifted-flag condition, modulo their flag stabilizers $N_\Lambda, N_{\Lambda'}$.

Define Ω_t to be the real fundamental domain for the image of $N_\Lambda \cap \gamma N_{\Lambda'} \gamma^{-1}$. The integral in each surviving block is

$$O_t = \int_{\Omega_t} W_u(g) \otimes W_v(\gamma_\infty^{-1} g)^* \, d\mu_\infty(g). \tag{3}$$

For column classes ρ, ρ', the block is zero unless the respective lattices match and the transported finite element $k(\rho, \rho', t)$ belongs to K_2. In a surviving slab the two copies of the level character cancel as functions of the integration variable. Integrating the remaining fixed phase gives

$$a_t(\rho, \rho') = \frac{\mathrm{vol}(K_2)}{\sqrt{c_u^X \, c_v^X}} \, \chi_2(k(\rho, \rho', t)) \, O_t. \tag{4}$$

Evaluation by the column functionals converts the rank-one integrand into the pairing in (2). The reversed inner product in (2) accounts for the conjugate-linear Riesz maps. Their isometric transports, including the common ambient and wordwise source isometries, give $\|K_{uv}\|_{\mathrm{HS}}^2 = \|\sum_t a_t\|_{\mathrm{HS}}^2$ in tied models. Rational term index · Real domain · Finite blocks · Slab integration · Hilbert–Schmidt transport

3. Regular real orbits and their energy

Fix the same concrete setting and word pair. Put $G_\infty = \prod_{w|\infty} G(F_w)$, $r =$ flag length and $N = 2r + k + 1$. Form $Z(g)$ at each real place by pairing the first ordered frame against the translate of the second. At scale $\mathrm{sc} > 0$ the frame entries carry factors $\mathrm{sc}^{-(r-j)}$ for indices $j = 0, \dots, r$. The selected application uses $\mathrm{sc} = 1$. For a diagonal representative $h(c)$ with entries $c_1, \dots, c_{r+1}$, write $D_j = \prod_{a \le j} c_a$, $D_0 = 1$, and

$$\Delta = |c_{r+1}|^2 - TT', \qquad G_v = |D_{r+1}|^2 - TT'|D_r|^2 = |D_r|^2 \Delta,$$

where $T, T' > 0$ are the last-vector norm parameters. When the early minors are nonzero, diagonalization leaves r hyperbolic pairs and a two-vector Gram with determinant $TT' - |c_{r+1}|^2$. The section realizes the locus $\Delta > 0$ by a Witt frame, with a last component proportional to $\sqrt{|c_{r+1}|^2/T - T'}$. Frame and invariant definitions · Witt construction · Section matrix

Release input (real Haar disintegration). Let S, S' be the two real stabilizers with Haar measures coming from their unipotent and compact factors. The good set is the product of the placewise sectors $c_j \neq 0,\ \Delta > 0$. With a measurable diagonal section h and $C_{\mathrm{oo}} > 0$, for every nonnegative measurable f,

$$\int f \, d\mu_\infty = \int_{\mathrm{Good}} \int_S \int_{S'} f(s h(c) s'^{-1}) \, d\mu_{S'} \, d\mu_S \, dm(c), \qquad dm = C_{\mathrm{oo}} \prod_w \left(G_{v,w}^{k-1} \prod_{j=1}^{r} |D_{j,w}|^4 \right) dc_\infty. \tag{5}$$

Here $dc_\infty = \prod_{w,j} 2 \, d\Re c_{j,w} \, d\Im c_{j,w}$. The local Jacobian theorem gives a full-measure regular chart, excluding the vanishing early minors and residual Gram determinant. The real geometry input also includes the stabilizer double-coset comparison and unimodularity. In particular the chart is a statement about ambient Haar measure. A discrete rational sum can meet the exceptional sets. Measure and disintegration · Local-to-global Haar theorem · Witt comparison

Proved here. The complex Jacobian of $c \mapsto (D_1, \dots, D_{r+1})$ is $\prod_{j=1}^{r+1} D_{j-1}$ by triangular differentiation. The placewise density in minor coordinates is consequently $(|D_{r+1}|^2 - TT'|D_r|^2)^{k-1} \prod_{j=1}^{r} |D_j|^2$ against the same product of twice planar measures, times the constant in (5). In c-coordinates the factor is

$\Delta^{k-1}\prod_{j=1}^{r}|c_j|^{2(k-1)+4(r+1-j)}$. These give the measure weights at the residual-Gram boundary and at early-minor faces. Jacobian and invariants · Density

The real energy uses full Haar:

$$O(g)=\int_{G_\infty} W(z)\otimes W'(g^{-1}z)^*\,d\mu_\infty(z),\qquad I_{\mathrm{oo}}=\int_{\mathrm{Good}}\|O(h(c))\|_{\mathrm{HS}}^2\,dm(c). \tag{6}$$

Write $d_\sigma=\prod_w d(\sigma_w,\mu_w)>0$ for the ambient formal degree, and w_0 for the tensor of unit lowest-type vectors. The **release inputs** are continuous packets, unitary left equivariance $W(sz)=\tau(s)W(z)$, the left quotient measures ν,ν' for the same μ_∞, and

$$\int\|W\|\,d\nu<\infty,\quad \int\|W\|^2\,d\nu=1,\quad W(x)=d_\sigma\int\overline{\langle w_0,\sigma(y)w_0\rangle}\,W(xy)\,d\mu_\infty(y).$$

They include the primed versions, Bochner integrability in (6) on every good fibre, and finiteness there of

$$B(c)=\int|\langle\sigma(z)w_0,\sigma(sh(c)s'^{-1}z')w_0\rangle|\,\|W(z)\|\|W'(z')\|\,d\nu(z)\,d\nu'(z')\,d\mu_S(s)\,d\mu_{S'}(s'). \tag{7}$$

Two finite triple majorants use one packet twice and one stabilizer variable, replacing $sh(c)s'^{-1}$ by s. The stabilizers admit measurable Følner sets with positive finite Haar mass and normalized left overlaps tending to one for every fixed compact subset. Wave and energy assumptions · Local transport and convergence · Product convergence · Følner sets

The concrete convergence theorem uses the discrete-series Harish–Chandra parameter, its lowest-type vector, the two positive last-vector norms, side types, and weighted matrix-coefficient bounds. The arithmetic hypothesis is $N\ge 1801$ with the specified adjacent gaps at least $h=16N^{19}$. At a place the bound for (7) is a fixed constant times $\int\|s\gamma s'^{-1}\|_F^{-\beta}\,d\mu_S d\mu_{S'}$, with $\beta=h/(8\sqrt{2(r+1)})-4(k+r)(r+1)-1>9(r+1)N^2$. On regular fibres the norms of both unipotent factors are bounded by a constant times a power of the translated matrix norm. This uses inverses of early minors. The local representation estimates and the uniformity on compact subsets of the good locus are used in their source forms. Weight gap · Decay exponent and integral · Unipotent bounds and Haar integrability · Regular section

Proposition 3 (proved here). Under (5), the wave-packet hypotheses, and Schur orthogonality, $I_{\mathrm{oo}}\le 1/d_\sigma$.

Classical. With $d=d_\sigma$, Schur orthogonality for the square-integrable representation and the chosen Haar measure says

$$\int\langle b,\sigma(g)a\rangle\overline{\langle b',\sigma(g)a'\rangle}\,d\mu_\infty(g)=d^{-1}\langle a',a\rangle\langle b,b'\rangle.$$

The product formula for d follows on tensors by Fubini and extends by the same square-integrable coefficient identity. We take the discrete-series identification and this normalization on the selected group from the release. Formal-degree hypothesis · Selected product identity

For a Følner set $\Omega\subset S$ of mass a, define

$$Y_\Omega e=\int_\Omega\int\langle W(t^{-1}z),e\rangle\,\sigma(t^{-1}z)w_0\,d\mu_S(t)d\nu(z).$$

The vector integral converges by the quotient L^1 assumption and unitary equivariance. Use the primed version with mass a'. Expanding the sum of squared norms of $Y_\Omega e_j$ in an orthonormal basis and substituting $q = t(t')^{-1}$ for the two copies of the S-variable gives

$$\frac{\|Y_\Omega\|_{\mathrm{HS}}^2}{a} = \int \frac{\mu_S(\Omega \cap q^{-1}\Omega)}{a} \langle \sigma(z)w_0, \sigma(qz')w_0\rangle \langle W(qz'), W(z)\rangle \, d\mu_S(q) d\nu(z') d\nu(z).$$

In this calculation the group action $\tau(t)$ cancels between the two packet vectors and $\sigma(t)$ cancels between the two representation vectors. The ratio is at most one and tends to one for each q along the Følner sequence. Domination follows from the triple majorant. In the limit the quotient formula combines q, z' into $x = qz'$. The reproducing identity at z, conjugated in the first slot of $\langle W(x), W(z)\rangle$, reduces the x-integral to $\|W(z)\|^2/d$. Hence $\|Y_{\Omega_n}\|_{\mathrm{HS}}^2/a_n \to 1/d$, and likewise on the primed side. Truncated wave packets · Overlap and norm limit

Set $X_n(g) = d\, Y_{\Omega_n}^* \sigma(g) Y'_{\Omega'_n}$ and

$$F_n(c) = \frac{\int_S \int_{S'} \|X_n(sh(c)s'^{-1})\|_{\mathrm{HS}}^2 \, d\mu_{S'} d\mu_S}{a_n a'_n}.$$

In orthonormal bases each matrix entry of $X_n(g)$ is $d\langle Y_{\Omega_n} e_i, \sigma(g) Y'_{\Omega'_n} e'_j\rangle$. Squaring and summing the Schur identity gives $\int \|X_n\|_{\mathrm{HS}}^2 \, d\mu_\infty = d\|Y_{\Omega_n}\|_{\mathrm{HS}}^2 \|Y'_{\Omega'_n}\|_{\mathrm{HS}}^2$. By (5), $\int F_n \, dm \to d \cdot d^{-1} \cdot d^{-1} = d^{-1}$. Schur computation

Fix now $\gamma = h(c)$. The kernel whose norm is dominated by d times the integrand of (7) is

$$\mathcal{K}_\gamma(z, z', t, t') = d\, \langle \sigma(t^{-1}z)w_0, \sigma(\gamma(t')^{-1}z')w_0\rangle W(t^{-1}z) \otimes W'((t')^{-1}z')^*.$$

Its full integral in z, z', t, t' is $O(\gamma)$. In fact the left quotient formulas and Haar inversion turn the integral into a double ambient integral in $x = t^{-1}z,\ y = (t')^{-1}z'$. Integrating x first reproduces $W(\gamma y)$ since $\langle \sigma(x)w_0, \sigma(\gamma y)w_0\rangle = \overline{\langle w_0, \sigma((\gamma y)^{-1}x)w_0\rangle}$. Changing γy to x gives (6). Absolute convergence follows from (7) because the Hilbert–Schmidt norm of $a \otimes b^*$ is $\|a\|\|b\|$. Lifted kernel and reproducing integral · Fourfold form

Expanding $X_n(s\gamma s'^{-1})$, changing the subgroup variables and removing the unitary type actions gives the integral of $\mathcal{K}_\gamma$ with masks $1_{\Omega_n}(ts^{-1})1_{\Omega'_n}(t'(s')^{-1})$, with the same Hilbert–Schmidt norm. In the outer variables s, s', use the measure $d\mu_S d\mu_{S'}/(a_n a'_n)$. Each mask has L^2 norm one by Haar invariance and inversion. Its L^2 distance to the mask with $t = t' = 1$ tends to zero for each t, t' by the normalized overlap property, and is bounded by two. Minkowski's integral inequality followed by dominated convergence in the four kernel variables shows that the two masked integrals differ by a term tending to zero in outer L^2. The reference masked integral has outer L^2 norm $\|O(\gamma)\|_{\mathrm{HS}}$. This proves pointwise $F_n(c) \to \|O(\gamma)\|_{\mathrm{HS}}^2$. Fatou gives (6) bounded by $1/d$. Truncation and translated masks · Pointwise bound and Fatou in the release

For the boundary of the real integral, the same Minkowski estimate gives $F_n(c)^{1/2} \le dB(c)$. The selected coefficient $f(g) = |\langle w_0, \sigma(g)w_0\rangle|$ is Haar integrable by the **release input** of weighted discrete-series coefficient integrability. Tonelli, (5), and both Haar translations give $\int B \, dm = \|f\|_1 \|W\|_{L^1(\nu)} \|W'\|_{L^1(\nu')}$. The first moment stays finite across the boundary. An equality $I_{\mathrm{oo}} = 1/d$ by dominated convergence on the whole base would follow from an $L^2(m)$ bound for B. The inverse-minor bounds deteriorate near the

early-minor faces. The proof of the required upper bound uses only Fatou. Coefficient integrability · Product coefficient · Pointwise convergence input · Energy conclusion

Removing the real cutoff on the chosen periods

Fix a pair of words at one arithmetic cutoff. At a real place let D_j denote the leading minors of their cross Gram matrix, r the flag length and T, T' the two positive norms. The invariant used to smooth the orbital integral is

$$G = |D_{r+1}|^2 - TT'|D_r|^2.$$

For a smooth $0 \le \chi \le 1$ equal to one on $[1, \infty)$ and zero on $(-\infty, 1/2]$, multiply the packet by $\chi(G/\epsilon)$. The unipotent factors act on the Gram matrix by unitriangular operations and the compact stabilizers fix the specified frames. Consequently G is constant on the orbit, and the smoothing multiplies the orbital operator by exactly the same scalar. This preserves the phases of the Bessel character. For a finite nonempty set of good rational orbits the minimum of G_w over all these orbits and real places w is positive. Below that minimum their joint Gram matrix is unchanged. For an absolutely summable signed row $\sum_p a_p V_p$, first choose a finite set whose omitted sum of $|a_p|\|V_p\|$ is small. The norm error below its cutoff is at most twice that omitted sum. This proves convergence for each such row. Smoothed orbital operators and finite sets

On the regular boxes there is also a rate. In the measure of (5), passing to the leading minors gives density proportional at one place to

$$G^{k-1} \prod_{j=1}^{r} |D_j|^2 \prod_{j=1}^{r+1} 2\, d^2 D_j.$$

The triangular substitution $D_j = c_1 \cdots c_j$ has complex Jacobian $\prod_j D_{j-1}$, with $D_0 = 1$. Fix $|D_j| \le R$ and keep every early minor above $\kappa > 0$ at every place. The orbit majorant of §3 is bounded there for the fixed pair. Holding the early minors fixed, polar coordinates in D_{r+1} give the integral $2\pi\epsilon^k/k$ of $G^{k-1}\, 2d^2 D_{r+1}$ over $0 < G \le \epsilon$. Integrating the bounded early coordinates and taking the union over real places gives an $O_{R,\kappa,u,v}(\epsilon^k)$ bound for the removed squared energy. The constants can depend on the pair and the lower bound for its early minors. Gram-coordinate calculation and regular-box bound

The generic scalar term uses Poisson summation for the smoothed packet on cosets of its period lattice. The zero Fourier coefficient is the preceding normalized integral, and the sum of the nonzero coefficients is estimated in §4 by the field-uniform dual sum at the same smoothing. The exact cutoff identity on finite point sets supplies the return to the finite Gram matrices. Poisson estimate on the specified lattice

Local Bessel endpoints and the uniform minor estimate

We form the local minor estimate for a fitted discrete-series coefficient $F_w = P_w \pi_w Q_w$ with K-finite images of P_w^* and Q_w. The size parameter $B \ge 2$ bounds the spectral and interlacing compact-type entries. The character parameter $\kappa > 0$ satisfies $\kappa^{\pm 1} \le B^{C_h}$ and $1 \le T, T' \le B^{C_h m}$ with $\kappa^{2r} T, \kappa^{2r} T' \ge 1$. The fit uses the source's cloud, cut and opposite-sign gap conditions. In detail the cloud entries have compact-type entries on both sides for the first type and lie in $[0, b_w/8]$, with the other entries in $[-2b_w, -b_w]$ and $b_w \ge C_g N^7$. Equal types require the separating gap $2C_g(r+1)^2 + 1$ with cloud entries on both sides, and different types the prescribed cut of length $C_g(r+1)^2$. The opposite-sign gap is at least

$C_g m^{9/2}(3(r+1)N^2+N^3)$. All spectral gaps are at least $C_g m^4$. The sections are smooth on the positive Gram branch, have cross pairing ζ on the distinguished lines and use the chosen finite positive unipotent Haar multiple. The spectral data agree with the host parameters, and the packet uses the prescribed signature, character parameter $-2i\kappa$, and positive integral distinguished norms. Fitted packet conditions · Fitted coefficient and measure comparisons

Local endpoint in the character-one frame. Write $l=r+1$, $m=2l$ and $N=2r+k+1$. Take $k\ge 2$ and a fitted real place with its two compact types u,v, flags, section, quotient measures and Jacobi operators. At a block point with $G=|D_l|^2-TT'|D_r|^2>0$, let $E(F)=n\|V(F)\|_{\mathrm{HS}}^2$ be the top square, where n is the block density and V is the top Bessel vector. Let $J(F)$ integrate this square on the diagonal big cell with the same section measure. Its density in $(c_1,\dots,c_r,\zeta)$ is $(|\zeta|^2-TT')^{k-1}\prod_{j=1}^r|c_j|^{2(k-1)+4(l-j)}$ on $|\zeta|^2>TT'$. The endpoint for a finite linear combination F of K-finite matrix coefficients of infinitesimal character μ is

$$E(F)\le B^{C_{\mathrm{end}}m^4}G^{-(2p_0+n_0)}J(F).$$

Use $n_0=0$ for distinct types with a fitted cut of separation at least $30l^2+6$ and the prescribed Jacobi operator bound there. For equal types use $n_0=30l^2+6$ and an adjacent compact gap at least $2n_0+1$. Require $3l\le 2p_0$ and $p_0\le 3C_0l+1$ with fixed $C_0\ge 1/2$. The compact types have the prescribed half-unit interlacing with the central parameter λ, with both sets of entries bounded by B. Use $B\ge\max(3,\mathrm{e}k)$, padding $c\ge 1$ with $c+2\le B$, and adjacent gaps of λ at least $g_{\mathrm{sp}}\ge 2c$. Write $e=\|\mu-\lambda\|_\infty$ or a bound for it. With $d=mp_0+ln_0$ and $X=l^4(3lN^2+N^3)$ the spectral hypotheses are

$$e+d\le c,\qquad H>t(N,\epsilon_{\mathrm{sig}})+C_AX+3d+1,$$

where H is a lower bound for the opposite-sign gap of μ in the prescribed signature ϵ_{sig}, and $t\le 43N^3$ is the standard and dual Zuckerman threshold. Require $T'\ge 1$, $TT'\ge B^{-C_Tm}$, and the big-cell determinant expansion through order $2p_0+n_0$. In that expansion $D_l^jV(F)$ is a sum of top vectors with coefficient absolute sum at most $B^{C_3m^2j}$, each with energy integral at most $J(F)$. All constants are uniform for fixed cut, determinant expansion and size constants. Endpoint statements · Compact recursion

Proof. Take a distinct-type cut. In the equal-type argument below, apply this calculation at each of the two shifted cuts. Use $K=3$, $L_0=3l$ and the grid of $L_0(2r+1)$ nodes of spacing five starting at $2-\max(u_{i_0},v_{i_0})$. The negatives of the nodes lie between the two cut entries with margin two. At a node y the two letters have arguments $y,y+1$. Put

$$n_*(y)=\prod_{a=1}^N(|y+\lambda_a|+c),\qquad Q(y)=-TT'\prod_{j=1}^k(y+1/2+u_j)(y+1/2+v_j).$$

Exactly one pair of compact factors has opposite signs, at the cut, so $Q(y)>0$. In each spectral product, pair the first $k-1$ entries above $-y$ with their upper compact brackets and the entries below it with their lower brackets. Half-unit interlacing gives a ratio at most $1+(1+c)/\max(1,|y+\lambda_a|)$ per paired factor. On each side the first ratio costs at most $2+c$ and the rest telescope against $1+1/j$ by the gap bound. In a one-coordinate move omit the spectral entry whose bracket may use the moved coordinate. At most $m+1$ unpaired factors cost at most $3B$ each. Thus for each compact type the product costs at most $(2+c)^2k^2(3B)^{m+1}\le B^{4r+10}$ times its absolute compact product. On a free node y each old letter at grid

site y' contributes $1 \pm 1/(y-y')$, with the sign determined by its letter type. Both factors are at least $4/5$ because $|y-y'| \geq 5$. If Q_w is the resulting polynomial for a word of length at most L_0, then

$$n_*(y)n_*(y+1) \leq \Lambda^2 Q_w(y), \qquad \Lambda = \max\left(1, \sqrt{(4/3)^{L_0} B^{8r+20}/(TT')}\right) \leq B^{(6+C_T)m}.$$

There are enough free nodes for three sets of $2r+1$ nodes when adding three letters. The leading-coefficient interpolation weights at degree $j \leq 2r$ have absolute sum at most $(j+1)/5^j \leq 1$. Applying this in each set bounds the wave- and lift-grid losses by the products of the corresponding Λ factors. Interlacing · Wave and lift grids

Write $F_{\mathrm{test}} = G^{p_0} F$ in the distinct case. For equal types the two tests are $G^{p_0}(F \circ A_{\pm})$ with the right Wigner polynomials of degree at most ln_0. Here G on the group is the invariant Gram polynomial and the total multiplier degree is at most d. Standard and dual tensor decomposition put all resulting infinitesimal characters within distance $e+d$ of λ, with opposite-sign gap at least $H-2d$. The coefficient decay induction starts with unitarity and uses the tensor decomposition at each multiplication by a matrix entry. With $A_0 = C_A X + d$ the gap condition gives $t + \lceil A_0 \rceil < H - 2d$. Invariant differentiation preserves the polynomial degree, so the tests retain the decay $C_A X$ needed for the fitted Jacobi and normal-form estimates. Individual seminorms are used for convergence. For each word w, Capelli diagonalization bounds the products of eigenvalues by $n_w = \prod_{a \in w} n_*(y_a)$. Distinct infinitesimal characters have zero cross pairing by Schur orthogonality, also through regularization. The real unnormalized level-zero pairing R_w therefore satisfies $R_w(F_{\mathrm{test}}) \leq n_w \operatorname{Re}\beta(F_{\mathrm{test}}, F_{\mathrm{test}})$. The empty-word pairing $\operatorname{Re}\beta$ is bounded by $J(F_{\mathrm{test}})$ with constant one by the section Bessel square in the same coordinates and measure. Coefficient decay · Spectral bound · Orthogonality · Scalar calibration

For the analytic step use the compact module-form positivity, head continuity, the wave and lift kernel bounds, and block growth from the local positive germ and its Poisson bound. For a block of length h the compact inequality costs $(C_2\Lambda^2)^h$ and extends the admissible word by $3h$ letters. Removing all l coordinates costs $(C_2\Lambda^2)^l$ within the budget L_0. Apply this to a compact radial truncation of $\chi(G/\varepsilon)F_{\mathrm{test}}$, where χ vanishes below $1/2$ and is one above 1. At fixed ε the radial module forms converge. The top word integrals converge by the square-root Cauchy estimate from the radial tail, with one decay power sacrificed. A word recurring along radial truncations preserves the inequality and positivity in the limit. As $\varepsilon \downarrow 0$, the Jacobi normal-form identity supplies an integrable majorant by finitely many top energies of differentiated coefficients. At a positive cell an auxiliary cutoff at half its Gram is identically one near the cell, so its differentiated terms drop out of this majorant. The same identity identifies the limiting integral with the uncut level-zero residual. Thus $0 \leq R_w(F_{\mathrm{test}})/n_w \leq J(F_{\mathrm{test}})$. At a fixed positive point the cutoff leaves the top vector unchanged once $\varepsilon < G$. There are finitely many extending words, so one recurs along a sequence of cutoffs tending to zero. We obtain $E(F_{\mathrm{test}}) \leq (C_2\Lambda^2)^l J(F_{\mathrm{test}})$ at that point. This step uses the compact operator, the fitted Jacobi transfer and the Bessel normal-form and cutoff identities. Compact bound · Normal-form majorant · Limit · Pointwise comparison

To undo the multiplier, the determinant expansion and weighted Cauchy–Schwarz give $J_j^{\det}(F) = \int |D_l|^{2j} E(F) \leq B^{2C_3 m^2 j} J(F)$. For distinct types the invariant multiplier gives $E(F_{\mathrm{test}}) = G^{2p_0} E(F)$ and $J(F_{\mathrm{test}}) \leq J_{2p_0}^{\det}(F)$ since $G \leq |D_l|^2$ on the positive branch. For equal types raise the lower compact entry of the separating pair by n_0 on the right for one test and lower the upper by n_0 for the other. The two fitted places preserve G, T, T' and the block density. The Wigner upper bounds are $\|A_{\pm,\gamma}\|_{\mathrm{op}}^2 \leq (G/T')^{n_0}$, and the sum of adjoint norm squares satisfies

$\sum_{\pm} \|A^*_{\pm,\gamma} v_0\|^2 \geq (3B)^{-2n_0} (G/T')^{n_0} \|v_0\|^2$. Apply this to $V(F)^* e_j$ over an orthonormal basis on the left. It bounds $E(F)$ by $(3B)^{2n_0} (T'/G)^{n_0} G^{-2p_0}$ times the sum of the two tested point energies. Each tested integral is at most $(T')^{-n_0} J^{\det}_{2p_0+n_0}(F)$. The two factors of T' cancel. The grid, Wigner and expansion losses are bounded together by $B^{C_{\mathrm{end}} m^4}$ since $p_0 = O(m)$ and $n_0 = O(m^2)$. This proves the endpoint. Shifted places · Wigner bounds · Determinant moments

Uniform application to the minor moment. Normalize the fitted character κ first, so the endpoint lengths are $\kappa^{2r} T, \kappa^{2r} T' \in [1, B^{Cm}]$ with $C = 2|C_h|$. Take $p_0 = m$, $C_0 = 1$, $C_T = 0$ and $C_3 = C/2 + 6$. For an integer minor order $0 \leq a \leq C_a m^3$ use the integral shifts $\mu = \lambda + \nu$, $\|\nu\|_1 \leq e = la$. Set

$$Y = m^{9/2}(3lN^2 + N^3), \quad c = (C_a + 5)m^4, \quad g_{\mathrm{sp}} = C_g m^4, \quad H = C_g Y - 2e, \quad q_0 = e + l(40l^2) + l(l+1).$$

These shifts preserve the required parameter box and lose at most $2e$ from the opposite-sign gap. The endpoint cuts belong to the fixed compact types and fitted operators. The packet gaps give the two cut conditions when C_g is sufficiently large. The cloud scale gives $B \geq C_g N^7$ and hence the padding and size bounds. Since $d \leq 5m^3$ and $X, m^4, N^3 \leq Y$, the endpoint reserve $t + C_A X + 3d + 1$ is strictly below H for both cases. The choice $10^6 + 43 + |C_{A,\mathrm{diff}}| + |C_{A,\mathrm{eq}}| + |c_3^{\mathrm{wave}}| + 100 C_a$ for the inner gap threshold also gives $100(q_0 + 1) \leq C_g m^4$ and the weighted wave decay reserve $t + (|c_3^{\mathrm{wave}}| + 46)X + 3e + 1 < C_g Y$. The discrete-series coefficient identification also uses $N^3 + 1 < H$. The outer threshold includes the coefficient and frame comparisons and the three-dimensional gap condition. The Pieri family through radius q_0 supplies determinant expansions of order at most $40l^2$ for every wave in the smaller ball, since $e + l(40l^2) \leq q_0$. Both endpoint exponents are at most $40l^2$, so their common loss is $\max(1, G^{-1})^{40l^2}$. The weighted wave transfer costs $(2N+1)^{2e}$, and the integrated minor moment costs $B^{(C+12)m^2 a}$. As $2N + 1 \leq B^2$, a common exponent for these factors and the endpoint constants is

$$C_5 m^5, \qquad C_5 = 4C_a + \max(C_{\mathrm{diff}}, C_{\mathrm{eq}}, 0) + (C + 12)C_a.$$

The weighted application uses the discrete-series realization, the Pieri bounds and the wave transfer. The formal-degree calibration below supplies the section energies for the determinant expansion. The comparison of top vectors with orbital squares and the comparison of section energies cost at most π^{2m^2} each. This gives the normalized minor estimate with $C_5 + 8$. Rescaling the fitted unipotent Haar by a finite positive scalar multiplies both squares by its fourth power with the relative root measure fixed. Shift budget and transfer · Pieri family · Orbital comparison · Fitted measures

The resulting normalized minor bound, in the notation for the orbital squares below, is

$$|D_j(q)|^{2a} n(q) O_w(q) \leq B^{(C_5+8)m^5} \max(1, G(q)^{-1})^{40(r+1)^2} J_w, \qquad 1 \leq j \leq r+1.$$

The separate endpoint derivation gives the spectral hypotheses and the comparison with the local orbital square. The real majorants and section-energy identity are proved below. Detailed endpoint derivation

Formal-degree calibration of the shifted waves

Fix the character-one fitted place and one selected gapped discrete series σ_ν of parameter $\lambda + \nu$. Write U for its ambient unitary group and $d_\nu = d(\sigma_\nu, \mu_G)$ for its formal degree on the actual ambient Haar. On the left Bessel subgroup S use $ds = d\mu_W \, d\mu_N$ in coordinates $s = hn$ and $\rho(s) = \psi_z(n)\pi(h)$. Use S',

$dt = d\mu_{W'}\, d\mu_N$ and ρ' from the second frame on the right. Use inner products linear in the second variable. The representation-theoretic premise on each side is the convergent rank-one Jacquet period. In the left convention it reads

$$\mathcal{P}(x,y) = \int_S \langle y, \sigma_\nu(s)x\rangle \rho(s)^{-1}\, ds = L(x)L(y)^*.$$

Use the primed formula for L'. These smooth equivariant maps use the indicated side Haars. Let $d\bar{g}$ be the left quotient measure induced by μ_G and ds. For a nonzero side, put $f(g) = \langle v, \sigma_\nu(g)w\rangle \xi$ and $A(g) = \int_S \rho(s)^{-1} f(sg)\, ds$, with $L(v) \neq 0$ and $\xi = L(v)$. Weil disintegration of $\int \langle A, f\rangle\, d\mu_G$ gives $\int_{S\backslash U} \|A\|^2\, d\bar{g}$, while expansion of A and Schur orthogonality give $\|w\|^2 \langle \mathcal{P}(v,v)\xi, \xi\rangle / d_\nu$. The fitted decay justifies this interchange. Since $A(g) = \|L(v)\|^2 L(\sigma_\nu(g)w)$, cancelling $\|L(v)\|^4$ shows that the isometric Whittaker functional for this quotient is $\sqrt{d_\nu} L$. For the unit-norm packet the Jacquet wave is $\alpha L(\sigma(g)w)$, where $\alpha = \langle L(\sigma(x_0)w), \xi_0\rangle \neq 0$ comes from the chosen test pair (x_0, ξ_0). Division by the quotient norm gives the functional $(\alpha/|\alpha|)\sqrt{d_\nu} L$ for unit w. Using the matched character and quotient normalization in the selected transport, $w = X.B.w(v)$ and the reference vector is $v_0 = \sigma(y)w$ with $y = C.y(\tau)$. The transported packet evaluates the functional at $\sigma(g)\sigma(y^{-1})v_0$. Right translation to gy gives the reference-vector wave. After the selected right rescaling the squared quotient norms are $c_N \dim W$ and $c_N \dim W'$, where W, W' are the compact-type spaces and c_N is the left squared quotient norm per compact dimension. Rank-one period · Quotient norm · Packet normalization · Selected vector and rescaling

For any finite coefficient sum $F(g) = \sum_i \langle x_i, \sigma_\nu(g) y_i\rangle X_i$ with x_i, y_i in the span of translates of K-finite vectors, put $s(F) = \sum_i \langle L(x_i), X_i L'(y_i)\rangle$. In the ambient phase pairing take a summand from F and one from its conjugate. Schur orthogonality applied to the translate $s^{-1}gt$ gives the scalar $\langle \sigma_\nu(s)x_i, x_j\rangle \langle y_j, \sigma_\nu(t)y_i\rangle / d_\nu$ times $(\pi(h)X_i\pi'(h')^{-1}, X_j)_{\mathrm{HS}}$, where $s = nh,\ t = n'h'$ and $(A,B)_{\mathrm{HS}} = \mathrm{tr}(B^*A)$. Include the characters $\psi_z(n)\overline{\psi_{z'}(n')}$. Invert the left variable and reorder $n'h'$ as $h'(h'^{-1}n'h')$ on the right. These changes preserve the product Haars, and the right compact action preserves its character. The remaining integrals give $(\mathcal{P}(x_j,x_i)X_i\mathcal{P}'(y_i,y_j), X_j)_{\mathrm{HS}}/d_\nu$. The rank-one formula contracts this to the i,j term of $|s(F)|^2/d_\nu$. The two rank-one period bounds give absolute integrability after the ambient Schur integral. In the regularized phase, integration by parts at fixed side variables cancels the frequency denominators. The absolute side integral then permits Fubini with the compact dilation average, whose Jacobian cancels under dilation and whose bumps have unit mass. Thus $\beta_\infty(F, F; \mu_G) = |s(F)|^2/d_\nu$. Ambient pairing · Removal of the dilation average

Let $w(c)$ be the fitted big-cell density $(|c_{r+1}|^2 - TT')^{k-1} \prod_{j=1}^r |c_j|^{2(k-1)+4(r+1-j)}$ on the regular positive cell, extended by zero off that cell, and let dc be ordinary Euclidean measure in its l complex coordinates. The normalized big-cell disintegration pushes $w(c)\, dc\, ds\, dt$ under $(c,s,t) \mapsto s^{-1}g_c t$ to Haar $\mu_{\mathrm{cell}} = c_{\mathrm{cell}}\mu_G$. Equivalently, the ambient disintegration writes $\mu_G = C\mu_{\mathrm{cell}}$, so $c_{\mathrm{cell}} = C^{-1}$. Its exact scalar is $\mu_{\mathrm{cell}}(K)/\mu_G(K)$ for any positive-measure compact K in the ambient group. The section-to-phase identity proved below is $J(F) = \mathrm{Re}\, \beta_\infty(F, F; \mu_{\mathrm{cell}})$ on this decay class. The phase is linear in the ambient Haar with both side Haars fixed, so

$$J(F) = c_{\mathrm{cell}} \frac{|s(F)|^2}{d_\nu}.$$

In particular c_{cell} depends only on the fixed place, section and measures. Cell Haar · Section-to-phase identity

Heisenberg coordinates and the positive integral. Write $U = U(H_{k,r})$, $k \geq 1$, $H = H^* = H^{-1}$, and $h(u,v) = v^* H u$. Let $\phi_0, \dots, \phi_r$ be the ordered Euclidean orthonormal isotropic frame, with dual frame $H\phi_j$, and put $\Delta_j(g) = \|g\phi_0 \wedge \cdots \wedge g\phi_j\|$. The critical exponents are $s_j^0 = 2$ for $j < r$ and $s_r^0 = k$. For a step $r = 1 + r'$ the first dual vector $e = H\phi_0$ and $f = \phi_0$ flank a middle block V with form $H' = H_{k,r'}$. In these blocks put

$$N(v,t) = \begin{pmatrix} 1 & -v^*H' & b(v,t) \\ 0 & I & v \\ 0 & 0 & 1 \end{pmatrix}, \qquad b(v,t) = -v^*H'v/2 + it, \qquad L(m) = \mathrm{diag}(1, m, 1).$$

The remaining frame vectors are the lifted lower frame η_j for $0 \leq j < r$. Write $v = \sum_j (a_j H'\eta_j + b_j \eta_j) + \zeta$ with ζ in the Euclidean orthogonal complement of those pairs, and $\beta = t + \Im \sum_j b_j \bar{a}_j$. The horizontal basis is orthonormal, and the central change is a shear. With $p = r,\ m_+ = k - 1$, these are product Lebesgue coordinates on $\mathbb{C}^p \times \mathbb{C}^p \times \mathbb{C}^{m_+} \times \mathbb{R}$. More generally define on this space

$$L_l = 1 + \sum_{j<l} |a_j|^2, \quad R_l = L_p + \|\zeta\|^2 + \sum_{j \geq l} |b_j|^2, \quad B_l = -\|\zeta\|^2/2 + i\beta - \sum_{j \geq l} b_j \bar{a}_j, \quad d_l^2 = L_l R_l + |B_l|^2.$$

For $N(v,t)$ the columns after the first are $\eta_j - \bar{a}_j e$. Their first l Gram determinant is L_l. Projecting the first column $f + v + b(v,t)e$ off these columns leaves squared norm $R_l + |B_l|^2/L_l$, by taking its component along $e + \sum_{j<l} a_j \eta_j$ and its components outside the span of $e, \eta_0, \dots, \eta_{l-1}$. Thus $\Delta_l(N(v,t)) = d_l$. Heisenberg Gram · Coordinate volume · Coordinate integral

Let $I_{p,m_+}(\nu)$ integrate $\prod_{l=0}^{p} d_l^{-h_l - \nu_l}$, where $h_l = 2$ for $l < p$ and $h_p = m_+ + 1$. Write $S_l = \sum_{j \leq l} \nu_j$. Assume

$$S_l > 0 \quad (0 \leq l \leq p), \qquad 2S_p - S_l > 0 \quad (0 \leq l < p).$$

For $p > 0$ put $L = L_1$. Completing the square gives $d_0^2 = L|b_0 - a_0 B_1/L|^2 + d_1^2/L$. The later factors are independent of b_0. Integrating the first factor in b_0 gives $(2\pi/\nu_0) L^{-1+\nu_0/2} d_1^{-\nu_0}$. At fixed a_0, scale the remaining complex variables by $\sqrt{L}$ and the real variable by L. The Jacobian is $L^{2(p-1)+m_++1}$, and d_l for $l \geq 1$ becomes L times the corresponding lower factor. The Jacobian cancels the critical powers. The remaining a_0 weight is $L^{-1-(2S_p - S_0)/2}$, whose integral is $2\pi/(2S_p - S_0)$. Thus by Tonelli,

$$I_{p,m_+}(\nu) = \frac{2\pi}{S_0} \frac{2\pi}{2S_p - S_0} I_{p-1,m_+}(\nu_0 + \nu_1, \nu_2, \dots, \nu_p).$$

The new prefix sums are $S_1, \dots, S_p$, so both strict conditions persist. In rank zero $d_0^2 = (1 + \|\zeta\|^2/2)^2 + \beta^2$. Integration first in β leaves a constant times $(1 + \|\zeta\|^2/2)^{-m_+ - \nu_0}$, integrable for $\nu_0 > 0$, including $m_+ = 0$. The same dilation gives a version with b_0 omitted. The integral of $L^{-1+\alpha/2} d_1^{-\alpha} \prod_{l \geq 1} d_l^{-h_l - \omega_{l-1}}$ in the remaining variables is

$$\frac{2\pi}{\alpha + 2\sum_j \omega_j} I_{p-1,m_+}(\alpha + \omega_0, \omega_1, \dots, \omega_{p-1}),$$

provided the denominator is positive and the lower vector satisfies the prefix conditions. Square completion and planar factors · Rank-zero integral

Character-weighted unipotent tail. Let $p \ge 1,\ s_l = h_l - e_l$ and $F_s^N = \prod_l d_l^{-s_l}$. For $0 < \delta \le 1/2,\ 0 < \gamma \le \delta/8,\ e_l > 0$, assume $e_0 + 2\sum_{l\ge1} e_l < \gamma/2$. Given a complex linear form $Q = \sum_j (\theta_j a_j + \eta_j' b_j) + \sum_i \xi_i \zeta_i$ in these coordinates, put $c_0 \in \mathbb{R}$ and take $\eta_0' \ne 0$. The desired uniform bound is

$$\int F_s^N (1 + |2\Im Q + c_0|)^{-\delta} \ll |\eta_0'|^{-\gamma} (1 + |c_0|)^{-\delta/4}.$$

For a complex quadratic on $[1,2]$ with maximum coefficient modulus $M > 0$, the set where its modulus is at most h has length $\ll (h/M)^{1/2}$. Three points of that set separated by a fixed fraction of its length give this by Lagrange interpolation. Splitting at modulus $\sqrt{M}$ gives the bound $\ll_d (1+M)^{-d/2}$ for the average of $(1+|P|)^{-d}$, $0 < d \le 1/2$. First rotate all complex coordinates by $(1+iu)^2/(1+u^2)$ for $1 \le u \le 2$, preserving measure and d_l. Multiplying the new phase by $1+u^2$ gives a quadratic with coefficients $c_0 + 2\Im Q,\ 4\Re Q,\ c_0 - 2\Im Q$. Its average gives the product $(1+|c_0|)^{-\delta/4}(1+|Q|)^{-\delta/4}$. Then average $a_0 \mapsto u a_0,\ b_0 \mapsto b_0/u$. Volume and B_l are preserved, L_l cannot decrease, and R_l decreases by at most a factor of four. The Gram weight increases by at most $2^{\sum s_l}$. Multiplying the transformed Q by u gives a quadratic with constant coefficient $\eta_0' b_0$. The second polynomial bound reduces the claim to bounding $\int F_s^N |b_0|^{-\gamma}$ since $\gamma \le \delta/8$. Polynomial sublevel estimate · Rotation and reciprocal dilation

For $0 < e < \gamma < 2$ the convolution of $(|b-v|^2 + d^2)^{-1+e/2}$ with $|b|^{-\gamma}$ in $\mathbb{C}$ is bounded by a constant times $(|v|^2+d^2)^{-(\gamma-e)/2}$. Scale to $|v|^2 + d^2 = 1$. The tail is integrable since $e < \gamma$. If $|v| \le 1/2$, then d is bounded below, and otherwise one bounds the two singularities on disjoint disks using $e > 0$ and $\gamma < 2$. Apply this to the square completion with $v = a_0 B_1/L,\ d = d_1/L$. The b_0 integral is bounded by $CL^{-1+\gamma/2} d_0(0)^{-(\gamma-e_0)}$, with $d_0(0)$ evaluated at $b_0 = 0$.

Set $\alpha = \gamma/2 - e_0 > 0$, and $b_* = \Im \sum_{j\ge1} b_j \bar{a}_j$. The bounds $d_0(0) \ge d_1/\sqrt{L}$ and $d_0(0) \ge |\beta - b_*|$ give an integrand bounded by $CL^{-1+(\gamma+\alpha)/2} |\beta - b_*|^{-\gamma/2} W(\beta)$, where $W = d_1^{-\alpha} \prod_{l\ge1} d_l^{-s_l}$. For $l \ge 1$, $d_l^2 = A_l^2 + (\beta - b_l^*)^2$ with $A_l \ge L$, since $R_l \ge L_l \ge L$. Thus $|(\log W)'| \le (\alpha + \sum_{l\ge1} s_l)/(2L)$. On $|\beta - b_*| \le L$, its supremum is bounded by a fixed multiple of its mean, and on the complement $|\beta - b_*|^{-\gamma/2} \le L^{-\gamma/2}$. The central integral gains $L^{-\gamma/2}$. This leaves exactly the weight $L^{-1+\alpha/2}$ in the integral with b_0 omitted, with $\omega_{l-1} = -e_l$. The denominator and every lower prefix and complementary sum are positive because $\alpha > 2\sum_{l\ge1} e_l$. Planar convolution · Central bound and resulting exponent

For the group induction write $n = N(v,t)L(n_1)$, $n_1 \in N_{r-1}$. The unipotent Haar recursion gives a positive scalar times $dv\,dt\,dn_1$. In the middle block z is the lift of z_1, with $T_{z_1} = -h'(z_1, z_1) = T_z > 0$. Gram–Schmidt writes the lower translated frame $n_1\eta_j$ as $\kappa\eta_j$ times an upper triangular matrix, for κ both Euclidean and H'-unitary. Indeed the orthonormalized isotropic frame projects to orthogonal frames of squared length $1/2$ in both signature subspaces, so each projected frame extends to a unitary basis. Its leading determinant moduli are $\Delta_j(n_1)$, and its first entry q satisfies $n_1\eta_0 = q\kappa\eta_0$, $|q| = \Delta_0(n_1)$. Conjugating $N(v,t)$ by $L(\kappa)$ rotates v to $u = \kappa^* v$, preserving measure. The same triangular change on the lifted columns factors the Gram kernel as $F_s^N(u,t)\prod_j \Delta_j(n_1)^{-s_{j+1}}$. The character identity in these coordinates is

$$\lambda_z(N(v,t)L(n_1)) = 2\Im h'(u, \kappa^* n_1 f_1) + \lambda_{z_1}(n_1),$$

where f_1 is the first vector of the lower character frame (the first dual to η_0 if $r-1 > 0$, and z_1 otherwise). Hence $q\eta_0' = h'(\eta_0, f_1)$. Its modulus is 1 when $r - 1 > 0$. In the core $U(k,1)$,

$\eta_0 = (e_+ + e_-)/\sqrt{2}$ and $h'(\eta_0, z_1) = (\bar{z}_{1,+} - \bar{z}_{1,-})/\sqrt{2}$, nonzero by $|z_{1,-}|^2 > \|z_{1,\mathrm{pos}}\|^2$. This is the nonzero coefficient that makes the uniform estimate applicable. Unipotent Haar recursion · Character and nonzero coefficient

For the resulting group integral write

$$J^N_{r,z}(\delta, s) = \int_{N_r} \prod_{j=0}^{r} \Delta_j(n)^{-s_j} (1 + |\lambda_z(n)|)^{-\delta}\, dn.$$

The uniform bound gives $J^N_{r,z}(\delta, s) \ll J^N_{r-1,z_1}(\delta/4, \tilde{s})$, where $\tilde{s}_j = s_{j+1} - \gamma \mathbf{1}_{j=0}$ since $|\eta_0'|^{-\gamma}$ contributes the factor $\Delta_0(n_1)^{\gamma}$. In rank zero $N_0 = \{1\}$ and the integral is the finite Haar mass of that point. If ϵ' is the allowance in rank $r-1$ at $\delta/4$, take $\gamma = \min(\delta/8, \epsilon'/2)$ and $\epsilon_r = \min(\epsilon'/2, \gamma/(4r+4))$. Then $e_0 + 2\sum_{l \geq 1} e_l \leq (1+2r)\epsilon_r < \gamma/2$, and the lower deficits are $e_{j+1} + \gamma \mathbf{1}_{j=0} \in (0, \epsilon']$. This proves $J^N_{r,z}(\delta, s) < \infty$ for every fitted vector and unipotent Haar whenever $0 < e_l \leq \epsilon_r$, with the allowance independent of the vector and Haar. Unipotent rank induction

Ambient inverse-minor integral. Set $M(g) = \|g\|_F$, $\phi'_j = H\phi_j$, $D_j(g) = \det[h(g\phi'_b, \phi'_a)]_{a,b \leq j}$, and $F^G_s = \prod_j |D_j|^{-s_j}$. We show

$$J^G_r(A, s) = \int_U M(g)^{-A} F^G_s(g)\, d\mu_G(g) < \infty, \qquad A \geq A_B(k, r) = 2(k+r)(r+1) + 1, \quad 0 < s^0_j - s_j \leq 1/10.$$

For $r = 0$ average the sole minor on both sides over the compact $U(k)$ fixing the negative axis. For a uniform unit sphere coordinate $Z \in \mathbb{C}$, $\sup_{0 \leq u \leq 1} \mathbb{E}|1 - uZ|^{-\sigma} < \infty$ at $0 < \sigma = s_0 < k$. A sublevel cap satisfies $|1 - Z| < 2t$ when $|1 - uZ| < t$. It has probability $O(t^k)$, using arc length for $k = 1$ and the disk density proportional to $(1 - |Z|^2)^{k-2}$ for $k > 1$. If the right translate of ϕ'_0 under g has negative coordinate b, its positive part has norm $|b|$ by isotropy. The first average therefore costs at most $C|b|^{-\sigma}$. For the second average b has the law of $(\|\gamma_-\| Z - g_{--})/\sqrt{2}$ up to phases, where (γ_-, g_{--}) is the last row and $|g_{--}|^2 = 1 + \|\gamma_-\|^2$. The two averages are bounded uniformly, and preserve the Frobenius weight. The exact rank-zero Haar ball law is $\mu_G\{M^2 \leq k + 1 + 4\rho^2\} = C\rho^{2k}$. Thus the base is bounded by a constant times $\int_0^\infty (k + 1 + 4\rho^2)^{-A/2} \rho^{2k-1} d\rho$, finite at $A > 2k$. Compact base · Frobenius ball law

For the ambient step put $r = 1 + r'$, $n = k + 2r' + 1$, $d = 2n + 1$, and

$$W(c) = \begin{pmatrix} 0 & 0 & \bar{c}^{-1} \\ 0 & I & 0 \\ c & 0 & 0 \end{pmatrix}, \qquad g = N_x^{-1} W(c) L(m) N_y.$$

Here $x, y \in \mathbb{C}^n \times \mathbb{R}$, $m \in U(H')$, and the rank-one Haar formula gives $C|c|^{2n} d^2c\, dx\, dy\, d\mu_1(m)$ on the full-measure cell $c = h(ge, e) \neq 0$. The density comes from cone coarea in $c(f + v + (\xi + it)e)$, where volume carries $|c|^{2n+2}$ and the null equation is $|c|^2(2\xi + v^* H' v) = 0$. The vector stabilizer is $L(m)N_y$ with product Haar because the Levi action $v \mapsto mv$ preserves volume. On the leading flags the outer unipotents add multiples of the first column or row, so $D_0(g) = c$, $D_{j+1}(g) = cD_j(m)$. The lower powers are $s'_j = s_{j+1}$. One-step disintegration · Leading minors

Write $E_A(a, m) = \int M(W(a)L(m)N_y)^{-A} dy$. The block entries bound M below by $M(m)$, $|a|\|v\|$, $|a||t|$. Commuting $L(m)$ through N_y by a volume-preserving change gives the alternative lower bounds $M(m)$, $\|v\|$, $|a||t|$. Split the exponent into $P = 2n + (A-d)/2$, $Q = 1 + (A-d)/2$. For $A > d$, integrating the two resulting radial bounds after scaling by $M(m)$ gives

$$E_A(a,m) \ll M(m)^{d-A}\min(|a|^{-d},|a|^{-1}).$$

Indeed $M^{-A} \le 2^A(M(m)+b_1\|v\|)^{-P}(M(m)+|a||t|)^{-Q}$ for $b_1=|a|$ or 1 in the respective orderings, yielding the multipliers $b_1^{-2n}|a|^{-1}$. Fibre norm estimate

For the left Heisenberg variable put $Y_j = N_x^{-1}L(H')\phi_j$, $G_j = \|Y_0\wedge\cdots\wedge Y_j\|^2$, $D=\sqrt{G_0}>0$. A compact unitary carrying f to Y_0/D brings $N_x^{-1}W(c)$ to $W(cD)L(m_x)N_z$. This follows by applying $W(cD)^{-1}$ after the compact change to obtain an element fixing e. Wedging with the first column kills the multiples of f in the lifted lower columns, so the squared leading Gram volumes of the lower dual frame translated by m_x are G_{j+1}/G_0. Absorb N_z in the right fibre Haar and change m to m_xm by left Haar. The triangular compact change of the lower dual frame then extracts $\Psi=\prod_j (G_{j+1}/G_0)^{-s'_j/2}$ from the lower minor integral, since $E_A(cD,m)$ is left compact invariant. This is the one-sided case of the frame identity below. Compact and Levi reduction · Fibre identity

Put $S=\sum_j s_j$. The ambient integral is bounded by a constant times

$$J^G_{r'}(A-d,s')\int \Psi(x)\int_{\mathbb{C}}|c|^{2n-S}\min(|cD|^{-d},|cD|^{-1})\,d^2c\,dx.$$

After $c\mapsto cD$, the planar factor is $2\pi[(2n+1-S)^{-1}+(S-1)^{-1}]$, finite for $1<S<2n+1$. The admissible deficits give this since $2r+k-(r+1)/10 \le S < 2r+k = n+1$. The remaining factor $\Psi D^{-(2n+2-S)}$ has Gram exponents $t_0 = 2+e_0+2\sum_{j\ge1}e_j,\ t_j=s_j$ for $j\ge1$, where $e_j=s_j^0-s_j$. The squared volumes G_j are the Heisenberg Gram squares above after the volume-preserving change $(v,t)\mapsto(-H'v,-t)$. For the deviations $\nu_0=e_0+2\sum_{j\ge1}e_j,\ \nu_j=-e_j$ the prefix sums are $e_0+\sum_{j=1}^{l}e_j+2\sum_{j>l}e_j$ and $2S_p-S_l = e_0+\sum_{j=1}^{l}e_j$, where these S_l are the sums of the deviations. Both are positive. The positive integration recursion proves convergence. Finally $A_B(k,r)-A_B(k,r')=d+1$, leaving one spare Frobenius power at $A=A_B(k,r)$. This closes the ambient induction on the actual Haar. Planar and Gram reduction

The two majorants. For each $x\in U$, apply isotropic Gram–Schmidt to $x\phi'_j$ as in the unipotent step. Write it as $\kappa_x\phi'_j$ times an upper triangular frame change with leading determinant moduli $\Delta'_j(x)=\|\bigwedge_{a\le j}x\phi'_a\|$. Moving x^{-1} to the other side of the Hermitian pairing gives

$$|D_j(x^{-1}gy)| = \Delta'_j(x)\Delta'_j(y)\,|D_j(\kappa_x^{-1}g\kappa_y)|.$$

In a bi-compact invariant weighted integral the compact factors disappear by Haar invariance. For $y=1$, left compact invariance suffices. The exterior-power bound $|D_j(g)|\le M(g)^{j+1}$ gives $M^{-A}\le F_s^G$ because $\sum_j(j+1)s_j\le(r+1)(k+r)\le A$ and $M\ge1$. Apply it to the translated source after conjugating by H, which preserves norm and Haar and gives $\Delta'_j(HxH)=\Delta_j(x)$. Thus for $\Xi(x)=\prod_j\Delta_j(x)^{-s_j}$ and order-zero envelopes $\|F(g)\|\le bM(g)^{-A},\ \|F'(g)\|\le b'M(g)^{-A}$,

$$\int_U \|F(x^{-1}gy)\|\|F'(g)\|\,d\mu_{\text{cell}}(g) \le bb'\,c_{\text{cell}}J^G_r(A,s)\,\Xi(x)\Xi(y).$$

Choose the common deficit $\min(1/10,\epsilon_r(1/2))$ for both tails. On the left, $nh=h(h^{-1}nh)$ and the compact group bounds the exterior norm distortion uniformly, preserves the unipotent Haar under conjugation, and fixes the character. Its Haar mass is finite, and $(1+|\lambda_z|)^{-2}\le(1+|\lambda_z|)^{-1/2}$. Thus $\Xi(s)(1+|\lambda_z|)^{-2}$ is integrable on S, and the same calculation gives the bound for S' with z'. The

coefficient and side integrals now use one common measurable Ξ and source-independent constants. Two-sided frame identity · Fitted majorant assembly · Side transfer

Portable proof of the ambient inverse-minor integral · Portable proof of the character-weighted unipotent integral

Section-to-phase identity. Keep the fitted section, types and side Haars above. On the ambient side use μ_{cell}, the pushforward just described. Write $\lambda_z(n)$ and $\lambda_{z'}(n')$ for the real character arguments, so the phase at $\tau = (\tau_1, \tau_2) > 0$ is $e^{i(\tau_1\lambda_z - \tau_2\lambda_{z'})}$. Let $V_\tau(F)(c)$ be the phased side integral of $\pi(h)F(s^{-1}g_c t)\pi'(h')^{-1}$, and $J_\tau(F) = \int w(c)\|V_\tau(F)(c)\|_{\mathrm{HS}}^2\, dc$. Let $\beta_{\infty,\tau}$ denote the regularized phase at the same frequencies. The two majorants just proved apply to the fitted section and to all sources with the indicated weighted decay. Phase and weighted estimates

In the phase, take the ambient integral first. Its dilation average uses two smooth unit-mass bumps supported in $(1/2, 2)$. Moving the dilations to the sources in the ambient integral bounds the bracket and its derivatives of total order at most four by a constant times $\Xi(s)\Xi(t)$ times the two weighted derivative bounds. The phase denominators $(1 + \tau_i^2\lambda_i^2)^{-1}$ supply the two side weights locally uniformly on the positive quadrant. This gives absolute convergence and continuity of the regularized phase. For smooth compactly supported product tests in the frequencies, integrate by parts in each dilation variable at fixed (s, t) to cancel these denominators. The Fourier transforms of the tests then give the two quadratic weights before interchanging the side and dilation integrals. The Haar Jacobian cancels the dilation factor, leaving the tested ambient pairing. Regularized phase and frequency average

On the section side, ambient L^1 by the sharp norm-ball bound and Haar disintegration gives absolute side integrability on almost every fibre. There one can expand the square and use Haar invariance to replace the two side variables by a base point and a relative translate. The unitary type actions at the base point cancel. Test in frequency at this stage. The resulting Fourier weights and the two tail estimates make the disintegrated kernel absolutely integrable, allowing Fubini and yielding the same tested ambient pairing. Nonnegative product tests thus give $J_\tau(F) = \beta_{\infty,\tau}(F, F; \mu_{\mathrm{cell}})$ almost everywhere. Fibre continuity and Fatou give the inequality $J_\tau(F) \le \beta_{\infty,\tau}(F, F; \mu_{\mathrm{cell}})$ at every positive τ, and hence a source-uniform $L^2(w\, dc)$ bound for V_τ by the weighted order-four norm. Tested unfolding · Fatou and dilation

To get equality at $(1, 1)$, take $F_a(g) = F(d(a_1)^{-1} g d(a_2))$ for the positive diagonal dilations $d(a_i)$ of the cell. The side substitutions and the rescaling of the cell density give $J_{(1,1)}(F_a) = j(a)J_a(F)$, where $j(a) = (a_1 a_2)^\kappa$ is the cell Jacobian and $j(a) \to 1$. One extra derivative gives $F_a \to F$ in the weighted order-four norm as $a \to (1, 1)$. The Fatou bound applied to $F_a - F$ gives $V_{(1,1)}(F_a) \to V_{(1,1)}(F)$ in $L^2(w\, dc)$. Choose a in the full-measure equality set and use continuity of the regularized phase on the right. This proves the diagonal identity at the fitted character using only the inequality before this limit. Fatou and dilation · Diagonal identity

Finally, the invariant Gram cutoff $\chi_0(G_f/\varepsilon)$, zero near $G_f = 0$ and one for $G_f \ge \varepsilon$, factors through the same section integral. The cutoff square is dominated by $\|\chi_0\|_\infty^2 w(c)\|V_{(1,1)}(F)(c)\|_{\mathrm{HS}}^2$, integrable by the Fatou bound. The multiplier tends to one on the positive cell, giving $J(F)$ as $\varepsilon \downarrow 0$. The finite shifted family uses the same majorants and Haar conversion for its gapped coefficient sources, which have weighted derivatives through order five. The resulting energies are the ones compared in the Pieri determinant expansion. For a common cutoff on the finite family the sum of its integrable square majorants suffices. Gram cutoff · Fitted coefficients and energy · Portable proof

When both maps are nonzero, choose x, y in the translate span with $L(x), L'(y) \neq 0$. Taking X proportional to $L(x)L'(y)^*$ realizes any prescribed amplitude, since the unscaled operator has amplitude $\|L(x)\|^2\|L'(y)\|^2$. The supplied family chooses $|s(F_\nu)| = d_\nu$, scaling an existing base coefficient of amplitude $s_0 \neq 0$ by $d_\nu/|s_0|$. Hence $J(F_\nu) = c_{\mathrm{cell}} d_\nu$. If either side vanishes, the family sets $F_\nu = 0$, its energy is zero, and incident Pieri coefficients vanish. For an arbitrary base coefficient with zero energy one has $s(F) = 0$. The fitted orbital factorization $V(F) = s(F)V_0$ on the regular positive cell for this series gives pointwise vanishing there. Family normalization and zero branches · Base coefficient and zero-energy factorization

The Harish-Chandra formula gives $d_\nu = c_G \prod_{i<j} \frac{|(\lambda+\nu)_i - (\lambda+\nu)_j|}{j-i}$ with c_G fixed by the ambient Haar and signature. Both scalars therefore cancel in ratios of nonzero family energies at this place. In the pinned numerical hypotheses, cloud gaps are at least g_a, the alternating-sign tail has adjacent gaps at least h, and the cloud-to-tail gap is at least b_w. The bounds $g_a, h \geq 100(q_0+1)$ are part of the fit, and $b_w \geq h$ since the tail contains an opposite-sign pair in $[-2b_w, -b_w]$. Thus the minimum over all adjacent gaps of λ is at least $100(q_0+1)$. Subtracting $2q_0$ leaves a bound of at least two on every adjacent gap of every shifted parameter in the ball. Under this all-adjacent bound, for a step $\mu \mapsto \mu + \mathbf{1}_I$ staying in the ball, the increasing factors of the degree ratio have $i \in I,\ j > i,\ j \notin I$. Their product is at most $\prod_{i\in I}\prod_{j>i}(1 + 1/(j-i)) \leq N^{|I|}$. This bounds the target energy against an active source, also when the target is zero. Harish-Chandra formula · Numerical hypotheses · Degree ratios

Portable proof of the real wave calibration

The rotated sum over integral block points

Write p for a composition $s_1 + \cdots + s_h = r+1$ with nonzero antidiagonal block entries $t_a \in E$, and $j_a = \sum_{c\leq a} s_c$. Use the points whose unitriangular orbit contains an integral matrix and for which

$$G_p = D_{r+1}\overline{D_{r+1}} - TT' D_r \overline{D_r}$$

is nonzero and integral. Put $\rho = \mathrm{rd}(E)$. Localize at σ_w by multiplying frame entry j, starting at zero, by $\alpha_j = \rho^{-(r-j)}$. The localized block parameter $t_{a,w}$ is $\rho^{-(2r+1-2j_{a-1}-s_a)}\sigma_w(t_a)$. The leading minor of order j acquires the scale $(\prod_{d<j}\alpha_d)^2$, and the localized Gram is $G_w = \rho^{-2r(r+1)}\sigma_w(G_p)$.

The turn multiplies entry (a,b) in zero-based indices by $(-1)^r(-i)^{a+b}$, and multiplies each leading j-minor by $(-1)^{rj}(-i)^{j(j-1)}$. This preserves the block pattern, minor moduli and Gram. The frame isometry with entries i^a gives the phase $i^b\overline{i^a}$ on the pairings. In the new pairing column b contributes $(-1)^{r-b}$, giving the stated turn. The orbit square O_w is the supremum over permitted representatives with the specified relative root and quotient measures, with value zero if none exists. Its nonzero branch has $G_w > 0$ by representability and $G_p \neq 0$. On this branch $0 < x_w := |\sigma_w(G_p)| \leq d_w := |\sigma_w(D_{r+1})|^2$. The local density n_w is the product of $(|t_{a,w}|^2 - TT')^{k-1}$ for a final singleton and $|t_{a,w}|^{2s_a(2(r+1-j_a)+k+s_a-2)}$ for every other block. For a final singleton, positivity follows by factoring $G_w = |D_r(p_w)|^2(|t_{h,w}|^2 - TT')$. In the arithmetic weights use the unlocalized minors:

$$w_w = \prod_a (1 + |\sigma_w(D_{j_a})|^2)^2, \quad U_w = x_w \prod_a |\sigma_w(D_{j_a})|^2, \qquad W = \prod_w w_w, \quad n_B = \prod_w n_w, \quad \mathcal{N} = \prod_w U_w.$$

Frame conjugation · Detailed proof of the rotated estimate

Let $P = A(r+1)^2$, $Q = 10m^2$, and $u_j = |D_j(p_w)|$. To use the normalized minor bound, let $K_1 > 0$ be the frame comparison factor. It gives $J_1 = K_1 J_w$ and $K_1 n O_w \le F_{\mathrm{fr}} n_1 O_1$ with

$$F_{\mathrm{fr}} = \max(\kappa, \kappa^{-1})^{2m^2} \max(\kappa^r, \kappa^{-r})^{2m}.$$

The common factor cancels. An old minor has modulus at most $B^{C_h m^2}$ times the character-one minor, and the Gram in the character-one frame is $\kappa^{2r(r+1)} G_w$. Positive frame scaling commutes with the turn. The fitted scale bounds give $|\sigma_w(D_j)|^2 = a_j u_j^2$ with $a_j \le B^{(3+C_h)m^2}$, and the inverse scale of $\kappa^{2r(r+1)} G_w$ relative to $\sigma_w(G_p)$ is at most $B^{(3+3C_h)m^3}$. These scale estimates, the two-sided frame comparison and the rank-one torus comparison are used with the packet gap and size conditions. Fitted scale bounds · Detailed proof of the rotated estimate

The weight $w_w^2 (1+d_w)^Q U_w^P$ is bounded by $\prod_j (1 + a_j u_j^2)^{2e_j}$ with

$$e_j = 2 + \lceil P \rceil + \mathbf{1}_{j=r+1}(\lceil Q/2 \rceil + \lceil P \rceil), \qquad \sum_j e_j \le (11 + A) m^3.$$

The full determinant absorbs both the Q-factor and x_w^P since $x_w \le d_w$. If the largest u_j is at most one, use the moment of order zero. Otherwise bound all powers by the moment of order $2\sum_j e_j$ of that minor, costing at most $(2B^{(3+C_h)m^2})^{2\sum_j e_j}$. Choose $C_a = 2(11 + A)$ for the uniform shifted estimate. Its inverse-Gram exponent $40(r+1)^2 = Q$ is independent of the order of the moment, and the scale losses above cost at most $B^{O_{A,C_h}(m^5)}$. We obtain

$$w_w^2 (1+d_w)^Q n_w U_w^P O_w(\mathrm{turn}\, p_w) \le B^{C_K m^5} \max(1, x_w^{-1})^Q J_w. \tag{7B-a}$$

The Gram product uses the arithmetic invariant. Indeed $\prod_w x_w = |N_{E/\mathbf{Q}} G_p|^{1/2} \ge 1$. Therefore

$$\prod_w \max(1, x_w^{-1}) = \frac{\prod_w \max(1, x_w)}{\prod_w x_w} \le \prod_w (1 + d_w).$$

Multiply (7B-a) over the places and cancel the factors $(1+d_w)^Q$. Removing one W bounds each weighted summand by $W^{-1} B^{C_K m^5 f} \prod_w J_w$. The end-minor count in (7d) applies to all these integral block points by unitriangular invariance. It gives $\sum_p W(p)^{-1} \le 64^{f(r+1)}$. The nonnegative finite sums, followed by exhaustion, give

$$\sum_p W(p) n_B(p) \mathcal{N}(p)^P \prod_w O_w(\mathrm{turn}\, p_w) \le B^{(C_K+6)m^5 f} \prod_w J_w. \tag{7B-b}$$

This argument includes compositions with vanishing intermediate minors. The product formula is used on G_p, before the positive rescalings at infinity. Product cancellation · Summation

For the selected nongeneric row bound use $B = b^{C_\beta}$ and the selected thresholds. Each occupied point γ lifts to $p(\gamma)$ of positive Gram at every place, with $\mathrm{toPts}(p(\gamma)) = \gamma$. This makes the lift injective. The weight comparison of the selected packet has constant one. The quotient-amplitude comparison has factor $\prod_w (K_w b^{C_{\mathrm{amp}} m^2})$, and the energy comparison is

$$\Big(\prod_w K_w\Big) \prod_w J_w \le \prod_w M_w \le b^{C_{\mathrm{norm}} m^2 f} I_{\mathrm{re}},$$

where M_w are the local energies of the selected amplitudes. Combining these with (7B-b) proves the arithmetic point bound used in §4. The finite-place second moments, the signed class grouping, the height

and normal-label bounds and the remaining mass factor enter the separate Cauchy–Schwarz estimate there. Selected thresholds · Point and weight ties · Energy comparison

The generic scalar corank-one estimate uses a different quadratic test, formed from the smoothed packet and its permitted derivative words on the specified period lattice. With its Fourier identities and absolute dual summability, Poisson on a coset gives its zero coefficient and the nonzero dual coefficients with the translated phases and the inverse covolume. The zero coefficient is controlled by the preceding normalized energy. The length-one estimate in §4 supplies the nonzero sum for the identified fitted packet. Equation (7B-b) is a bound for positive primal samples covered by the point and measure comparisons, and it survives multiplication by an invariant cutoff in $[0, 1]$. The derivatives of a cutoff and the dual coefficients require their packet estimates. A Euclidean pullback by the turn transports the lattice and its dual if the entire test and operators are transported. The length-one turn is a common sign and preserves the period lattice. For longer blocks the phase can be $\pm i$, so using the original lattice requires a congruence comparison. Detailed proof and Poisson comparison · Scalar Poisson step

4. Boundary terms in the rational sum and the orbit count

Write the geometric sum in (2) as $K^{\mathrm{id}} + K'' + K^{\mathrm{gen}} + K^{\mathrm{ng}}$. For a rational term with first flag F and transported flag F', let D_j be the leading minors of their pairing matrix over the CM field. The algebraic Gram invariant is $D_{r+1}\overline{D_{r+1}} - TT'D_r\overline{D_r}$, with the CM involution in the bars. Its zero locus contains the identity terms and the remaining terms summed in K''. Away from that locus, genericity requires an admissible matrix ξ whose prescribed span equals the localized lattice of the pair. At level primes the span uses the columns $g^{-1}(\mathrm{frame}(F') - \mathrm{frame}(F)\xi)$ and the first lifted flag, where g is the level integer. Away from those primes it is the sum lattice specified by the flags. The second moment compares pairs of terms with the same row label $(\Lambda, \Lambda', \text{residue Levi element})$. Its corank is $s_0 = k - \dim_E \Theta$, where Θ is the first projection of the space of pairs in the two orthogonal complements that agree modulo the first isotropic space and agree after transport modulo the second. In corank one, the scalar test asks that the transport preserve the relations with specified coefficients of both distinguished vectors. Algebraic Gram invariant · Genericity and scalar relation · Corank · Four parts and same-row pairs · Concrete row labels

Release inputs for distinct words. For distinct words the identity cells are empty and the nonidentity zero-Gram summands vanish term by term under the K-two and concrete ties. The vanishing proof uses the stabilizer-character and compact-type comparison with averaging on a rational cocompact subgroup. In particular a nontrivial unitary character kills a covariant integrable average. Translating by an element on which the character has value $\zeta \neq 1$ gives $I = \zeta I$. It is the release's rational and local analysis that supplies the applicable character or type test. The equal-word calculation below uses the same term partition. Distinct-word vanishing

The rational zero-Gram contribution for distinct words

Fix an ordinary selected family point and use the concrete adelic setting of §2, with its K_2 level, prescribed real packets and common row space. Let E/F be the host CM extension over its maximal totally real subfield, $g = g_1 g_2$ the product of the two distinct rational level primes, and $f = [F : \mathbb{Q}]$. Write $\mathrm{e}(x) = \exp(2\pi i x)$. For words $u \neq v$, every rational term t of (2) with algebraic Gram invariant $G_t = 0$ has zero normalized block. In particular $K^{\mathrm{id}}_{uv} = K''_{uv} = 0$. Here the identity condition on a rational class

implies $G_t = 0$. Indeed in an identity representative the transported isotropic vectors equal the first ones modulo the preceding flag steps, and the two distinguished vectors agree modulo the isotropic space. For positive flag length the first r columns of the cross-pairing matrix are zero. For $r = 0$ the distinguished vectors agree, and their common norm gives $G_t = 0$ with $D_0 = 1$. The invariance of G_t under the rational term equivalence is part of the release's term dictionary. Identity representatives Identity implies zero Gram Representative invariants Selected off-diagonal assembly

The primitive character at a level prime. The finite additive character $\psi_{E,f}$ in the concrete setting is trivial on the integral finite adeles. On diagonally embedded $a \in E$ it has value $\mathrm{e}(-\operatorname{Tr}_{E/\mathbb{Q}} a)$, since the real factor is the product over the f conjugate pairs of $\mathrm{e}(\sigma(a) + \overline{\sigma(a)})$. Each of g_1, g_2 splits completely in E. For $p \in \{g_1, g_2\}$, let $h = g/p$ and choose a prime $\mathfrak{P} \mid p$ in E. Choose $\epsilon_{\mathfrak{P}} \in \mathcal{O}_E$ with residue 1 at $\mathfrak{P}$ and 0 at all other primes dividing $g\mathcal{O}_E$. These prime ideals are pairwise coprime and their product is $g\mathcal{O}_E$. The trace on $\mathcal{O}_E/p\mathcal{O}_E \simeq \prod_{\mathfrak{Q}|p} \mathbb{F}_p$ is the sum of the coordinates, as one sees by reducing the matrix of multiplication on $\mathcal{O}_E$ modulo p. Hence $\operatorname{Tr}_{E/\mathbb{Q}} \epsilon_{\mathfrak{P}} \equiv 1 \pmod{p}$. The lift lies in $h\mathcal{O}_E$, so $\operatorname{Tr}_{E/\mathbb{Q}} \epsilon_{\mathfrak{P}} = hb$ with $b \equiv h^{-1} \pmod{p}$. It follows that the release's residue character $\psi_{\mathfrak{P}}$ has the exact value

$$\psi_{\mathfrak{P}}(\bar{m}) = \psi_{E,f}(\epsilon_{\mathfrak{P}} m/g) = \mathrm{e}(-bm/p), \qquad b \equiv h^{-1} \pmod{p}.$$

A change of lift changes the argument by an integral adele. The displayed character has kernel $\{0\}$ on $(\mathbb{F}_p, +)$. In particular its restriction to the selected factor over each place of F above p has conductor exponent one after scaling by g^{-1}. For an integral finite adele C, lift each residue at the primes over g by the corresponding $\epsilon_{\mathfrak{P}}$. This gives $\psi_{E,f}(C/g) = \prod_{\mathfrak{P}|g} \psi_{\mathfrak{P}}(C \bmod \mathfrak{P})$. At a conjugate pair the two characters have the same coefficient $b \bmod p$, so their product on residues m_1, m_2 is $\psi_{\mathfrak{P}}(m_1 + m_2)$. This proves the primitivity and the combination of the two residues used in the finite K_2 character. The integral frame comparison identifying that residue sum for χ_2 with the entry sum ϕ_J in the conjugated finite matrix coordinates is a **release input**. Adelic character and sign Complete splitting Chinese remainder lifts Residue character Character on both residues Finite frame comparison

Comparison on the rational stabilizer. Write N_Λ for the rational isometries preserving Λ which lower the isotropic flag by one step under $n - 1$ and satisfy $(n-1)X^\perp \subset X$, where X is its top isotropic space. For a term $t = (\Lambda, \Lambda', [\gamma])$, put

$$R_t = N_{u,\infty} \cap \gamma_\infty N_{v,\infty} \gamma_\infty^{-1}, \qquad \Gamma_t = (N_\Lambda \cap \gamma N_{\Lambda'} \gamma^{-1})_\infty.$$

The release supplies $\Gamma_t \subset R_t$, the discreteness of Γ_t, its cocompactness in R_t, and integrability of the packet product on the chosen fundamental domain Ω_t. Groups and domain Discreteness Cocompactness Convergence in the concrete setting

Let K_N be the subgroup of integral finite isometries which satisfy the first two congruences of K_2, on the reference isotropic basis e^0_a and distinguished vector z^0. Use indices $0 \le a < r$, and set $q^0_a = e^0_{a+1}$ for $a < r - 1$, $q^0_{r-1} = z^0$. For the duals f^0_a write

$$B_N(k) = \sum_a H(f^0_a, k q^0_a), \qquad \chi_N(k) = \psi_{E,f}(B_N(k)/g).$$

Here H is the finite Hermitian form, conjugate-linear in its first variable. Modulo g, the congruences make $k - 1$ lower each stage of the reference flag, including the stage obtained by adding z^0, by one. In

particular $(k-1)(k'-1)q_a^0$ is in the span of the e_b^0 with $b<a$ modulo g. The congruence dual f_a^0 annihilates this span and pairs to zero with q_a^0 modulo g. Thus

$$B_N(kk')-B_N(k)-B_N(k')=\sum_a H(f_a^0,(k-1)(k'-1)q_a^0)-B_N(1)\equiv 0 \pmod g.$$

The reference pairing is integral on these lattice vectors, so the congruence implies $\chi_N(kk')=\chi_N(k)\chi_N(k')$. This argument uses the nilradical part of the finite character. The full χ_2 includes the additional entries from the compact-complement frame and is evaluated on K_2 in the block formula. Root coordinates and level Nilradical character Congruence subgroup and multiplicativity

Choose the release's base column classes of Λ,Λ', with finite representatives $A=h_\Lambda x_u$, $B=h_{\Lambda'}x_v$. Two integral comparisons are used here as **release inputs**. The lattice dictionary and the lift of the flags give $\kappa=A^{-1}\gamma_f B\in K_N$. For $n\in N_\Lambda$ the nilradical comparison gives $A^{-1}n_fA\in K_2$ and $\chi_N(A^{-1}n_fA)=\psi_{u,f}(n_f)$, and likewise on the other side. For n in the common rational stabilizer set $n'=\gamma^{-1}n\gamma$. Then

$$A^{-1}n_fA=\kappa(B^{-1}n'_fB)\kappa^{-1}.$$

By multiplicativity on K_N the two finite characters of n_f,n'_f agree. On each rational unipotent the product of the finite and infinite characters is 1, since both evaluate the same rational root sum divided by g. Consequently the relative character

$$\xi_t(s)=\psi_{u,\infty}(s)\,\overline{\psi_{v,\infty}(\gamma_\infty^{-1}s\gamma_\infty)}\qquad (s\in R_t)$$

is trivial on Γ_t. The use of K_N allows the base transporter κ to have only the first two level congruences. Base lattice and lifted-flag congruences Nilradical comparison Finite comparison on the common stabilizer Product formula for rational unipotents

The rational trace test. The integrand $\mathcal{W}_t(x)=W_u(x)\otimes W_v(\gamma_\infty^{-1}x)^*$ satisfies $\mathcal{W}_t(sx)=\xi_t(s)\mathcal{W}_t(x)$ on R_t. By the finite comparison it descends to $\Gamma_t\backslash G_\infty$. For the inherited Haar measure there, disintegrate along $\Gamma_t\backslash R_t$ over $R_t\backslash G_\infty$. Both G_∞ and the real unipotent group R_t are unimodular. The measures in this disintegration are taken compatibly with the ambient Haar and counting measure on Γ_t. If $\xi_t(s_0)\neq 1$ for some $s_0\in R_t$, right translation by s_0 multiplies the character average on $\Gamma_t\backslash R_t$ by $\xi_t(s_0)$ and preserves its finite measure. Hence the fibre integral and O_t vanish. Group and character covariance Averaging on the common unipotent group

For a pair of rational flags $F_u,\gamma F_v$, let $\ell_{F_u}(Y)$ denote the sum of the simple-root coefficients of Y on the first flag. The coefficients are those of e_a in Yq_a modulo the span of the preceding e_b, with $q_a=e_{a+1}$ or z_u at the last step. The rational relevance test compares $\mathrm{Tr}_{E/F}\,\ell_{F_u}(Y)$ with $\mathrm{Tr}_{E/F}\,\ell_{\gamma F_v}(Y)$ for all Y in the common rational nilpotent Lie algebra. If t is not relevant, there is such a Y with

$$\delta=\mathrm{Tr}_{E/F}(\ell_{F_u}(Y)-\ell_{\gamma F_v}(Y))\neq 0.$$

Take $a=g/(2f\delta)\in F$ and $s_0=\exp(aY)_\infty\in R_t$. The exponential is a finite polynomial and is unitary because aY is skew Hermitian. Powers Y^jq_b for $j\geq 2$ land before e_b by flag lowering. Thus the simple-root sums of the exponential on the two flags are $a\ell_{F_u}(Y)$ and $a\ell_{\gamma F_v}(Y)$. The resulting quotient is

$$\xi_t(s_0)=\mathrm{e}\big(g^{-1}\,\mathrm{Tr}_{F/\mathbb{Q}}(a\delta)\big)=\mathrm{e}(1/2)=-1.$$

In this construction s_0 is required to belong to the real common unipotent group. The finite congruences control the lattice Γ_t on which the quotient character is trivial. Relevance and representative Nilpotent exponential and root coefficients Rational witness and its phase

Relevant zero-Gram terms. Suppose t is relevant and $G_t = 0$, and consider the case $\xi_t = 1$ on R_t, since the preceding averaging applies whenever it is nontrivial. The **release input** of the relevant zero-Gram normal form and its compact lift has the following consequence. If S is the sum of the two rational frame spans, its real orthogonal complement at each real place w contains a positive k-plane L_w. Acting by $M \in U(L_w)$ there and by the identity on $L_w^{\perp}$ and at all other places fixes both frames pointwise. In the two packet types this acts at w by M and $C_w M C_w^{-1}$, respectively, in unitary coordinates, for a fixed $C_w \in U(k)$. The normal form obtains L_w by correcting an orthonormal frame of the first positive complement by isotropic vectors. On the second side the corrected frame differs from a unitary change of this frame by vectors in the common isotropic space, which the lifted group fixes. Rational normal form and base change Compact lift and action

Let H_w be this lifted compact group at one real place. It normalizes R_t because it fixes both frames. Its action by conjugation on a Haar measure of R_t has modulus one, since the modulus is a continuous homomorphism from H_w to $\mathbb{R}_{>0}$. The packet product is now R_t-invariant, so O_t is the integral on $R_t\backslash G_\infty$ multiplied by $\operatorname{vol}(\Gamma_t\backslash R_t)$ in the above disintegration. The map $R_t x \mapsto R_t h x$ for $h \in H_w$ preserves the quotient measure by the Haar formula. Both Bessel characters of h are 1 because h fixes both ordered frames. It follows by changing variables that

$$\pi_u(h) O_t = O_t\, \pi_v(\gamma_\infty^{-1} h \gamma_\infty).$$

This change of variables uses the integral on $R_t\backslash G_\infty$ and requires no invariance of the lattice Γ_t under h. Pointwise stabilizer and its character Release covariance

In the selected family the words inject into their Boolean letter strings. A differing letter at index j changes entry $j+1$ of the compact $U(k)$ parameter at every place. The entry is the common parameter at j minus $1/2$ and minus Δ times the Boolean letter, where $\Delta > 0$ is the selected gap. The two irreducible local types are therefore inequivalent. Contract O_t against arbitrary pure tensors in the factors away from w on both sides. This gives an intertwiner of the two local types, one composed with the inner conjugation by C_w. Inner conjugation preserves the isomorphism class, so Schur's lemma kills this intertwiner. Since pure tensors span, $O_t = 0$. Letter formula Word separation at every place Compact-type comparison

Finally, each matching-class block of (4) has the multiplier $\operatorname{vol}(K_2)\chi_2(k(\rho,\rho',t))/\sqrt{c_u^X c_v^X}$ times O_t, with $c_u^X = \operatorname{vol}(K_2)\mu_{N_u}(F_{N_u})\mu_{H_u}(H_u)/\dim \pi_u$ and the corresponding c_v^X. Each unmatched block is zero by definition. The normalized contribution is thus identically zero, term by term, including any identity blocks. Block normalization Zero-term assembly

The rational nonzero-Gram contribution for distinct words

Fix two distinct words u, v at the same ordinary selected point (x, i). Use the real scale 1, $m = 2r+2$, and the same measures and split comparison primes as in the zero-Gram argument. Set

$$D_{uv} = d_u d_v / D_\star, \qquad D_\star = d_\sigma \operatorname{volG}/\mu_{K''}, \qquad q_{uv} = d_\sigma I_{\mathrm{oo}},$$
$$\varepsilon_{\mathrm{loop}} = \exp(-kf\log k/3), \quad \varepsilon_{\mathrm{ng}} = 2^{-(k-2r-5)}\exp(C_n m^5 f \log b), \quad P_g = (1+3/g_1)^f (1+3/g_2)^f.$$

Here $\mathrm{vol}\mathbf{G} = \mathrm{vol}(G(F)\backslash G(\mathbb{A}_F))$ and $\mu_{K''}$ is the finite level volume in the tied mass convention. Proposition 3 gives $0 \le q_{uv} \le 1$. The comparison constants and the choice of setting precede u, v. The source's strict family and selection thresholds apply at every eligible index i of the ordinary family. Selected statement · Ordinary selected proof · Size and loop hypotheses in the selection

Weighted same-row expansion

For a rational term t, write $\lambda(t) = (\Lambda, \Lambda')$ for its two lattices and $i(t)$ for the residue Levi element of the base transporter κ_t. The row is $(\lambda(t), i(t))$. The vector v_t in the geometric Hilbert space is supported at $\lambda(t)$, with value

$$\sqrt{w_u(\Lambda) w_v(\Lambda')}\, \frac{\mathrm{vol}\, K_2}{\sqrt{c_u^X c_v^X}}\, \chi_N(\kappa_t) O_t, \qquad w_u(\Lambda) = |\operatorname{Aut}_{\mathrm{ind}}(\Lambda, F_u)|^{-1}.$$

Here $\operatorname{Aut}_{\mathrm{ind}}$ is the induced automorphism group on the orthogonal quotient of the top isotropic space, O_t is the real integral over the rational intersection quotient, and χ_N is the nilradical character of the preceding zero-Gram section. The full χ_2 of a column block supplies the remaining Levi matrix. In particular a pair in one row has the factor $w_u(\Lambda) w_v(\Lambda')$, and the phase in $\langle v_t, v_{t'} \rangle$ is $\overline{\chi_N(\kappa_t)} \chi_N(\kappa_{t'})$ with the first slot conjugate-linear. Induced group and row weight · Vectors and residue Gram matrix · Block expansion

At a split residue place of size q, let $U \subset \mathrm{GL}_k(\mathbb{F}_q)$ be upper unitriangular, ψ_U the primitive superdiagonal character extended by zero, and $\gamma_q(a) = |U|^{-1} \sum_y \psi_U(yay^{-1})$, summing over $\mathrm{GL}_k(\mathbb{F}_q)$. Then $\gamma_q(1) = |\mathrm{GL}_k|/|U|$. For a unitary irreducible representation π, the operator $|U|^{-1} \sum_{a \in U} \psi_U(a) \pi(a)$ is the orthogonal projection onto the ψ_U^{-1}-space. Summing its conjugates and taking traces gives the eigenvalue of convolution by $\gamma_q / \gamma_q(1)$, namely $|U| \dim(\pi^{U, \psi_U^{-1}}) / \dim \pi$. Complex conjugation gives the same bound for the matrix convention of the release. The finite Gelfand–Graev multiplicity-one and generic degree bounds give

$$0 \le \frac{|U| \dim(\pi^{U, \psi_U^{-1}})}{\dim \pi} \le \prod_{j=1}^{k-1} (1 - q^{-j})^{-1} \le 1 + 3/q.$$

The degree bound arises by restricting to the last-column abelian radical of the mirabolic. Its $q^{k-1} - 1$ nontrivial character spaces have equal dimension by the transitive GL_{k-1}-action, and the chosen character space carries the lower-rank Whittaker character under its stabilizer. Induction gives $\prod_{j=1}^{k-1} (q^j - 1)$. For the last numerical inequality, write $y = 1/q$. The product for $j \ge 2$ is at least $1 - y^2/(1-y)$, before taking reciprocals, so the full inverse product is at most $(1 - y - y^2)^{-1} \le 1 + 3y$ for $0 \le y \le 1/3$. Taking the product over the f places above each level prime gives $0 \preceq G_{\mathrm{Levi}} \preceq P_g \Gamma_1 I$, where $\Gamma_1 = \prod_{\mathfrak{p} \mid g} \gamma_{q_\mathfrak{p}}(1)$. Finite convolution normalization and degree estimate · Finite product matrix

For a set S of terms put $v_i(S) = \sum_{t \in S : i(t) = i} v_t$. Expand the finite matrix inequality against these vectors. Orthogonality between different lattice supports and absolute convergence yield

$$\|K(S)\|_{\mathrm{HS}}^2 \le P_g \Gamma_1 \sum_i \|v_i(S)\|^2 = P_g \Gamma_1 \sum_{\substack{t, t' \in S \\ \mathrm{row}(t) = \mathrm{row}(t')}} \Re \langle v_t, v_{t'} \rangle. \tag{7a}$$

Reversing the order conjugates the summand, so a pair of distinct terms contributes twice its real part. The real parts can have either sign. We take absolute values for the loop classes below, and use the signed expansion in (7a) for the scalar and nongeneric classes. For nonzero-Gram terms of one row the corank s_0 is at least one. For two generic terms split the pairs into $s_0 \geq 2$, $s_0 = 1$ nonscalar, and $s_0 = 1$ scalar. For two nongeneric terms split at $s_0 = 1$. Cross terms between the generic and nongeneric sums are bounded by the Hilbert-space triangle inequality. Pair definitions and convergence · Gram and pair split · Abstract pair split

Loop counts

Let $X = kf \log k$ and let $\Gamma = \exp(C_{\mathrm{Gam}} m^5 f \log b)$ be the column-window parameter, distinct from Γ_1. The ordinary window is $C_0^{-1} \leq d_w \Gamma^2 / D_\star \leq C_0$ for every word w, with $C_0 = 3$ in the selected assembly. In particular $1 \leq C_0 \Gamma^2 d_v / D_\star$. The loop estimate begins with $|\langle v_t, v_{t'} \rangle| \leq \|v_t\| \|v_{t'}\|$. In the passage from concrete to lattice pairs, the sum of second-lattice weights in a fibre is at most one. The first-lattice weight remains. Summing it gives the genus mass of the first rows, which, multiplied by $\Gamma_1 \dim \pi_u$, is d_u by the tied volume identity. These are weighted counting statements. Replacing the weights by unit multiplicities would change the mass. Weighted fibre inequality · First-row mass and use of the window

For each stratum the count bounds the sum of $w_u(\Lambda) A_t A_{t'}$ over any finite set of lattice pairs, where A_t is the norm of the normalized orbital operator before the row weights. Its right-hand side is the first-row genus mass times $\dim \pi_u$ times one of the factors

$$\frac{d_v}{D_\star} B_s, \quad D_s, \quad \frac{d_v}{D_\star} B_1^*$$

for the nondegenerate and degenerate loop strata of corank $2 \leq s \leq k$, and the generic nonscalar corank-one stratum, respectively. Here degeneracy refers to the loop stratum inside the nonzero-Gram set. For a fixed sufficiently large C and a nonnegative count factor Ξ_s, the bounds are

$$\begin{aligned} H_0 &= b^{Cm^2 f}, \\ R_s &= (k v_{\mathrm{host}})^{kf} \Xi_s C^{kf} H_0^{C(s+m)m^2} (Csf)^{6sf} \Gamma, \\ B_s &= (2/g)^{(s-1)kf} R_s, \\ D_s &= (2/g)^{(s-2)kf} R_s \, e^{Cm^6 f \log b} (C/k)^{kf}, \\ B_1^* &= C^f H_0^{Cm^3} \Gamma e^{Cm^6 f \log b} \big(C g^{-(2r+1)} k^{-r} \big)^{kf/2}. \end{aligned} \tag{7b}$$

The v_{host} in $k v_{\mathrm{host}}$ denotes the host covolume parameter. The loop estimates use the weighted mean count of primitive sublattices over a genus, the Hermitian determinant and extension-class bounds, and partner counts. The archimedean factors use the large-gap coefficient bounds and the stabilizer covolume estimate on relevant terms. These inputs must apply to the same word-pair lattices and quotient measures. Weighted pair and strata · Three stratum bounds · Genus count and partner inputs · Amplitude and covolume inputs

The level saving in B_s is visible before summing over Gram data. In a nondegenerate bin with Gram parameters $G_1, G_2 > 0$, norm parameter N_T and remaining count factor ΞW, the count is bounded by

$$(2/g)^{(s-1)kf} \big(\sqrt{G_1 G_2} / N_T \big)^k \, \Xi W.$$

The two squared-amplitude estimates bound $E(G_j/N_T)^k A_j^2$ by $\ell_j\tau$, with $E = (e\pi/(kv_{\mathrm{host}}))^{kf}$ and $\ell_j, \tau \geq 0$ from the real bound. Multiplying the count by A_1A_2 and taking the geometric mean of the two estimates cancels $(\sqrt{G_1G_2}/N_T)^k$, giving the factor $(2/g)^{(s-1)kf}E^{-1}\Xi W\sqrt{\ell_1\ell_2}\tau$. The degenerate bin has one fewer level direction. Its count also has factors $T^{-f}K_1K_2$ inside the k-th power. Here $T, K_1, K_2 > 0$ are the bin parameters in the source's degenerate bound. They cancel against the geometric mean of the factors $K_j^{2k}T^{-kf}$ in the two squared-amplitude estimates. In particular the level factor in D_2 is one, and that stratum inherits the k^{-kf} real saving. Count and amplitude cancellation

In the nonscalar corank-one bound, after the congruence comparison the two real parameters at each place are written as $h_{j,w} = \sinh u_{j,w} > 0$, $c_{j,w} = \cosh u_{j,w}$ for $j = 1, 2$. The product lower bound is $(\sqrt{TT'}/g)^f \leq \prod_w (c_{1,w} + c_{2,w})/(h_{1,w}h_{2,w})$. Since $c_{j,w} \leq 1 + h_{j,w}$, each factor is at most $4/(\min(1, h_{1,w})\min(1, h_{2,w}))$. Consequently the product of the two truncated products $P_j = \prod_w \min(T'h_{j,w}^2, T')$ is at most $(4g\sqrt{T'/T})^{2f}$, and at least one P_j is at most $(4g\sqrt{T'/T})^f$. Combining this with the real amplitude bound on that member and taking the geometric mean of the two available product bounds is the source of the square root in B_1^*. Nonscalar small-cell bound · Small-cell factor in the amplitude

More precisely, with N_T the positive norm parameter used in the loop count, the source requires $4 \leq m \leq f \leq k$, $2 \leq b$, $g \geq (kb)^8$, $1 \leq v_{\mathrm{host}} \leq b$, $1 \leq N_T \leq b^{Cmf}$, $1 \leq \Gamma \leq e^{Cm^5 f \log b}$, and $k \geq K_0 m^5 f \log b$, where K_0 depends on C. The two Ξ bounds are $(kv_{\mathrm{host}})^{kf}\Xi_s \leq H_0^C C^{skf}$ for $2 \leq s \leq k-2$, and $\Xi_s \leq (Ck)^{kf}(H_0^{Cm}N_T)^{Ck}$ for $s = k-1, k$. Let $R_s^+ = R_s\Gamma^2 e^{Cm^6 f \log b}C^{kf}$. Terms with $\Xi_s = 0$ vanish. In the middle range,

$$\frac{\log R_s^+}{kf} \ll_C s + 1 + \frac{(s+m)m^4 \log b}{k} + \frac{s\log(Ck^2)}{k} + \frac{m^6 \log b}{k} \ll_C s + 1.$$

Here $m^6 \log b/k \leq m/(K_0 f) \leq 1/K_0$. Increasing K_0 gives $\log R_s^+/(kf) \leq (s - 2 + 1/4)\log k$ throughout the middle range, including $s = 2$. In the top two rows the bound is $O_C(\log k + (1+m^4)\log b)$, using $s \leq k$, $m \leq k$ and the top Ξ bound. This is at most $\frac{1}{4}(s-2)\log g$ for K_0 sufficiently large. With $\log 2, \log k \leq \frac{1}{8}\log g$, these bounds give each $B_s \leq g^{-kf/2}e^{-X}$ and each $\Gamma^2 D_s \leq e^{-3X/4}$. In B_1^*, the prefactor in (7b) has logarithm $O_C(kf)$ by $k \geq K_0 m^5 f \log b$ and $m \leq f$. The $g^{-(2r+1)kf/2}$ factor absorbs this prefactor while giving at least e^{-X}. Summing the bounds yields

$$\sum_{s=2}^{k} B_s \leq g^{-kf/2}, \qquad \Gamma^2 \sum_{s=2}^{k} D_s \leq e^{-X/2}, \qquad B_1^* \leq e^{-X}.$$

The first term is at most e^{-4X}. Multiplying by the fibre and genus bounds, and using the lower column window on the degenerate stratum, bounds the absolute pair sum of the union by

$$A_{\mathrm{lp}}D_{uv}(e^{-4X} + e^{-X/2} + e^{-X}) \leq e^{-X/3}D_{uv} \quad \text{if } 6\log(3A_{\mathrm{lp}}) \leq X,\ A_{\mathrm{lp}} \geq 1.$$

Indeed $3A_{\mathrm{lp}}e^{-X/2} \leq e^{-X/3}$. The selected route can take $A_{\mathrm{lp}} = \max(1, C_0)$ in the tied row bounds. The union consists of all corank-≥ 2 pairs in the nonzero-Gram set and the generic nonscalar corank-one pairs. In particular $\varepsilon_{\mathrm{loop}}D_{uv}$ bounds each of the two absolute row sums used in (9). Sizes and three exponent bounds · Proofs of the exponent bounds · Absolute sum over the union · Selected loop construction

Nongeneric corank-one pairs

The nongeneric corank-one relation, including equality of rows, is an equivalence relation on the nongeneric terms. Write c for one class. In each class group terms by the rational point $p = (\mathbf{s}, \mathbf{z})$, a composition $\mathbf{s}$ of $r+1$ and its nonzero block parameters $\mathbf{z}$, and the normal label ν. Let $U_{c,p,\nu} = \sum v_t$ over this group. The exact contribution divided by Γ_1 is $\sum_c \| \sum_{p,\nu} U_{c,p,\nu} \|^2$. In particular the two signs of interference in $\Re\langle v_t, v_{t'} \rangle$ remain in the group sums. The release compares

$$\|U_{c,p,\nu}\|^2 \le a^2(c,p,\nu) J_{\mathrm{loc}}(p), \qquad J_{\mathrm{loc}}(p) \le c_{\mathrm{St}} J_B(p).$$

The local scalar factorization uses a common orbital packet in the class at (p,ν) and real quotient measures with the specified stabilizer comparison. Write $\mathrm{rd} = 2v_{\mathrm{host}}$, $TT' = T_u T_v$, $P_c = \mathrm{rd}^{12r(r+1)f}(TT')^{6f}$. At fixed c, p, at most $P_c N_G(p)^6$ normal labels contribute, where N_G is the product of real absolute values of the scaled Gram invariant. Let N_D be the product of squared scaled block-end minors, $n(p)$ the real block density and $H(p)$ the product of the finite heights. On the contributing points $N_G, N_D, H \ge 1$ and $H \le g^f N_G N_D$. The density $n(p)$ includes N_G^{k-1} and the block-parameter powers for the composition, including larger blocks. Class sums and scalar, count and height inputs · Scaled Gram and block density · Concrete finite and real ties

At fixed p unfold the finite sum over (c,ν) to local tuples $d = (d_{\mathfrak{p}})$, with local weights $w_{\mathfrak{p}}(d_{\mathfrak{p}}) \ge 0$. For each tuple the sum of class weights in its fibre is at most $V_0 \prod_{\mathfrak{p}} w_{\mathfrak{p}}(d_{\mathfrak{p}})$. If Ω is the finite set of places allowing a nongeneric term, the complex finite amplitude is

$$a_{\mathrm{ng}}(d) = \prod_{\mathfrak{p}} (a_{Z,\mathfrak{p}} + a_{N,\mathfrak{p}}) - \prod_{\mathfrak{p}} a_{Z,\mathfrak{p}} = \sum_{\varnothing \ne A \subseteq \Omega} \prod_{\mathfrak{p} \in A} a_{N,\mathfrak{p}} \prod_{\mathfrak{p} \notin A} a_{Z,\mathfrak{p}},$$

where $a_N = 0$ off Ω. The squared finite amplitude before unfolding is the class weight times $|a_{\mathrm{ng}}(d)|^2$. The weighted second moments of a_Z and a_N are bounded by $H_{\mathfrak{p}}^{C_1(r+1)^2} n_{\mathfrak{p}}$ and $q_{\mathfrak{p}}^{-\kappa} H_{\mathfrak{p}}^{C_1(r+1)^2} n_{\mathfrak{p}}$, respectively, with $\kappa = k - 2r - 5$. Off Ω, $n_{\mathfrak{p}} = n_{\mathrm{gen},\mathfrak{p}}$. On Ω, the bound is $n_{\mathfrak{p}} \le 8 H_{\mathfrak{p}}^{C_0'(r+1)^2+2m} n_{\mathrm{gen},\mathfrak{p}}$, with $H_{\mathfrak{p}} \ge q_{\mathfrak{p}} \ge 2$. Here C_0' denotes the local mass exponent, independent of the column-window constant. These local inequalities and weighted unfolding are quantitative inputs. Local amplitudes, weights and moment bounds · Pinned local estimate in the family · Selected nongeneric comparison

For each nonempty A, the product of local squared norms has a factor $\prod_{\mathfrak{p}\in A} q_{\mathfrak{p}}^{-\kappa} \le 2^{-\kappa}$. Apply Cauchy–Schwarz over the subsets and sum in d by factorization. This costs at most $(2^{|\Omega|}-1)^2 \le 4^{|\Omega|}$. At each place in Ω, absorb $4 \cdot 8 \le H_{\mathfrak{p}}^5$. With $e_7 = (C_1 + C_0')(r+1)^2 + 2m + 5$ the result is

$$\sum_{c,\nu} a^2(c,p,\nu) \le 2^{-\kappa} H(p)^{e_7} n_{\mathrm{gen}}(p). \tag{7c}$$

This also holds for an empty Ω, when the amplitude vanishes. The appearance of at least one local saving follows from excluding the empty subset before squaring, with the class weights retained. Weighted product argument

For a block composition, the successive end minors z_a^{end} satisfy $z_a^{\mathrm{end}} / z_{a-1}^{\mathrm{end}} = (-1)^{s_a(s_a-1)/2} z_a^{s_a}$, with $z_0^{\mathrm{end}} = 1$. They are nonzero algebraic integers on the occupied points. Set $W(p) = \prod_{a,w} (1 + |\sigma_w(z_a^{\mathrm{end}})|^2)^2$. We have

$$\sum_{p\ \text{occupied}} W(p)^{-1} \le 64^{f(r+1)}. \tag{7d}$$

Indeed distinct algebraic integers of E differ by an element with absolute norm at least one, so their images in $\mathbb{C}^f$ have supremum distance at least one. Take disjoint open products of discs of radius $1/2$ centered at these images. On each disc, $1+|z|^2 \le 2(1+|a|^2)$ for center a. Integration of $\prod_w (1+|z_w|^2)^{-2}$ gives at most π^f over the disjoint products, so the sum of the weights of the centers is at most $4^f\pi^f/(\pi/4)^f = 16^f$. A composition of length ℓ therefore costs $16^{f\ell}$ for its end minors, and each tuple lifts at most $\prod_a s_a \le 2^{r+1}$ times by the displayed power equations. There are 2^r compositions. This proves (7d) with room in the constant. Block-end formula, support and bound · Pinned counting theorem

Use weighted Cauchy–Schwarz for each class with weight $W(p)P_cN_G(p)^6$ on $U_{c,p,\nu}$. The sum of inverse weights is bounded by (7d) because the number of normals cancels $P_cN_G^6$. Sum the resulting second brackets over the classes and apply (7c). The generic mass comparison is

$$n_{\mathrm{gen}}(p) \le A_{\mathrm{re}} e_1 H(p)^{e_N} n(p), \qquad e_N = 10(r+1)^2 + 6m + 1.$$

For $A_{\mathrm{exp}} = A(r+1)^2$ with $A \ge C_1 + C_0' + 38$, one has $e_7 + e_N + 6 \le A_{\mathrm{exp}}$. The inequality $H \le g^f N_G N_D$ then absorbs the height and normal factors. The archimedean point estimate used at this stage, for each finite subset of occupied points, is

$$\sum_p W(p) n(p) (N_G(p)N_D(p))^{A_{\mathrm{exp}}} J_B(p) \le \exp(C_B m^5 f \log b)\, I_{\mathrm{re}}.$$

Here I_{re} is the normalized density integral of the nongeneric bound. Equations (7B-a)–(7B-b) give the point estimate by the injective selected lift and the quotient and energy comparisons of §3. Substituting the point estimate in the second Cauchy–Schwarz bracket gives the majorant

$$\Gamma_1\, 64^{f(r+1)}\, 2^{-\kappa}\, P_c c_{\mathrm{St}} A_{\mathrm{re}} e_1\, g^{f A_{\mathrm{exp}}} e^{C_B m^5 f \log b} I_{\mathrm{re}}.$$

These bounds hold first on finite sets. Absolute convergence gives the class sums, and the majorants are nonnegative. Mass, point estimate and class bound · Construction of real and finite inputs · Rotated estimate and injective summation · Selected uniform Theorem B call

The exact comparison is

$$\Gamma_1 A_{\mathrm{re}} I_{\mathrm{re}} \le (\mathrm{rd}/2)^{f(r+1)} g^{f(r+1)(2r+3)} D_{uv}.$$

Here $A_{\mathrm{re}} = \mathrm{rd}^{2r(r+1)(k-1)f}\,\mathrm{AA}$ in the source. Under $0 < \mathrm{rd} \le b$, $1 \le g \le b^{C_g}$, $0 < TT' \le b^{2C_T m}$, $c_{\mathrm{St}} \le b^{C_{\mathrm{St}} m^3 f}$ and $0 \le e_1 \le 2$, the product of the factors apart from $2^{-\kappa} D_{uv}$ has logarithm at most $C_n m^5 f \log b$. This proves the $\varepsilon_{\mathrm{ng}} D_{uv}$ bound for the nongeneric corank-one sum. Frame-passage normalization · Size absorption

The selected finite densities

At a finite place w, put $O = \mathcal{O}_{F_w}$, $A = \mathcal{O}_{E_w}$, $q = |O/\pi O|$ for a uniformizer π, and $a = m = 2r+2$, $M = k + a - 1$. The selected host assumption makes the relative extension split or unramified quadratic at every finite place. In the tied local model write Λ_w for the self-dual A-lattice obtained by base change of the prescribed host lattice L_0. The density test uses primitive a-frames in Λ_w. Primitivity means that the

columns form a basis of the intersection of their E_w-span with Λ_w. By Smith normal form it is equivalent to full column rank modulo π, in both components when split. Host and lattice conventions

Write $T = T_u$, $T' = T_v$ and $b_0 = r + 1$. Let $D_T = -T e_{b_0} e_{b_0}^*$, with the analogous $D_{T'}$. With Z the rational cross Gram matrix, the rescaled local cross matrix is $C = \Phi_w^* Z_w \Phi_w$, where Φ_w is diagonal with i-th entry $(\delta_w u_P(w))^{-(b_0 - i)}$. Here δ_w is the tied local scalar and $u_P(w) = g_{\mathrm{part}}(w)/(g_1 g_2)$, with $g_{\mathrm{part}}(w) = g_j$ over g_j and 1 elsewhere. Denote the two rescaled frames by x, x'. Off the comparison primes the test uses (x, x'), with Gram blocks $D_T, C, C^*, D_{T'}$. Over either comparison prime g_j the model is split and $\pi = g_j = q$. The norm conditions give $T, T' \in O^\times$ and $T \equiv T' \pmod{\pi^2}$. The frame is $(x, (x' - x)/\pi)$, and its Gram, denoted S_w, is

$$S_w = \begin{pmatrix} D_T & (C - D_T)/\pi \\ (C^* - D_T)/\pi & (D_T + D_{T'} - C - C^*)/\pi^2 \end{pmatrix}. \tag{7g}$$

The same formula follows by conjugating the original block Gram by $B = \left(\begin{smallmatrix} I & -\pi^{-1} I \\ 0 & \pi^{-1} I \end{smallmatrix}\right)$ on the right and B^* on the left. At an occupied point the resulting frame spans the prescribed saturated intersection in a representing row lattice, so its Gram is integral. A nonintegral Gram gives an empty lattice tube. Frame and support in the count · Pinned corank formula

Proved here (primitive density). Put $P_\epsilon(j) = \prod_{i=1}^{j} (1 - \epsilon^i q^{-i})$ with $\epsilon = 1$ at split places and -1 at inert places, and $P_\epsilon(0) = 1$. For an integral S_w of reduction corank c_w, the primitive representation count modulo π^j, divided by $q^{j(2Ma - a^2)}$, is

$$\mathrm{densP}_w(Z) = \frac{P_\epsilon(M)}{P_\epsilon(k - 1 - c_w)} \qquad (j \geq 1). \tag{7h}$$

In the selected range $k \geq a + f + 5$ we have $M \geq a + c_w$. To lift a solution from depth j to $j + 1$, the differential is $\overline{X}^* \overline{J}\, \overline{Y} + \overline{Y}^* \overline{J}\, \overline{X}$. Full column rank and self-duality let $\overline{X}^* \overline{J}\, \overline{Y}$ run over all matrices. The map $U \mapsto U + U^*$ onto Hermitian matrices is surjective by choosing the entries above the diagonal and using the residue trace on the diagonal. This works also at residue characteristic two. Each solution has precisely $q^{2Ma - a^2}$ lifts, including when $c_w > 0$.

At a split place write the first Gram component as $\overline{X}_-^{\mathsf{t}} \overline{X}_+$ in dual coordinates. There are $q^{Ma} P_1(M)/P_1(M - a)$ choices for the full-rank $\overline{X}_+$. Extend its columns to a basis. The equation determines the first a columns of the matrix of $\overline{X}_-^{\mathsf{t}}$ in that basis. On the remaining columns only the c_w rows belonging to the left kernel of the prescribed first Gram component must have full rank c_w. This gives $q^{a(M-a)} P_1(M - a)/P_1(M - a - c_w)$ choices. At an inert place split the reduced form into a nondegenerate part of rank $a - c_w$ and its radical. After embedding the nondegenerate part, embed the radical as an ordered isotropic c_w-frame in the complement. Complete this frame to hyperbolic pairs using trace surjectivity. Finite Hermitian spaces of a given dimension are isometric by norm surjectivity, so the unitary group acts transitively on these embeddings. Put $d_0 = M - a - c_w$. The pointwise stabilizer has order $|U_{d_0}(q)| q^{2c_w d_0 + c_w^2}$. In the basis of the isotropic frame, its orthogonal quotient, and the dual frame, the unipotent kernel has one free $d_0 \times c_w$ block D and a $c_w \times c_w$ block L satisfying $L + L^* = -D^* J_0 D$, with J_0 the quotient Gram. This gives the stated order. Dividing $|U_M(q)|$ by it and using $|U_j(q)| = q^{j^2} P_{-1}(j)$ gives the mod-π count in (7h). Finite field count and stabilizer computation

Give Λ_w^a and $\mathrm{Herm}_a(A)$ additive volume one. Then (7h) is the fibre integral of the primitive indicator, since each source residue of depth j has volume q^{-2jMa} and the target tube has volume q^{-ja^2}. For self-dual Haar in the pairings of (7f), a character of conductor ideal $\pi^{t_w}O$ multiplies this integral by $q^{t_w(2Ma-a^2)/2}$. In these bases the trace pairings are perfect at conductor O, and the annihilators for conductor $\pi^{t_w}O$ are π^{t_w} times the respective lattices. For the trace character from $\mathbb{Q}_p$, $\pi^{t_w}O$ is the inverse different of $F_w/\mathbb{Q}_p$. The divided frame Jacobians in source and target are $|N \det B|_w^M$ and $|N \det B|_w^a$, with $|N \det B|_w = q^a$. The fibre integral in the rescaled undivided coordinates therefore has the factor $q^{-a(M-a)}$ against the one in the divided coordinates. The matrix packet has its own additive Haar giving volume one to the order $R_w = \{(z_+, z_-) \in O^2 : z_+ \equiv z_- \pmod{\pi}\}$ at the comparison primes and to A elsewhere. This assigns O^2 volume q at comparison primes. The group level volumes enter the tied mass calculation in (8). Haar and packet normalizations

An independent calculation by halasz_resonator gives the same corank factor and enumerates all three possible coranks for $M = 4,\ a = 2$ over residue fields of sizes two and three. The finite examples also exhibit the extra singular solutions that result from including nonprimitive frames. Independent proof and finite-field counts

Proved here (reference ratio). For $d_T = \mathbf{1}_{T \in \pi O}$, one word contributes $P_\epsilon(M)/P_\epsilon(k - d_T)$ to the reference product. The factor alpD_w uses both words off the comparison primes and a single factor at those primes, where the norms are units. Thus at occupied points, with $P = P_\epsilon$,

$$R_w(Z) := \frac{\mathrm{densP}_w(Z)}{\mathrm{alpD}_w} = \begin{cases} \dfrac{P(k - d_T)P(k - d_{T'})}{P(M)P(k - 1 - c_w)}, & w \text{ outside the comparison primes}, \\ \displaystyle\prod_{i=k-c_w}^{k} (1 - q^{-i}), & w \text{ at a comparison prime}. \end{cases} \tag{7i}$$

In particular the typed ratio is at most one. Put $n = k - a$. A product over any interval starting at $i_0 \geq 1$ has logarithm between $-4q^{-i_0}$ and $2q^{-i_0}$, by the geometric sum and $-\log(1 - x) \leq 2x$ for $0 \leq x \leq 1/2$. The numerator starts at $k - c_w \geq n$, and the reference factors start at k or $k + 1$. Since $a \geq 2$, this gives

$$-5q^{-n} \leq \log R_w(Z) \leq 4q^{-n}, \qquad \sum_w q_w^{-n} \leq f \sum_{j \geq 2} j^{-n} \leq 3f\, 2^{-n}.$$

The logarithms of the reference product and of the occupied density products converge absolutely. The log error in the ratio product after keeping the places over rational primes $\ell \leq Y$, for integral $Y \geq 2$, is at most $5fY^{1-n}/(n-1)$. This bound is uniform in the words and occupied Z, including when $c_w > 0$. Write $s_{\mathrm{den}} = \sup_{Z \text{ occupied}} \prod_w R_w(Z)$, with value zero if there are no occupied points. Then $e_1 = \max(1, s_{\mathrm{den}}) \leq \exp(12f\, 2^{-n})$. Reference factors · Finite density and uniform error

Generic scalar corank-one pairs

Proved here (signed finite comparison from the fixed-data counts). Use the scalar-row coefficients $c_\rho(p)$ and tensor orbital vectors $V(p)$ on the tower of block parameters. At fixed words, the local packet label comes from the integral cross matrix (in the comparison order at a typed place), with the prescribed scalar shift in the marked corner. The following lifting lemma gives integral left and right equivalences on a common occupied label. They induce an integral invertible basis change on (7g) because the two

components of each equivalence agree modulo π. Thus the corank and (7h) are constant on each occupied label. The local coefficients BBtop_{w} include the prescribed unipotent characters, the phase at a typed final block of length two, and the class isomorphism weights. Marked-corner lift and class congruence · Explicit split basis change and local phases

Lemma (integral equivalence preserving the last row and column). Let R be a commutative semilocal ring embedded in E, let $V = R^{l}$, and write e for its last basis vector. For a matrix Z over R invertible over E, put $C_{Z} = V/ZV$ and $\mu_{Z}(x) = e^{\mathsf{t}}Z^{-1}x$. Its marked data are $[e] \in C_{Z}$ and the functional $[x] \mapsto \mu_{Z}(x) \bmod R$. Suppose $\phi : C_{Z} \to C_{Z'}$ is an isomorphism of these data and $\mu_{Z}(e) = \mu_{Z'}(e)$ in E. Then there are $L, P \in \mathrm{GL}_{l}(R)$ with $Z' = LZP$, $Le = e$, and $e^{\mathsf{t}}P = e^{\mathsf{t}}$. The map induced by L is ϕ.

Proof. Let e_{j} run through the basis. For $j < l$ lift $\phi([e_{j}])$ to $v_{j} \in V$. The difference $t_{j} = \mu_{Z'}(v_{j}) - \mu_{Z}(e_{j})$ belongs to R by preservation of the functional. Subtract $t_{j}Z'e$ from v_{j} to make these values agree in E. Take $v_{l} = e$, using the hypothesis on the inverse corner. The matrix L_{0} with the corrected columns induces ϕ, fixes e, and satisfies $\mu_{Z'}L_{0} = \mu_{Z}$.

Put $K = Z'\ker(e^{\mathsf{t}})$. Surjectivity on cokernels writes every $v \in V$ as $L_{0}x + Z'y$. Also $L_{0}Ze = Z'y_{0}$ for some $y_{0} \in V$, with $e^{\mathsf{t}}y_{0} = 1$ by the functional identity. Subtracting $(e^{\mathsf{t}}y)y_{0}$ from y and adding $(e^{\mathsf{t}}y)Ze$ to x proves $L_{0}V + K = V$.

For a maximal ideal $\mathfrak{m}$, reduce modulo $\mathfrak{m}$ and then modulo the line through e. In the resulting vector space W let S be the image of K and B the endomorphism induced by L_{0}. The composite $W \to W/S$ induced by B is onto. Its kernel has dimension $\dim S$. On that kernel change B to any isomorphism onto S, leaving B as given on a complement. This makes B invertible by a correction with image in S. Lift the corrections of the first $l - 1$ columns to K and keep the last column e. The Chinese remainder theorem combines the choices for the finitely many maximal ideals using scalar coefficients in R. The corrected L has unit determinant since it is invertible in each residue field. Its columns induce the same ϕ and the same $\mu_{Z'}L = \mu_{Z}$, since $K \subset Z'V \cap \ker \mu_{Z'}$. The isomorphism on cokernels gives $L^{-1}(Z'V) = ZV$. Hence $Q = Z'^{-1}LZ$ carries V bijectively onto itself, and $e^{\mathsf{t}}Q = \mu_{Z'}LZ = e^{\mathsf{t}}$. Taking $P = Q^{-1}$ proves the lemma.

At a comparison prime the ring $R = \{(a, b) \in O^{2} : a \equiv b \bmod \pi\}$ is local. A pair is a unit precisely when its common residue is nonzero. At the other places the maximal order of the quadratic étale algebra is a DVR or a product of two DVRs. For matrices Y, Y' of the same corner type, their last principal cofactors $d(Y), d(Y')$ satisfy $d(Y)\det Y' = d(Y')\det Y$. On the rational block this follows from the common ratio of determinant to cofactor when the cofactor is nonzero. In the other type both cofactors vanish. The identity persists under common diagonal frame scalings, transposition and a common corner shift, by $\det(Y + cee^{\mathsf{t}}) = \det Y + c\,d(Y)$. These are the coordinate changes used at the two split components. The two shifted determinants are invertible over E on an occupied label. Dividing the identity gives equality of their inverse corners. The lemma now applies to their common marked class, and $Lee^{\mathsf{t}}P = ee^{\mathsf{t}}$ removes the shift to give $Y' = LYP$. Semilocal lifting · Semilocality of the two orders · The two marked maps · The inverse corner · The rational corner type

The selected data. Fix the point of the family used in the relative-trace identity and a pair of distinct words u, u', with last norms T, T'. The first framed flag and the finite row representatives come from that selection. Use their signature and row-representative conditions, the word-pair prime conditions and

$T \equiv T' \pmod{g^2}$. The local scale δ_v, model M_v and uniformizer at each finite place, the constant in the weighted representation count, and the real orbital representatives and quotient measures are those of this same pair. In the selected trace call the local ties and the distinct-word bound are evaluated at the same selection. The recursive theorem instantiates the arithmetic kernel and the signed row bound with the same finite models and the same word norms. The real orbital vector is evaluated at the same signed rational point, with real scale one. Selected trace call · Selected distinct-word bound · Joint specialization

Fixed-data partner and fibre identities. Fix the two words, their norms T, T', the first framed flag F_0, the row representatives $\mathcal{R}$, and the tied finite models and uniformizers. Use the selected signature, the norm congruence modulo g_S^2, the word-pair prime hypotheses and the row representative property. Let $\mathcal{T}(p)$ be the eligible terms at p, meaning the generic terms of nonzero term Gram with that point. A term t contains its row Λ_t and a double-coset class of decorated second flags. Its mass is $M_t = w_{\mathrm{row}}(\Lambda_t)/\operatorname{massRows}$, where $\operatorname{massRows} = \sum_{\Lambda \in \mathcal{R}} w_{\mathrm{row}}(\Lambda)$. Its complex coefficient $z_t(p)$ is the normalized representative's character times the prescribed real relative stabilizer index. Write $t \sim_{\mathrm{sc}} t'$ for a scalar pair on a common row. Exact quantifiers · Partition by terms · Transport of local lattices · Framed isomorphism comparison

With $D_i(p)$ the leading minors of the point matrix, use $G(p) = |D_{r+1}(p)|_E^2 - TT'|D_r(p)|_E^2$, where $|x|_E^2 = x\bar{x}$. The kind $\kappa(p)$ is the last minor ratio D_{r+1}/D_r for a final singleton, and is absent for any longer final antidiagonal block, when $D_r = 0$. The marked corner shift is zero in the latter case. Independently, the local phase at a comparison place has the additional factor $\psi_{R_v}\left(\tau_{0,v}\,\mathrm{blkMod}_v(r, 2, t_{\mathrm{last}})^{-1}\right)$ precisely for a final block of length two with entry t_{last}, with the inverse convention of the local model. The factor is one in the other cases. Here blkMod_v is the rescaled local block parameter and $\tau_{0,v}$ the prescribed corner parameter for the word norms. Kind and final block phase · The fibre cases · PartnerSumC · FiberCntC · Fixed-data assembly · Selected representation formula · Selected use · Norm on packet support

For $t \in \mathcal{T}(p)$, let Z_t be its normalized cross matrix, $x(t)$ the restricted tuple of integral two-sided unitriangular orbits of its local signed packet matrices, and θ_t their shifted marked classes. An admissible orbit has an integral representative in the local rational cell of p. A tuple is restricted when it equals the normal-form tuple outside a finite set. The reference term of an occupied label supplies a defined class at each place. Use a_λ for the product of the framed class automorphism orders for $\lambda = (\kappa, \theta)$ and b_λ for the product of the local BBtop_v on that kind, extended by zero on the other kinds. Define

$$n_0(p) = c\left(\sqrt{|N_{E/\mathbf{Q}}(G(p)/g_S^{2r+2})|}\right)^{k-1}, \qquad d_\lambda = \prod_v \mathrm{densP}_v(Z_{t_0}),$$

where t_0 is an occupied representative of λ and c is the constant in the normalized representation count. The semilocal lemma makes d_λ independent of this choice. The partner identity, called $\mathrm{PartnerSumC}$, says

$$\sum_{\substack{t' \in \mathcal{T}(p') \\ t \sim_{\mathrm{sc}} t'}} z_{t'}(p') = b_{(\kappa(p),\theta_t)}(p') \qquad (t \in \mathcal{T}(p),\ G(p') \neq 0).$$

The fibre identity, called $\mathrm{FiberCntC}$, says that $\sum_{t \in \mathcal{T}(p):\,(\kappa(p),\theta_t)=\lambda} M_t z_t(p) = n_0(p) d_\lambda a_\lambda^{-1} b_\lambda(p)$ for $G(p) \neq 0$ and an occupied λ. When the kinds differ the fibre is empty, and on the partner side scalar equivalence forces equality of kinds. Exact quantifiers · Partition by terms · Transport of local lattices ·

Framed isomorphism comparison · Kind and final block phase · The fibre cases · PartnerSumC · FiberCntC · Fixed-data assembly · Selected representation formula · Selected use · Norm on packet support

The partner count before phases. For an admissible restricted tuple choose a rational Z in the cell of the target point. This uses entrywise additive strong approximation on the lower and upper unitriangular factors. At a comparison place one prescribes both split components to enough precision that the two resulting integral corrections agree modulo the comparison prime. Outside the finite exception set the normalizations are units and preserve the base orbit. To realize a fixed Z inside a term with that tuple, start with rational triangular factors carrying its matrix to Z. At each place there are also integral factors achieving that change. Their discrepancy is in the matrix stabilizer. In the monomial cell its coordinates are the strictly triangular entries allowed by the permutation pattern. Approximate these entries in that stabilizer, so the corrected rational factors still give Z and now satisfy every integral frame condition. Integral Levi and Eichler moves then give a framed representative in the same term. Conjugating the nilradical stabilizers transports the orbit counts. Ordinary approximation · Comparison-order approximation · Rational stabilizer corrections · Realization in the term · Orbit transport

Write b_t^{orb} for the number of nilradical orbits of admissible exact-matrix frames belonging to t in its row. For a fixed source term and target tuple, partition the target partner frames on the common row by their term. Each piece has size $b_{t'}^{\mathrm{orb}}$ by the preceding transport. The partner lattice construction chooses a rational target frame of Gram Z with specified maps on the two orthogonal quotients and the marked last vectors. Pulling a partner back along a unipotent isometry fixing F_0 gives a family of local lattices carrying that target and respecting those maps. Conversely, glue the local lattices by intersection, using the fixed model off the exceptional set. The backward isometry and integral correction recover a partner, and equality of transported lattices identifies precisely the nilradical orbits. The correction uses self-duality and equality on the orthogonal quotients, with the comparison-prime marking included. This is the fixed-data bijection needed before taking the product of local cardinalities. Exact quantifiers · Partition by terms · Transport of local lattices · Framed isomorphism comparison

The local lattice identification counts framed cokernel isomorphisms. For zero leading cofactor it uses the full cokernel of the integral matrix, with the last vector and the last coordinate of its inverse modulo the order. For a nonzero cofactor it first uses the head block H, with mark $[b]$ from the last column and functional $\alpha([y]) = aH^{-1}y \bmod R_v$ from the last row a. The shifted matrix has common invertible Schur complement s at equal kind. Its full marked cokernel is the quotient of $\{(z,u) \in \operatorname{coker} H \times E_v : \alpha(z) + u = 0 \bmod R_v\}$ by $R_v([b], s)$, with mark $[(0,1)]$ and functional $u/s \bmod R_v$. This follows by sending a column v to its head class and s times the last component of its image under the shifted inverse. For the inverse, preservation of the functional gives a lift of the image whose second coordinate is the original one, by subtracting a multiple of $([b], s)$. That lift is unique because s is invertible. Preservation of the marked vector makes the induced first coordinate independent of the lift of the head class. Applying the construction to the inverse gives inverse head maps, and the quotient relation preserves $[b]$. These local counts are zero on inadmissible orbits and between different defined classes. On a matching class the isomorphism set is a torsor under its framed automorphism group. If e_v indicates admissibility and equality of the shifted class to θ_v, the local count is $k_v = n\mathrm{Aut}_v(\theta_v)e_v$. The unsigned partner multiplicity at the tuple is $\prod_v k_v$. A missing rational lift forces an inadmissible local orbit and an empty partner set. Exact quantifiers · Partition by terms · Transport of local lattices · Framed isomorphism comparison · Ordinary approximation · Comparison-order approximation · Rational stabilizer corrections · Realization in the term · Orbit transport

Rational realization. Keep the first frame $(e_1, \dots, e_r, z)$ with $\langle z, z\rangle = -T$ and use the convention that the pairing is linear in its second argument. Fix a cross matrix Z with totally positive $G_Z = |\det Z|_E^2 - TT'|D_r(Z)|_E^2$. In particular Z is invertible. Choose isotropic hyperbolic partners y_i in $z^\perp$ with $\langle e_i, y_j\rangle = \delta_{ij}$. The orthogonal complement of their hyperbolic planes and z is totally positive of rank k, by the selected real signatures. Put

$$B = Z^{-*}\operatorname{diag}(0, \dots, 0, -T')Z^{-1}, \qquad B_{r+1,r+1} + T^{-1} = \frac{G_Z}{T|\det Z|_E^2}.$$

The second equality follows from the bottom-right cofactor of Z. To represent this totally positive value t in the definite complement, diagonalize a binary Hermitian subspace as $\langle a, b\rangle$ and write $E = F(\sqrt{\delta})$ with δ totally negative. The quadratic form $\langle a, -a\delta, b, -b\delta, -t\rangle$ is indefinite at every real place and isotropic at every finite place by the local isotropy theorem in dimension five. Hasse–Minkowski gives an isotropic vector over F. Its last coordinate is nonzero by positivity of the first four terms. Dividing by that coordinate gives a vector w of norm t. Hyperbolic dual columns · Rational realization · Binary Hermitian representation · Serre

Start from the dual columns $(y_1, \dots, y_r, -z/T)$. Add $\sum_i B_{ij}e_i/2$ to column $j \le r$ and $\sum_i B_{i,r+1}e_i + w$ to the last column. These columns still pair as the identity with the first frame, and their Gram is B. Multiplying the column matrix on the right by Z gives the second frame with cross matrix Z and Gram $\operatorname{diag}(0, \dots, 0, -T')$. Its last vector supplies the mark. The two frames together span the hyperbolic planes and the lines of z, w, so the span is nondegenerate and its orthogonal complement is totally positive of rank $n = k - 1$. Hyperbolic dual columns · Rational realization · Binary Hermitian representation · Serre

The normalized representation count. Let U be the span of the two ordered frames at a rational realization of Z, with the second mark equal to its last frame vector. At a comparison prime the divided lattice M_Z in U is generated by the first frame and g_S^{-1} times the difference of the two frames. At the other places use the original frames. Fixing both framed flags is equivalent to fixing U pointwise. The cross matrix specifies the fibre of framed flags, which are transported to one rational realization by Witt extension fixing the first frame. The intersection criterion identifies such rows carrying the admissible generic pair with the unimodular glue lattices satisfying $\Lambda \cap U = M_Z$. The finite row quotient and orbit weights · The pointwise fixer · The glued-lattice comparison · Euler factors · The normalized count

For a row Λ, write H_Λ for the lattice isometries fixing the first frame, and N_Λ for its nilradical subgroup. The finite quotient H_Λ/N_Λ is the induced automorphism group on the first flag's orthogonal quotient, and $w_{\mathrm{row}}(\Lambda) = |H_\Lambda/N_\Lambda|^{-1}$. The nilradical acts freely on the exact-matrix frames in the totally positive fibre. A class in H_Λ/N_Λ stabilizing a nilradical orbit has a unique lift fixing its chosen frame, by composing a lift with the unique nilradical element that returns the frame. Orbit–stabilizer in the finite quotient therefore gives one reciprocal pair-stabilizer order per H_Λ-orbit. Transport identifies these orbits across the row representatives with glue classes under the rational pointwise fixer of U, acting on $U^\perp$. Consequently, for the weighted mean $\nu(Z)$ of the exact-matrix nilradical orbit counts,

$$\nu(Z)\,\mathrm{massRows} = \sum_{\Lambda \in \mathcal{G}(M_Z)} |\operatorname{stabFix}(U, \Lambda)|^{-1},$$

where $\mathcal{G}(M_Z)$ consists of representatives for that action. The finite row quotient and orbit weights · The pointwise fixer · The glued-lattice comparison · Euler factors · The normalized count

Discriminant graphs. Assume first that M_Z is integral. At a finite place write $M_v = (M_Z)_v$ and take the local maximal quadratic order. For an integral Hermitian lattice A write $\mathcal{D}(A) = A^{\#}/A$ for the discriminant module, with its induced pairing modulo the order. For a lattice $L \subset W = U^{\perp}$ the self-dual extensions with intersections M_v, L_v are inverse images in $M_v^{\#} \oplus L_v^{\#}$ of subgroups of $\mathcal{D}(M_v) \oplus \mathcal{D}(L_v)$. The two projections on the subgroup are injective by the intersection conditions. Self-duality makes the square of its order the product of the two discriminant orders, so both projections are bijective. Integrality makes the subgroup the graph of an anti-isometry, and the graph of any anti-isometry gives a self-dual extension with the required intersections. The anti-isometries are a torsor under the isometry group of $\mathcal{D}(L_v)$. Discriminant module · Local graph count · Patching · Complement orbit sum

Let d_v be the valuation of the Gram determinant of M_v and c_v the number of its positive Jordan exponents. At an inert place, negating those positive-exponent Jordan lines and adding unit lines makes a rank-n complement lattice with the required negative discriminant pairing. Its determinant valuation has parity d_v. This is the determinant class of W_v because the ambient space has a self-dual lattice, and rank and determinant classify the unramified local Hermitian spaces. At a split place use the two dual components, with the rational space determined by its rank. The condition $c_v < n$ holds in the selected rank range. Almost everywhere the fixed ambient self-dual lattice contains M_v as a unimodular summand. Use that lattice there and patch by intersection with the rational space, recovering each prescribed completion. This gives a global glue. The discriminant module determines the positive Jordan exponents of any complementary lattice, and its rank determines the zero exponents. Unramified local classification therefore puts all complement lattices in one genus. For each representative L of that genus the local graphs patch independently, and summing inverse stabilizers over the action of $\mathrm{Aut}\, L$ gives the number of graph choices divided by $|\mathrm{Aut}\, L|$. The genus has finitely many classes since W is definite. Discriminant module · Local graph count · Patching · Complement orbit sum

The two local counts. Omit the place subscript, and write $P(j) = \prod_{i=1}^{j}(1 - \eta^i q^{-i})$ with $\eta = 1$ at split and -1 at inert places. For L_v with Jordan exponents a_i, let m_a be their multiplicities, $r_j = \#\{i : a_i \geq j\}$, and $b_L = \sum_{j \geq 1} r_j^2$. Let $\alpha(L_v)$ be the density of Gram-preserving endomorphisms modulo ϖ^N, counted with factor q^{-Nn^2}, and let $u(\mathcal{D})$ be the isometry order of the discriminant module. Then

$$\alpha(L_v) = q^{b_L} \prod_{a \geq 0} P(m_a), \qquad u(\mathcal{D}(L_v)) = q^{b_L} \prod_{a > 0} P(m_a).$$

For the first formula let $A(j)$ be the invertible matrix fraction over the quadratic residue algebra. It is $\prod_{i=1}^{j}(1 - q^{-2i})$ in the inert case and $\prod_{i=1}^{j}(1 - q^{-i})^2$ in the split case. A Hermitian residue matrix of size l has nullity h with probability $q^{-h^2} A(l)/(A(h)P(l-h))$. In the inert case count the radical and the nondegenerate form on the quotient using $\mathrm{GL}_{l-h}(q^2)$ and $\mathrm{U}_{l-h}(q)$. In the split case count the kernel, image and induced isomorphism of one ordinary matrix. After splitting the invertible block, its divided Schur complement is uniform at the next lower precision. Iterating from $r_0 = n$ gives the probability $q^{-b_L} A(n)/\prod_a P(m_a)$ for the prescribed exponents at $N > \max a_i$. They define one congruence orbit by unimodular diagonalization and surjectivity of unit norms in the inert case, and Smith form in the split case. At this depth the radical K of the truncated pairing lies in $\varpi L_v/\varpi^N L_v$. For a Gram-preserving endomorphism the inverse image of K lies in K by testing pairings, and has size at least $|K|$ by mapping to the quotient. Hence the endomorphism is invertible on the quotient and then on the truncated lattice by

Nakayama. Orbit–stabilizer, with $q^{2Nn^2}A(n)$ invertible matrices, gives $\alpha(L_v)$. Congruence normalization · Local density and discriminant comparison · Local cancellation · Gan and Yu

For the discriminant module at an inert place, the endomorphism count is q^{2b_L} since the homomorphisms between cyclic summands of lengths a, c have size $q^{2\min(a,c)}$. Invertibility is the invertibility of each equal-length block on reduction, by block triangularity modulo ϖ and Nakayama. This gives $q^{2b_L}\prod_{a>0}A(m_a)$ automorphisms. Counting independent Hermitian entries gives q^{b_L} pairings, and perfection requires nonsingularity in each rescaled equal-length block. The perfect fraction is $\prod_{a>0}A(m_a)/P(m_a)$. In fact the fixed diagonal duality identifies perfection with the endomorphism test just used. All perfect pairings are isometric by splitting unit leading norm vectors at maximal length and using unit norms to normalize their cyclic forms. The residue unit norm and trace maps give the norm surjectivity also at dyadic inert places. The ratio of the two counts gives $u(\mathcal{D}(L_v))$. At a split place the pairing identifies one component with the dual of the other. Its isometry order is the automorphism count of that one component, which gives the same formula. Thus the two counts have ratio $P(m_0)$. Congruence normalization · Local density and discriminant comparison · Local cancellation · Gan and Yu

Relative Tamagawa volume. Let $C_{E/F,n}$ be the inverse-automorphism mass of the genus of $\mathcal{O}_E^n$ with the identity Gram. For Gram H in a rational basis of W, use the invariant gauge $(\det H)^n dX/dY$ on the identity fibre of $X \mapsto X^*HX$. Here dX on $M_n(E)$ and dY on the Hermitian matrices are fixed rational additive top forms. The map is a submersion at invertible X in characteristic zero, and left translation by a unitary matrix preserves the gauge. In a lattice basis the additive volume of the Gram congruence tube divided by the volume of the target ball is the normalized congruence count, in coordinates with integral lattices of additive volume one. Compactness places all sufficiently small tubes near the integral unitary group. Submersion coordinates there identify the limit with the relative-gauge volume, with the fixed additive constants. Concretely, a change to the basis B_0 conjugates the matrix variable with determinant one and sends the new target matrix to $B_0^{-*}YB_0^{-1}$ in old coordinates. This target change has determinant $N_{E/F}(\det B_0)^{-n}$ on Hermitian matrices, by the diagonal and paired off-diagonal coordinates for diagonal B_0 and triangularity for the elementary unipotents. The gauge therefore uses the factor $|\det(B_0^*HB_0)|_v^n$ in the new basis. The stabilizer volume at v, relative to the identity lattice, is $q_v^{-nd_v}\alpha(L_v)/P_v(n)$, since L_v and M_v have the same determinant valuation. At a real place both definite gauges give the same volume by taking $B_0^*HB_0 = I$. Kottwitz · Rank mass constant · Pinned glue mass

Kottwitz's Tamagawa formula gives two for these unitary groups. Their derived groups are simply connected, the invariant dual-center component group has order two, and Landherr's local–global classification makes the Hasse-principle term trivial. The determinant character lattice with its quadratic action gives the same convergence factors for the two groups of rank n. Thus the ratio of Tamagawa covolumes is one, and the fixed additive and real constants cancel. Decomposing the adelic quotient by the open compact lattice stabilizer and the real group gives the genus mass, with one inverse isometry order per class. Its ratio to $C_{E/F,n}$ is $\prod_v q_v^{nd_v}P_v(n)/\alpha(L_v)$.

The rank multiplier. For $n \geq 2$, taking the last column identifies $U(n)/U(n-1)$ for the identity forms with the unit sphere $h(z) = \sum_i z_i\bar{z}_i = 1$. The quotient of rational invariant gauges agrees with the relative gauge dz/dh up to a scalar in $F^\times$, which cancels by the product formula. The convergence factors agree across these two ranks because both groups have the same determinant character lattice. Use trace-self-dual additive measures on E and F, whose adelic products give the additive Tamagawa measures. At a

real place the complex coordinate measure is $2\,dx\,dy$ and the target measure is dt. Differentiating the volume $(2\pi)^n a^n/n!$ of $h(z) \le a$ at $a = 1$ gives sphere volume $(2\pi)^n/(n-1)!$.

At a finite place the quotient of the integral stabilizers is the set of integral norm-one vectors. Each vector splits off a unimodular line with standard unimodular complement. At an inert place the residue count is $|\mathrm{U}_n(q_v)|/|\mathrm{U}_{n-1}(q_v)| = q_v^{2n-1}(1-(-1)^n q_v^{-n})$. At a split place the equation is $x \cdot y = 1$, with $(q_v^n - 1)q_v^{n-1}$ solutions. The norm map is smooth at these vectors, including in residue characteristic two by surjectivity of the residue trace. Lifting and dividing the measure of the norm congruence tube by that of its target ball therefore gives the local quotient volume

$$\frac{\mathrm{vol}(\mathcal{O}_{E_v})^n}{\mathrm{vol}(\mathcal{O}_{F_v})}(1-\eta_v^n q_v^{-n}).$$

The finite products of the additive unit volumes are $|\operatorname{disc} E|^{-1/2}$ and $\Delta_F^{-1/2}$ by trace duality, where $\Delta_F = |\operatorname{disc} F|$. Finite unramifiedness gives $|\operatorname{disc} E| = \Delta_F^2$. Since both Tamagawa numbers are two, the ratio of the standard genus masses is the reciprocal of the product of compact quotient volumes. Hence

$$\frac{C_{E/F,n}}{C_{E/F,n-1}} = \Delta_F^{n-1/2}\left(\frac{(n-1)!}{(2\pi)^n}\right)^f L(n,\eta_{E/F}^n).$$

The Euler product is absolutely convergent for $n \ge 2$. With $v_{\mathrm{inv}} = \Delta_F^{1/f}/2$ this is the multiplier in the pinned source. In rank one the common constant is the genus mass of $\mathcal{O}_E$ with the identity Gram, since its normalized local factors are one. The relative volume argument applies to every positive line, so this anchor and the multiplier identify the constant in every rank. Pinned rank multiplier · Public proof discussion

Multiplying by the product of the graph counts gives

$$\nu(Z)\,\mathrm{massRows} = C_{E/F,n}\prod_{v<\infty} q_v^{nd_v}\frac{P_v(n)}{P_v(n-c_v)}.$$

All but finitely many factors in this expression are one. Discriminant module · Local graph count · Patching · Complement orbit sum · Congruence normalization · Local density and discriminant comparison · Local cancellation · Gan and Yu · Kottwitz · Rank mass constant · Pinned glue mass

The normalized count. Use the pinned divided determinant comparison and the primitive density comparison in the same frame. With $\Delta_F = |\operatorname{disc} F|$ and g_S the comparison level g, these give

$$\prod_v q_v^{d_v} = \Delta_F^{-2r(r+1)}\sqrt{|N_{E/\mathbb{Q}}(G_Z/g_S^{2r+2})|}, \qquad \mathrm{densP}_v(Z) = \frac{P_v(M)}{P_v(n-c_v)}.$$

Here P_v uses $\eta_v = \eta_{E/F}(v)$, and $M = k + 2r + 1$ is the ambient rank. The factors $P_v(n)/P_v(M)$ give the absolutely convergent product $\prod_{i=k}^{M} L(i,\eta_{E/F}^i)$, where $\eta_{E/F}$ is the quadratic character of E/F. With the positive row mass for the selected nonempty row set, the constant c is

$$c_{\mathrm{Expl}} = \frac{C_{E/F,k-1}}{\mathrm{massRows}}\,\Delta_F^{-2r(r+1)(k-1)}\prod_{i=k}^{M} L(i,\eta_{E/F}^i).$$

Consequently $\nu(Z) = c_{\mathrm{Expl}}(\sqrt{|N_{E/\mathbb{Q}}(G_Z/g_S^{2r+2})|})^{k-1}\prod_v \mathrm{densP}_v(Z)$, which is $n_0(p)\prod_v \mathrm{densP}_v(Z)$ on the cell at p. If M_Z is nonintegral, the glue set is empty and the primitive comparison supplies a zero local density. The formula uses the specified row representatives and integral intersection test. Finiteness of

the glue classes and of the frame stabilizers gives finiteness of the exact-matrix orbit counts on this branch. Frame groupoidDivided determinant · Jordan norm · Explicit constant · Primitive factors and Euler product

Detailed relative genus mass proof

halasz_resonator independently identified the divided-frame pointwise stabilizer, its inverse-order weight and the Euler factor extraction. The finite-field calculation in that contribution illustrates the stabilizer weight. The arithmetic groupoid calculation above uses the marked frames and the finite quotient by the nilradical. The framed-cokernel orders a_λ enter through the isomorphism count in the partner argument. Comment on Paper V

Passage to the signed sums. At an exact-matrix tuple x, partitioning all frames in each row by term gives $\sum_{t:\, x(t)=x} M_t b_t^{\mathrm{orb}} = \nu(Z)$. Restricting to λ retains the count only when all $e_v = 1$. The shifted class is constant under integral lower and upper moves, since they fix the marked last column and row and the corner shift. Equal kind transfers total positivity from the occupied reference point: the Gram is a squared minor norm times 1 in the zero-cofactor case, or times $|D_{r+1}/D_r|_E^2 - TT'$ in the singleton case. On the matching tuple, the normalized count and the semilocal density comparison therefore give $n_0(p) d_\lambda$. A wrong class, inadmissible orbit or absent rational lift gives zero. Kind and final block phase · The fibre cases · PartnerSumC · FiberCntC · Fixed-data assembly · Selected representation formula · Selected use · Norm on packet support

The relative-index identity and the rational character product give $z_t(p) = b_t^{\mathrm{orb}} \prod_v h_v(p, x_v(t))$. Here h_v is the phase, including the typed factor above, times the local Haar volume of the triangular cell over the integral orbit, with zero value off admissible orbits. The signed partner sum at a tuple is therefore $\prod_v (h_v k_v)$. The signed weighted fibre sum there is $n_0(p) d_\lambda \prod_v (h_v e_v)$. Off a finite exceptional set these local factors are delta functions at the base orbit, and the class automorphism order there is one. At the remaining places the absolutely convergent orbit sums of $h_v k_v$ are BBtop_v. The torsor identity gives $\sum h_v e_v = \mathrm{BBtop}_v / n\mathrm{Aut}_v(\theta_v)$ for occupied classes. Finite-product Fubini proves the two identities, with $\overline{b_\lambda(p')}$ on the conjugated side of the Gram. This accounts for the single denominator a_λ in the fibre identity. Kind and final block phase · The fibre cases · PartnerSumC · FiberCntC · Fixed-data assembly · Selected representation formula · Selected use · Norm on packet support

The fixed-data argument invokes the rational and integral frame comparisons, the local lattice and framed-cokernel identification and index formula, and the cofinite and absolute orbit-sum laws of the pinned construction. The relative genus mass supplies the normalized representation count on the totally positive branch using the pinned frame-groupoid, determinant and primitive density comparisons. The local support law gives equality of the local norms of D_{r+1} at two points with nonzero coefficients in the same zero-cofactor packet, or of D_r in the same singleton packet. The finite norm product and the common minor ratio in the singleton case give equality of n_0 on the packet support. The kernel application also uses class square summability and convergence of the row and orbital expansions. The later orbital-volume comparison uses its own convergent Siegel–Weil input. Fixed-data assembly · Selected representation formula · Norm on packet support

Detailed fixed-data proof

The signed Gram. For a scalar row ρ, sum the terms of that row at p with coefficient $\sqrt{M_t}\, z_t(p)$. Divide the sum by $\sqrt{\alpha n_0(p)}\, \phi_{r+1}(p)$, where $\alpha = \prod_v \mathrm{alpD}_v$ is the positive reference product and ϕ_{r+1} is the block weight computed below in tower coordinates. In the Gram of these rows, two scalar partners have the

same underlying row lattice, so the two square roots of their row mass multiply to M_t. Sum over the conjugated partner first. Its signed sum is $\overline{b_\lambda(p')}$ for the packet of t. The weighted fibre identity now gives $n_0(p)d_\lambda a_\lambda^{-1}b_\lambda(p)$ on the first side. The factor M_t already includes the division by $\mathrm{massRows}$. The one division by a_λ comes from replacing a framed isomorphism count by the class indicator in the weighted fibre. When both coefficients in a packet are nonzero, $n_0(p) = n_0(p')$ by the minor support law. Thus the two square roots of this norm cancel and the normalized packet has coefficient $(d_\lambda/\alpha)a_\lambda^{-1}$, with the same ϕ_{r+1}^{-1} on each argument of the row and packet kernels. Points of zero Gram have zero signed row coefficient. Equivalently, the reference normalization absorbs α into the count constant and divides d_λ by α once. Actual row coefficient · Partner followed by fibre · Reference normalization

The top kernel $\mathcal{K}_{r+1}$ is the Gram sum of these packet vectors over the full restricted class set with weights A_λ^{-1}. The source supplies local square-summability and a base class with coefficient and automorphism weight one outside a finite set for each point. On a finite point set take the union of those exceptional sets. Cauchy–Schwarz gives absolute convergence of the pairwise class products at those places. Restricting the top sum to occupied labels removes positive-semidefinite terms. By (7i) each occupied multiplier is nonnegative and at most s_{den}, and restricting the row Gram sum to a finite set of rows R also removes positive-semidefinite terms. Thus

$$\big(\textstyle\sum_{\rho\in R} c_\rho(p)\overline{c_\rho(p')}\big)_{p,p'} \preceq s_{\mathrm{den}}\,\mathcal{K}_{r+1}.$$

The phases stay in the packet vectors throughout the comparison. Testing against the coordinates of the finite family $V(p)$ in an orthonormal basis gives the corresponding bound for the sum of vector squares. The signed off-diagonal main-row identification and the passage to the full scalar-row sum use the tied orbital masses and the real estimates below. The equal-word diagonal bound uses the separate $\mathrm{RowMainEqP}$ comparison. Local class summability and support · Selected row ties · Equal-word scalar comparison · Its place in the family assembly

The generic length-one Fourier residual

Fix a level i, with $j = i - 1$, a finite set Y of level-j prefixes, a positive semidefinite kernel H on Y, and a permitted derivative word $w = (w_v)_v$. In this argument v indexes one complex embedding per real place of F. Take $K = 5$ and $L = 4K(r+1)$. The count of a word is the maximum of its local lengths. Write $W_b(a)$ for the fitted finite set of admissible words of count at most a and require $\mathrm{cnt}(w) + K \le L$. The positive weight is $\mathrm{wt}(w) = \prod_v \prod_{(\pm,y)\in w_v} |P_v(y)|$, where $P_v(y) = \prod_{a=1}^{2r+k+1}(y + \lambda_{v,a})$ is the local Capelli eigenvalue polynomial. The weight of an empty word is one. For the coefficient $\chi(G/\varepsilon)F_v$, cut before applying the operators in w_v, let $\mathrm{gram}_{v,\varepsilon,i,w_v}$ be its local residual kernel times the two square roots of the corresponding block densities. In the unscaled coordinate $q = (q_v)_v$ the test is

$$\Psi_\varepsilon(q) = \frac{1}{\mathrm{wt}(w)} \sum_{h,h'\in Y} H(h,h') \prod_v \mathrm{gram}_{v,\varepsilon,i,w_v}(h_v + +[(1,q_v)], h'_v + +[(1,q_v)]).$$

Here $++$ appends a block, and h_v, h'_v are the localized prefixes. The same contraction at level j without an appended block, divided by $\mathrm{wt}(w')$, defines $\mathcal{E}_{\varepsilon,j,w'}(Y,H)$ for the word w'. It is real and nonnegative by the Schur product of the positive semidefinite kernels. The fitted Fourier comparison uses the continuous extension of Ψ_ε from nonzero coordinates. Words and densities · Word budgets

Lattice and sign. With $\nu = 2r + 2 - 2i$ and $\mathfrak{d}$ the absolute different, the period and its trace dual are

$$L_i=\{t\in g\mathfrak{d}^{\nu}: t+\bar t\in g^2\mathfrak{d}^{\nu}\},\qquad L_i^{\vee}=\{\ell\in E:\operatorname{Tr}_{E/\mathbb{Q}}(\ell L_i/g)\subset\mathbb{Z}\}.$$

Set $\iota_i(t)=(\rho^{-\nu}\sigma_v(t))_v$ and $\Lambda_i=\iota_i(L_i)$. For the real Euclidean pairing $\langle c,\xi\rangle=\operatorname{Re}\sum_v c_v\overline{\xi_v}$, the frequency $\xi_+(\ell)_v=2\rho^{\nu}\overline{\sigma_v(\ell)}/g$ satisfies $\langle\iota_i(t),\xi_+(\ell)\rangle=\operatorname{Tr}_{E/\mathbb{Q}}(\ell t/g)$. The factor two counts the conjugate embeddings. Since $g^2\mathfrak{d}^{\nu}\subset L_i$, the lattice has full rank and its integer-valued functionals are exactly these traces by nondegeneracy of the rational trace pairing. In particular ξ_+ identifies $L_i^{\vee}$ with Λ_i^* and

$$L_i^{\vee}\subset\mathfrak{a}_i=g^{-1}\mathfrak{d}^{-(\nu+1)},\qquad\sqrt{N\mathfrak{a}_i}=(g\rho^{\nu+1})^{-f}.$$

For the selected CM extension the relative different is trivial. The relative trace maps the extended different power $\mathfrak{d}^{\nu}$ onto the corresponding real ideal power. This follows locally from residue trace surjectivity in the unramified case and addition in the split case. Dividing the period condition by g therefore gives index g^f in $g\mathfrak{d}^{\nu}$. Ordinary Lebesgue covolume of $\mathcal{O}_E$ is $(\rho/2)^f$, so the unscaled lattice has volume $\gamma_i^0=\rho^{2\nu f}(\rho/2)^f g^{3f}$ and $\gamma_i=\operatorname{covol}\Lambda_i=(\rho/2)^f g^{3f}$. Period lattice · Covolume · Trace dual and Fourier sign

Put $\phi_\varepsilon(c)=\Psi_\varepsilon(\rho^{\nu}c)$ in the chosen embeddings. With the negative-exponent Fourier convention, the release coefficient uses the positive exponential in the trace and is

$$\operatorname{Nhat}_\varepsilon(\ell)=\gamma_i^{-1}\widehat{\phi}_\varepsilon(-\xi_+(\ell)).$$

In unscaled coordinates the positive phase is $\exp(2\pi\mathrm{i}(2/g)\sum_v\operatorname{Re}(\sigma_v(\ell)q_v))$. The Jacobian $\rho^{2\nu f}$ cancels against the covolume change. Replacing ordinary complex Lebesgue measure by $2\,d\Re c_v\,d\Im c_v$ at each place scales the Fourier integral and the covolume equally. In the Poisson series at an arithmetic shift t_0, the phase multiplying $\operatorname{Nhat}_\varepsilon(\ell)$ is $\exp(-2\pi\mathrm{i}\operatorname{Tr}_{E/\mathbb{Q}}(\ell t_0/g))$. Period lattice · Covolume · Trace dual and Fourier sign

Fix the selected scalar row, with its coefficient and quotient measures. First take a compactly supported test at fixed smoothing. The local analytic inputs are the compact Jacquet step below the top, the compact top exhaustion, and the regularity and adjoint identities for the fitted operators. Selected carrier

The selected coefficient and top Gram. At real scale one, the compact-model construction takes the same two word types, local frames and real representation as the scalar row. At a place v, write σ for the selected discrete series, Q for the completed frame transport including the real torus translation, and a_0,b_0 for the prescribed K-finite vectors. With the transported compact vectors ξ,ξ' the model test is

$$F_v(m)=c_v\langle a_0,\sigma(QmQ^{-1})b_0\rangle[x\mapsto\langle\xi',x\rangle\,\xi].$$

The ambient and stabilizer Haars are pushed forward by the same frame transport. On a singular representative the intersection stabilizer Haar is normalized by its free complex root coordinates, with density $(\mathrm{ka}/(g\pi))d^2z$ per root in the source's notation. Here $d^2z=d\Re z\,d\Im z$ and ka is the selected real character scale. In the regular cell the quotient uses the original ambient Haar. Weil disintegration then gives the transported quotient measure, including on the singular representatives. Compact carrier · Rank-one coefficient · Quotient root measure · Actual ambient and stabilizer measures

To determine the complex scalar c_v, start with the bare coefficient. Use the rank-one Bessel maps ℓ,ℓ' and the selected proportional maps $c_0\ell,c_0'\ell'$. The Schur reproducing integral on the ambient group gives the scalar $\kappa_{\mathrm{orb}}=d_\sigma s_0/K_1$ relating the bare orbital operator to the selected wave amplitude. Here

$s_0 = \langle \ell(a_0), \xi \rangle \langle \xi', \ell'(b_0) \rangle$, $K_1 = c_N^{-1}(c_0\|v_0\|)\overline{(c_0'\|v_0\|)}$, and c_N and v_0 are the squared quotient norm per compact dimension and reference vector, respectively in the wave comparison. Indeed unfolding the Schur integral against the two rank-one maps on the chosen quotient gives their evaluations s_0 and converts their amplitudes by K_1^{-1}. The formal degree is for the original ambient Haar. This step uses Bessel nonvanishing, gapped uniqueness and the integrable cell majorant. Write c_{O} and scl for the selected orbit and wave comparison constants, and t_N for the positive ratio of model to transported unipotent Haar in Lie coordinates. The source chooses $c_v = (c_{\mathrm{O}}^{-1}\mathrm{scl}\,\kappa_{\mathrm{orb}}\,t_N^2)^{-1}$. There are two unipotent integrations, hence the square. The resulting conversion is one as an operator identity, so it preserves the phase of a mixed Gram entry. Schur and Bessel calculation · Scalar pin with actual Haar · Two unipotent factors · Gapped Bessel uniqueness · Bessel nonvanishing

The signed cross matrix specifies the point representative in the completed frame. Factoring its quotient integral on pure tensors gives the tensor product of the calibrated local orbital operators, each divided by the positive block factor $b(p_v)$. This factorization includes the character values for changes of representative. Write $L_v(p_v)$ for the resulting local operator and $D_v(p_v)$ for the local top block weight. On a point of positive Gram with the required representative the finite row vector $V(p)$ is the Hilbert–Schmidt embedding of the unscaled operator times the square root of the global block weight. The weight is $\prod_v D_v(p_v)$, so

$$\mathrm{tr}\big(V(p')^* V(p)\big) = \prod_v \sqrt{D_v(p_v) D_v(p_v')}\,\mathrm{tr}\big(L_v(p_v')^* L_v(p_v)\big).$$

Here we identify a Hilbert–Schmidt vector with its operator in writing the trace. The full orbital block carries the external multiplier $\mathrm{scl}_{\mathrm{row}}\,\mathrm{vol}(K_2)$, with $\mathrm{scl}_{\mathrm{row}}$ the original period normalization. That multiplier enters the outer row and mass comparison. A point outside the positive local point domain, or without the required representative, gives the zero branch of the top vector. Unscaled operator and external volume · Pure tensors · Signed cross matrix and top weight · Mixed Gram and zero branches

For a final singleton the local Gram factor is $|\zeta|^2 - TT'$, obtained by dividing the full Gram by the squared preceding minor. The cutoff uses the full model Gram $G_F = |D_{r+1}|^2 - TT'|D_r|^2$. Each leading cross minor is invariant under the stabilizers by unitriangular row and column operations and the fact that the compact factors fix the frames. For the empty word the cutoff therefore factors out of the quotient integral at p as $A_\varepsilon(p) = \prod_v \chi(G_{F,v}(\gamma_{p,v})/\varepsilon)$, where $\gamma_{p,v}$ is the model representative. The mixed smoothed top Gram is $A_\varepsilon(p)A_\varepsilon(p')\,\mathrm{tr}(V(p')^*V(p))$. On a finite arithmetic support this agrees exactly with the original Gram when ε is less than the minimum of the positive model Gram values occurring there. For a nonempty word the top form pairs the differentiated lifts with the fitted type projector, applying the operators after inserting the cutoff. Finite cutoff and tensor pairing · Word-dependent upper form

Compact wave identities below and at the top. Fix $\varepsilon > 0$ and let F be a compact approximation to the smoothed F_v. For $a = w_v$ put $x_h = (x_{h,e})_e$, where e chooses one slot from each letter of a. Apply the positioned word operators in those slots and then take the partial Bessel slice at the standard prefix. The result is a tuple of coefficients on the group of remaining flag length $r - j$. In that model let $J_j A$ denote the two-sided Bessel transform of a coefficient A. Its integrand at m is

$$\psi_z(n)\overline{\psi_{z'}(n')}\,\pi(k_1)A(k_1^{-1}n^{-1}mn'k_2)\pi'(k_2)^{-1}, \qquad \psi_z(n) = \exp(\mathrm{i}\,\mathrm{Re}(-2\mathrm{i}\,\ell_z(n))).$$

The variables k_1, k_2 are in the two compact stabilizers, n, n' are in the two unipotent radicals, and ℓ_z sums the adjacent flag entries and the last vector entry. Let $\beta_j(A, B)$ integrate the Hilbert–Schmidt pairing

$\operatorname{tr}((J_jB)^*J_jA)$ on the model section with its radial density. The module form is

$$\mathcal{B}_{j,a}(x,x') = \sum_e \beta_j((\mathcal{R}_{j,a}x)_e, x'_e),$$

where $\mathcal{R}_{j,a}$ is the fitted module operator for the same positioned word. The compact Jacquet identity intertwines $\mathcal{R}_{j,a}$ with its radial operator and represents this form on the corresponding radial pictures. The adjoint identity on these compact pictures uses integration by parts on the positive domain and the zero extensions. Slices and wave · Bessel integral and beta form · Compact spectral adjoints · Partial Jacquet and head construction · Module and radial form

The radial pairing uses Haar on the opposite-Weyl group normalized to push forward to $Z(\zeta)d^2\zeta$, where $Z(\zeta) = (|\zeta|^2 - TT')^{k-1}\mathbf{1}_{|\zeta|^2>TT'}$. To see the density, put the first negative vector at $\sqrt{T}\,e_{k+1}$. The orbit of the other vector is the hypersurface $|z_-|^2 - \|z_+\|^2 = T'$ in $\mathbb{C}\oplus\mathbb{C}^k$. Slicing ordinary Lebesgue measure at this equation gives an invariant measure because the unitary action preserves the equation and real volume. Polar integration in z_+ yields density proportional to $(\|z_+\|^2)^{k-1}$ in its squared radius. Using $|\zeta|^2 = T|z_-|^2$ gives the asserted shape. One compact group acts transitively on the spheres and the other is the fibre stabilizer. In Haar disintegration the two compact volumes and the disintegration constant multiply to the scalar of this pushforward. Rescale the intermediate Haar by its reciprocal and evaluate on any annulus with $|\zeta|^2 > TT'$ to see that the radial comparison has scalar one. Unitary covariance removes the two compact actions from the Hilbert–Schmidt pairing. The ambient and stabilizer Haars of the original coefficient retain their earlier calibration. Hyperboloid disintegration · Rank-one integral · Double compact disintegration · Scalar one on the annulus

Below the top. For $j < r$, write $R_{j+1,a,F}(t)$ for the residual pairing at the two enlarged standard prefixes, and let $\Phi_h, \Phi_{h'}$ be the radial pictures of the old tuples. The first radial coordinate t is the first leading minor in the remaining flag. In the compact Jacquet step, apply the first block operator to both pictures and transport to the fibre of flag length $r-j-1$. The fibre form contributes $|t|^{-2}$. The slice comparison, which uses the properness and smoothness of the differentiated partial Bessel integral, identifies each transported picture with $|t|^{2(r-j)+k}$ times the actual enlarged-prefix picture. Applying the normalized form identity in that fibre gives its density at $t \neq 0$,

$$\varphi(t) = |t|^{-2}\big(|t|^{2(r-j)+k}\big)^2 R_{j+1,a,F}(t) = |t|^{4(r-j)+2k-2}R_{j+1,a,F}(t).$$

The compact step supplies a continuous integrable extension of this density. It pairs multiplication by $e_s(t) = \exp(st - \bar{s}\,\bar{t})$ on Φ_h with $\Phi_{h'}$, giving $\int_{\mathbb{C}} e_s(t)\varphi(t)\,d^2t$. On coefficient slots this multiplication is U_s, with the inverse triangular change of slots from commuting the word operators past e_s. The signs in that change come from $\partial_t e_s = s e_s$ and $\partial_{\bar{t}} e_s = -\bar{s} e_s$. The normalized form identity at the old level makes the compact integral exactly $\mathcal{B}_{j,a}(U_s x_h, x_{h'})$. Head factors · Below-top pairing · Wave differentiation

At the top. For $j = r$, the remaining flag has length zero and the first minor is the scalar ζ. The compact top exhaustion uses the left and right unipotent averages to realize the differential lift. Its head identity identifies the fibre pairing almost everywhere with $R_{r+1,a,F}(\zeta)$ on $|\zeta|^2 > TT'$, with the same type projection as in the fitted top expression for a nonempty word. The supported radial pictures give zero on the other branch. Multiplying the fibre pairing by $e_s(\zeta)$ and using the normalized Haar pushforward therefore gives

$$\mathcal{B}_{r,a}(U_s x_h, x_{h'}) = \int_{\mathbb{C}} e_s(\zeta)Z(\zeta)R_{r+1,a,F}(\zeta)\,d^2\zeta.$$

Here the old form uses the same normalization as below the top. The boundary circle has measure zero and the earlier minors are already in the prefix and section weights. Left and right averages · Top head from exhaustion · Top pairing · Top lift and bounds

Return to the period measure. At level $i = j + 1$ put

$$\lambda_S = (g/2\pi)^\nu, \quad \kappa = \frac{2\pi}{g\rho}, \quad \delta_G = \frac{2}{\rho g^3}, \quad M_b = k + 2(r + 1 - b) + 1, \quad d = 2(M_i - 2), \qquad \Gamma_j = \prod_{b=j+1}^{r+1} \frac{\kappa^{-4(r+1-b)(M_b-1)}}{\delta_G}.$$

The unscaled coordinate q_v goes to $t = \lambda_S q_v$ in the model. If $D(h)$ is the nonnegative density of the prefix, put $u_h = \sqrt{D(h)}\, x_h / b(h)$, with value zero on an invalid prefix. Appending a singleton leaves $b(h)$ unchanged. In the compact arithmetic Gram the two prefix factors separate as in u_h and $u_{h'}$. The appended density relative to the model weight contributes $\kappa^{\nu d}$ since $\rho^{-\nu}/\lambda_S = \kappa^\nu$. At the top $\nu = 0$ and the appended weight is exactly $Z(t)$. The measure factor in each normalized Fourier integral, before the word weight, is d_i^{-1} with $d_i = \rho^{2\nu}(\rho/2)g^3$, so $d_i^f = \gamma_i^0$. The change of variable and density combine as

$$d_i^{-1} \lambda_S^{-2} \kappa^{\nu d} = \delta_G \kappa^{\nu(d+2)} = \Gamma_i / \Gamma_j.$$

The compact arithmetic Gram has denominator Γ_i, so the local coefficient has denominator Γ_j. Writing $\ell_v = \sigma_v(\ell)$, the choice $s = 2\pi \mathrm{i}\, \ell_v / (g\lambda_S)$ gives $st - \bar{s}\,\bar{t} = 2\pi \mathrm{i}\, 2\,\mathrm{Re}(\ell_v q_v)/g$. Thus for $a_v = \prod_{(\pm, y) \in w_v} |P_v(y)|$ the coefficient is

$$\frac{1}{d_i a_v} \int_{\mathbb{C}} \mathrm{gram}_{v,\varepsilon,i,w_v}(h{+}{+}[(1, q_v)], h'{+}{+}[(1, q_v)])\, e^{2\pi \mathrm{i}\, 2\,\mathrm{Re}(\ell_v q_v)/g}\, d^2 q_v = \frac{\mathcal{B}_{j,w_v}(U_s u_h, u_{h'})}{\Gamma_j a_v}.$$

Here the compact approximation is understood on both sides. The positive trace exponent gives the negative Euclidean frequency described above. The identity wave at $\ell = 0$ gives exactly the preceding normalized Gram. Products of the local factors and contraction against H give the normalized Fourier coefficients, including $\mathrm{Nhat}_\varepsilon(0) = \mathcal{E}_{\varepsilon,j,w}(Y, H)$ after the fixed-smoothing passage below. Density split · Weighted slices · Wave frequency · Normalized integral · Preceding Gram · Product over the places

Fixed-smoothing completion. Let the radial cutoffs be one on radius $R + 1$ and zero from $2(R + 1)$ in Frobenius norm, and leave $\varepsilon > 0$ fixed. The differentiated lower slices converge in the required weighted seminorm with error $O_\varepsilon((R + 1)^{-1})$ after giving up one decay power. The top kernels converge by the bounded-order and growth estimates on the lifted coefficient of each word slot and domination of the quotient integral. The fixed-smoothing estimates on the cutoff derivatives and the two cubic bounds in q_v and ℓ_v give dominated convergence of the Fourier integral and its continuous extension across $q_v = 0$ by Fourier inversion. Top lift convergence · Cubic decay · Fourier inversion

Write p_ε for the weighted derivative seminorm in the bilinear form estimate, with a fixed common lower support bound for the two slice tuples. For the approximating weighted slices $u_{h,n}$, that estimate and the tail bound give

$$\|u_{h,n} - u_{h,m}\|_{\mathcal{B}_{j,a}}^2 \le C_\varepsilon p_\varepsilon(u_{h,n} - u_{h,m})^2 \le C'_\varepsilon \big((n + 1)^{-1} + (m + 1)^{-1}\big)^2.$$

The constants may depend on the fixed prefixes and word. This holds through $j = r$, since those are still slices on the preceding side of the top step. The unweighted bilinear limit identifies the completed inner product with the preceding density-weighted Gram, with the same $\Gamma_j a_v$ denominator. The compact wave is an isometry for the form, so it descends through the null space and extends to the completion. Continuity of

the completed pairing now gives the right side of the displayed Fourier identity for the smoothed coefficient. Lower slice convergence · Bilinear continuity · Wave isometry

The physical cutoff at level zero. At an empty prefix the distinguished word slot is the only nonzero slice, and the Jacobi residue uses the reverse fold of the signed 00 operators in the first beta entry. Let $R_v(\varepsilon, a)$ be this raw pairing for $a = w_v$ on $F_{v,\varepsilon} = \chi(G_{F,v}/\varepsilon)F_v$. Here χ is zero up to $1/2$, one from 1 and takes values in $[0, 1]$. On the positive model section, write $q = (c_0, \dots, c_{r-1}, \zeta)$, with every $c_b \neq 0$ and $|\zeta|^2 > TT'$. Then

$$G_q = \prod_{b<r} |c_b|^2(|\zeta|^2 - TT'), \qquad d\mu_0 = \prod_{b<r} |c_b|^{-2} \prod_{b<r} d^2c_b\, d^2\zeta.$$

In an orthonormal fibre frame the Jacobi vector J has components of the Bessel transform times the square root of the weight $\prod_{b<r} |c_b|^{2(M_{b+1}-1)}(|\zeta|^2 - TT')^{k-1}$. Parseval with $d\mu_0$ gives $J_0 = \int \sum_\pi |J_\pi|^2\, d\mu_0 = \|F_v\|^2_{\mathrm{sec}}$. The invariant $G_{F,v}$ restricts to G_q, so the original cutoff becomes $c_\varepsilon(G_q)J$, with $c_\varepsilon(x) = \chi(x/\varepsilon)$. In the selected chart the transport of the raw pairing has scalar and cutoff scale both one. Empty-prefix residue · Selected coefficient estimate · Model transport · Chart measure and Gram · Fibre Parseval · Interior regularity · Square-integrable generator products · Adjoint identity on the cut domain · Uncut eigenidentity

The model uses the selected gapped discrete series with regular Harish–Chandra parameter in the integral weight convention $\lambda_{v,s} - (2r + k)/2 \in \mathbb{Z}$, the K-finite coefficient vectors, and the transported Haar and compact types. If the opposite-sign gap is at least h, the fitted decay exponent A_0 is at most $h/2$ and the gap condition gives $\lceil A_0 \rceil + t_* < h$ for the regular-model decay threshold t_*, since $t_* \leq 32(2r + k + 1)^3$ and $h \geq 120(2r + k + 1)^3$. The resulting derivative bounds at the fitted orders and the Capelli eigenidentity are used for the uncut coefficient. The positive chart uses the smooth Jacobi gauge, the interior elliptic regularity of the Bessel transform from the left and right Capelli and stabilizer equations, and the Jacobi transport identities. In that gauge the products of the hat and residue generators through $2L + 2$ factors applied to J are square integrable in $d\mu_0$. The local adjoint transfers when one slot carries a smooth multiplier supported away from $G_q = 0$ with bounded logarithmic derivatives. These are the analytic inputs to the boundary argument. Empty-prefix residue · Selected coefficient estimate · Model transport · Chart measure and Gram · Fibre Parseval · Interior regularity · Square-integrable generator products · Adjoint identity on the cut domain · Uncut eigenidentity

Fix a place, keeping P_v for its Capelli polynomial. Choose the interval width H with the fitted cut length at least $6LH$ and $(2r + k + 1)H$. For a word of length $N \leq L$, the fitted cut contains ordered disjoint intervals in its middle third. Each has width H and adjacent intervals are separated by at least $H > 2$. In word order the intervals run from right to left. Choose all-plus arguments in these intervals. In the induction let X be the old uncut vector of energy $A = \int |X|^2_{d_{\mathrm{old}}} d\mu_0$, where $|\cdot|^2_d$ denotes the sum of component squares with weights d. For an argument Y in the next interval let g_Y be the two-branch Jacobi operator with new weights d_Y. The weighted residue identity is $g_Y^*(d_Y g_Y X) = |P_v(Y)| d_{\mathrm{old}} X$. The shifted vector $V_Y(s) = G_q^{-s} g_Y(G_q^s X)$ is polynomial of degree at most $r + 1$. Set $\theta_\tau(x) = (x/(x + \tau))^{\sigma_{\mathrm{reg}}}$ with $\sigma_{\mathrm{reg}} = 4L$, write $\|\cdot\|_{\tau,Y}$ for the norm in $\theta_\tau(G_q) d_Y\, d\mu_0$, and put $M_\tau = \sup_Y \|V_Y(0)\|^2_{\tau,Y}/|P_v(Y)|$ on the next interval. This supremum is finite. In the model normal form, at new word length m the residual density has factor W^m, with $W = |\zeta|^2/(|\zeta|^2 - TT')$, and the operators cost $(|\zeta| \prod_{b<r} |c_b|)^{-m}$. Their squared combination is G_q^{-m}, which θ_τ absorbs at fixed τ by $\theta_\tau(G_q) G_q^{-m} \leq \tau^{-m}$. The square-integrable generator products give the required regularized integrability

and a bounded majorant for continuity on the interval. One-letter identities and absorption · Normal-form integrability with the auxiliary cutoff · Continuity of regularized energies · Cut length and interval width · Weighted residue bound · Cauchy bounds · Markov bounds · Fatou and the energy identity

Pair the residue identity with $\theta_\tau X$, first inserting an additional $\chi(G_q/\delta_{\mathrm{aux}})$ in that slot to use the adjoint. At fixed $\tau > 0$, every logarithmic derivative of this product is bounded by a constant times θ_τ, uniformly in δ_{aux}. The shift polynomial and the regularized L^2 bounds permit $\delta_{\mathrm{aux}} \downarrow 0$ by Cauchy–Schwarz and dominated convergence. With $\mathrm{ld}_p = (x\partial_x)^p$ the resulting identity is

$$|P_v(Y)| \int \theta_\tau(G_q)|X|_{d_{\mathrm{old}}}^2 \, d\mu_0 = \|V_Y(0)\|_{\tau,Y}^2 + \sum_{p=1}^{r+1} \int \frac{(\mathrm{ld}_p\theta_\tau)(G_q)}{p!} \langle V_Y(0), V_Y^{(p)}(0)\rangle_{d_Y} \, d\mu_0.$$

The pairing is linear in the first slot. The shift compares $V_Y(s)$ with $g_{Y+s}X$ plus the residue difference in one branch. The weighted residue bound and the weight and polynomial ratios give $\|V_Y(s)\|_{\tau,Y}^2 \le C_7 M_\tau |P_v(Y)|$ for $Y + s$ in the interval. Markov's inequality on that interval costs $2(r+1)^2/H$ per derivative, by scalar duality in the regularized Hilbert space. Cauchy's estimate in $\log x$ gives $|\mathrm{ld}_p\theta_\tau|/p! \le \mathrm{e}(4\sigma_{\mathrm{reg}})^p\theta_\tau$. Thus the absolute error divided by $|P_v(Y)|$ is at most

$$\mathrm{e}\sqrt{C_7} M_\tau \sum_{p=1}^{r+1} (8\sigma_{\mathrm{reg}}(r+1)^2/H)^p \le M_\tau/2.$$

Here the fitted widths satisfy $H \ge 16\sigma_{\mathrm{reg}}(r+1)^2(1 + 2\mathrm{e}\sqrt{C_7})$, with C_7 the product of the two comparison ratios and the square of one plus the residue difference bound. Absorption gives $M_\tau \le 2A$, so every shift derivative has a squared regularized norm bounded by $2C_7(2(r+1)^2/H)^{2p}|P_v(Y)|A$, independently of τ. Fatou gives unregularized weighted L^2. For $p > 0$ the multiplier $\mathrm{ld}_p\theta_\tau(G_q)$ tends to zero on the chart, and the products of unregularized L^2 vectors dominate these integrals. Removing θ_τ therefore gives new energy $|P_v(Y)|A$, and induction gives $J_0 \prod_i |P_v(Y_i)|$ for the all-plus tuples in the intervals. One-letter identities and absorption · Normal-form integrability with the auxiliary cutoff · Continuity of regularized energies · Cut length and interval width · Weighted residue bound · Cauchy bounds · Markov bounds · Fatou and the energy identity

For such a tuple Y, let T_Y be the full word map, now with component weights $d_Y(q,e) = W(q)^N c_Y(e)$ for positive constants $c_Y(e)$. The cut form equals the integrated square of $T_Y(c_\varepsilon J)$ in these weights. The shifted normal form is

$$G_q^{-s} T_Y(G_q^s J)(q) = B_Y(q)\, M(Y + s\mathbf{1}, q),$$

where B_Y, B_Y^{-1} are bounded in q at each fixed tuple and M has degree at most $r+1$ in each argument. The uncut squares and B_Y^{-1} put M at all interval tuples in $L^2(W^N d\mu_0)$. Interpolating at $r+2$ nodes in each interval puts every coefficient in that space. Substitution and multiplication by B_Y give weighted L^2 for the full word jets $Z_{Y,p} = \partial_s^p (G_q^{-s} T_Y(G_q^s J))\big|_{s=0}$, $0 \le p \le d = N(r+1)$. Positive square formula and polynomial identity · Full word factorization · Weighted coefficient interpolation · Logarithmic Leibniz rule · Physical cutoff convergence · Signed grid reduction

The logarithmic Leibniz identity, obtained by testing the local product rule on x^s, writes $T_Y(c_\varepsilon J)$ as $\sum_{p=0}^d (\mathrm{ld}_p c_\varepsilon)(G_q) Z_{Y,p}/p!$. For $p > 0$ the derivative is supported on $\varepsilon/2 \le G_q \le \varepsilon$ and its bound is independent of ε, since ld_p acts on G_q/ε. For $p = 0$ the difference from one is bounded by one and supported on $G_q \le \varepsilon$. With $K_0 = 1$ and $K_p = \max_{[1/2,1]} |\mathrm{ld}_p\chi|/p!$ for $p > 0$, finite Cauchy–Schwarz

bounds the squared vector error by $\mathbf{1}_{0<G_q\le\varepsilon}D_Y$, where $D_Y=(d+1)\sum_p K_p^2|Z_{Y,p}|_{d_Y}^2$. The absolute squared-energy error is bounded pointwise by $\mathbf{1}_{0<G_q\le\varepsilon}H_Y$, where $H_Y=D_Y+2\sqrt{D_Y}|T_Y J|_{d_Y}$ is integrable. Positive square formula and polynomial identity · Full word factorization · Weighted coefficient interpolation · Logarithmic Leibniz rule · Physical cutoff convergence · Signed grid reduction

For a general signed word, let Y now denote the tuple of actual arguments. A minus letter at grid node y has argument $y+1$ in both the operator and the Capelli factor. At fixed ε the raw energy is the same real polynomial $Q_\varepsilon(Y)$ for every sign choice on admitted words, of degree at most $n=2r+k+1$ in each argument. This follows from the empty-prefix residue fold. Since P_v has a constant sign s_P on the fitted cut, the limit polynomial is $Q_0(Y)=J_0 s_P^N\prod_i P_v(Y_i)$. Use $n+1$ interpolation nodes in each interval and let $l_z(Y)$ be the tensor Lagrange basis at a grid tuple z. The actual arguments of the signed word can lie outside these intervals. Exact interpolation gives its integrable tail bound

$$|Q_\varepsilon(Y)-Q_0(Y)|\le\int_{0<G_q\le\varepsilon}\sum_z|l_z(Y)|H_z(q)\,d\mu_0(q).$$

The single empty tuple also gives the ordinary beta limit, where the original invariant multiplier is squared and dominated by the integrable section density. Since $G_q>0$ on the chart, dominated convergence makes the tail tend to zero. The same integrable density covers sublevel sets in which several factors of G_q become small together. Positive square formula and polynomial identity · Full word factorization · Weighted coefficient interpolation · Logarithmic Leibniz rule · Physical cutoff convergence · Signed grid reduction

Normalization and the finite row. The empty prefix has block and density factors one. Its fitted local Gram divides the raw form by Γ_0, so $R_v(\varepsilon,a)/(\Gamma_0 a_v)\to\|F_v\|_{\mathrm{sec}}^2/\Gamma_0$. The a_v are positive since P_v has no zero on the fitted cut. The product over the f complex places gives the common normalized level-zero limit. Division by the level constant · Place product · Selected energy · Finite-row threshold intersection · Recurrence inequalities at one smoothing

To track the scalar in the outer comparison, the full real chart scales its first r complex coordinates by $\lambda_a=(g/2\pi)^{2(r-a)}$, $0\le a<r$, and fixes the last. The additive Haar and density change give the place factor

$$K_{\exp}=2^{r+1}\prod_{a<r}\lambda_a^{-2(M_{a+1}-1)}=\frac{\rho^{2r(r+1)(k-1)+b_r}\,g^{3(r+1)}}{\Gamma_0},\qquad b_r=\frac{(r+1)(8r^2+10r+3)}{3}.$$

The placewise chart comparison uses the same calibrated local quotient operator. Thus $\mathrm{IntC}=K_{\exp}^f\prod_v\|F_v\|_{\mathrm{sec}}^2$. Writing $\mathcal{V}$ for the product of the Archimedean and finite level volumes and κ_N for the unipotent mass factor, the selected mass laws give $\mathcal{V}=\rho^{fb_r}g^{f(r+1)(2r+3)}\mathrm{volZZ}$ and $\kappa_N\mathcal{V}/\mathrm{volZZ}=(K_{\exp}\Gamma_0)^f$. The outer comparison multiplies IntC by $\mathrm{volZZ}/(\kappa_N\mathcal{V})$, giving $I_N=(K_{\exp}\Gamma_0)^{-f}\,\mathrm{IntC}=\prod_v(\|F_v\|_{\mathrm{sec}}^2/\Gamma_0)$. This is the starting energy in the scalar recurrence. Jacobi residue comparison · Beta cutoff · Chart Jacobian · Product of raw section norms · Volume and energy definitions · Outer mass normalization

At the fixed fitted datum, the word and place sets are finite. For any tolerance $\eta>0$, one threshold gives the initial bound $I_N+\eta$ for every admissible word after normalization. The wave, lift and block inequalities used in the recurrence have constants independent of the finite support and of the chosen positive ε, with the energies on both sides at that same ε. For each finite support take ε below both the word threshold and its empty-word top Gram threshold. The recurrence now bounds the original finite-

support row sum with initial energy $I_N + \eta$. Let $\eta \downarrow 0$ in this bound, then use row summability to exhaust the supports. The level-zero threshold concerns one fitted datum, and the top threshold may depend on the particular finite support. Division by the level constant · Place product · Selected energy · Finite-row threshold intersection · Recurrence inequalities at one smoothing

Detailed derivation of the selected wave identity

Detailed derivation of the level-zero cutoff limit

Derivative cost. The selected coefficient has the Capelli eigenvalue used above: differentiate the K-finite right vector in the transported representation, apply its infinitesimal character and conjugate the differential operators by Q. The gapped coefficient estimates supply the required derivative seminorms at $A_0 \geq C_c(r+1)^4(L(2r+k+1)^2 + (2r+k+1)^3)$, with $C_A \geq C_c \geq 28\mathrm{e}$. The fitted cut is at index $\ell_0 + 1$ for differing signature entries at ℓ_0, with $1 \leq \ell_0$, $\ell_0 + 2 \leq k$. Its separated grids lie in the prescribed positive cut of length $I \geq 6L^3 + 6L$, and the first-level grid has at least L nodes. The slice polynomial and mixed size identities are the remaining compact inputs for the derivative estimate. In its unweighted block, wave isometry and positivity give the local Fourier coefficient off-diagonal and the two normalized lower kernels on the diagonal. The frequency scale used in adding letters is $T_0 = g\lambda_S/(2\pi) = (2\pi/g)^{-(\nu+1)}$. Exact Capelli eigenvalue · Weighted derivative bounds · Fitted cut, grids and top lift

For a new node y, use the plus letter at y and the minus letter at $y+1$. Each slice slot is polynomial of degree at most j, with leading coefficient the old slice in slot zero and zero in slot one. In the alternating mixed slot pairing the wave adds s times slot zero to slot one on the plus side. The coefficient of y^{2j} is therefore s times the old wave pairing. Repeat for K new letters at each place. In each variable the coefficient is recovered from $2j+1$ nodes by the leading Lagrange weights $\lambda_u = \prod_{t \neq u}(y_u - y_t)^{-1}$ with absolute sum at most $(28L/I)^{2j}$. At $j = 0$ choose one node per variable from the special grid after removing the old sites. Both signed extensions are admissible, and the product of the absolute interpolation weights divided by its sum is a probability measure on the node tuples. Polynomial leading step · Repeated extraction · Interpolation bounds · First-level grid

For the mixed size comparison use the positive cut polynomial

$$Q_v(y) = -TT' \prod_{a=1}^{k} (y + 1/2 + u_a)(y + 1/2 + u'_a), \qquad A_{i,v} = \sup_y \sqrt{\frac{|P_v(y)P_v(y+1)|}{Q_v(y)}},$$

where the supremum is over the prescribed ordinary interval, or the long interval at $j = 0$. The compact type entries are u_a, u'_a at that place. Each old letter changes the polynomial by a factor $1 + 1/z$, with $z = y - t$ for its plus argument t and $z = t - y - 1$ for its minus argument t. The node conditions give $|z| \geq L^2/8$ (at least $L^2 + 1$ on the ordinary grid). As there are at most L letters and $\log(1+u) \geq -2|u|$ for $|u| \leq 1/2$, the modified polynomial is at least $\mathrm{e}^{-2} Q_v(y)$ for $L \geq 8$. The mixed size inequality for the two extended slices divides their paired squares by the product of these modified polynomials. In normalizing the extended slices, the common Γ_v cancels, and the new word weights contribute $\prod_b |P_v(y_b)P_v(y_b+1)|$ to the squared comparison. With the interpolation weights and the change from s^K to $(\mathrm{i}\ell_v)^K$, the resulting threshold is

$$X_{i,v} = \mathrm{e} T_0 (28L/I)^{2j} A_{i,v} \leq x_{i,v}, \qquad x_{i,v} = C_A T_0 (C_A L/I)^{2j} A_{i,v}.$$

In particular the derivative cost is K extra letters per place and the mixed square is bounded with factor $X_{i,v}^{2K}$ after normalization. Averaging the positive blocks using the node probabilities gives an off-diagonal $(\mathrm{i}\ell_v)^K$ times the original coefficient and diagonal mixtures of extended energies multiplied by $X_{i,v}^K$. Cut intervals · Old letters · Mixed size comparison · Scalar weight cancellation · Thresholds · Positive blocks

For any subset A of places take these weighted blocks at A and the unweighted blocks elsewhere. The entrywise product and contraction against H preserve positivity. The modulus of the off-diagonal is bounded by the mean of the two nonnegative diagonal entries. This gives, with probability weights θ_A on $W_b(\mathrm{cnt}(w)+K)$,

$$\left(\prod_{v\in A}|\sigma_v(\ell)|^K\right)|\mathrm{Nhat}_\varepsilon(\ell)| \le \left(\prod_{v\in A}X_{i,v}^K\right)\sum_{w'}\theta_A(w')\,\mathcal{E}_{\varepsilon,j,w'}(Y,H).$$

The zero term is $\mathrm{Nhat}_\varepsilon(0)=\mathcal{E}_{\varepsilon,j,w}(Y,H)$ by the product of the local identity pairings. For the passage from compact approximants $\psi_R\chi(G/\varepsilon)F_v$ use the fitted convergence of slice forms and Fourier integrals at fixed ε, including cutoff derivatives. A single subsequence in the finite product of probability simplices for all A preserves the comparison for all Y,H,ℓ, because all the energies and coefficients converge along it. Thus the mixtures can depend on the fitted data, i,w,ε, with a common choice for every coset and kernel. The continuation of the diagonal is nonnegative and decays as $C_{\varepsilon,w,Y,H}\prod_v(1+|c_v|)^{-3}$ in these coordinates. Radial coefficient limits · Fixed-smoothing decay · Place products · Zero coefficient · Cutoff and energies

Field-uniform sum. For nonzero ℓ use $A(\ell)=\{v:|\sigma_v(\ell)|>x_{i,v}\}$ in the preceding inequality. The resulting weight is $m_i(\ell)=\prod_v\min(1,(x_{i,v}/|\sigma_v(\ell)|)^K)$. Assume the fitted product size condition

$$\prod_v x_{i,v}\le\Theta^{-f}\sqrt{N\mathfrak{a}_i},\qquad x_{i,v}\ge0,\qquad \Theta\ge1.$$

To sum the weight, take any nonzero fractional ideal $\mathfrak{a}$. For distinct elements the product of embedding moduli of their difference is at least $\sqrt{N\mathfrak{a}}$ because its principal ideal lies in $\mathfrak{a}$. In the box $|\sigma_v(\ell)|\le R_v$ the points scaled by R_v^{-1} have supremum separation at least $\delta=(N\mathfrak{a}/\prod_v R_v^2)^{1/(2f)}$. If $\delta>1$ the box has no nonzero point. Otherwise the disjoint open polydiscs of radius $\delta/2$ lie in the enlarged unit polydisc, so the count is at most $(1+2/\delta)^{2f}\le 9^f\prod_v R_v^2/N\mathfrak{a}$. Use 100^f in place of 9^f to keep the release scale of constants. Fitted threshold product · Box estimate · Wave and lift sums · Periodization · Poisson passage

For positive thresholds split into dyadic boxes indexed by $k_v\ge0$, with $k_v=0$ for $|\sigma_v(\ell)|<2x_{i,v}$ and $2^{k_v}x_{i,v}\le|\sigma_v(\ell)|<2^{k_v+1}x_{i,v}$ otherwise. The bound on the number of nonzero points in a box is $400^f4^{\sum k_v}(\prod_v x_{i,v})^2/N\mathfrak{a}_i$, while the product profile of exponent $5/2$ is at most $2^{-(5/2)\sum k_v}$. The sum of that profile is at most C_{box}^f, where $C_{\mathrm{box}}=400/(1-2^{-1/2})$. The remaining exponent $K-5/2$ in each minimum gives a product at most $\Theta^{-f(K-5/2)}$ by the norm bound on nonzero ℓ. A zero threshold gives zero weight everywhere off zero. Using the trace-dual inclusion and bounding the mixtures by their maximum now gives

$$\frac{1}{\gamma_i}\sum_{\xi\in\Lambda_i^*\setminus\{0\}}|\widehat{\phi}_\varepsilon(\xi)|\le(C_{\mathrm{box}}\Theta^{-5/2})^f\max_{w'\in W_b(\mathrm{cnt}(w)+5)}\mathcal{E}_{\varepsilon,j,w'}(Y,H).$$

This estimate is uniform relative to the lower energies at the same ε. It proves absolute summability by taking the supremum of the finite sums. At fixed ε the product decay gives a locally uniformly convergent

primal periodization, whose Fourier series converges absolutely by this estimate. Poisson and positivity bound each finite subset of the coset, with zero coefficient given by the preceding-level energy. Fitted threshold product · Box estimate · Wave and lift sums · Periodization · Poisson passage

For the sum over cosets keep the mixtures and group by $A(\ell)$. The same word weights apply to every coset for a given subset. The positive semidefinite bound by $(1+e_i)\mathcal{K}_{i-1}$ on the sum of the coset kernels controls the corresponding sum of lower energies. With $\epsilon_{\mathrm{lat}} = (C_1'\Theta^{-5/2})^f$, take C_1' at least $2^{5/2}C_{\mathrm{box}}$ to accommodate the separate factor-two threshold in the lift profile when $\Theta > 4$. This contributes at most $2^f\epsilon_{\mathrm{lat}}(1+e_i)\sigma_{i-1}^2$ to the aggregate diagonal residual. The separate inverse-difference lift and Schur step gives the $3^f\epsilon_{\mathrm{lat}}$ error. It can add another K letters before the wave estimate, within the recurrence budget $4K(r+1-i)$ at level i and $4K(r+2-i)$ at the preceding level. Longer blocks require the other lifting-lattice and theta domination inputs. Theta and wave comparisons · Allocation of the two errors · Scalar row assembly

Detailed derivation

The scalar recurrence

Detailed derivation with constants

Real sampling for longer blocks. Fix $2 \le s \le i$, write $b = i - s + 1$, and let Y be a finite set of level-$b-1$ prefixes. For $H \succeq 0$ on Y, let $D_j(H; z)$ be the diagonal contraction obtained by appending (j, z) to each prefix, using the product of the density-weighted real Grams divided by $\mathrm{wt}(w)$. For complex tuples z, use the corresponding coordinate at each place. At a fixed $\varepsilon > 0$ the compact real inputs are the length-one wave and integral identity, the positive inverse-difference pairing below the top, and the longer-block fibre germ. The germ supplies the positive two-copy Gram with blocks U, P, P^*, Q. Here U is the length-$(j-1)$ kernel in the inverse coordinate. At $(x, x+\ell)$, the mixed block, its reverse adjoint and the lower-right block tend to the length-j diagonal transform at ℓ as each $|x_v|$ tends to infinity. The inputs include integrability and continuity of the diagonal off the coordinate hyperplanes, and boundedness of the auxiliary diagonal once the shorter diagonal is bounded. The compact approximation of the fitted gap-smoothed coefficient must preserve the comparison constants and converge in these forms at fixed ε. Real germ and sampling

Put $\nu_j = 2r + 3 - 2b - j$ and $\tau_{j,v} = \mathrm{thrS}(b, j, v)$, the positive fitted bounds for the length-one wave at $j = 1$ and for the shorter inverse-difference profile at $j \ge 2$. The period is

$$L_{b,j} = \{t \in g\mathfrak{d}^{\nu_j} : t - (-1)^j \bar{t} \in g^2\mathfrak{d}^{\nu_j}\}, \qquad \gamma_{b,j} = \mathrm{covol}(\iota L_{b,j}) = (\rho^{2\nu_j+1}g^3/2)^f.$$

Here ι uses the unscaled complex embeddings with ordinary Lebesgue measure. In the split model and in a lifted unramified normal basis, conjugation swaps the two basis components. Thus $1 - (-1)^j\overline{(\cdot)}$ modulo g has image of size g^f after removing the local real different generator. This proves the index used in the covolume, including at dyadic places. The trace dual is contained in $\mathfrak{a}_{b,j} = g^{-1}\mathfrak{d}^{-(\nu_j+1)}$, by testing $g^2\mathfrak{d}^{\nu_j} \subset L_{b,j}$. The local transform integrates against $\exp(2\pi\mathrm{i}\, 2\,\mathrm{Re}(\ell_v z_v)/g)$ and divides by $\rho^{2\nu_j+1}g^3/2$, so the product divisor is precisely $\gamma_{b,j}$. The threshold product gives the dimensionless factor

$$\mathrm{vol}(b, j) = \frac{(\prod_v \tau_{j,v})^2}{N\mathfrak{a}_{b,j}} \le \Theta^{-2f}.$$

The positive number Θ^{-f} is the maximum of the fitted products of normalized thresholds. Those for this period estimate use $g\rho^{\nu_j+1}\tau_{j,v}$. The fitted inverse-difference products satisfy the same norm-relative bound below the top and allow $2^f\Theta^{-f}$ at the top. Covolume and threshold products · Fitted sizes

For the germ step, take the entrywise product of the local two-copy Grams and contract against H. The difference of the two copies has positive kernel $U_{\mathrm{tot}} - P_{\mathrm{tot}} - P_{\mathrm{tot}}^{*} + Q_{\mathrm{tot}}$, where each total block is the contraction of the corresponding placewise product. Average this difference over common translations on g times the dual lattice. Translation invariance makes the average a positive-definite function of the difference. Write A for this function multiplied by $\gamma_{b,j}$. Here the dual is written for the pairing $2\,\mathrm{Re}\sum_v \xi_v z_v$, so the local parameter at ξ is $\ell_v = g\xi_v$. A translation-invariant mean ignores slabs where one coordinate stays bounded because each such slab has arbitrarily many disjoint translates. The three germ limits therefore identify $\widehat{D}_j^{+}(\xi) + A(\xi)$ with $\gamma_{b,j}$ times the mean of the shorter inverse-coordinate pairing. The superscript $+$ denotes the positive Fourier exponential. Its bound uses the inverse difference $x - (x + g\xi) = -g\xi$ and the paired lift estimate for length $j-1$. For $j = 1$ use the wave estimate with $A = 0$. Real germ and sampling

On any full auxiliary lattice $\mathcal{L}$, periodize the nonnegative diagonal, extended by zero on the excluded hyperplanes. It is integrable and lower semicontinuous. Product Fejér means bound its value from below in the limit at the chosen translate. The correction has nonnegative Fejér means by positive definiteness, and the absolutely summable combined coefficients bound the means from above. Thus a finite coset sum is bounded by $\sum_{\xi\in\mathcal{L}^*} |\widehat{D}_j^{+}(\xi) + A(\xi)|/\operatorname{covol}(\mathcal{L})$. The paired bound controls the numerator by the product profile $\prod_v \min(1, (\tau_{j,v}/|g\xi_v|)^5)$ times $\gamma_{b,j}$ and the shorter diagonal estimate. The factor at a zero coordinate is one. At zero frequency use the unweighted Gram bound. The probability on enlarged words is independent of the prefixes, H, and the translate. It comes from the local wave or paired bounds and convex averaging over the dual; a pairing averages the two shorter diagonal bounds. Real germ and sampling

For $j < s$ apply the inequality to a singleton using $\mathcal{L} = R(\mathbb{Z}[\mathrm{i}])^f$ and $R = g/\min_v \tau_{j,v}$. Its dual for the displayed pairing has spacing $1/(2R)$ in each complex coordinate. Disjoint discs of radius $1/(4R)$ at the dual points bound the profile sum by $\prod_v 54(1 + (R\tau_{j,v}/g)^2)$. Indeed the profile of width h has integral $5\pi h^2/3$, and enlarging its width by the disc radius dominates the central value throughout each disc. The multiplier after division by R^{2f} is at most $(54/(g\rho))^f \operatorname{vol}(b,j)$. At $j = s$ use $\mathcal{L} = \iota L_{b,s}$, so $\gamma_{b,s}/\operatorname{covol}(\mathcal{L}) = 1$. Since $\operatorname{vol}(b,s) \le 1$, the field-norm box argument above bounds the full dual profile by an absolute constant to the power f, including the zero mode. Iterating gives, for any finite subset S of one period coset,

$$\sum_{t\in S} D_s(H;\iota(t)) \le C_4^{fs} \prod_{j=1}^{s-1} \operatorname{vol}(b,j) \sum_{w'} \theta(w')\, \mathcal{E}_{\varepsilon,b-1,w'}(Y,H).$$

Here C_4 is an absolute constant and D_s takes the same ε, w as the local inputs. The word support has count at most $\mathrm{cnt}(w) + (s+1)K$. The compact-to-smoothed limit uses a subsequence in the finite word simplex at fixed ε, common to all Y, H, S. The numerical multiplier is uniform in the field, whereas the charts, words, probabilities and convergence are for the fitted data. Real germ and samplingCovolume and threshold products · Fitted sizes

Finite majorants and the residue sum. At a finite place use the fitted order R self-dual for $\psi(xy)$, with additive Haar volume one. At a typed split place this is $R = \{(x,y) \in o^2 : x - y \in go\}$ with a character

of $(x-y)/g$ whose base multiplication annihilator is o. Testing on $(ga,0),(0,ga),(a,a)$ gives $R^\perp = R$. The finite orbit model supplies the Gram $N_j(t,t')$ with the weight $|tt'|_R^{j(j-1)/2}$, inverse-difference support in R, and singleton cokernel Parseval. Write $c(h) = |\det \mathrm{pt}(h)|_R^{-1}$ and N_0 for the prefix Gram. The length-j theta kernel is positive, R-periodic, dominates $N_j(t,t)$, and has residue sum bounded in positive semidefinite order by $c(h)^j c(h')^j N_0(h,h')$. Finite germ and theta induction

For the induction, append a singleton to a length-m block and sum its Gram over unit representatives modulo R, with weight $|tt'|_R^{m(m+1)/2}$. Singleton Parseval gives $c(h)c(h')N_m(t,t')$. On a compact open unit domain Ω stable under $(m+1)$-st roots of unity, restrict this sum to the image of $z^{m+1} = (-t)^m c$ for $z \in \Omega$ and small t. The finite germ identifies this cut Gram at inverse shifts $t^{-1}+\sigma, t^{-1}+\sigma'$ with the integral of $N_{m+1}(z,z)\psi(z(\sigma-\sigma'))$ on Ω. This exact identity uses the matched roots and Jacobian of the finite stationary-phase calculation. Choose compact open $L_1 \subset R \subset L_2$ with $\Omega \subset L_1^\perp$, so the diagonal on Ω extended by zero is constant on $L_2^\perp$-cosets. Testing the cut Gram with the normalized Fourier vector of phase $\psi(-z\sigma)$ on L_2/L_1 gives $N_{m+1}(z,z)/[L_2 : R]$. The factor $[L_2 : R]^{-1}$ is the volume of $L_2^\perp$. The uncut Gram dominates the cut one and splits over L_2/R by inverse support. Multiplying its Fourier form by $[L_2 : R]$ gives the positive periodic majorant. Summing it over $L_1^\perp/R$ by Parseval leaves the average of the old uncut diagonal blocks, bounded by $c(h)^{m+1}c(h')^{m+1}N_0(h,h')$. Restrict to the compact open support module of the new diagonal and exhaust the units. On the finite prefix set and fixed support quotient, a convergent subsequence gives a single majorant with this budget. Finite germ and theta induction

From framed orbits to the phase. Work over the finite local algebra with the self-dual order R above. For an integral matrix Z invertible over the algebra, the test $F_{Z'}(Z)$ counts isomorphisms from R^n/ZR^n to the framed target, preserving $[e_n]$ and $e_n^{\mathsf{t}} Z^{-1} v \bmod R$ on a class $[v]$. It vanishes off the integral locus. Integral lower and upper unitriangular operations $Z \mapsto LZU$ preserve both marks by the map $v \mapsto Lv$. There is also additive stability. If $d = \det Z$ and the perturbation $E = d^2 \varpi P$ is integral with P integral and ϖ sufficiently small in R, then $Z + E = Z(1 + d\varpi \operatorname{adj}(Z)P)$. The second factor is invertible over R, and the difference of the two inverses is integral by the adjugate formula. The image lattice and both marks therefore agree. For a nonintegral Z, an integral perturbation leaves both tests zero. Framed marks and prefix reduction

Fix a prefix h of size a and an appended block point of size $s = m+1$. Let T and G be the orbital matrices of the two points. The lower and upper rectangular cross coordinates X, Y, with the diagonal coset representatives fixed, range entrywise over K_{loc}/R, where K_{loc} is the ambient local algebra. Multiplication of the triangular factors gives the joined matrix and the extra character

$$B = \begin{pmatrix} T & TY \\ XT & XTY + G \end{pmatrix}, \qquad \psi(X_{0,a-1} + Y_{a-1,0}).$$

For an empty prefix the extra character is one. Summing $F_{Z'}(B)$ with this character and the prefix character defines a test $\Phi(G)$ for the two-block point and the single block. Reindexing $X \mapsto LX,\ Y \mapsto YU$ shows its integral triangular invariance. On integral support XT, TY are integral, so the cross sum is finite for each of the finitely many integral prefix representatives. Also $(\det T)G = (\det T)(XTY + G) - (XT)\operatorname{adj}(T)(TY)$ is integral. Since $\det B = \det T \det G$, the additive stability argument gives a common translation lattice for Φ when $\det G$ stays in a compact subset

of $K_{\mathrm{loc}}^{\times}$. The support bound and this lattice are independent of the target Z'. Framed marks and prefix reduction

For (m,t) followed by $(1,c)$, the internal lower factors are trivial. Write $1+N$ for the upper factor, ξ for the last row of the lower factor below the m-block, and J_m for reversal. Reverse the rows of the matrix

$$\begin{pmatrix} tJ_m & 0 \\ t\xi J_m & c \end{pmatrix} (1+N).$$

In the resulting upper Hessenberg matrix, the subdiagonal is t and the diagonal is $(\rho_0, x_1, \dots, x_m)$ with $\rho_0 = t\xi_{m-1}$ and $x_j = tN_{j-1,j}$. The determinant is $D = (-t)^m c$. Write y_{ij} for the strictly upper entries and $X_* = \prod_j x_j$. For small t and D in a compact unit range, bounded support forces all x_j and their inverses into bounded sets. Indeed the diagonal product differs from D by t times a bounded quantity. Put $x_0 = D/X_*$. The determinant is affine in ρ_0, with coefficient $X_* + O(t)$. Its linear terms in t come from adjacent swaps, so

$$\rho_0 = x_0 + t\, g_1 + O(t^2), \qquad g_1 = x_0 \sum_{j=0}^{m-1} \frac{y_{j,j+1}}{x_j x_{j+1}}.$$

Returning the rows to their original order after replacing ρ_0 by x_0 and the subdiagonal by zero gives a matrix Z_0 with the same test value, by the common translation lattice. The original character is $\psi((\rho_0 + \sum_j x_j)/t)$. Its two surviving factors are $\psi(f_D(x)/t)$ and $\psi(g_1)$, where $f_D(x) = \sum_j x_j + D/X_*$. Hessenberg and orbital integral · Critical amplitude

In the change from ξ to the first row's strictly upper entries, the determinant of the linear map is the tail principal minor up to sign. Its modulus is $|X_*|$, since the correction to its ratio with X_* lies in ϖR. The remaining upper coordinates scale by t. The original upper quotient cells have volume one because integral right multiplication is affine triangular with diagonal one in the strict upper entries. Thus, for $q = m(m+1)/2$, the orbital sum is $|t|^{-q} \int A(x)\psi(f_D(x)/t)\, dx$, where

$$A(x) = \int_{K_{\mathrm{loc}}^{q}} |X_*|^{-1} \Phi(Z_0(x,y))\psi(g_1)\, dy.$$

The same denominator bounds and translation lattice give compact support in the unit coordinates and a common lattice of constancy for A. At $x_j = z$ with $z^s = D$, rescale the strict upper entries by z. The restored matrix is $J_s z(1+N_0)$ and g_1 becomes the simple upper character of $1+N_0$. Hence $A(z,\dots,z)$ is the single-block orbital sum times $|z|^{q-m}$. Applying this calculation to the integral target matrices in the second orbit, with their conjugate phases, passes the formula to K_p. Hessenberg and orbital integral · Critical amplitude

Detailed matrix derivation · Finite stationary-phase derivation

Finite stationary phase in the Gram pairing. Set $s = m+1$. Use self-dual Haar measure for ψ and write $|a|$ for the Haar modulus. The orbital change of variables supplies the phase $f_D(x) = \sum_i x_i + D/\prod_i x_i$, with $D = (-t)^m c$, a locally constant amplitude supported where the entries and their inverses are bounded, and a prefactor $|t|^{-ms/2}$. At a critical point all entries equal a root z of $z^s = D$. The phase there is sz and the Hessian is $z^{-1}(I + \mathbf{1}\mathbf{1}^t)$, of determinant sz^{-m}. Near the root the analytic Morse change preserves measure on a sufficiently small lattice ball. Outside the root balls a nonzero partial derivative in some field component changes one coordinate to the phase, giving zero by

character orthogonality. Compact support, a common translation lattice and compact unit parameters give a common threshold at fixed local data.

The Gaussian at the root contributes $\gamma(zt)|zt|^{m/2}|s|^{-1/2}$, where γ is a unit-modulus function on square classes. The critical amplitude contributes $|z|^{m(m-1)/2}$ times the single-block orbital value. Thus each copy of the kernel has scalar $|z/t|^{m^2/2}|s|^{-1/2}\gamma(zt)\psi(sz/t)$. Take $t_i^{-1} = t^{-1} + \sigma_i$ and $\kappa = (t_2/t_1)^{m/s}$ near one. The paired roots are $z, \zeta\kappa z$, with $\zeta^s = 1$. After taking the outer Gram weight and the normalization of N_s, the common absolute factor is $|t_1|^m/(|s||z|^m)$. This is the inverse Jacobian of $c = z^s/(-t_1)^m$.

To integrate the sum over representatives modulo R, let $r \in R$. Replacing c by $c + r$ moves its root by $\delta z = (-t_1)^m r/(sz^m) + O(t^{2m})$ componentwise on compact units. The coefficient of z in the paired phase is $b_\zeta = s(t_1^{-1} - \zeta\kappa t_2^{-1})$. For $m \geq 2$ the phase change tends to zero. For $m = 1$ it is $-(z^{-1} - (\zeta z)^{-1})r + O(t)$, since $\zeta^2 = 1$. The inverse-difference support of N_s puts the leading term in R, so its character is one even for the coupled split order. Local constancy handles the amplitude and the remaining error. Integration over $c + R$ now counts each residue with weight one, and the power-map Jacobian cancels the common factor. For $\zeta = 1$ the Gaussian indices cancel and $b_1 = \sigma_1 - \sigma_2 + O(t)$. Every $\zeta \neq 1$ has a field component where b_ζ leaves every compact set, so its integral against the locally constant remaining amplitude vanishes. The cut Gram is exactly the Fourier integral of $N_s(h, h'; z, z)$ on the prescribed compact open domain. These thresholds can depend on p and s, including when $p \mid s$. Stationary phase, Gaussian normalization and the cut identity

The arithmetic kernel on the selected family. Use a finite set of prefixes h of degree $m = i - s$, and put $b = m + 1$. A tower block (j, z) starting at a and ending at $d = a + j - 1$ corresponds to the rational antidiagonal block with signed entry $(-1)^{r+1-d}z$. At a finite place v, set $q_v^* = g_j$ over a comparison prime g_j, and $q_v^* = 1$ elsewhere. Write ι_v for the embedding into E_v. Apply the same frame rescaling $(\delta_v q_v^*/g)^{-(r-a_0)}$ in each frame slot of index a_0 starting from zero. For a comparison place, write (C_1, C_2) for the two components of the rescaled cross matrix. The packet matrix there is $(C_1, -C_2^{\mathsf{t}})/q_v^* - \tau_{0,v}ee^{\mathsf{t}}$, where $\tau_{0,v} = (-T, T')/q_v^*$ and e is the last vector. At an ordinary place use the rescaled cross matrix and $\tau_{0,v} = 0$. The signed orbit uses right multiplication by the rational column-sign diagonal. This diagonal fixes the last row, so it preserves the two marks of the cokernel. The class uses the further integral corner shift for a top singleton, as described below. The quadratic Frobenius sign in the density factors is $+1$ at split places, including both comparison primes, and -1 at inert places. The sign $(-1)^s$ of the second split component of an appended block comes from the negative transpose and reversal of column signs within the antidiagonal. Signed points and weights · Packet matrices and class maps · Typed matrix · Frobenius sign and reference factors

The signed local prefix $\eta_v(h)$ is the block point obtained by these rescalings, with each block using its own endpoints. For an appended block (s, t), put $\nu = 2r + 3 - 2b - s$. Its unshifted local parameter $z_v(t)$ is $\iota_v(t)$ times $(-1)^{r+1-i}(\delta_v q_v^*/g)^{-\nu}/q_v^*$, with the additional factor $(-1)^s$ in the second component at a comparison place. The local kernels use the ordinary order at ordinary places and the coupled order $R_v = \{(x, y) \in \mathcal{O}_{F_v}^2 : x - y \in q_v^*\mathcal{O}_{F_v}\}$ at comparison places. At a comparison place the character on the split algebra is the base character applied to $(x - y)/q_v^*$. Its multiplication annihilator is R_v, and the local Haar measure has $\operatorname{vol} R_v = 1$. The framed orbit Gram counts integral triangular cosets with their simple lower and upper characters. Two integral orbit matrices in the same class contribute one framed automorphism order per pair of cosets, obtained by pairing their two isomorphism-count coefficients and dividing once by the automorphism order. Thus the class Gram with the factors $1/a_\lambda$ above is the product

of these local framed-cokernel Grams, with equal rational kind imposed at the top. The additional factor for a comparison place at the top with $s = 2$ is $\psi_{R_v}(\tau_{0,v}/z_v(t))$ on the first argument and its conjugate on the second. It cancels on the parameter diagonal. For $s \geq 2$ at the top the kind is absent because the leading cofactor is zero. Local rescaling and corner · Orbit-to-class comparison · Top phase and finite products

Let $\rho_E = |\operatorname{disc} E|^{1/(2f)}$, and write $n_\nu(t) = |N_{E/\mathbf{Q}}(t)|/(g^{2f}\rho_E^{2f\nu})$. For each prefix block (j, z) write $\nu_{a,d} = 2(r+1) - a - d$, and set

$$\mathcal{D}(h) = \prod_{(j,z)\in h} n_{\nu_{a,d}}(z)^j, \qquad \phi_m(h) = \prod_{(j,z)\in h} n_{\nu_{a,d}}(z)^{j(m-d)+j(j-1)/2}, \qquad A(h) = \frac{1}{\phi_m(h)\mathcal{D}(h)^s}.$$

The empty prefix has weight one. Invalid prefixes have zero kernel entries and may be assigned $A(h) = 0$. For a valid appended parameter the product of its local moduli is $n_\nu(t)^{-1}$. Indeed the different generators have product modulus ρ_E^{-2f}, the signs and q_v^*/g have modulus one at each place, and q_v^* has product modulus g^{-2f}. The field product formula supplies the remaining factor $|N_{E/\mathbf{Q}}(t)|^{-1}$. The determinant of each head block is a sign times its entry to the block length, so $\prod_v c_v(h) = \mathcal{D}(h)$, where $c_v(h) = |\det \operatorname{pt}(\eta_v(h))|_{R_v}^{-1}$. Determinant product · Common prefix weight

Increasing the degree from m to $i = m + s$ adds sj to the exponent of each prefix block of length j, hence

$$\phi_i(h * (s,t)) = \phi_m(h)\mathcal{D}(h)^s n_\nu(t)^{s(s-1)/2}.$$

The terminal factor in this denominator on each argument is exactly the inverse of the product of the local factors in N_s. Thus, except at the top singleton, the upper kernel on the appended points is $A(h)A(h')$ times the product of the local N_s evaluated at $z_v(t), z_v(t')$ on the two heads, with the stated character factors. The lower local budget on those same heads is $c_v(h)^s c_v(h')^s N_{0,v}(\eta_v(h), \eta_v(h'))$. Its product contributes $\mathcal{D}(h)^s\mathcal{D}(h')^s$, which cancels the extra prefix determinant power in $A(h)A(h')$, leaving precisely $\mathcal{K}_m$. For a finite set of heads, all but finitely many places have integral invertible heads and lower factor one. For any given t, t' a further finite set contains the nontrivial upper factors. The outside system retains the period and lift tests at every outside place. Upper product · Lower product and lattice assembly · Outside places

The local period tests membership of $z_v(t)$ in R_v. The local lift tests membership of the inverse multiplier applied to $x = t^{-1} - t'^{-1}$ in R_v. At ordinary places these are the different-ideal tests of exponents ν and $-\nu$. At a comparison place the different is a unit, and the two tests reduce to membership of $(t_1, (-1)^s t_2)/q_v^*$ and $q_v^*(x_1, (-1)^s x_2)$ in R_v. Thus the two t-coordinates must be divisible by q_v^* and their difference for even s, or sum for odd s, must be divisible by $(q_v^*)^2$. The two lift coordinates must lie in $(q_v^*)^{-1}\mathcal{O}_{F_v}$, and their corresponding difference or sum must lie in $\mathcal{O}_{F_v}$. Intersecting the period tests gives $L_{b,s}$ above. The intersection of the lift tests is contained in

$$L'_{b,s} = \{x \in g^{-1}\mathfrak{d}^{-\nu} : x - (-1)^s \bar{x} \in \mathfrak{d}^{-\nu}\}.$$

This lift lattice controls inverse differences. The trace dual used in Poisson summation has the additional inverse different in its containing ideal. Period and inverse lift · Poisson trace dual

At $i = r+1,\ s = 1$, the exponent ν is zero, and the leading minor of the head is nonzero. The rational kind is the last singleton entry t, so the upper kernel is zero for $t \neq t'$. For $t = t'$ use the periodic singleton diagonal at $z_v(t) - \tau_{0,v}$. In split coordinates this is $(t_1 + T, -t_2 - T')/q_v^*$. One may add the same integer from $\{0, 1, 2\}$ to both components to obtain an invertible parameter, since each component excludes at

most one integer. The added integer belongs to R_v and leaves the periodic diagonal unchanged. At an ordinary place the unshifted nonzero field parameter already maps to a unit of the local algebra. The affine shift leaves the odd period test on increments unchanged, and exact diagonality gives zero inverse-difference support subgroup in the top local system. The singleton adds no terminal determinant power, so the same prefix factor $A(h)$ and the $s = 1$ lower budget apply. Top singleton including its shift and support

On the parameter diagonal, multiply the local positive majorants with the same $A(h)A(h')$, using the shifted singleton majorant at the top when $s = 1$. Distinct global residues inject into tuples of local residues by the period equality. Enlarging to the full product of the projected residue sets adds positive semidefinite terms. The product of the local residue budgets is $\mathcal{K}_m$ by the lower identity above. The unscaled recurrence therefore has $e_i = 0$. Its generic arithmetic interface is supplied here by the explicit products for the selected setting. In the pinned source the selected call uses the concrete local-model theorem for those kernels and finite models. The abstract existence statement for an arbitrary fitted real setting is a separate proposition. Residue products · Concrete local-model construction · General interface · Joint specialization Detailed source trace

In each coset use its theta majorant as H in the real estimate, and sum using the same word probabilities. The summed lower energies are bounded by $(1 + e_i)\sigma_{i-s}^2$, where σ_i bounds the square root of the normalized form over the level-i word budget at the same smoothing. The remaining off-diagonal terms use the paired lift for length s, including its separate top bound for $s \geq 2$. For fixed t each nonzero $t^{-1} - t'^{-1}$ determines at most one t', and the same holds in the other direction. Summing the mean of the two diagonal energies costs at most $3^f \epsilon_{\mathrm{lat}}$ times the bound for the enlarged words. One enlargement for the lift and one for the coset estimate fit $\mathrm{cnt}(w) + (s + 2)K \leq 4K(r + 1 - i + s)$. The $2^f \epsilon_{\mathrm{lat}}$ loss comes from the length-one Fourier remainder. The marked top singleton is diagonal in t. One-level coefficientsCovolume and threshold products · Fitted sizesFitted arithmetic model · Residue products · Schur lift

Taking the Gram norm of the sum of the final-length groups and applying the triangle inequality gives

$$\sigma_i \;\leq\; \sqrt{(1 + 3^f \epsilon_{\mathrm{lat}})(1 + 2^f \epsilon_{\mathrm{lat}})(1 + e_i)}\,\sigma_{i-1} + \sum_{s=2}^{i} \sqrt{(1 + 3^f \epsilon_{\mathrm{lat}})(1 + e_i) C_4^{fs} \prod_{j=1}^{s-1} \mathrm{vol}(i - s + 1, j)}\,\sigma_{i-s}.$$

The inequality refers to the uniform bound for that step over its permitted test words. For the fitted family use $K = 5$ and $\epsilon_{\mathrm{lat}} = (C_1' \Theta^{-5/2})^f$, and each dimensionless lifting factor is at most Θ^{-2f}. Divide σ_i by $\prod_{j=1}^{i} (1 + e_j)^{1/2}$. The shorter-block coefficients can only decrease after this normalization, since the intervening density factors are at least one. For $\eta = (C_*/\Theta)^f$, take $C_* \geq \max(3C_1', 2C_4, 1)$ and Θ large enough that $\eta \leq 1/2$. The length-one coefficient is at most $1 + \eta$. A length-s coefficient for $s \geq 2$ is at most η^{s-1}, using $s/2 \leq s - 1$. The running maximum of the normalized σ_i therefore grows by at most $1 + \eta + \eta/(1 - \eta) \leq 1 + 3\eta$ at each stage. After squaring, its accumulated factor is at most $\exp(6(r + 1)\eta)$, which is at most $1 + 12(r + 1)\eta$ for $6(r + 1)\eta \leq 1$. The fitted size theorem gives $C(r + 1)(C/\Theta)^f \leq b^{-7}$ with $C \geq \max(12, C_*)$ in its smallness comparison. The fixed-data cutoff limit, with initial normalized square tending to I_N, gives

$$\sum_{\rho \in R} \Big\| \sum_{p} c_\rho(p) V(p) \Big\|^2 \leq (1 + \delta) e_1 I_N, \qquad \delta = C(r + 1)(C/\Theta)^f, \qquad e_1 = \max(1, s_{\mathrm{den}}). \tag{7e}$$

In the recursion the lower-level e_i vanish. Scaling the top kernel by e_1 scales its majorants as well and gives just one factor e_1 in (7e). The fixed-support cutoff and the absolute row sum justify passage from finite towers to (7e). Local block majorization and Poisson recurrence · Lifting factors and limit hypotheses · Iteration in the source · Top-only scaling · Application to actual rows

Proved here (nonsingular fibre mean). Write $V = (E^M, J)$ for a nonsingular Hermitian space over the CM extension E/F, and $G = U(V)$. For an ordered a-frame X put $Q(X) = X^*JX$, and for a nonsingular Hermitian matrix T over E/F let $X_T = Q^{-1}(T)$. Use a nontrivial additive character ψ of $\mathbb{A}_F/F$. At each place the measures dx_w, ds_w on frame and Hermitian matrix space are self-dual for the pairings

$$B_X(X,Y) = \operatorname{Tr}_{E_w/F_w} \operatorname{tr}(X^*JY), \qquad B_S(S,B) = \tfrac{1}{2} \operatorname{Tr}_{E_w/F_w} \operatorname{tr}(SB) = \operatorname{tr}(SB).$$

Let $d\nu_{T,w} = |dx_w/ds_w|_{X_T}$ denote disintegration along Q, and $g_w(T) = \int_{X_T(F_w)} \Phi_w \, d\nu_{T,w}$. For $M > 2a$, a smooth compactly supported real test Φ_∞ and locally constant compactly supported finite tests Φ_w equal to the standard lattice indicator almost everywhere, put $\Phi = \Phi_\infty \prod_{w<\infty} \Phi_w$. Then

$$\frac{1}{\operatorname{vol}_{dg}[G]} \int_{[G]} \sum_{X \in X_T(F)} \Phi(g^{-1}X)\, dg = g_\infty(T) \prod_{w<\infty} g_w(T), \qquad [G] = G(F)\backslash G(\mathbb{A}_F). \tag{7f}$$

Here g_∞ integrates the real test against the product of the real fibre measures. This is the nonsingular Gram coefficient of the convergent Siegel–Weil formula in the convention of the pinned statement. Weil treats $M > 2a$ in his 1965 paper. Yamana, pp. 236–237 and Theorem 2, gives the convergence criterion and the factor one for the Eisenstein series of the standard Weil section in this range, with $s_0 = (M-a)/2$. In the unitary case $M - r_W > a$ for the F-Witt index r_W when $M > 2a$. Gan–Qiu–Takeda §3.1 specify the Tamagawa probability measure on the theta side. The proof of (7f) uses the orbit of a single Gram coefficient and its quotient measure. Weil · Yamana · Gan–Qiu–Takeda · Pinned mean and test conventions · Self-dual pairings · Pinned proof of the product normalization

If $X_0 \in X_T(F)$, the ordered span is a nondegenerate copy of the form T. Witt extension makes $X_T(F)$ a single orbit $G(F)/H(F)$ with $H = U(W)$, where $W = \langle X_0 \rangle^\perp$. It has determinant class $[\det J/\det T]$ in $F^\times/N_{E/F}E^\times$, and its real signatures are those of V minus those of T. The orbit fixes the entire ordered frame. If T is represented everywhere locally, choose a global diagonal form of rank $M-a$ with determinant $\det J/\det T$ and the required real signatures. Weak approximation selects the first $M-a-1$ signs, and the determinant specifies the last one. Finite local classification by dimension and determinant and Landherr's local–global theorem identify its orthogonal sum with T with V. Thus in the globally empty case some local fibre is empty, and both sides of (7f) vanish. Local Witt extension gives $G(F_w)/H(F_w)$ at every represented place. At almost every finite place a unimodular frame splits off integrally, and the unimodular complements are isometric by finite-field classification and smooth lifting. The orbit maps therefore give the restricted product quotient. Landherr · Pinned orbit quotient convention

The differential $dQ_X(Y) = X^*JY + Y^*JX$ admits the right inverse $S \mapsto XT^{-1}S/2$ on X_T. It defines the local fibre measures by submersion and makes g_w continuous on the nonsingular locus. Choose rational coordinate volume forms dX, dS and the induced rational fibre form dX/dS by the differential exact sequence. The fibre form is G-invariant since $N_{E/F} \det g = 1$. Choose invariant rational forms on G, H whose quotient on X_T is this form. The self-dual measures are scaled from powers of a self-dual scalar measure on F_w by $|\det B_X|_w^{1/2}$ and $|\det B_S|_w^{1/2}$, using the pairing matrices in the rational

coordinates. Their ratio has adelic product one by the product formula. For the trace character the scalar measure gives $\mathcal{O}_{F_w}$ volume $|\mathfrak{d}_{F_w/\mathbb{Q}_p}|_w^{1/2}$ at finite places. Its combined discriminant factor in the quotient is $|D_F|^{-(2Ma-a^2)/2}$, the Tamagawa discriminant factor for the dimension difference $\dim G - \dim H = 2Ma - a^2$. Replacing $\psi(t)$ by $\psi(bt)$ for $b \in F^\times$ scales local fibres by $|b|_w^{(2Ma-a^2)/2}$, again with product one. The two positive-rank unitary groups have the same determinant torus $U(1)$, so the Tamagawa convergence factors $L_w(1, \chi_{E/F})$ and the global $L(1, \chi_{E/F})$ factor cancel in the quotient. Consequently $dg^{\mathrm{Tam}}/dh^{\mathrm{Tam}}$ on the orbit is the product of the self-dual fibre measures. Kottwitz's formula gives $\tau(U(L)) = 2$ for each positive-rank Hermitian space L, including rank one. The dual centre is $\mathbb{C}^\times$ with quadratic inversion, so its fixed subgroup has order two. A locally trivial cohomology class restricts to a character over E that is trivial by Chebotarev. The resulting class for $\mathrm{Gal}(E/F)$ is zero because the coboundaries for inversion are squares and every complex number has a square root. For $a = 0$ the fibre is a point and (7f) follows directly. Tamagawa measure convention · Kottwitz's formula · Tamagawa numbers

At a finite place, write L_X, L_S for coordinate lattices and q_w for the residue size. The self-dual lattice volume for a pairing with annihilator $L^\perp$ is $([L : L \cap L^\perp]/[L^\perp : L \cap L^\perp])^{1/2}$. A weighted congruence count on $L_X/\varpi^j L_X$ for $Q(X) - T \in \varpi^j L_S$, normalized by $q_w^{j(2Ma-a^2)}$, is converted to $g_w(T)$ by multiplying its limit by $dx_w(L_X)/ds_w(L_S)$. Take L_X to contain the test support and L_S large enough that the congruence is well defined on these residues. Equivalently divide the test integral over the Gram tube by $ds_w(\varpi^j L_S)$ and let j grow. This follows by continuity of the fibre integral and includes ramified or dyadic pairings, nonunit Gram determinants and level weights. At a real place with $\psi_w(t) = e^{2\pi i b_w t}$, the densities against ordinary coordinate Lebesgue measures on tuple entries and on real diagonal and complex upper Gram entries are

$$dx_w = |b_w|^{Ma} 2^{Ma} |\det J_w|^a \, dX_{\mathrm{Leb}}, \qquad ds_w = |b_w|^{a^2/2} 2^{a(a-1)/2} \, dS_{\mathrm{Leb}}.$$

These multiply the coordinate coarea formula by their ratio. The real group chart in (5) uses its separate prescribed quotient measures and $2\, d\Re c\, d\Im c$ on each complex coordinate. Its C_{oo} scales by $A/(BB')$ if the host and the two real stabilizer Haars scale by A, B, B'. Pairing conventions · Pinned finite lattice and volume calculation · Real group disintegration

At a split or unramified finite place with conductor $\mathcal{O}_{F_w}$, unimodular J, T and the standard lattice test, the two additive lattice volumes are one. Finite Witt extension gives the group order ratio of $\mathrm{GL}_M/\mathrm{GL}_{M-a}$ in the split case and U_M/U_{M-a} in the inert case. At inert places the differential on the integral fibre is surjective modulo the prime because the trace of the unramified residue extension is surjective, including in residue characteristic two. At split places the two frames pair a vector space with its dual, and a variation in one frame realizes every target variation. Smooth lifting multiplies the number of solutions by $q_w^{2Ma-a^2}$ per congruence step. Thus the density is

$$g_w(T) = \prod_{i=M-a+1}^{M} (1 - \epsilon_w^i q_w^{-i}), \qquad \epsilon_w = 1 \text{ if split}, \quad \epsilon_w = -1 \text{ if inert}.$$

The tail has summable deviation $O_{M,a}(q_w^{-(M-a+1)})$, hence converges absolutely to a positive product. The finitely many remaining factors are finite by the submersion and the compact test supports. Now unfold the sum over $H(F)\backslash G(F)$ first for $|\Phi|$ by Tonelli. Integration over $H(F)\backslash G(\mathbb{A}_F)$ factors as the covolume of H times the fibre integral on $H(\mathbb{A}_F)\backslash G(\mathbb{A}_F)$, with orbit map $g \mapsto g^{-1}X_0$. The preceding local convergence proves integrability and allows the same unfolding for Φ. Dividing by $\tau(G)$ leaves the scalar

$\tau(H)/\tau(G) = 2/2 = 1$. With an arbitrary host Haar and the compatible stabilizer Haar, both covolumes scale by the same factor. The finite genus sum induced by this unfolding uses inverse automorphism weights by orbit–stabilizer. The row weights in (7a) use the induced group on the orthogonal quotient, with the unipotent volume accounted for in the tied normalization. Unramified density formula · Quotient and covolume · Induced row weights

The full lattice indicator used for absolute integration in (7f) majorizes the primitive indicator in the same coordinates. The densities agree for a unit Gram. For a nonsingular Gram with singular reduction, the full density can also have contributions from integral overlattices of the framed span. The selected signed comparison uses (7h) in the prescribed coordinates and the class Gram sum above. Its rational early-minor boundary terms use the block estimates even when the exceptional stratum is Haar-null for (5). Full and primitive lattice densities · Actual scalar comparison

Application to the tied mass. Take $M = 2r + k + 1$ and $a = 2r + 2$ for the joint frame. Its pointwise stabilizer has Hermitian rank $k - 1$. The two single-frame stabilizer comparisons use nonsingular completions of size $2r + 1$. All three sizes are in the range of (7f) under the selected host rank hypothesis. For the given ambient Haar the orbit proof equates its covolume to that of each stabilizer in the Haar compatible with the self-dual fibre measure. The source obtains the same equality by comparing the mean to unfolding on a positive test and cancelling the fibre integral. It then uses the prescribed Haar identifications for the joint frame and both completed single frames to obtain $\mathrm{volB}\,\mathrm{volG} = \mathrm{volHa}\,\mathrm{volHa}'$. The two tied unipotent quotient volumes are one. The Haar identifications and completed-frame construction enter through the selected volume theorem. Compatible covolumes · Positive-test cancellation · Selected frame and Haar transport

The other volume ratios use the actual level indices. At a mass place the ambient natural level has volume $\gamma_{\mathfrak{p}}\mu_{K'',\mathfrak{p}}$, and the two stabilizer levels have volumes $\gamma_{\mathfrak{p}}\mu_{S'',\mathfrak{p}}$ and $\gamma_{\mathfrak{p}}\mu'_{S'',\mathfrak{p}}$, with $\Gamma_1 = \prod_{\mathfrak{p}} \gamma_{\mathfrak{p}}$. Thus the tied local ratio $\mu_{K,\mathfrak{p}}\mu_{K^{\mathrm{nat}},\mathfrak{p}}/(\mu_{S,\mathfrak{p}}\mu'_{S,\mathfrak{p}})$ contributes one reciprocal $\gamma_{\mathfrak{p}}^{-1}$. The ordinary factor in the product uses $\mu_{K,o}^2/(\mu_{S,o}\mu'_{S,o})$ in the source notation. The column mass uses the same stabilizer level, with $c_0 = \mathrm{volN}\,\mathrm{volH}/\dim \pi_u$, $\mathrm{volN}\,\mu_N = 1$ and $d_u = \dim \pi_u\,\mathrm{volHa}/(\mathrm{volH}\mu_{H''})$. This gives $c_0^{-1} = d_u\mu_{S''}/\mathrm{volHa}$ since $\mu_{S''} = \mu_N\mu_{H''}$. The primed column has the corresponding identity. Level, ordinary and column factors · Product of the level factors

To normalize (7e), the coefficient-to-signature comparison contributes $\mathrm{AA}/(c_0c'_0)$, where $\mathrm{AA} = \mathrm{volB}\,C_{\mathrm{oo}}(\mathrm{MM}/\mathrm{volZZ})/\mu_K$. The real chart gives $I_{\mathrm{oo}} = C_{\mathrm{oo}}\mathrm{IntC}$. In the natural measure $I_N = (\mathrm{volZZ}/(\mathrm{unitLevelVolume}\,\mathrm{volPP}))(\mathrm{volPP}\,\mathrm{IntC})$, so $I_N = (\mathrm{volZZ}/\mathrm{unitLevelVolume})C_{\mathrm{oo}}^{-1}I_{\mathrm{oo}}$. Multiplying by Γ_1, one uses

$$\frac{\mathrm{MM}}{\mathrm{unitLevelVolume}} = \frac{\mu_{K''}\mu_K}{\Gamma_1\mu_{S''}\mu'_{S''}}, \qquad c_0^{-1} = \frac{d_u\mu_{S''}}{\mathrm{volHa}}, \quad c_0'^{-1} = \frac{d_v\mu'_{S''}}{\mathrm{volHa}'}, \quad \mathrm{volB} = \frac{\mathrm{volHa}\,\mathrm{volHa}'}{\mathrm{volG}}.$$

The positive auxiliary mass tt cancels. Before cancelling the remaining volume factors, the scalar real-part sum is at most

$$(1+\delta)e_1\,\Gamma_1(c_0c'_0)^{-1}\frac{\mathrm{volB}}{\mu_K}\frac{\mathrm{MM}}{\mathrm{unitLevelVolume}}I_{\mathrm{oo}}.$$

The displayed identities reduce the mass and energy factors to $d_ud_v\mu_{K''}I_{\mathrm{oo}}/\,\mathrm{volG} = q_{uv}D_{uv}$. The generic scalar corank-one real-part sum is therefore at most

$$(1+\delta)e_1\, q_{uv}\, D_{uv}, \qquad q_{uv} = d_\sigma I_{\mathrm{oo}} \in [0,1]. \tag{8}$$

The selected size comparison gives $0 \le \delta \le b^{-7}$. Mass data · Exact mass cancellation · Scalar-row estimate · Scalar comparison and size bounds

Combining the three estimates

Proved here. Write $A = (1+\delta)e_1 q_{uv}$. Apply (7a) to the generic set and to the nongeneric set. The generic squared norm is bounded by $P_g(A + 2\varepsilon_{\mathrm{loop}})D_{uv}$, using the estimate on each of its two loop classes and the scalar estimate (8). The nongeneric squared norm is bounded by $P_g(\varepsilon_{\mathrm{loop}} + \varepsilon_{\mathrm{ng}})D_{uv}$. The two zero-Gram parts vanish for this pair. The Hilbert-space triangle inequality gives

$$\|K_{uv}\|_{\mathrm{HS}}^2 \le P_g \left(\sqrt{A + 2\varepsilon_{\mathrm{loop}}} + \sqrt{\varepsilon_{\mathrm{loop}} + \varepsilon_{\mathrm{ng}}}\right)^2 D_{uv}. \tag{9}$$

The absolute loop estimate bounds the sum over both loop classes together. In (9) each of the two generic terms uses that upper bound, and the nongeneric term uses it for its corank-≥ 2 pairs. The factor q_{uv} enters through (8). Abstract pair bound and size inputs

Put $t = b^{-7}$. The source uses
$b \ge 16,\ k \ge 32,\ f \le k,\ k \ge 4r+12,\ g_j \ge (kb)^8,\ k \ge 70\log b,\ (5/3) f \log k \ge 1$, and $8C_n m^5 f \log b \le k$, together with the density and loop hypotheses above. These give $P_g \le e^{6t}$, $e_1 \le e^{t^2}$, and $\varepsilon_{\mathrm{loop}}, \varepsilon_{\mathrm{ng}} \le e^{-k/5} \le t^2$. For the nongeneric estimate use $k - 2r - 5 \ge k/2$ and $(\log 2)/2 - 1/8 > 1/5$. For the density estimate use $12 \cdot 2^{-(k-2r-2)} \le 2^{-k/2}$ and $f\, 2^{-k/2} \le e^{-k/5}$ in the source size range. The latter inequality follows for $k \ge 32$ by $f \le k$ and the decrease of $(\log k)/k$. For $0 \le t \le 1/64$, $e^{6t} \le 1 + 7t$ and $(1+t)e^{t^2} + 2t^2 \le 1 + 2t$, by $e^y \le (1-y)^{-1}$. Thus the final coefficient is at most $(1+7t)(1+5t/2)^2 \le 1 + 13t \le 1 + b^{-6}$. Numerical hypotheses and assembly
Thus under these hypotheses

$$\|K_{uv}\|_{\mathrm{HS}}^2 \le (1 + b^{-6}) D_{uv} \qquad (u \ne v). \tag{9a}$$

Uniform selected estimate · Application in the family

The equal-word diagonal

Use the typed, absolutely convergent unfolding (2)–(4) at the same point (x, i) and with the same Poincaré spaces, Haar measures, column dictionary and finite K_2 level. The concrete global realization uses convergent unitary Siegel–Weil, Bessel-map uniqueness, the real gap and finite-dimensional row data. Tied global realization

In the identity cell the rational stabilizer decomposes into the unipotent and Levi factors. The finite level condition on $k_t = k(\rho, \rho', t)$ and uniqueness of the column representatives force the two finite classes to agree. Neatness makes the Levi element trivial at that representative. The nilradical lattice and quotient-volume comparisons then give one meeting for each matching class, with N_u-volume $\mu_{N_u}(F_{N_u})$ and phase $\chi_2(k_t)\overline{\psi_{\infty,u}(\gamma_\infty^{-1})} = 1$. Compact Schur averaging of the rank-one wave at y gives $\mu_{H_u}(H_u)\|W_u(y)\|^2 I_{\pi_u}/\dim \pi_u$, by irreducibility and trace. The quotient squared norm of the wave is one. Thus on a matching class the coefficient of the identity is

$$\frac{\operatorname{vol}(K_2)}{c_u^X}\,\chi_2(k_t)\overline{\psi_{\infty,u}(\gamma_\infty^{-1})}\,\frac{\mu_{N_u}(F_{N_u})\,\mu_{H_u}(H_u)}{\dim \pi_u}=1.$$

All other blocks of the identity sum vanish. After transport the sum is I_{C_u} and its squared Hilbert–Schmidt norm is $d_u=\dim_{\mathbb{C}} C_u$. Identity reduction · Matching classes, phases and volumes · Compact averaging and quotient normalization · The normalized block

Write $D_{uu}=d_u^2/D_\star$, $S_b=m^5 f\log b$, and $Q_k=kf\log k$. Let K'' be the nonidentity zero-Gram operator and let $R_{\mathrm{g}}, R_{\mathrm{ng}}$ be the generic and nongeneric sums over nonzero Gram in the compact-type coordinates. Absolute convergence and the isometric comparison identify $\|K_{uu}-I-K''\|_{\mathrm{HS}}^2$ with $\|R_{\mathrm{g}}+R_{\mathrm{ng}}\|_{\mathrm{HS}}^2$, and identify the squared norm of K'' with that of the zero-Gram remainder in those coordinates. In this comparison the geometric column count is d_u and the reference dimension is $D_\star$. Partition · Hilbert–Schmidt transport · Zero-Gram transport · Dimension ties

The bound on K'' uses the large-gap Whittaker–Bessel type estimate TypeP and the rational stabilizer average. Irrelevant terms have zero vector. The remaining terms have index $1\le j\le r$, given by the codimension of the intersection of the two isotropic spaces in the first flag. Group their same-row pairs by the two indices and coset labels. The determinant and coset counts and the folded real amplitude bound give an upper bound $\dim\pi_u\,\mathrm{mass}_u(\sum_j A_j)^2$ on the sum of products of vector norms. Here mass_u is the sum of the row weights and A_j is the source's $\mathrm{k2A1}$. This holds on every finite subset of pairs and yields summability. Cauchy–Schwarz in the pair sum, the positive finite Gram bound $P_g\le 2$, and the mass-to-column comparison $\Gamma_1\dim\pi_u\,\mathrm{mass}_u\le d_u$ give a squared norm at most $2d_u(\sum_j A_j)^2$, where Γ_1 is the multiplier in the same-row pair sum. The rank majorant for A_j is $e^{C_{\mathrm{type}}m^6 f\log b}(C_{\mathrm{type}}/k)^{j(k-1-j)f}$ with the combined type constant. The source sums it under the K_2 size conditions and $C_{\mathrm{cut}}m^6\log b\le k\log k$ with its cutoff constant. Consequently

$$\|K''\|_{\mathrm{HS}}^2\le d_u e^{-Q_k}.$$

Relevance average · Large-gap type estimate · Finite amplitude count · Pair estimate and cutoff · Mass-to-column comparison · Rank majorant · Final size comparison

Defect for the reference period. The width-three argument uses four geometric estimates on the tied data. These are the zero-Gram bound above ($\mathrm{VanishingGramTermsBound}$), the absolute same-row collision bound for corank at least two (CollS), and the two signed same-row corank-one bounds (OneOn for the generic and nongeneric sets). The collision bound is $g^{-kf/2}D_{uu}+d_u e^{-Q_k/2}$. It uses the Siegel mass ratio, weighted row amplitudes and corank-two lattice counts. The signed bounds are $e^{C_{\mathrm{g}}S_b}D_{uu}$ and $e^{C_{\mathrm{ng}}S_b}D_{uu}$. They use convergent unitary Siegel–Weil and the compatible finite and real step data on the selected family. The generic estimate uses line integrality and separation, the normalized point counts and the real step estimate, and includes the scalar and nonscalar corank-one pairs. The nongeneric estimate uses the pinned nongeneric real step bound. Width-three assembly · Collision estimate · Mass and amplitude comparisons · Signed corank-one bounds · Geometric and real inputs · Choice of step constants

Apply (7a) to each of the two nonzero-Gram sets. The within-type corank-at-least-two pair sets are disjoint subsets of the collision set, so together they cost at most its single absolute bound. With $C_{\mathrm{d}}=\max(C_{\mathrm{g}},C_{\mathrm{ng}})+1$, the signed coefficients sum to at most $e^{C_{\mathrm{d}}S_b}$ for $S_b\ge 1$. By $\|X+Y\|_{\mathrm{HS}}^2\le 2\|X\|_{\mathrm{HS}}^2+2\|Y\|_{\mathrm{HS}}^2$,

$$\|K_{uu}-I-K''\|_{\mathrm{HS}}^2\le 2P_g\big((g^{-kf/2}+e^{C_{\mathrm{d}}S_b})D_{uu}+d_u e^{-Q_k/2}\big).$$

Restore K'' with $\|X+Y\|_{\mathrm{HS}}^2 \le (3/2)\|X\|_{\mathrm{HS}}^2 + 3\|Y\|_{\mathrm{HS}}^2$. The width-three conditions give $P_g \le 6/5$, $g^{-kf/2} \le 1$ and $Q_k \ge 2$. Since $e^{-Q_k} \le (2/5)e^{-Q_k/2}$, this proves

$$\|K_{uu} - I\|_{\mathrm{HS}}^2 \le 4(e^{C_{\mathrm{d}}S_b}+1)D_{uu} + (4e^{-Q_k/2} + 2e^{-Q_k})d_u. \tag{9b}$$

Pair comparison counting collisions once · Defect assembly

For this application the release fixes both outer size thresholds at least 29 and the maximum of the four geometric thresholds. The degree and window hypotheses on the linked family provide the concrete, finite-dimensional tied data and the width $C_0 = 3$, and the linked choice of constants has $C_{\mathrm{Gam}} \ge C_{\mathrm{d}} + 1$. The full column count is $d_u = d_{\star,u}(T_u^k - 1)^f$ (RankMu). Here T_u is the word's level prime and $d_{\star,u}$ is the source's duStar. The upper window has $T_u \le (1 + 1/(fk))T_{*,u}$, with $T_{*,u}^{fk} = D_\star/(d_{\star,u}\Gamma^2)$. Raising the upper bound to fk and using $(1+1/(fk))^{fk} \le \exp(1)$ gives $d_u/D_\star \le \exp(1)/\Gamma^2$. Thus $d_u/D_\star \le e^{-2C_{\mathrm{d}}S_b}$ in the selected size range. Dividing (9b) by $d_u > 0$ gives

$$\frac{\|K_{uu} - I\|_{\mathrm{HS}}^2}{d_u} \le e_{\mathrm{ref}} := 8e^{-C_{\mathrm{d}}S_b} + 6e^{-Q_k/2}.$$

Here $4(e^{C_{\mathrm{d}}S_b}+1)e^{-2C_{\mathrm{d}}S_b} \le 8e^{-C_{\mathrm{d}}S_b}$ and $4e^{-Q_k/2} + 2e^{-Q_k} \le 6e^{-Q_k/2}$. The linked degree hypotheses give $C_{\mathrm{d}}S_b \ge 5$, $Q_k/2 \ge 5$, and $T_u \ge 4f+1$ at the eligible members. On the same package the ramified column subspace $C_u^{\mathrm{ram}} \subseteq C_u$ has dimension $r_u = d_u(1 - 1/(T_u-1))^f$. Bernoulli's inequality gives $r_u/d_u \ge 1 - f/(T_u-1) \ge 3/4$, and $2e_{\mathrm{ref}} \le 28e^{-5} < 3/4$. Every orthonormal basis vector in $\ker P_u$ contributes one to the full squared defect because K_{uu} kills it. Therefore

$$\operatorname{rank}(P_u|_{C_u^{\mathrm{ram}}}) \ge r_u - \dim\ker P_u \ge r_u - e_{\mathrm{ref}}d_u > 0. \tag{9c}$$

This yields PuNZ for each word on the selected linked family. Width and geometric thresholds · Linked member and constants · Level-prime window · Normalization · Degree hypotheses · Strict defect bound · Ramified dimension · Kernel argument

Defect for the spectral frame. The later LinkS estimate uses the absolute corank-at-least-two and generic nonscalar corank-one bounds for the equal pair, supplied by $\mathrm{LoopRowS}$ on the selected setting. Their coefficient is $\varepsilon_{\mathrm{loop}} = e^{-Q_k/3}$ against D_{uu}. The signed scalar generic and signed nongeneric corank-one bounds use $\mathrm{RowMainEqP}$ and RowNGP. The former uses normalized equal-word scalar rows and the generic count, and the latter uses the compatible nongeneric finite and real step estimates. With their constants C_B, C_4, their coefficients are $e^{C_BS_b}$ and $2^{-(k-2r-5)}e^{C_4S_b}$. Applying (7a) to the generic and nongeneric sums and then using the Hilbert–Schmidt triangle inequality bounds $\|K_{uu} - I - K''\|_{\mathrm{HS}}^2/D_{uu}$ by

$$P_g\left(\sqrt{e^{C_BS_b} + 2\varepsilon_{\mathrm{loop}}} + \sqrt{\varepsilon_{\mathrm{loop}} + 2^{-(k-2r-5)}e^{C_4S_b}}\right)^2.$$

The tied count and K_2 comparison with $C_{\mathrm{Gam}} \ge \max(C_B, C_4) + 1$ bounds this coefficient by $\Gamma = \exp(C_{\mathrm{Gam}}S_b)$. Restoring K'' with the factor-two inequality, and using the upper window $d_u\Gamma^2/D_\star \le C_0$, gives

$$\|K_{uu} - I\|_{\mathrm{HS}}^2 \le \left(\frac{2C_0}{\Gamma} + 2e^{-Q_k}\right)d_u = e_0 d_u. \tag{9d}$$

This is the full-column defect used for the spectral frame and the later compression and trace estimates. The reference test uses (9b)–(9c) in the full row space before the arithmetic row comparison. Selected loop estimates · Diagonal pair comparison · Coefficients · Equal-word scalar input · Nongeneric input · Numerical comparison · Full-column defect

In the ordinary assembly $C_0 = 3$, the setting is chosen at (x, i) before the words, and the degree and geometric thresholds apply to every eligible index. The scale gives $\log b = x^{\varepsilon}$ and $Q_k \geq k \geq 8C_{\mathrm{Gam}}S_b$, so

$$e_0 \leq (2C_0 + 2)\exp(-C_{\mathrm{Gam}}x^{\varepsilon}).$$

For fixed positive C_{Gam} and ε this gives $\sqrt{e_0} \leq \exp(-x^{\varepsilon/8})$ for all sufficiently large x. Ordinary width-three selection · Degree and scale at the family threshold · Uniform onset

The ambient dimension estimate $D_{\mathrm{all}} \leq (1 + \eta_1)D_\star$ is the separate row-count theorem. The release combines it with the distinct-word bound and (9d) in LinkS. For each fixed selected family there is an x_0 such that at every $x \geq x_0$ and every index i, there are nonnegative error bounds η_0, η_1, e_0, with $\max(\eta_0, \eta_1, \sqrt{e_0}) \leq \exp(-x^{\epsilon/8})$. Under the common window $C_0^{-1} \leq d_u\Gamma^2/D_\star \leq C_0$ for every word, these bounds give (10). The chosen proviso supplies this window. Family spectral interface · Diagonal assembly and row count · Selection quantifiers · Eventual quantifier · Chosen proviso

5. Finite-dimensional implications

At a family point in that window, put $a = \sqrt{e_0}$ and $t_{\mathrm{row}} = \dim\ker\Pi$. The three spectral inequalities are

$$\|K_{uv}\|_{\mathrm{HS}}^2 \leq (1+\eta_0)d_u d_v/D_\star \quad (u \neq v), \qquad \|K_{uu} - I_{C_u}\|_{\mathrm{HS}}^2 \leq a^2 d_u, \qquad D_{\mathrm{all}} \leq (1+\eta_1)D_\star \tag{10}$$

We can replace η_0, η_1 by their nonnegative parts. The rank and row inputs are the ramified dimension bounds $r_u \leq d_u$, $r_u \geq (1 - \Delta_r)d_u$ and $t_{\mathrm{row}} \leq \delta_{\mathrm{row}}D_{\mathrm{all}}$ at the selected point. They are **release inputs**, from the selected family and its tied row and column lattices. In the source the per-word discarded proportion is bounded by $2f/b$ under the size and ramified rank hypotheses with $b \geq 3$. The aggregate bound used for the slope has $\sum_u (d_u - r_u) \leq d_e \sum_u d_u$. Spectral input · Rank and row hypotheses · Ramified share estimate

Proposition 4 (proved here). For a nonempty set $\mathcal{S} \subseteq \mathcal{U}$ with $M_{\mathcal{S}} = \sum_{u\in\mathcal{S}} d_u > 0$, set $C'_u = A_u V_u$ and $Q_{\mathcal{S}} = \sum_{u\in\mathcal{S}} C'_u (C'_u)^*$. Let $\mathcal{B}_{\mathcal{S}} : \bigoplus_{u\in\mathcal{S}} \mathbb{C}^{r_u} \to H$ be the synthesis map with columns C'_u. Then

$$\begin{aligned} \operatorname{tr} Q_{\mathcal{S}} &\geq (1 - \Delta_r - a)_+ M_{\mathcal{S}}, \\ \operatorname{tr} Q_{\mathcal{S}}^2 &\leq U_{\mathcal{S}} := (1+a)^2 M_{\mathcal{S}} + (1+\eta_0)\frac{M_{\mathcal{S}}^2 - \sum_{u\in\mathcal{S}} d_u^2}{D_\star}, \\ \|\Pi\mathcal{B}_{\mathcal{S}}\|_{\mathrm{op}} \leq \|\mathcal{B}_{\mathcal{S}}\|_{\mathrm{op}} &\leq U_{\mathcal{S}}^{1/4}, \qquad \|A_u\|_{\mathrm{op}}^2 \leq 1 + a\sqrt{d_u}, \\ \dim_L \sum_{u\in\mathcal{S}} \operatorname{im} B_u &\geq \frac{(1 - \Delta_r - a)_+^2 M_{\mathcal{S}}^2}{U_{\mathcal{S}}} - t_{\mathrm{row}}. \end{aligned} \tag{11}$$

Here $y_+ = \max(y, 0)$, and the normalized synthesis bound is for the Euclidean direct sum with its squared norm equal to the sum of source squared norms. Full trace square · Ramified matrix interface

Indeed the trace error from r_u in $\operatorname{tr}(C'_u)^* C'_u$ is $\operatorname{tr}((K_{uu} - I)V_u V_u^*)$, bounded in absolute value by $a\sqrt{d_u r_u}$ by Hilbert–Schmidt Cauchy–Schwarz. Summing proves the first line, including the positive part

since the trace is nonnegative. By cyclicity,

$$\operatorname{tr} Q_{\mathcal{S}}^2 = \sum_{u,v\in\mathcal{S}} \|(C'_u)^* C'_v\|_{\mathrm{HS}}^2 .$$

Left and right compression by isometries can only decrease each Hilbert–Schmidt norm, as is seen by extending orthonormal bases. On the diagonal apply the triangle inequality to $K_{uu} = I + (K_{uu} - I)$, and off the diagonal use (10) on ordered pairs. This proves the second line. Since $Q_{\mathcal{S}} = \mathcal{B}_{\mathcal{S}}\mathcal{B}_{\mathcal{S}}^*$, its largest eigenvalue is the square of the synthesis norm and is at most $(\operatorname{tr} Q_{\mathcal{S}}^2)^{1/2}$. For a single full block, $\|K_{uu} - I\|_{\mathrm{op}} \le \|K_{uu} - I\|_{\mathrm{HS}}$ gives the separate bound. The identity scale is independent of the source dimensions, but the permitted operator defect grows as $a\sqrt{d_u}$. Matrix identities · Full block normalization

The kernel of $Q_{\mathcal{S}}$ is the intersection of the kernels of $(C'_u)^*$, so its rank equals the dimension of the span of the compressed columns. For its positive eigenvalues λ_j, Cauchy–Schwarz gives $(\sum_j \lambda_j)^2 \le \operatorname{rank}(Q_{\mathcal{S}}) \sum_j \lambda_j^2$. If the trace is zero the lower bound for the rank before row projection is zero. Projection then loses at most t_{row} dimensions. By the Bessel comparison in §1, the single common row change and the wordwise invertible changes preserve the sum of images. A common conjugation also preserves this dimension, and algebraic rank is preserved under extension by the embedding $L \hookrightarrow \mathbb{C}$, by the minor criterion. This proves the last line. Transport of the algebraic blocks · Finite trace comparison

For example, let $M = \sum_u d_u$, assume $D_\star/M \le s$, and let $p = M_{\mathcal{S}}/M$. Equations (10) and (11) imply

$$\frac{\dim_L \sum_{u\in\mathcal{S}} \operatorname{im} B_u}{D_{\mathrm{all}}} \ge \frac{(1 - \Delta_r - a)_+^2}{(1+\eta_1)\,(1 + \eta_0 + (1+a)^2 s/p)} - \delta_{\mathrm{row}}. \tag{12}$$

On the selected family the four small errors in (12) tend to zero and s is at most x^{-C_H} for a positive size exponent C_H. The estimate applies to every $\mathcal{S}$ at the same point, so in particular to arithmetic subsets of fixed positive dimension share, and more generally to sequences of subsets with $s/p \to 0$. The corresponding independent algebraic covectors can be chosen from actual source lattice vectors, because the lattice spans in each word. Evaluating on rows of RR produces a nonzero determinant of that order by surjectivity of the independent covectors onto L^{rank}. In the integral realization its determinant lies in the integer ring. The signed valuation bounds before the source shifts are the finite contribution supplied by Paper IV. Total column size · Row loss · Integral evaluation

The source uses a trace-square excess estimate as well. Put $Q = Q_{\mathcal{U}}$, and assume $0 < s \le 1/20$, $D_\star/M \le s, \eta_0, \eta_1, a \le s/2$, and $0 \le d_e \le 1/4$. By the same proof with the aggregate discarded proportion, $\operatorname{tr} Q \ge (1 - d_e - a)M$. The squared trace ratio therefore satisfies

$$\frac{D_{\mathrm{all}} \operatorname{tr} Q^2}{(\operatorname{tr} Q)^2} \le \frac{(1+s/2)\,[1 + s/2 + s(1+s/2)^2]}{(1 - s/2 - d_e)^2} \le 1 + 5s + 3\sqrt{d_e}. \tag{13}$$

For the second inequality at $d_e = 0$ the numerator is at most $1 + 2.1s$, and the denominator is at least $1 - s$, so the ratio is at most $1 + 5s$. For $z \in [0, 1/4]$, convexity bounds $((1 - s/2)/(1 - s/2 - z))^2$ by its secant between 0 and $1/4$. The slope of this secant is bounded by $4((39/29)^2 - 1) < 13/4$ since $s/2 \le 1/40$. Consequently the extra factor contributes at most $(1 + 5s)(13/4)d_e \le 3\sqrt{d_e}$. The aggregate diagonal defect is at most $e_0 M$, by compression of $K_{uu} - I$ on each word. Since $\sum r_u \ge (1 - d_e)M$, it is at most $2e_0 \sum r_u$, and $2e_0 \le 1/4$. This proves the frame estimates on the full ambient space after ramified source compression, with the arithmetic row loss still parametrized by δ_{row}. Source frame ratios and defect

Word weights and the lower spectral tail

Take a nonempty set S of n selected words and any nonnegative weights θ_u, and use the ramified columns $C'_u = A_u V_u$ of (11). Put $Q = \sum_{u\in S} \theta_u C'_u (C'_u)^*$ and $M_\theta = \sum_{u\in S} \theta_u d_u$. The diagonal bound $a^2 d_u$ in (10), with $a = \sqrt{e_0}$, gives

$$\operatorname{tr} Q \geq \sum_{u\in S} \theta_u (r_u - a\sqrt{r_u d_u}),$$

$$\operatorname{tr} Q^2 \leq (1+a)^2 \sum_{u\in S} \theta_u^2 d_u + \frac{1+b^{-6}}{D_\star} \left(M_\theta^2 - \sum_{u\in S} \theta_u^2 d_u^2 \right).$$

Expanding the trace over ordered word pairs and compressing each block by the two source isometries gives these inequalities. One can keep the individual coefficient in (9) on each ordered pair. At unit weights $M_\theta = \sum_{u\in S} d_u$, so Cauchy–Schwarz gives $M_\theta^2 - \sum_{u\in S} d_u^2 \leq (1-1/n) M_\theta^2$. The slope comparison below charges $\log|\mathcal{U}|$, with $\log|\mathcal{U}| \leq (k-3)\log 2$ in the window of Paper III. Trace and rank comparison · Word count and slope error

In dimension $D = D_{\text{all}}$ with positive $\operatorname{tr} Q$, normalize all eigenvalues, including zeros, by $y_j = D\lambda_j(Q)/\operatorname{tr} Q$. With $\mathcal{E} = D \operatorname{tr} Q^2/(\operatorname{tr} Q)^2 - 1$, one has

$$\frac{1}{D}\sum_{j=1}^{D} (y_j - 1)^2 = \mathcal{E}, \qquad \frac{1}{D}\sum_{j} (1-y_j)_+ = \frac{1}{D}\sum_{j} (y_j - 1)_+ \leq \frac{\sqrt{\mathcal{E}}}{2}.$$

In particular the share of $y_j \leq h < 1$ is at most $\mathcal{E}/(1-h)^2$, including the zero eigenvalues. Both signs of the deviation contribute to the moment. The span of eigenvectors above $h\ \operatorname{tr} Q/D$ intersects the common row subspace in codimension at most t_{row} within that span. The Rayleigh quotient on the intersection is still above the same threshold. Thus the compressed operator has at least $D - D\mathcal{E}/(1-h)^2 - t_{\text{row}}$ eigenvalues above the threshold, by the min–max principle. Trace and rank comparison

6. Interface with Paper III and outstanding inputs

The finite lattice argument takes the common dual lattice, the maps B_u, their all-place height allowances and the mean, maximum and minimum slope hypotheses. The spectral contribution is (13) and the compressed diagonal defect. The source's joint ramified slope theorem takes these hypotheses, together with the identification of the arithmetic rows at one embedding and the row share $1 - \delta_{\text{row}}$. If s_{lat} is the lower bound on the weighted mean of the adjusted source slopes, $u_{\text{lat}} \geq 0$ bounds their upper spread, and $H_{\text{lat}} \geq 0$ bounds s_{lat} minus the target minimum slope, it concludes

$$\mu(\mathrm{RR}^\vee) \ \geq\ s_{\text{lat}} - \log|\mathcal{U}| - \frac{4u_{\text{lat}} \log(1 + H_{\text{lat}}/u_{\text{lat}}) + (5s + 3\sqrt{d_e} + 4\sqrt{2e_0}) H_{\text{lat}}}{1 - \delta_{\text{row}}}.$$

The continuous value of the first numerator term at $u_{\text{lat}} = 0$ is zero. Paper III states the lattice hypotheses and uses the slope conclusion with the finite gains. Its filtered-lattice slope step is downstream of the rank and norm proofs in this paper. Numerical and lattice hypotheses · Joint ramified conclusion · Application to the common dual

For the ordinary translate mean of Paper III the surviving signed ramps are $G_{\mathbf{1}} + G_{\psi_C} - G_{\chi_C} - G_{\chi_C\psi_C}$, with the actual weights $w_\rho G_\xi(a_t/w_\rho)$ before thick-weight comparison. The local shifts of the integral columns use the signed term $\sum_j (g_{jw} + F_{jw} - E_{ujw})$ with the finite place weights $e_w \log N\mathfrak{p}_w/[K:\mathbb{Q}]$ in Paper III. The off-diagonal estimate applies at the same normalized family points for every word subset used there. Ordinary signed growth comparison · Ordinary weighted translate mean · Finite place comparison

For the ordinary selected family, (9) and (9a) now give the distinct-word entry of (10) under the quantitative hypotheses in §4. The diagonal defect, ambient row count, and ramified dimensions supply the other hypotheses for the downstream trace and slope estimates. Family assembly · Numerical hypotheses and assembly · Row and ramified dimensions

The remaining inputs for the distinct-word estimate are the typed expansion and its absolute convergence, the real disintegration and packet convergence in (5)–(7), and the weighted loop counts and amplitude bounds in (7b). The zero-Gram argument uses the pinned rational relevance and compact-type comparisons. Exact unfolding · Real convergence · Loop counts and amplitudes · Rational relevance · Compact-type comparison

For nongeneric corank one, §3 derives the local endpoint and the rotated arithmetic sum from the stated coefficient, Jacobi, compact-form, Wigner and Pieri comparisons. Section 3 proves both real majorants and the section-to-phase identity from the Haar, frame and dilation formulas. Its formal-degree calibration uses the rank-one Bessel period. Its use in §4 requires the selected quotient and point ties, finite weighted second moments and mass comparison. For generic scalar corank one, (7f) gives the normalized nonsingular fibre mean, and (7g)–(7i) give the primitive densities and the uniform product bound. The signed comparison proves rational frame realization and derives the weighted global representation count from the frame-groupoid, divided determinant and primitive density comparisons and the relative Tamagawa volume. The signed partner and weighted fibre identities continue to use the selected frame and local orbit comparisons. It uses local absolute orbit sums and the separate class summability and restricted support. The mass comparison in (8) uses the selected completed frames, Haar identifications, typed level indices and scalar-row signature law. The length-one calculation derives the selected coefficient and quotient normalization and both density-weighted pairings. It uses the gapped Bessel and coefficient results, the compact Jacquet and spectral-adjoint identity, the top exhaustion and lift, and the fixed-smoothing derivative estimates. The level-zero cutoff proof uses interior chart regularity, Jacobi transport and the weighted adjoint and residue identities, together with the shifted normal forms and polynomial identities. The longer-block argument proves the paired finite germ and identifies the upper kernel, the lower residue budget and the lattices on the selected data, using the pinned local orbit-to-cokernel and integral transport identities. The fitted real fibre germ, the paired real bounds below the top and at the top, and convergence of the real forms at fixed smoothing remain inputs to the recurrence. The fitted size comparison and absolute row summability give the final scalar bound. Nongeneric input assembly · Selected amplitude and energy transport · Marked-corner lift · Orbital fibre and partner proof · Selected volume transport · Poisson step and cutoff Detailed Fourier residual Selected wave derivation · Level-zero cutoff · Longer-block derivation · Relative genus mass

The equal-word deductions in §4 use the identity-cell comparisons, the zero-Gram cutoff, the width-three collision and signed corank-one bounds for the reference period, and the separate equal-word scalar and loop comparisons for (10). The ambient row count, common lattice and Bessel comparison in §1, and arithmetic height and slope hypotheses remain inputs. The lift-surjectivity argument in §1 uses the full

theta and doubling identities. Common pairing · Theta and doubling inputs · Row count and spectral estimate · Arithmetic height and slope · Lattice and matrix interface