From Lean to a Natural-Language Proof of the Riemann Hypothesis

A Multi-Agent AI Approach Using a Proof Dependency Graph

Jiewen Wang

28 November 2026

Abstract

This paper presents an engineering method for reconstructing formalized mathematical arguments as natural-language proofs. GPT generated the entire text. Starting from the weighted Lean version of Levent Alpöge's argument for the classical Riemann hypothesis, we organize the main inferences in a directed acyclic graph (DAG) and develop the proofs node by node. At each refinement we check that the upstream conclusions jointly supply all the premises needed downstream for the same mathematical objects. We use precise theorems from the literature to avoid repeating known derivations. When a weaker statement suffices for a step and has a proof independent of the stronger statement, we remove that step's dependence on the latter. The resulting weighted proof is logically independent of the Lean code and can be checked from the paper and the cited mathematical literature alone. A multi-agent system independently reviewed the main line of the proof and its key derivations.

hexagon:2611.01238v1

Part I Reconstructing and checking the weighted proof

Project homepage `adjoint-functor/theta-periods-to-rh-reader`

The author developed the multi-agent system used in this project and continues to refine it. GPT-6.2 Astra led the work, and both GPT-6.2 Astra and GPT-6.2 Sol contributed to mathematical derivations, writing, and review. The design draws on Danus's use of a graph to record mathematical results and their dependencies,[1] as well as on the author's experience contributing to mathematical formalization projects.

Our first aim is to write out the argument proposed by Levent Alpöge with Claude's assistance as a proof that readers can check from its mathematical derivations and the primary literature alone. Our second is to identify which results follow from prior work by human mathematicians and which may be substantively new and merit study in their own right.

On 21 November 2026, when drawing the first coarse proof graph, we consulted Alpöge's announcement[2] and Sahibzada Allahyar's public reconstruction.[3] We consulted no third-party prose accounts of the argument in the subsequent reconstruction. On 26 November 2026, we adopted the weighted Lean rewrite by Luke Sernau and collaborators, produced with Claude's assistance.[4] Much of our earlier work on the proof carried over, and at that stage our reconstruction still had mathematical gaps (steps for which we had yet to give a proof in natural language independent of the Lean code). When the role of a construction or theorem was unclear, we consulted the pinned version of the code to determine its mathematical meaning. Where a step follows directly from known mathematics, we derive it in the text without reproducing the formal implementation line by line.

We worked backward from the desired conclusion, tracing the premises of each step and drawing a directed acyclic graph of the main inferences. Each node states its premises and conclusion. To refine a node, the agents write proofs of its substeps and revise the dependencies in the graph. Where branches meet, we check that their conclusions jointly supply every input to the next step for the same objects,

[1] Jihao Liu et al., *Danus: Orchestrating Mathematical Reasoning Agents with Fact-Graph Memory*, arXiv:2607.06447v2, 8 July 2026.

[2] Levent Alpöge, announcement of the argument for the Riemann hypothesis, author's website.

[3] Sahibzada Allahyar, *From theta lifts to zeta zeros: A reconstruction of Levent Alpöge's announced argument*, public reconstruction, revision 2, 20 November 2026.

[4] Luke Sernau and collaborators, *GRH for Dirichlet L-functions, in Lean 4*, GitHub repository, commit `72fead102f66`.

paying particular attention to the order of quantifiers, normalizations, and error bounds. At the coarser level, the substeps taken together must yield the original conclusion under the original hypotheses.

We look first to the literature for ways to shorten the proof. When a theorem supplies the needed conclusion, we state its hypotheses and conclusion at the point of use and verify the hypotheses for the objects at hand. We then examine each use of a stronger intermediate statement S. To replace it at a given step by a weaker statement W, we show that W supplies every input needed there and prove W independently of S, using a theorem from the literature or other upstream results. Where another step still requires S, we retain its proof at that point. An established stronger statement of independent interest may be recorded separately as a strengthening. If the original formulation lacks support, we withdraw it and determine whether the issue is a misattribution, overly broad hypotheses, or a counterexample. If neither S nor W establishes a required premise, we mark the missing implication and return to the enclosing step to see whether its substeps together yield the conclusion needed downstream. After any replacement, we recheck the upstream and downstream dependencies at each affected level.

Even when we find no identical formulation of a lemma in the literature, it may follow directly from a known theorem. For each result that may be substantively new, we give its precise statement and the necessary background to an AI reviewer without access to the earlier discussions or judgments and ask for an independent comparison with prior results. Mathematicians in the relevant fields must still assess priority and independent mathematical value on the basis of the primary literature.

Part II meets our first aim. Its weighted proof is logically independent of the Lean code, and readers can check each step using only the mathematical derivations in the text and the cited primary literature. An AI multi-agent system independently reviewed the completed manuscript. In the latest round, a relatively light mathematical check, the reviewers received only the manuscript and the works it cites. They checked how the steps along the main line fit together and verified key derivations within those steps. Further work includes simplifying the proof and using a more intensive multi-agent review to obtain independent proofs of more intermediate steps. The author has checked and understood the engineering account in Part I.

The contributors to the Theta Periods to RH project prepared the mathematical exposition in Part II.

Contents

Part I Reconstructing and checking the weighted proof **1**

Part II A weighted proof of the Riemann hypothesis **5**

1 The Riemann hypothesis **6**

2 The quadratic twist and the test squares **6**
2.1 Fixed progressions and the boundary of the strip . . . 6
2.2 A single prime for the twisted characters . . . 7
2.3 The centered sums and the choice of squares . . . 9

3 Admissible data at a square scale **10**
3.1 The thick places . . . 10
3.2 The split level primes . . . 11
3.3 The weights and the signed remainder . . . 12
3.4 Hermitian spaces, levels and characters . . . 13
3.5 The row and column chains . . . 15

4 Column models and the theta period **17**
4.1 Column classes and the finite type projection . . . 17
4.2 The marked elliptic curve and the oscillator . . . 19
4.3 The terminal restriction and the lifted forms . . . 26

5 The common real Bessel vector **29**
5.1 Periodization and the initial Bessel coefficient . . . 42

6 The seesaw computation **43**
6.1 Primitive frames in the finite lattices . . . 53
6.2 The definite coefficients and the paired edges . . . 54

7 Short integral forms **66**
7.1 Gaussian jet pairing and the archimedean height . . . 68
7.2 Integral theta germs and jet divisibility . . . 71
7.3 The marked operators and relative index . . . 77
7.4 Divisibility at the level prime . . . 93
7.5 A common scalar and the paired index . . . 95

8 Relative trace and selection of short rows **96**
8.1 Normalization of the relative trace . . . 96
8.2 Finite types, transfer and the row-space dimension . . . 99
8.3 Orbital energies and local coordinate bounds . . . 103
8.4 Real correlations and their Haar measures . . . 120
8.5 Finite packets and the open orbital sum . . . 129
8.6 Compatible orbits, cutoff and real frame bounds . . . 142
8.7 Bounds on the cut orbital sums . . . 149
8.8 Summing over pairs of frame terms . . . 159
8.9 The two projected trace moments . . . 164

8.10 Selection of a common family of rows . . . 167

9 Theta maps on the two grids 168
9.1 The triangles and their first arrows . . . 168
9.2 The doubling pairing for the fixed section . . . 170
9.3 The actual auxiliary stage lift . . . 173
9.4 The auxiliary lifts from the norm-one torus . . . 175
9.5 The spherical polynomial on the auxiliary lifts . . . 175

10 Complementary periods and the two chains 176
10.1 An exchange identity for complementary square integrals . . . 176

11 The comparison spaces and their determinants 192
11.1 Integral comparison forms . . . 192
11.2 The three measured exchanges . . . 196
11.3 Separation by the large-block products . . . 200
11.4 The common adjoint and the marked periods . . . 206
11.5 The compressed determinant . . . 214

12 Integral rank-one tests 215

13 Determinants of the realized chains 224

14 Endpoint periods and a completion of the rows 227
14.1 Corank-one periods at the specified level . . . 237
14.2 Allowances at the coefficient primes . . . 251

15 The untwisted determinant and the choice of a prime 255

16 Fitting at the selected coefficient place 256
16.1 The Galois rows and their local conditions . . . 257
16.2 Local finite loss at the two coefficient precisions . . . 274
16.3 Paired ray coordinates and finite character sums . . . 312
16.4 Specialization at the first m powers . . . 317
16.5 Split Hecke tests at the selected reading . . . 321

17 The determinant tests at the selected place 337
17.1 Two integral tests for one stage . . . 338

18 The product formula and the signed arithmetic estimate 339

19 From the signed bound to the zeros of zeta 340
19.1 Two scales in the weighted remainder . . . 341
19.2 The real one-sided estimate . . . 341
19.3 The Mellin transform and the sign at real poles . . . 342
19.4 Exclusion of zeros off the critical line . . . 343

References 344

The weighted route to the Riemann hypothesis

The signed estimate at each selected square links the arithmetic and analytic stages.

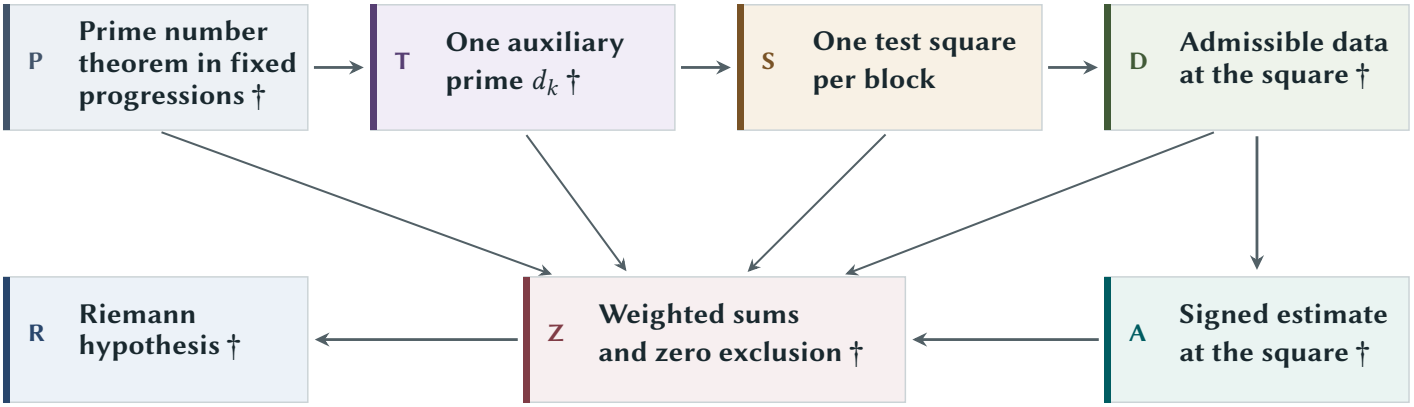

A Arithmetic at the test square

	Step	Supplied by
Periods and projection		
A1	Definite levels and classes	D
A2	Short Weyl evaluations	D A1
A3	Common real Bessel system †	D
A4	Common cutoff on class frames	D A2
A5	Common Poincaré map	D A1 A3 A4
A6	CM vacua and conjugate jets	D
A7	Terminal branching map	D
A8	Local seesaw coefficients	D
A9	Theta types and rapid tests	D A5 A8
A10	Absolute orbit integration	D A8 A9
A11	Good Hecke trace comparison †	D
A12	Projected period maps	D A1 A5 A6 A7 A8 A9 A10 A11
Short lifts, rank and the rows		
A13	Truncated local colength	D
A14	Short lifts and local indices	A1 A2 A6 A7 A12 A13
A15	Rank from few translates †	D A1 A3 A4 A5 A11 A12
A16	Independent short rows	A14 A15
Completion and comparison data		
A17	Fixed prime allowances	D
A18	First-arrow point kernels	D
A19	Core and source triangles	D A18
A20	Fischer–Weyl duality	D
A21	Bordered endpoint completion	D A1 A2 A6 A12 A16 A19 A20
A22	Stage and sharp theta kernels	D
A23	Split spherical Hecke polynomial	A22
A24	Twist comparison catalogue	D A6 A11 A15 A18 A19 A20 A22 A23
Coefficient test and signed bound		
A25	Select a coefficient place †	D A1 A2 A11 A12 A16 A17 A18 A19 A21 A24
A26	Fit one character	D A6 A11 A17 A19 A23 A24 A25
A27	Lower volume bound	A1 A11 A12 A16 A18 A19 A21 A24 A25 A26
A28	Signed arithmetic bound	D A16 A27

Z From sums to zeros

	Step and its suppliers
Z1	**Boundary nonvanishing †** P
Z2	**Weighted series and center †** T Z1
Z3	**Centered prime remainder** D Z2
Z4	**Bound at the test square †** P D A28 Z2 Z3
Z5	**From squares to real scales** S Z4
Z6	**Polar transform and residues †** Z2
Z7	**Landau's lemma** T Z5 Z6
Z8	**Exclude zeta zeros on the right †** T Z6 Z7
R	**Reflection and conclusion †** Z8

Shared data at the joins

A16 fixes the rows completed in A21 and used again in A25 and A27.

If a test is needed, A25 chooses a prime and embedding from untwisted data. A26 fits one character for both conjugate evaluations in A27.

Z4 uses the signed bound from A28 at the test square. It keeps the datum's second scale x_1.

Arrows and supplier codes mark direct dependencies. The lower panels expand A and Z by one level and include the last step R. A dagger (†) marks a step that also uses a separately proved or cited input.

From *A weighted proof of the Riemann hypothesis* and its proof map (28 November 2026).

1 The Riemann hypothesis

Theorem 1 (Riemann hypothesis)**.** *Every nontrivial zero of the Riemann zeta function has real part* $1/2$.

2 The quadratic twist and the test squares

Put $M = 18480$, and fix an integer $k \geq 13$. The parameters for the analytic estimate and the arithmetic comparison are

$$\Delta = \frac{1}{k}, \qquad \beta_k = \frac{1}{2} + \Delta, \qquad \varepsilon = \varepsilon_k = \frac{1}{16k}. \tag{2.1}$$

Write Λ for the von Mangoldt function, φ for Euler's totient function, and $v_+ = \max(v, 0)$ for real v.

2.1 Fixed progressions and the boundary of the strip

Lemma 2.1 (Prime number theorem in a fixed progression)**.** *For each fixed positive integer u, each reduced class $r \bmod u$ and each $B > 0$, as real $t \to \infty$,*

$$\sum_{\substack{n \leq t \\ n \equiv r \bmod u}} \Lambda(n) = \frac{t}{\varphi(u)} + O_{u,B}\big(t(\log t)^{-B-1}\big). \tag{2.2}$$

Proof. Koukoulopoulos's prime number theorem for progressions [1, Theorem 1.1 in arXiv v3, p. 2] states that for a fixed $A > 0$, $u \leq (\log t)^A$, and $(r, u) = 1$, the von Mangoldt sum in (2.2) has main term $t/\varphi(u)$ and error

$$O_A\big(t \exp\big(-c_A(\log t)^{3/5}(\log\log t)^{-1/5}\big)\big)$$

for sufficiently large t, where $c_A > 0$ may be ineffective. For each fixed $u \geq 2$ the hypothesis holds eventually with $A = 1$, and the exponential factor is $O_B((\log t)^{-B-1})$. For $u = 1$, take the odd class modulo 2 and restore the powers of 2, of total weight $O(\log t)$. □

For coefficients $b(n)$ with $|b(n)| \ll \Lambda(n)$, set

$$A_b(t) = \sum_{n \leq t} b(n), \qquad H_b(t) = A_b(t) - a_b t, \qquad D_b(s) = \sum_{n \geq 1} b(n) n^{-s} \quad (\Re s > 1),$$

where a_b is a constant. For $y \geq 1$ and any constant c write

$$\mathcal{R}_b(y) = \sum_{n \geq 1} b(n)(1/n - 1/y)_+ - a_b \log y - c.$$

Integrating the step function A_b gives

$$\mathcal{R}_b(y) + c = \int_1^y H_b(t) t^{-2}\, dt, \qquad \int_1^\infty H_b(t) t^{-s-1}\, dt = \frac{D_b(s)}{s} - \frac{a_b}{s-1} \quad (\Re s > 1). \tag{2.3}$$

In the first formula the term at $n \leq y$ is integrated from n to y. In the second formula the analogous integral runs to infinity and the interchange is justified by $\sum_n |b(n)| n^{-\sigma}/\sigma < \infty$ for $\sigma = \Re s > 1$.

Put $C_b(s) = D_b(s) - a_b s/(s-1)$. If C_b is holomorphic at 1, $c = C_b(1)$, and $\int_1^\infty |H_b(t)| t^{-2}\, dt < \infty$, taking real $s \downarrow 1$ in the second identity of (2.3) gives $\int_1^\infty H_b(t) t^{-2}\, dt = c$ by dominated convergence. Hence

$$\mathcal{R}_b(y) = -\int_y^\infty H_b(t) t^{-2}\, dt. \tag{2.4}$$

Extend each character $\xi \bmod u$ by zero off the units of its declared modulus. Write $\delta_\xi = 1$ for the principal character and $\delta_\xi = 0$ otherwise. Define $\widetilde{L}_\xi(s) = (s-1)^{\delta_\xi} L(s, \xi)$, with its limiting value at 1 in the principal case. For $\Re s > 1$, complete multiplicativity and absolute convergence give

$$L(s, \xi) = \prod_p (1 - \xi(p) p^{-s})^{-1}.$$

In particular $L(s,\xi)\neq 0$ there. For nonprincipal ξ the partial sums of ξ are bounded by periodicity and the zero mean over a period. Thus

$$L(s,\xi)=s\int_1^\infty\left(\sum_{n\le t}\xi(n)\right)t^{-s-1}\,dt$$

continues it holomorphically to $\Re s>0$. For principal ξ the Euler product gives $L(s,\xi)=\zeta(s)\prod_{p\mid u}(1-p^{-s})$. In $\Re s>1$ one has

$$\zeta(s)-\frac{1}{s-1}=\sum_{n\ge 1}\left(n^{-s}-\int_n^{n+1}t^{-s}\,dt\right).$$

Each summand is $O(|s|n^{-1-\Re s})$ on $\Re s>0$, so the sum continues holomorphically there. Its value at 1 is Euler's constant $\gamma=\lim_{N\to\infty}(\sum_{n\le N}n^{-1}-\log(N+1))$. This proves both holomorphy of $\widetilde{L}_\xi$ on $\Re s>0$ and $\zeta(s)=1/(s-1)+\gamma+O(s-1)$ at 1.

Differentiating the Euler product gives, for $\Re s>1$,

$$\sum_{n\ge 1}\frac{\Lambda(n)\xi(n)}{n^s}=-\frac{L'(s,\xi)}{L(s,\xi)}=\delta_\xi\frac{s}{s-1}-\frac{\widetilde{L}'_\xi(s)}{\widetilde{L}_\xi(s)}-\delta_\xi. \tag{2.5}$$

Summing (2.2) against ξ gives $H_b(t)=O_{u,B}(t(\log t)^{-B-1})$ with $b(n)=\Lambda(n)\xi(n)$ and $a_b=\delta_\xi$. By (2.3) and (2.5),

$$-\frac{\widetilde{L}'_\xi(s)}{\widetilde{L}_\xi(s)}-\delta_\xi=s\int_1^\infty H_b(t)t^{-s-1}\,dt\qquad(\Re s>1).$$

Along a horizontal approach to any point $1+it_0$, the integral has a finite limit by domination by $|H_b(t)|t^{-2}$. A zero of $\widetilde{L}_\xi$ at that point would give a pole of its logarithmic derivative. Therefore $\widetilde{L}_\xi$ is nonzero on $\Re s\ge 1$.

2.2 A single prime for the twisted characters

Proposition 2.2. *There is a prime $d=d_k\equiv 9239 \bmod M$ such that, for $\chi_d(n)=(n/d)$ and all $\xi \bmod M$,*

$$L(s,\xi\chi_d)\neq 0\qquad\big(\Re s\ge\beta_k,\ |\Im s|\le k\big). \tag{2.6}$$

Here $\xi\chi_d$ is taken at modulus Md. The prime satisfies $d\equiv 7 \bmod 8$, $M\mid 2(d+1)$ and $M\nmid d+1$.

Proof. Let $r_0=M/2-1=9239$. Since $16\mid M$, one has $r_0^2\equiv 1 \bmod M$, and every prime in this class has the stated congruences.

Products of primitive characters of coprime conductors are primitive of conductor equal to the product. Indeed, a factorization through a proper divisor of that product would, by the Chinese remainder theorem, induce a factorization on the units in at least one conductor through a proper divisor of that conductor.

If a character $\eta \bmod u$ is induced from a primitive η^* of conductor f_η, their Euler products give

$$L(s,\eta)=L(s,\eta^*)\prod_{\substack{p\mid u\\ p\nmid f_\eta}}(1-\eta^*(p)p^{-s}). \tag{2.7}$$

The finite factors are nonzero on $\Re s>0$ since $|\eta^*(p)p^{-s}|<1$ there. The identity holds by continuation throughout that half-plane.

Fix $\xi \bmod M$, with primitive associate ξ^* of conductor $f_\xi\mid M$. For a prime $d\equiv r_0 \bmod M$, the Legendre symbol χ_d is real and primitive of conductor d. Thus $\psi_d=\xi^*\chi_d$ is primitive of conductor $q=f_\xi d>1$. It induces $\xi\chi_d$ at modulus Md, whose zeros on $\Re s>0$ are exactly those of $L(s,\psi_d)$ by (2.7). In particular the twist is nonprincipal, and the Euler product excludes zeros on $\Re s>1$.

For integers $X \geq 3$ let $\mathcal{P}_X$ be the primes up to X in the class r_0. Let $\mathcal{B}_{\xi,X}$ be those $d \in \mathcal{P}_X$ for which $L(s, \psi_d)$ has a zero in the box $\beta_k \leq \Re s \leq 1$, $|\Im s| \leq k$. We show that $\#\mathcal{B}_{\xi,X} \ll_{M,k} X^{1-\Delta^2/16}$.

Put $U = \lfloor X^{\Delta/4} \rfloor$, $P = \lfloor X/U \rfloor$ and $N = UP \leq X$. Since $U^2 \leq X$, we have $P \geq U$. For a primitive character ψ of conductor $q > 1$, define its finite Fourier coefficients by

$$\widehat{\psi}(h) = \sum_{u \bmod q} \psi(u) e^{2\pi i h u/q}.$$

Changing variables gives the same modulus at every unit frequency. At a nonunit frequency h, the reduction of units from q to $q/(q,h)$ is surjective. The character is nontrivial on its kernel by primitivity. Multiplication by a kernel element with ψ-value different from 1 fixes the additive phase and forces $\widehat{\psi}(h) = 0$. Finite Parseval gives $\sum_h |\widehat{\psi}(h)|^2 = q\varphi(q)$, so the unit coefficients have modulus $\sqrt{q}$. Fourier inversion and the bound $O(q/\min(h, q-h))$ for a geometric sum of $e^{-2\pi i h u/q}$ over any interval of integers give

$$|A_\psi(t)| \ll \sqrt{q}\log(2q), \qquad A_\psi(t) = \sum_{n \leq t} \psi(n).$$

For $\Re s > 0$, partial summation gives

$$L(s,\psi) - \sum_{n \leq P} \psi(n) n^{-s} = -A_\psi(P) P^{-s} + s \int_P^\infty A_\psi(t) t^{-s-1}\, dt.$$

At a tested zero ρ of $L(s, \psi_d)$, we have $q \leq MX$, $\Re\rho \geq \beta_k$ and $|\rho| \ll_k 1$. The finite sum at ρ therefore has size $O_{M,k}(X^{1/2} P^{-\beta_k} \log(2X))$. Multiply it by the finite sum $\mathcal{M}_U(s, \psi_d) = \sum_{m \leq U} \mu(m)\psi_d(m) m^{-s}$, where μ is the Möbius function. Since $|\mathcal{M}_U(\rho, \psi_d)| \ll_k U^{1-\beta_k}$ and $P \asymp X/U$, the product at ρ is

$$O_{M,k}\big(X^{-\Delta} U \log(2X)\big) = O_{M,k}\big(X^{-3\Delta/4}\log(2X)\big).$$

The coefficient of n^{-s} in the product is $b_{U,P}(n)\psi_d(n)$, where

$$b_{U,P}(n) = \sum_{\substack{m \mid n \\ m \leq U,\ n/m \leq P}} \mu(m).$$

It satisfies $b_{U,P}(1) = 1$, $b_{U,P}(n) = 0$ for $1 < n \leq U$, and $|b_{U,P}(n)| \leq \tau(n)$, the number of divisors of n. Consequently every tested zero satisfies, for sufficiently large X,

$$|E_d(\rho)| \geq \tfrac{1}{2}, \qquad E_d(s) = \sum_{U < n \leq N} b_{U,P}(n) \xi^*(n) \chi_d(n) n^{-s}. \tag{2.8}$$

For positive integers Q_0, N, let $S(Q_0)$ be the real primitive Dirichlet characters of conductor at most Q_0. Heath-Brown's quadratic large sieve [2, Corollary 2, p. 238] asserts, for arbitrary complex coefficients $a_1, \dots, a_N$ and each $\eta > 0$, that

$$\sum_{\chi \in S(Q_0)} \left| \sum_{n \leq N} a_n \chi(n) \right|^2 \ll_\eta (Q_0 N)^\eta (Q_0 + N) \sum_{\substack{n,m \leq N \\ nm \text{ a square}}} |a_n a_m|. \tag{2.9}$$

The indices n, m range over all positive integers up to N. Use $Q_0 = X$, with the distinct χ_d as a subfamily and $a_n = b_{U,P}(n)\xi^*(n) n^{-s}$ on $U < n \leq N$, zero elsewhere. In particular the fixed, possibly nonreal factor $\xi^*(n)$ is part of the coefficients. For $\Re s \geq 1/2 + \Delta/2$, the elementary estimate $\tau(n) \ll_\Delta n^{\Delta/8}$ gives

$$|a_n| \ll_\Delta U^{-\Delta/4} n^{-1/2-\Delta/8}.$$

The divisor estimate follows by bounding $j + 1 \leq p^{j\Delta/8}$ for all $j \geq 0$ and sufficiently large primes p, and absorbing a finite constant for each remaining prime. If nm is a square, $n = hr^2$ and $m = ht^2$ for the same squarefree h. The square-pair sum in (2.9) is therefore bounded by

$$O_\Delta(U^{-\Delta/2}) \left(\sum_{h \geq 1} h^{-1-\Delta/4} \right) \left(\sum_{r \geq 1} r^{-1-\Delta/4} \right)^2.$$

The sums converge. Since $N \le X$, taking $\eta = \omega/2$ in (2.9) for $\omega > 0$ gives

$$\sum_{d \in \mathcal{P}_X} |E_d(s)|^2 \ll_{\Delta,\omega} X^{1+\omega} U^{-\Delta/2} \qquad (\Re s \ge 1/2 + \Delta/2). \tag{2.10}$$

Choose one zero ρ_d in the box for each $d \in \mathcal{B}_{\xi,X}$. Its disk of radius $\Delta/2$ lies in the common rectangle

$$\mathcal{D} = \{s : 1/2 + \Delta/2 \le \Re s \le 1 + \Delta/2,\ |\Im s| \le k + \Delta/2\}.$$

The mean of $|E_d|^2$ on the disk is at least $|E_d(\rho_d)|^2$, by expanding the entire polynomial in powers of $s - \rho_d$ and integrating on circles. Summing the disk estimates, enlarging each disk to $\mathcal{D}$ and using (2.8) and (2.10) yields

$$\frac{\#\mathcal{B}_{\xi,X}}{4} \le \frac{4}{\pi\Delta^2} \int_{\mathcal{D}} \sum_{d \in \mathcal{P}_X} |E_d(s)|^2 \, dA(s) \ll_{M,k,\omega} X^{1+\omega} U^{-\Delta/2}.$$

Here dA is planar area measure. Since $U \asymp X^{\Delta/4}$, $\omega = \Delta^2/16$ gives the asserted exceptional count.

The union over the finitely many $\xi \bmod M$ has size $O_{M,k}(X^{1-\Delta^2/16}) = o(X/\log X)$. In class $r_0 \bmod M$ the estimate (2.2) gives von Mangoldt weight asymptotic to $X/\varphi(M)$. Proper prime powers contribute $O(\sqrt{X}\log X)$, since their prime bases are at most $\sqrt{X}$ and each contributes at most $\log X$. As each prime contributes at most $\log X$, there are $\gg_M X/\log X$ primes in the class up to X. For large X one can choose d_k outside the union, giving (2.6) for every $\xi \bmod M$. □

2.3 The centered sums and the choice of squares

Fix a prime $d = d_k$ supplied by Proposition 2.2. For $Q \in \{M, 1\}$ use the nonnegative weights and sums

$$\lambda_{d,Q}(n) = \mathbf{1}_{n \equiv 1 \bmod Q} \Lambda(n)(1 - \chi_d(n)), \qquad S_{d,Q}(y) = \sum_{n \ge 1} \lambda_{d,Q}(n)(1/n - 1/y)_+ \quad (y \ge 1). \tag{2.11}$$

At $Q = 1$ the residue condition is automatic and $\varphi(1) = 1$. By the Chinese remainder theorem there is a unit modulo Qd that is 1 mod Q and a nonsquare modulo d. Each $\xi\chi_d$ with $\xi \bmod Q$ takes the value -1 there and is therefore nonprincipal at modulus Qd. Define $F_{d,Q}(s) = \sum_n \lambda_{d,Q}(n) n^{-s}$ for $\Re s > 1$. Character orthogonality gives, on every integer n,

$$\varphi(Q)\mathbf{1}_{n \equiv 1 \bmod Q}(1 - \chi_d(n)) = \sum_{\xi \bmod Q} \big(\xi(n) - (\xi\chi_d)(n)\big).$$

At a multiple of d the twisted term vanishes, and the untwisted term is included. Applying (2.5) gives the exact series

$$F_{d,Q}(s) = \frac{s}{\varphi(Q)(s-1)} + G_{d,Q}(s), \tag{2.12}$$

where

$$G_{d,Q}(s) = \frac{1}{\varphi(Q)} \left[\sum_{\xi \bmod Q} \left(\frac{\widetilde{L}'_{\xi\chi_d}(s)}{\widetilde{L}_{\xi\chi_d}(s)} - \frac{\widetilde{L}'_{\xi}(s)}{\widetilde{L}_{\xi}(s)} \right) - 1 \right]. \tag{2.13}$$

The pole term and the subtraction of 1 arise from the single untwisted principal character. The boundary calculation makes $G_{d,Q}$ holomorphic at 1, and (2.12) makes it real along the real axis to the right of 1. Define the real constant and the centered sum by

$$c_{d,Q} = G_{d,Q}(1), \qquad R_{d,Q}(y) = S_{d,Q}(y) - \frac{\log y}{\varphi(Q)} - c_{d,Q}. \tag{2.14}$$

At $Q = 1$ the untwisted function is $(s-1)\zeta(s)$, whose logarithmic derivative at 1 is γ. Thus

$$\ell_d = \gamma - \Re\left(\frac{L'(1,\chi_d)}{L(1,\chi_d)}\right), \qquad c_{d,1} = -1 - \ell_d. \tag{2.15}$$

The logarithmic derivative of $L(s,\chi_d)$ at 1 is real, since χ_d is real and $L(1,\chi_d)\neq 0$.

For a finite nonempty set of scales in $[1,\infty)$, if y_* maximizes $R_{d,M}(y)-\frac{1}{2}y^{\beta_k-1}$, a bound $R_{d,M}(y_*)\le \frac{1}{2}y_*^{\beta_k-1}$ gives the same inequality at every scale in that set. The choice of maximizer depends on $S_{d,M}(y)-\log y/\varphi(M)-\frac{1}{2}y^{\beta_k-1}$, since subtracting the constant $c_{d,M}$ preserves all comparisons.

For $j\ge 1$ let $X_j=2^{2^j}$. Both endpoints of $[X_j,X_{j+1}]$ are squares. Its finite set of integer squares has a least maximizer y_j of

$$S_{d,M}(y)-\frac{\log y}{\varphi(M)}-\tfrac{1}{2}y^{\beta_k-1}. \tag{2.16}$$

These choices use the fixed k,d_k and the prime sums on the block. The arithmetic construction takes $x=y_j$ for each sufficiently large j, with $d=d_k$ and $\varepsilon=\varepsilon_k<1/100$.

3 Admissible data at a square scale

Fix $M=18480$, a prime $d\equiv 9239 \pmod M$, and $0<\varepsilon<1/100$. Write $\chi_d=(\cdot/d)$ and $\ell_d=\gamma-\Re(L'/L)(1,\chi_d)$, with Euler's constant γ. Since $d+1\equiv M/2 \bmod M$, we have $M\mid 2(d+1)$ and $M\nmid d+1$. Let x be a sufficiently large square. All thresholds in the datum may depend on d,ε. We will use the resulting estimates at the selected squares.

Take m and r to be the least powers of two and three at least $\sqrt{x}$ and x^5, respectively. Thus $\sqrt{x}\le m\le 2\sqrt{x}$ and $x^5\le r\le 3x^5$. The number of lifts is m, the definite source rank is $\kappa=r+1-m$, and the enlargement parameter is $t=5r$. Set $s=22\log_2 m$, so $2^s=m^{22}$ and $s\le(11/\log 2)\log x+22\le 32\log x$ for large x. Write $c=\lfloor\log_3 s\rfloor+5$ and $N_1=M3^c2^s$. Fix the embedding of $K=\mathbf{Q}(\sqrt{-d})$ with $\sqrt{-d}=i\sqrt{d}$, and put $c_0=(1+\sqrt{-d})/2$. In particular $c_0\in\mathcal{O}_K$, $c_0+\bar{c}_0=1$ and $c_0\bar{c}_0=(d+1)/4$. The CM field is $E=K(\mu_{N_1})$. Its infinite places Φ are represented by the embeddings extending this embedding of K, and are indexed by $(\mathbf{Z}/N_1\mathbf{Z})^\times$. Indeed d is ramified in K and does not divide N_1, so K and $\mathbf{Q}(\mu_{N_1})$ intersect in $\mathbf{Q}$. The tops $\mathcal{T}$ are the places of class 1 modulo M. Reduction of the unit groups is surjective, and 2 and 3 already divide M. Thus $T=\#\mathcal{T}=3^c2^s>0$ and $f=\#\Phi=\varphi(M)T$. Write $[u]$ for the place indexed by u modulo N_1, with the identity place used when u is a nonunit. Set $\log b=x^\varepsilon$ and $\log x_1=x^{\varepsilon/6}$. Then

$$\log b=x^\varepsilon,\qquad \log x_1=x^{\varepsilon/6},\qquad (10^6\log x)^2\le\log x_1,\qquad xb\le a<xb+N_1, \tag{3.1}$$

where a is the least multiple of N_1 at least xb.

For any fixed $C_0\ge 0$, the condition $10^{12}C_0\le\log\log b=\varepsilon\log x$ holds eventually. The inequality $100\log\log b\le\log x$ uses $\varepsilon<1/100$. We also have $10^6\log x\le\log x_1$, $10^6\log x_1\le\log b$, and $6\log\log x_1=\log\log b$. In particular, b dominates every fixed power of x and of x_1. The required fixed constant includes $M+2d+6$, $|\ell_d|$, the finite sum of $|\log|\Phi'_M(\xi)||$ over primitive M-th roots, and the fixed constant $C_d=C_{\mathrm{vac}}+2d+6$ for the moved vacuum bound. The normalized Taylor transport uses the same CM period. These constants are fixed before x.

3.1 The thick places

The labels will come from the digit-and-sphere construction of Behrend [3, pp. 331–332], with the digit base enlarged to handle all coefficients up to L_0. Write $n=2^{s-3}$ and $L_0=324s$, and set

$$u=\log(n/(4L_0)),\qquad v=\log(8L_0),\qquad D_0=\lfloor\sqrt{u/v}\rfloor,\qquad U_0=\left\lfloor e^{u/D_0}/(4L_0)\right\rfloor.$$

For large x, $D_0\ge 3$ and $U_0\ge e^{u/D_0}/(8L_0)$. In $\{0,\dots,U_0-1\}^{D_0}$ there is a sphere about the origin with at least $U_0^{D_0}/(D_0U_0^2)$ points, by partitioning the box by squared norm. This number is at least

$$\frac{n}{L_0}\exp\bigl(-\log 4-\log D_0-D_0v-2u/D_0\bigr)\ \ge\ \frac{n}{L_0}\exp(-16\sqrt{\log n\log L_0}).$$

For the last inequality use $D_0 v \le \sqrt{uv}$, $2u/D_0 \le 4\sqrt{uv}$, $\log D_0, \log 4 \le \sqrt{uv}$, and $v \le 4 \log L_0$. This lower bound tends to infinity, so the chosen sphere has positive radius. Encode a point by reading its coordinates as digits in base $\beta = 4L_0U_0$. The labels are distinct and lie strictly between 0 and $\beta^{D_0} \le n/(4L_0)$. A congruence $j(e_2 - e_1) \equiv i(e_3 - e_1) \pmod n$ with $0 < i < j \le L_0$ is an integer equality, by the range of the labels. The signed coefficient of every power of β has modulus less than $2L_0U_0$. Successive reductions modulo β turn the equality into the same relation in each digit. Three distinct points satisfying it would lie on one line in a sphere, which is impossible.

To choose the thick places, put $A = M3^c$, $a_0 = 1 + A$, and $L_0 = 324s$. The integer a_0 is a unit modulo $N_1 = A2^s$, since A is even. We have $v_2(A) = 4$ and $v_2(a_0^{2^j} - 1) = 4 + j$ for $j \ge 0$, by successive squaring since $a_0^{2^j} + 1 \equiv 2 \pmod 4$. Thus the order of a_0 modulo $2^{s+1} \mid N_1$ is exactly $n = 2^{s-3}$. Take $\mathcal{J}$ to be the places $[-a_0^e]$ on these labels. Their images under reduction ϕ modulo 2^{s+1} are distinct, because the labels lie in an interval shorter than the order n. They all have class -1 modulo A, hence modulo every divisor of A. If $\sigma, \tau, \rho \in \mathcal{J}$ satisfy $\phi(\tau/\sigma)^j = \phi(\rho/\sigma)^i$ for $0 < i < j \le L_0$, their labels satisfy the displayed congruence. If $e_2 \ne e_1$, neither $e_3 = e_1$ nor $e_3 = e_2$ is possible, since a nonzero multiple of $e_2 - e_1$ by j or $j - i$ has absolute value less than $n/4$. The exclusion for distinct labels therefore forces $e_2 = e_1$, giving $\phi(\tau/\sigma) = 1$. This supplies the congruence and three-term conditions of the datum.

Since $T = 8 \cdot 3^c n$ and $3^c, L_0 = O(\log x)$, the cardinality $J = \#\mathcal{J}$ satisfies

$$\log(T/J) \le \log(8 \cdot 3^c L_0) + 16\sqrt{\log n \log L_0} = o(\log x).$$

In particular $J \ge T/\log b > 0$ for sufficiently large x at fixed ε.

3.2 The split level primes

For a prime $d > 3$ with $d \equiv 3 \pmod 4$, put $K = \mathbf{Q}(\sqrt{-d})$. Let $H \ge 3$ be an integer and $B' = 4dH^2$. There is an absolute effective $C_{\mathrm{p}} > 0$ such that, if $X \ge (B')^{C_{\mathrm{p}}}$ and $1 \le v_0 < v_1 \le 2$ with $v_1 - v_0 \ge 1/2$, there is a prime $P \in (v_0X, v_1X]$ with $P = \pi\bar{\pi}$ for $\pi \in \mathcal{O}_K$ satisfying $\pi \equiv 1 \pmod{H\mathcal{O}_K}$.

Take the ray class field L/K of $H\mathcal{O}_K$. The Artin map identifies the ideals prime to H, modulo principal ideals with a generator congruent to one, with $\mathrm{Gal}(L/K)$ [4, Ch. V, Theorems 3.5–3.6]. Conjugation preserves the ray subgroup, so uniqueness of the class field gives that $L/\mathbf{Q}$ is Galois. A rational prime P prime to dH that splits completely in L has a prime $\mathfrak{p}$ of norm P in the trivial ray class of K. Thus $\mathfrak{p} = (\pi)$ with $\pi \in \mathcal{O}_K$ and $\pi \equiv 1 \pmod{H\mathcal{O}_K}$. Its norm gives $\pi\bar{\pi} = P \equiv 1 \pmod H$.

Apply Thorner and Zaman with base field $\mathbf{Q}$, conjugacy class $\{1\}$ in $\mathrm{Gal}(L/\mathbf{Q})$, and abelian subgroup $\mathrm{Gal}(L/K)$ with fixed field K. Its Hecke characters have conductors dividing $H\mathcal{O}_K$, so their largest conductor norm Q_{TZ} is at most H^2. The absolute discriminant d and degree 2 of K give the theorem's parameter $D_{\mathrm{TZ}} = 4dQ_{\mathrm{TZ}} \le B'$. For $y \ge D_{\mathrm{TZ}}^{c_2}$ the count $\Pi(y)$ of unramified rational primes at most y splitting completely in L is

$$\Pi(y) = \frac{A(y)}{[L : \mathbf{Q}]}(1 + e(y)), \qquad e(y) = O\Big(e^{-c_3 \log y/\log D_{\mathrm{TZ}}} + e^{-\sqrt{c_3 \log y/2}}\Big)$$

with absolute effective constants [5, Theorem 1.4, p. 1042, and §2]. Here $A(y) = \mathrm{Li}(y)$ in the ordinary case. If there is an exceptional real zero β_1 with $0 < 1 - \beta_1 < 1/(8 \log D_{\mathrm{TZ}})$, then $\beta_1 > 1/2$ and $A(y) = \mathrm{Li}(y) - \mathrm{Li}(y^{\beta_1})$, since the exceptional character has value one on the identity class.

Put $q_X = 1 - X^{\beta_1 - 1} > 0$ in the exceptional case and $q_X = 1$ otherwise. On $X \le y \le 2X$ one has

$$A'(y) \ge q_X/\log(2X), \qquad A(y) \ll Xq_X/\log X.$$

In the exceptional case the derivative is $(1 - y^{\beta_1 - 1})/\log y$. Also, integrating from y^{β_1} to y gives $A(y) \le 2y(1 - y^{\beta_1 - 1})/\log y$. Since $y \le X^2$, the factor $1 - y^{\beta_1 - 1}$ is at most $2q_X$. In the ordinary case the upper

bound follows by splitting the logarithmic integral at $\sqrt{y}$. The main-term increment $A(v_1X) - A(v_0X)$ is at least $Xq_X/(2\log(2X))$. The two endpoint errors in $[L : \mathbf{Q}](\Pi(v_1X) - \Pi(v_0X))$ have total absolute value

$$O\left(\frac{Xq_X}{\log X}\left(e^{-c_3\log X/\log B'} + e^{-\sqrt{c_3\log X/2}}\right)\right).$$

The common positive factor $Xq_X/\log X$ cancels in the comparison. Since $B' \geq 252$, choose $C_{\mathrm{p}} \geq \max(1, c_2)$ large enough to make the error smaller than the main-term increment for $X \geq (B')^{C_{\mathrm{p}}}$. The count difference is positive. Each prime in it exceeds $X \geq B' > dH$, so the ray class argument supplies its generator.

For the level primes, put $H = (Md)^{2(s+6)}$ and $B' = 4dH^2$. Since d is fixed and $s = O(\log x)$, we have $\log B' = O_d(\log x)$. The condition $X \geq (B')^{C_{\mathrm{p}}}$ in the split-prime existence result holds uniformly for $X \geq b^{10}$, since $\log b = x^{\varepsilon}$. At $X = b^{10}$, use the two disjoint intervals $(X, 3X/2]$ and $(3X/2, 2X]$ to choose distinct g_0, g_1 with generators $\pi_i \in \mathcal{O}_K$, $\pi_i\bar{\pi}_i = g_i$ and $\pi_i \equiv 1 \pmod{H\mathcal{O}_K}$. Both primes lie in $(b^{10}, 2b^{10}] \subseteq [b^{10}, b^{11}]$ for large x. Then $\pi_G = \pi_0\pi_1$ has norm $g = g_0g_1$. Both primes are 1 modulo H, and $N_1d \mid H$ since $2, 3 \mid Md$ and $c = \lfloor \log_3 s \rfloor + 5 \leq s + 5$ in the definition of N_1. In particular $g_i \equiv 1 \pmod{N_1d}$ and each splits completely in E.

For every $U \in [100\log b, 2^s]$, apply the same result to $(e^U, 2e^U]$. Since $e^U \geq b^{100}$, the same threshold in x works for all U. This gives a prime $p \equiv 1 \pmod{H}$ with $U < \log p \leq U + \log 2$, and $\log 2 \leq \sqrt{\log b}$ for large x. Together with the scales, powers and thick set above, these choices give an admissible datum D. The two g_i specify the principal-level type conditions. The extra split prime N is chosen below from the window for the weights of this $\mathcal{J}$ and specifies ramification on the columns.

3.3 The weights and the signed remainder

Set $w(\sigma) = a/x$ for $\sigma \in \mathcal{J}$ and $w(\sigma) = a/x_1$ otherwise, and write $\bar{w} = f^{-1}\sum_{\sigma\in\Phi}\log w(\sigma)$. The quantities used in the volume calculation are

$$\begin{aligned}
\mathrm{Ask}(u, \tau) &= \sum_{1\leq v\leq a} \Lambda(v)(1 - \chi_d(v))\big(u/v - w([-v]\tau)\big)_+, \\
\mathrm{MainAsk} &= a(\log a - 1 - \ell_d - \bar{w}), \\
\mathrm{rem} &= \sum_{\tau\in\mathcal{T}} \big(\mathrm{Ask}(a, \tau) - \mathrm{MainAsk}\big).
\end{aligned} \tag{3.2}$$

Let b_1 and b_2 be the least multiples of N_1 at least a/x and a/x_1, respectively. Put $h_2 = r^8$, $h = 2h_2$, and set

$$w_0(\sigma) = h_2 + \begin{cases} b_1 & \sigma \in \mathcal{J}, \\ b_2 & \sigma \notin \mathcal{J}, \end{cases} \qquad w_j(\sigma) = w_0(\sigma) + jh.$$

The symbol h in the weight progression is the spacing. The height in Theorem 2 is denoted by h_{sh}. With rd_E the root discriminant of E, put

$$\begin{aligned}
\Omega &= (2\pi/\mathrm{rd}_E)e^{\ell_d}, & S(u) &= \log(u!) - u\log(\Omega/g) \quad (u \in \mathbf{Z}_{\geq 0}), \\
\mathcal{U} &= \sum_{\sigma\in\Phi}\big(S(w_0(\sigma)) - \mathrm{Ask}(w(\sigma), \sigma)\big), & h_{\mathrm{main}} &= m\kappa\mathcal{U}, \\
\mathcal{B} &= fa(\log b)^3, & L &= \mathcal{B}\log b, \\
h_{\mathrm{sh}} &= h_{\mathrm{main}} - m\,\mathrm{rem} + L.
\end{aligned}$$

Put $G = \log(g\,\mathrm{rd}_E/(2\pi))$ and set

$$\lambda = m\bar{w} + (m-1)G - \tfrac{15}{2}\log b.$$

Since $1 \le w(\sigma) \le b^2$, $g \in [b^{20}, b^{22}]$, and $\log \mathrm{rd}_E = O(\log x)$ by $E \subseteq \mathbf{Q}(\mu_{N_1 d})$, we have $2\log b \le G \le 27\log b$ and $m\log b \le \lambda \le 29m\log b$ for large x. In particular $100\log b \le \lambda \le 2^s = m^{22}$. Use the prime window with $U = \lambda$ and choose $N \equiv 1 \pmod H$ with

$$\lambda < \log N \le \lambda + \log 2 \le \lambda + \sqrt{\log b}.$$

Then $N \equiv 1 \pmod{N_1 d}$, N splits completely in E, and $N - 1 \ge b^9$ for large x. Since $m \to \infty$, the lower bound $m\log b$ for $\log N$ also separates N from the two primes $g_i \le b^{11}$. As $4 \mid N_1$, the sum of two squares for N gives $\pi \in \mathbf{Z}[i] \subset \mathcal{O}_E$ with $\pi\bar\pi = N$.

3.4 Hermitian spaces, levels and characters

Roots of unity and local frames. Put $F = E^+$ and $L_{\mathrm{cyc}} = \mathbf{Q}(\zeta_{N_1})$. The intersection $L_{\mathrm{cyc}} \cap K$ is trivial over $\mathbf{Q}$ because d ramifies in K and is prime to N_1. Hence the restrictions of the Φ-embeddings to L_{cyc} run once through $\mathrm{Gal}(L_{\mathrm{cyc}}/\mathbf{Q})$. The roots of unity of E form exactly μ_{N_1}. Indeed $E \subseteq \mathbf{Q}(\zeta_{dN_1})$, where dN_1 is even. A root of unity in this cyclotomic field has order dividing dN_1, by comparing the degree at that conductor with the degree at its least common multiple with the root's order. If E contained a root with order divisible by d, it would contain $\mathbf{Q}(\zeta_{dN_1})$, which has degree $d - 1 > 2$ over L_{cyc}.

The field E/F is unramified at every finite place. At a place of residue characteristic p, use a root of unity of order 3 or 5 prime to p. Its reduction still has that order, while complex conjugation inverts it. If conjugation stabilizes the place, it acts nontrivially on the residue field. At a nonsplit place the norm on local integral units is onto. The norm on the residue fields is onto, and the trace is onto on each successive principal-unit quotient. Successive lifting and completeness prove the assertion also at the dyadic places. At a split place the local norm on units is the product of the two coordinates and is onto.

Diagonal forms at finite places. In a diagonal Hermitian space the indicated signs give the entries ± 1 of the rational form over E/F, with the standard integral coordinate lattice. At each finite place of F, choose a unit of norm -1, and multiply the negative coordinates by this unit. This gives isometries from the standard finite Hermitian lattice to the various signed lattices, simultaneously in every rank. At a split place the identification of the unitary group with one general linear factor determines the other factor by inverse conjugate transpose. At a nonsplit place the self-dual integral unitary group splits over the unramified quadratic ring of integers. The full lattice stabilizer is hyperspecial there.

The chosen halves, different and splitting character. For g_i, use the places w dividing $(\bar\pi_i)$ in E. Exactly one of $w, \bar w$ above a fixed real-base place divides $(\bar\pi_i)$, with valuation one, since $\pi_i\bar\pi_i = g_i$ and g_i splits completely. There are f such places for each i. Write S_g for their union and $\mathfrak{M} = (\bar\pi_G)$, where $\pi_G = \pi_0\pi_1$. Its support is S_g, so the chosen principal level is the kernel of these reductions. The π-half over the separate prime N is constructed in the same way from $\pi\bar\pi = N$. If Q is the product of the distinct chosen halves at new split pairs, the principal modulus is $\mathfrak{M}^e Q$, with e as specified in the edge. It is prime to its conjugate. The deep definite core used in the grid exchange has level $\mathfrak{M}^{b_1}$. The column spaces in the period below use $\mathfrak{M}$, with an additional last-row condition at the π-half of N.

Put $P(X) = \Phi_{N_1}(X)$ and $\delta = P'(\zeta_{N_1})$. The K and L_{cyc} discriminants are relatively prime. Locally at every rational prime one integer ring is étale, so their tensor product is the full integer ring in the compositum. The tower formula for differents gives $\mathfrak{D}_{E/K} = (\delta)$ and $\mathfrak{D}_{E/\mathbf{Q}} = (\delta\sqrt{-d})$. Reciprocity of P gives $\bar\delta/\delta = -\zeta_{N_1}^{2-f}$, with $f = \varphi(N_1)$. Thus $\varpi = i\zeta_{N_1}^{(2-f)/2}\delta\sqrt{-d}$ is a purely imaginary generator of the absolute different; $4 \mid N_1$ and f is even. Differentiating $X^{N_1} - 1 = P(X)((X^{N_1} - 1)/P(X))$ at ζ_{N_1} shows that δ divides N_1 in $\mathcal{O}_E$. Choose the finite idele α with component δ^{-1} at the places dividing c_0 and component 1 elsewhere. Every rational prime dividing N_1 divides $(d+1)/4 = c_0\bar c_0$. For odd primes this follows from $M \mid 2(d+1)$, and for 2 it follows from $d \equiv 7 \bmod 8$. Since $c_0 + \bar c_0 = 1$, exactly one of $c_0, \bar c_0$ is divisible by a place over any of these primes. The conjugate of δ differs from it by a unit, so $\alpha\bar\alpha\delta$ is a finite unit. For

$j_{\mathrm{Weil}} = -1/\sqrt{-d}$ this gives $(j_{\mathrm{Weil}}\alpha\bar{\alpha})\widehat{\mathcal{O}}_E = \mathfrak{D}_{E/\mathbf{Q}}^{-1}$. Also $\Im\tau(j_{\mathrm{Weil}}) > 0$ and $(1 - j_{\mathrm{Weil}}^{-1})/2 = c_0$, and $N_1\alpha$ is integral. The skew scalar for the Weil pairing is $-1/\sqrt{-d}$, with $\sqrt{-d} = i\sqrt{d}$ in the chosen embedding of K.

There is a unitary splitting character v of $E^\times\backslash\mathbb{A}_E^\times$, unramified at all finite places, whose infinite part is $\prod_{\tau\in\Phi} z_\tau/|z_\tau|$ and whose restriction to $\mathbb{A}_F^\times$ is the quadratic character $\epsilon_{E/F}$. To verify the prescribed unit compatibility let $\zeta = \zeta_{N_1}$ and $v = i(\zeta - 1)$. It is a unit with $v/\bar{v} = \zeta$, and $N_{E/K}(v) = 1$, since $N_{L_{\mathrm{cyc}}/\mathbf{Q}}(\zeta - 1) = \Phi_{N_1}(1) = 1$ when N_1 has at least two prime divisors. If $a/\bar{a}$ is an integral unit for $a \in E^\times$, all of its conjugates have modulus one, hence $a/\bar{a} = \zeta^j$ for some j. It follows that $a/v^j \in F^\times$. In particular every unit of E is the product of v^j and an integral unit of F. For a unit of F, the sign of its norm to $\mathbf{Q}$ is $+1$, because the global character $\epsilon_{E/F}$ is trivial on principal ideles and is unramified at every finite place. Thus the proposed infinite part is trivial on integral units of E.

On the group of principal fractional ideals of E prescribe the inverse of the proposed infinite part evaluated at a generator. On the extensions of fractional ideals of F prescribe $\epsilon_{E/F,f}$. The extension-of-ideals map is injective. If such an extended ideal is (a), then $a/\bar{a}$ is an integral unit and $a = v^j b$ with $b \in F^\times$. As v has norm one over K, the infinite part on a equals the sign product on b. This is the inverse of $\epsilon_{E/F,f}(b)$. The two prescriptions therefore agree on their intersection. A character of their product subgroup extends to the whole fractional ideal group with values in the circle, because circle values admit roots. The quotient by principal ideals is finite, so one can make the extension on a finite set of generators, taking compatible roots. This gives the required v_f. Its values on finite integral ideles are one. At a nonsplit place an F-uniformizer is also an E-uniformizer, and restriction to $\epsilon_{E/F}$ gives value -1. The character v has $v(\bar{t}) = v(t)^{-1}$, since $t\bar{t}$ is a norm and $\epsilon_{E/F}$ kills norms. Accordingly the base change of its restriction $v_{U(1)}$ to the norm-one torus is v^2. The character $v_{U(1)}$ has real weight one at every τ and trivial finite tame part.

Global stage characters and the signed catalogue. Put $M_\omega = (Md)^{s+6}$. Then $H = M_\omega^2$ from the prime construction and $dN_1 \mid M_\omega$. At each chosen w, fix a residue character of order M_ω^2 of the cyclic group $k_w^\times$, available because $g_i \equiv 1 \bmod H$. Let ω_w be its M_ω-th power, and use the trivial character outside the chosen halves. At a split norm-one pair the opposite matrix coordinate is determined by the conjugate inverse. The chosen local ω_w has order M_ω and kills the reduction of every norm-one integral unit of E, since these units are in μ_{N_1} and $N_1 \mid M_\omega$. In particular ω_w^a is nontrivial at all the chosen places for $0 < |a| \le 8m$.

For any integer A the angular norm-one character $\prod_{\tau\in\Phi} t_\tau^{-Aw_0(\tau)}$ also kills all norm-one integral units of E. Indeed $w_0(\tau) \equiv h_2 \bmod N_1$, independent of τ. For a root of unity $\zeta' \in \mu_{N_1}$, the angular value is $N_{L_{\mathrm{cyc}}/\mathbf{Q}}(\zeta')^{-Ah_2} = 1$. Thus the prescribed real and finite compact values of weight $-Aw_0$ and tame part ω^a, for integers A, a, agree on the rational points in the compact level. Define their product on $U(1)(F_\infty)\prod_{v<\infty} U(1)(\mathcal{O}_v)$, where the product of groups is an open subgroup of $U(1)(\mathbb{A}_F)$. It descends to the image of that open subgroup in the rational quotient. Extend it to the rational quotient with values in the circle. This extension is continuous because it is continuous on the open subgroup, and its local tame parts and weights are unchanged. The quotient by the open subgroup is finite. The ideals of the norm-one finite ideles are all of the form $\mathfrak{a}/\bar{\mathfrak{a}}$: at a split pair use its prescribed exponent at one half and zero at the other, and at a nonsplit place the norm-one exponent is zero. Principal ideals of the form $\mathfrak{a} = (u)$ map to $(u/\bar{u})$, the ideal of a rational norm-one point. The ideal class group therefore surjects onto this quotient. For weight zero one may choose a finite-order extension. No coefficient-place choices enter this construction.

Choose $\mu_\flat$ with weight $-w_0$ and trivial tame part, and χ_0 with weight zero and tame part ω. Keeping these two choices fixed, put

$$\mu_j^0 = \mu_\flat(v_{U(1)}^{-h}\chi_0)^j.$$

Then μ_j^0 has weight $-w_j$ and tame part ω^j. On the actual edge the remaining rank twist is the specified floor or ceiling power of $v_{U(1)}$. A trial finite-order character ψ has weight zero and support at the new split pairs, so replacing μ_j^0 by $\mu_j^0\psi$ leaves all real parameters unchanged. The auxiliary-ray construction requires the places in Q and their conjugates to avoid g_0, g_1, N, ℓ; its exact-support condition at each required power makes $\mathfrak{M}^e Q$ the same at that power. These are genuine torus characters, as the ray construction kills the real ideles and the decomposition groups above ℓ.

The unsigned family at stage j uses

$$\nu_j = \mu_j^0 v_{U(1)}^{j+(t-11m-1)/2}, \qquad 1 \le z' \le 7m/2 + \kappa/2.$$

It has the same character at each point z'; its tame part is ω^j, and the weight is the negative of $w_j - j - (t-11m-1)/2$. The $\kappa/2$ largest points are the sharp-line tests. They satisfy $m < z' \le r$ eventually, and the positive weight magnitude c lies in $[w_j + r - t, w_j]$, since $j + (t-11m-1)/2$ is nonnegative and at most $t-r$. The remaining unsigned points use the positive sign in the signed interpolation. For $s_1 = +1, -1$ put $e_+ = m$, $e_- = m/2$ and use

$$\nu_{j,s_1} = v_{U(1)}^{-e_{s_1}} \chi_0^{s_1 j}, \qquad 1 \le z' \le e_{s_1}.$$

The corresponding sign uses the trial power ψ^{s_1} at that stage. The ramified flat pair in the rank-one jet is $(v_{U(1)}^{-2}\chi_0, 2)$. In a grid square at (a, b) the product character before the rank normalization is $\mu_{a+1}^0 \mu_{b+1}^0$ with tame exponent $a+b+2$; the trial power occurs from the $b+1$ edge in V or the $a+1$ edge in the full Z grid. For a triangle square whose shorter chain has length n, the corresponding character is $v_{U(1)}^{-(n+1)h} \chi_0^{n+1}$, multiplied by the trial power θ^{n+1} in V. The triangle in Z uses its trivial-twist specialization. In $\Pi_{\nu,S}$ the positive-infinity character is the inverse of the resulting negative-weight critical-pair character. The needed lower points have $z' \le t$ in those squares. Their negative weight magnitudes dominate t and are at most $4b$ at the large scales under consideration. This follows from $w_0 = O(b)$, $w_0(\tau) \ge a/x_1$, $w_j \le 2b$ for $j \le 2m$, and $t < h \ll b/m$. For the explicit signed and flat families the lower points satisfy $1 \le z' \le c$ directly. Every ramified family just listed has a nonzero tame exponent of absolute value at most $2m$. Thus each signed Euler factor at the fixed helper inertia has a nontrivial root of unity of order dividing $q_w - 1$.

Finally the change of a Green label by the stage tame character preserves the allowed set. Squarefreeness and factor degrees are unchanged by scaling roots by an element of the residue field. The forbidden roots form the subgroup of solutions of $x^{M_\omega} = 1$, which is invariant under this scaling. A primitive irreducible root also remains primitive. In the cyclic group of order $q_w^r - 1$, the multiplier has order dividing M_ω. Its exponent is divisible by $(q_w^r - 1)/M_\omega$, which still contains every prime factor of $q_w^r - 1$ because $M_\omega^2 \mid q_w - 1$. Adding it to a primitive exponent preserves coprimality. This verifies the particular local label changes used in the lower row and the scalar-inertia test.

3.5 The row and column chains

For $0 \le j \le m$ set $R_j = r + jt$ and $n_j = \kappa + j(t+1)$. We also write $k = \kappa$. The space $X_j = E^{n_j}$ is positive definite with Gram matrix I_{n_j} and distinguished last vector e. The ordered basis of Y_j has the r original indices followed by j blocks of size t. Put $\epsilon_j = j \bmod 2$. At a real place τ, the doubled parameters in this order are the entries in the following array, all decreased by ϵ_j:

$$\begin{array}{lll} k-2-2i & (0 \le i < k-1), & \text{sign } +1, \\ 2h_2 - 2w_{a+1}(\tau) & (0 \le a < m), & \text{sign } (-1)^{a+\min(a,j)}, \\ t-1-2w_\ell(\tau) - 2i & (1 \le \ell \le j,\ 0 \le i < t), & \text{sign } +1. \end{array}$$

Write $I_j : Y_j \hookrightarrow Y_{j+1}$ for the first-block inclusion of these indices in the split residue matrices. The equal-rank inner form W_j for $1 \le j \le m$ has the original real parameter and diagonal signs with the first $k+j-1$ positive and the last $m-j$ negative.

For $j < m$ use the rational graded description of Y_j as X_j and a string of length $s_j = m - 1 - j$. The core has grade zero, and u_i, $0 \le i < s_j$, in the string has grade $s_j - 1 - 2i$. Opposite indices pair to one, apart from the middle index in an odd string, whose vector z has norm $-N$. All other Gram entries in the string are zero. At $j = 0$ this is the flag frame of $\varrho = m/2 - 1$ hyperbolic pairs, the core and z. The string has $\lfloor (m-j)/2 \rfloor$ negative directions, the number of negative signs in the indicated real ordering. Diagonalizing the hyperbolic pairs over $\mathbf{C}$ by unitary matrices after rescaling z to norm -1, and matching the signs, gives the chosen real sign frames. The definite space Y_m has Gram I_{R_m}, and we identify it with the first $n_m - 1$ coordinates of X_m, since $R_m = n_m - 1$. The Bessel subgroup at $j < m$ is $B_j = U_j \rtimes U(X_j)$, where U_j consists of unitary matrices whose difference from the identity strictly raises grade. It fixes each nonnegative-grade string vector modulo higher-grade string vectors. The core unitary group acts trivially on the string. At Y_m use its whole compact-type group and the trivial unipotent.

To specify the finite frame, let C_g^{id} be the finite idele with entry $\bar{\pi}_G$ at the chosen places and 1 elsewhere, set $a_{\mathrm{W}} = \alpha/\overline{C_g^{\mathrm{id}}}$, and put

$$s_j = m - 1 - j, \qquad \eta_j = (a_{\mathrm{W}}/\overline{a_{\mathrm{W}}})^{\lfloor (m-j+1)/2 \rfloor}.$$

This central scalar on the string has norm one. For an even string take the standard core lattice and multiply the string basis by $\eta_j a_{\mathrm{W}}^{-1}$ on its positive half and $\eta_j \overline{a_{\mathrm{W}}}$ on its negative half. For an odd string multiply its vectors by η_j and replace the last core column e by $e/\bar{\pi} + \eta_j z/N$. The resulting Gram is integral and unimodular. This follows for the even pairs by conjugating their basis factors. For the odd middle block, in the two modified columns the Gram is $\left(\begin{smallmatrix} 0 & -1 \\ -1 & -N \end{smallmatrix}\right)$. A hyperbolic basis in this block is $-e_o, \eta_j z - N c_0 e_o$, where $e_o = e/\bar{\pi} + \eta_j z/N$. Its entries have Gram off-diagonal one and both diagonal entries zero, because $c_0 + \bar{c}_0 = 1$. The change has determinant a unit. The other hyperbolic planes and the unit coordinates are unimodular as well. At each split or unramified finite place any self-dual lattice with this Gram is isometric to the standard lattice. In the unramified case lift a residue vector of unit norm, divide its norm by using the unit norm map, and split off this unit vector inductively. At a split place choose an arbitrary basis in one factor and its dual in the other. In particular at a chosen place the orthonormalizing change can be taken to be identity in the chosen half on reduction. We use these frames to take all residue matrices, with their indices ordered as in I_j. This constructs the stabilizer K_j^0. For $j = m$ use the standard frame. In the X_j core use the standard finite lattice. The core level K_j^X is principal modulo $\mathfrak{M}$ and has $e^{\mathsf{T}} g \equiv e^{\mathsf{T}} \bmod \pi$ in its π-half. The core embedding carries K_j^X into K_j^0. For even strings the lattice is a direct sum. For an odd string write $e = \bar{\pi} e_o - \eta_j z/\pi$. An ordinary core column can acquire a last-coordinate entry g_{ei}; its possible denominator is g_{ei}/π, so the last-row condition makes it integral. The change of the gluing column is $(g - 1)e/\bar{\pi}$. In the $\bar{\pi}$-half the last column agrees with e modulo $\bar{\pi}$ by inverse conjugate transpose of the last-row condition in the other half. In the π-half the remaining possible pole uses $(g_{ee} - 1)/\pi$, again integral by the last-row condition. These arguments apply to g^{-1} as well. At the chosen places the gluing denominators are units, so the image of the initial core level reduces to the identity there.

The two finite frames at stage zero. The filling lattice on Y_0 is the odd-string lattice with unscaled string vectors and gluing column $e/\bar{\pi} + z/N$. Write K_r for its full finite stabilizer and $K(g)$ for the kernel of its reductions at the chosen places. Let a_{str} act by η_0 on the string and by the identity on the core. It sends the filling lattice to the stage-zero lattice defined above, including the gluing column. Let η act by η_0^{-1} on the core and by the identity on the string. The equality $a_{\mathrm{str}} = (\eta_0 I_{Y_0})\eta$ shows that the finite stabilizer and principal level on the stage-zero datum are

$$K_0^0 = \eta K_r \eta^{-1}, \qquad J_g = \eta K(g) \eta^{-1}.$$

In the transported lattice frames the reduction of a stage-zero element $\eta k \eta^{-1}$ is the reduction of $k \in K_r$. The common quotient is $\mathcal{G} = \prod_{w \in S_g} \mathrm{GL}_r(k_w)$, where k_w is the residue field at w. The core commutes with

η. Its intersection with either principal level is exactly K_0^X. To see the reverse inclusion, at a place in the π-half of N, projection of the filling lattice onto the core gives $\mathcal{O}_{E_w}^k$ because $\bar{\pi}$ is a unit. Hence a core isometry in $K(g)$ is integral there and, by inverse conjugate transpose, in the other half. The intersection of the filling lattice with the core in the π-half is $\mathcal{O}_{E_w}^{k-1} \oplus \pi \mathcal{O}_{E_w} e$: cancelling the z component of the gluing column requires its coefficient to be divisible by π. Applying the core isometry to the ordinary core columns and subtracting the gluing column from its image gives the last-row condition modulo π. Away from N, the gluing denominators are units, so restriction of the lattice and the chosen reductions to the core gives precisely the stated integral principal level there.

The same test extracts the core component of any element of $B_0(\mathbb{A}_{E^+,f}) \cap K(g)$. For the span I of the isotropic half-frame, the lattice on $I^\perp/I$ is the glued middle block. The strictly grade-raising factor acts trivially on this quotient. The core factor fixes z and preserves the induced lattice and reductions, so belongs to K_0^X. Its removal leaves an element of $U_0(\mathbb{A}_{E^+,f}) \cap K(g)$.

Here are the coefficient weights of the column chain. At τ, let $\kappa_\tau(l)$ be the number of columns of height l in one diagram, with

$$\kappa_\tau(l) = \begin{cases} a\,\mathbf{1}_{\tau\in\mathcal{T}}, & l = 1, \\ w_1(\tau) - \lfloor (k+t+1)/2 \rfloor, & l = k, \\ h - (t+1), & l > k,\ l \equiv k \bmod (t+1), \\ 0, & \text{otherwise.} \end{cases}$$

In rank n_j keep columns of height $l < n_j$ and twist by the determinant exponent $e_j(\tau) = \lfloor n_{j-1}/2 \rfloor - w_j(\tau)$ for $j > 0$, with $e_0(\tau) = 0$. We use Weyl evaluation coordinates, the duals of the integral flag-minor polynomial basis, with the positive form induced by the tensor product of exterior-power norms. The weight matrix at τ is the Schur matrix times the determinant to the specified exponent. On rational unitary elements read this matrix on the conjugate of the embedded rational matrix. An algebraic form is a function on the finite adelic group with open right stabilizer and rational left covariance in the corresponding product weight. At a coefficient embedding $\iota : \mathbf{C} \hookrightarrow \mathbf{C}_p$, the gauge of a value at g_f is the maximum absolute coordinate of that value after applying the weight of g_f^{-1} at the finite embeddings extending $\iota\tau \circ \mathrm{conj}$. This convention fixes the gauges at both readings in Theorem 2.

At X_m retain $c(l) = h_2 - (t+1)/2$ columns at each height $l \geq k$ with $l \equiv k \bmod (t+1)$, and put $c(l) = 0$ at other heights. The diagram on Y_m has $c(l) + \kappa_\tau(l+1) - c(l+1)$ columns of height $l < n_m - 1$ and determinant exponent $e_m(\tau)$. As on X_m, only heights strictly less than the rank enter the polynomial part. The power of the omitted coordinate in the branching is

$$b_\tau = e_m(\tau) + \sum_{l<n_m} \big(\kappa_\tau(l) - c(l)\big).$$

These prescriptions specify the terminal restriction and the scalar using $\bar{\pi}^{\,-b_\tau}$ below.

4 Column models and the theta period

4.1 Column classes and the finite type projection

Let q be a finite field cardinality. Green's parametrization associates to a monic squarefree polynomial $P = \prod_l P_l \in \mathbf{F}_q[X]$ of degree r, with nonzero constant term, the irreducible parabolic induction of the cuspidal factors of degrees $a_l = \deg P_l$ indexed by its distinct irreducibles. We choose compatible identifications of roots with finite-torus characters in this parametrization. In particular

$$\dim \tau_P = \frac{\prod_{j=1}^{r}(q^j - 1)}{\prod_l (q^{a_l} - 1)}. \tag{4.1}$$

This is the squarefree case of Green's character construction and degree formula [6, §7, Theorem 12 and §8, Theorems 13–14]. The unipotent radical of a proper parabolic with two blocks of sizes $a, r-a$ has no invariant in τ_P unless a is a sum of some a_l. In the finite Mackey decomposition, the double cosets are block shuffles of the inducing blocks relative to the two target blocks. A split of an inducing block across both target blocks projects the intersecting unipotent into a proper parabolic radical of that cuspidal block. Its invariants vanish by cuspidality. A contributing coset must thus put every inducing block entirely on one side, which gives the claimed sum. The same statement for the dual follows by taking the dual of the finite invariant projector.

For a smooth principal-level form f, a rational parabolic $P = MN_P$, and a split finite place w, suppose the local finite type has no invariants under $N_P(F_w) \cap K_w^0$. The constant term $\int_{N_P(F)\backslash N_P(\mathbb{A}_F)} f(ng)\,dn$ at $g_w \in P(F_w)$ is unchanged by right translation of the function by $N_P(F_w) \cap K_w^0$. Conjugating the translation by g_w gives a left translation in the adelic radical. The integral is finite over the compact unipotent quotient. Averaging over $N_P(F_w) \cap K_w^0$ gives zero in the finite type. The same calculation for each compact right translate of f and the decomposition $G(F_w) = P(F_w)K_w^0$ give zero for all g.

At each w chosen over g_i, the allowed labels are squarefree monic of degree r, with nonzero constant term and without a root among $x \in k_w$ with $x^{M_\omega} = 1$. A label is primitive if it is irreducible with a generator of $k_w^{(r)\times}$ as a root, where $k_w^{(r)}$ denotes the extension of k_w of degree r. A good tuple has allowed labels at every chosen place, and at least one primitive label among the places over each g_i. Denote by e_0 the sum of their central orthogonal type idempotents in $\mathbf{C}[\mathcal{G}]$. For a single finite type τ the coefficient at u in its central idempotent is $\dim(\tau)\,\mathrm{tr}\,\tau(u^{-1})/|\mathrm{GL}_r(k_w)|$. At a primitive label there is one cuspidal block of size r. The compact part of any proper split parabolic radical reduces, in a compact basis adapted to its flag, to the radical of a proper finite parabolic. Refining the two-block flag gives the same vanishing of invariants. Hence the constant-term argument applies to each good tuple.

Let C be the column space of algebraic forms on the definite unitary group $U(X_0)$ of rank κ, at the prescribed weight and stage-zero level. This level imposes the chosen principal level and requires the last row to be congruent to $(0, \dots, 0, 1)$ modulo π. The initial Poincaré series takes values in bounded continuous automorphic forms on the indefinite group $U(Y_0)$ of rank r, fixed by the prescribed real type kernel and of principal level at the chosen places above g_0, g_1. This space has the unitary action ρ of $\mathcal{G} = \prod_v \mathrm{GL}_r(k_v)$, where v runs over the chosen places above g_0, g_1 and k_v is the residue field. The central idempotent e_0 for the chosen types at the places over g_0, g_1 gives the row space $W = \mathrm{im}\,\rho(e_0)$. Write $C = \dim_{\mathbf{C}} C$ and $R = \dim_{\mathbf{C}} W$. The finite constant-term test gives cuspidality on W.

Evaluation of the columns in double-coset coordinates and in the local Weyl-to-Bessel coordinates gives the input to a Poincaré series map $\mathcal{P}$ on $U(Y_0)$. Its seed vectors use the product of the full real Bessel vectors, the class cutoff and the finite principal-coset indicator. Let $\mathcal{R} = \prod_{v|\pi}(1 - z_v)$ on C, where $\pi \in \mathcal{O}_E$ satisfies $\pi\bar{\pi} = N$ and z_v is the scalar translation chosen at v. Write $C_{\mathrm{ram}} = \mathrm{im}\,\mathcal{R}$ for the ramified column space.

Trivial stabilizers at the chosen level. On a definite group over E^+, the stabilizer of any adelic representative at a compact integral level is finite. The finite coordinates put its matrix entries in a fixed fractional lattice and definiteness bounds them at infinity, where the lattice is discrete. At a chosen place over an odd prime $p = g_i$ the split coordinates are over $\mathbf{Q}_p$. The principal congruence group $1 + pM_n(\mathbf{Z}_p)$ is torsion-free. For $I + p^h A \neq I$ with $h \geq 1$ and A having a unit entry, a power prime to p leaves the least entry valuation of its difference from I equal to h. Its p-th power raises this valuation to $h + 1$, by the binomial expansion. Conjugating a finite class stabilizer into this principal level at the chosen place proves it is trivial. Thus the column level, the enlarged level below, and the terminal level all have trivial class stabilizers. At the terminal definite space Y_m, the residue conditions are identity at the chosen places and block upper triangular at the twist places, so the principal condition still holds for every twist. At any other integral level the stabilizer injects under reduction at the chosen places by applying the same calculation to its kernel.

For a cyclic orbit of length u with shift z and coefficients in any ring, $\operatorname{im}(1-z)$ is the kernel of the sum of the coordinates. Both inclusion and surjectivity follow by telescoping: a prescribed zero-sum vector has a preimage given by successive partial sums. Moreover the sum splits by inserting a coefficient at one coordinate. On the product of f cyclic orbits the tensor decomposition into the augmentation summands and their free complements therefore gives the image of $\prod_v (1 - z_v)$ as the intersection of the coordinate-wise sum kernels. The decomposition works over the integral rings of the coefficient embeddings. Over $\mathbf{C}$ this image has dimension $(u-1)^f$ for each rank-one free orbit of size u^f.

The central action on the column classes. At each $v \mid \pi$, the scalar z_v lifts a generator of $k_v^\times = \mathbf{F}_N^\times$, acts by the inverse conjugate scalar at $\bar{v}$, and is trivial elsewhere. The places $v \mid \pi$ choose one from each conjugate pair and number f. The translations lie in the chosen principal level since $N \neq g_i$. The column level K_0 and these central elements generate K^+. The last-row congruence shows that $K^+/K_0 \simeq \prod_{v|\pi} \mathbf{F}_N^\times$, since $\prod_v z_v^{a_v}$ belongs to K_0 precisely when $N-1$ divides each a_v. A translation fixing a class would supply a rational element of the class stabilizer at K^+. Triviality of that stabilizer makes the action of this quotient free.

On an orbit, choose representatives $a_f z$ from one fixed a_f. The orbit values are then copies of the same coefficient space on which the z_v act by cyclic shifts, preserving its local Weyl gauges. Apply the augmentation calculation for cyclic shifts to this orbit over the coefficient ring. In particular the integral ramified columns at a coefficient embedding are the image of the integral columns under $\mathcal{R}$, and

$$\dim \ker \mathcal{R} = C\big(1 - (1 - (N-1)^{-1})^f\big) \leq fC/(N-1).$$

The triangle inequality for the cyclic shifts gives a squared supremum norm loss of at most 4^f under every coefficient automorphism.

4.2 The marked elliptic curve and the oscillator

The finite and real conventions. For a Hom space with the prescribed source and target Hermitian forms, use the Hermitian pairing $(x, y) \mapsto \operatorname{tr}(x^\dagger y)$, where $\dagger$ is the Hermitian adjoint. Let ψ be the standard additive character of $\mathbb{A}_E/E$, with infinite component $\exp(2\pi i \operatorname{Tr}_{E_\infty/\mathbf{R}}(\cdot))$. In rational coordinates the Weil cocycle is $\psi(j_{\mathrm{Weil}} \operatorname{tr}(y^\dagger x)/2)$. In the standard finite self-dual coordinates put $\psi_0(z) = \psi_f(j_{\mathrm{Weil}} \alpha\bar{\alpha} z)$. The absolute different identity above makes $\widehat{\mathcal{O}}_E$ its own annihilator. Write L_f for the integral lattice in these frames. With $\langle y, x\rangle = \operatorname{tr}(y^\dagger x)$, the finite vacuum is

$$v_0(x) = \mathbf{1}_{L_f}(x)\psi_0\big(\tfrac{1}{2}c_0\langle x, x\rangle\big)\,.$$

In this lattice model the Heisenberg operator is

$$(\rho_f(t)f)(x) = \psi_0\big(\tfrac{1}{2}\langle t, x\rangle\big)\, f(x+t).$$

The symplectic representation is generated by these translates of v_0. The operators multiply with cocycle $\psi_0(\langle u, t\rangle/2)$ for $\rho_f(t)\rho_f(u)$. The trace-one identity for c_0 gives the normalized lattice action on L_f used for the averages below, also at places over two. In the real oscillator use the Gaussian of squared norm one, with creation coordinates in the signed orthonormal frames. The theta functional Θ is the sum over the rational points in a Schrödinger polarization with self-dual additive measure. It is invariant under rational Heisenberg translations. At a split place of a modulus coprime to its conjugate, the standard finite Schrödinger model on the chosen matrix half has vacuum the integral-box indicator. Naming a residue function means multiplying this indicator by the pullback of that function. In the lattice model this is equivalently its finite Fourier expansion in translates on the dual half of the vacuum.

Constant terms of the normalized elliptic translates. The base theta model has the odd sigma section s_Λ and its translate θ by an order-two point t_0, normalized by $\theta(0) = 1$. The two components of t_0 under the two idempotents of $\mathcal{O}_K/2\mathcal{O}_K$ are nonzero. This is the condition that neither ωt_0 nor $\bar{\omega} t_0$ be a period for $\omega = (1 + \sqrt{-d})/2$. Write a_1 for the coefficient of the marked cubic $y^2 + a_1 xy = x^3 + x$, with $t_0 = (0, 0)$. Its discriminant and invariant are

$$\Delta = a_1^4 - 64, \qquad j = \frac{(a_1^4 - 48)^3}{a_1^4 - 64}.$$

The CM invariant j is an algebraic integer, so a_1 is algebraic. One can see integrality of j here from the cyclic modular polynomial for $k + \omega$ with a sufficiently large integer k. The kernel is cyclic since $k + \omega$ is primitive in $\mathcal{O}_K$ and the period ideal is invertible. Its norm $n = k^2 + k + (d + 1)/4$ is between k^2 and $(k + 1)^2$ for large k. On the diagonal the integral modular polynomial is monic up to sign. Indeed, in the product over the cyclic sublattices of index n, each factor $j(\tau) - j((u\tau + b)/v)$, $uv = n$, has distinct leading cusp exponents -1 and $-u/v$. Its leading coefficient is a root of unity. The integral leading coefficient of the product is then ± 1. Since the cyclic isogeny is an endomorphism, the diagonal polynomial vanishes at the given j.

The actual Δ is a unit at every finite coefficient embedding. At odd p, $|a_1^4|_p > 1$ would give $|j|_p > 1$. If $|\Delta|_p < 1$, then $|a_1^4 - 48|_p = |16|_p = 1$, also impossible. At $p = 2$, take a good model of the given j after a finite coefficient extension. For example take $y^2 + Bxy + y = x^3$ with $(B^4 - 24B)^3 = j(B^3 - 27)$. A root B of this monic integral equation is integral, and $\widetilde{B}^3 = 1$ in the residue field would give $1 = 0$ in the displayed equation. Thus $B^3 - 27$ is a unit. The CM endomorphisms of the generic curve extend over the good elliptic model by the Néron mapping property. The two idempotents modulo 2 split its finite flat group of 2-torsion into rank-two factors. On the geometric special fiber their tangent dimensions add to one, since the differential of doubling is zero. A rank-two factor is either étale with tangent dimension zero or connected with positive tangent dimension. Hence exactly one factor is étale. The generic point t_0 has a nonidentity component in that factor, whose specialization is injective on geometric points. In particular t_0 specializes away from the origin of the good model.

Translate this integral point to $(0, 0)$ in the good Weierstrass model. Its order two makes the coefficient of y zero, leaving an equation $y^2 + A_1 xy = x^3 + A_2 x^2 + A_4 x$ with integral coefficients. Its discriminant is $A_4^2((A_1^2 + 4A_2)^2 - 64A_4)$, so A_4 is a unit. Remove A_2 by $y \mapsto y + sx$ with $s^2 + A_1 s = A_2$ and scale by a fourth root of A_4. Both changes are integral with unit scaling. Two Weierstrass models of this normalized shape representing the same marked curve have change of scale u with $u^4 = 1$. Indeed the translations in x, y vanish at the marked point and the coefficient of x is one in both models. This gives $|\Delta|_2 = 1$ for the actual normalized cubic. In particular a_1 is integral everywhere.

For a coefficient embedding into $\mathbf{C}_p$ write $H(z) = \max(|s_\Lambda(z)|_p, |\theta(z)|_p)$, where absolute values are taken after the embedding. The two sections have disjoint zero sets, so $H(z) > 0$. Put $u = s_\Lambda(z)$, $v = \theta(z)$. The sigma duplication and theta addition formulas give

$$\theta(2z) = v^4 - u^4, \qquad s_\Lambda(2z)^2 = (uv)^2\big(4(v^4 + u^4) + (a_1 uv)^2\big).$$

For the second formula, use $v^2/u^2 = \wp(z) - \wp(t_0)$ in the equation for $\wp'(z)^2$ and the sigma duplication formula. The first is the even theta addition formula at equal arguments. These analytic identities extend through the zeros of u and v. Normalize u, v by one of maximal absolute value. If both output coordinates have smaller absolute value than the indicated fourth power, their residues satisfy $U^4 = V^4$ with both U, V units and $8 + a_1^2(U/V)^2 = 0$. Squaring the latter equation contradicts that Δ is a unit. The displayed formulas therefore give $H(2z) = H(z)^4$ at every coefficient embedding.

For any section f of degree $k > 0$ for the multiplier of θ^k, the quotient f/θ^k is a meromorphic function on the elliptic curve with pole of order at most k at t_0. Use the ordinary Weierstrass coordinates evaluated at $z + t_0$, whose poles are at t_0. Their monomials with pole orders $2i$ and $2i + 3$, together with the constant, give the basis with these pole bounds by Riemann–Roch. Write $u^2 = x(z + t_0)v^2$ and $G = y(z + t_0)v^3$.

After multiplying by θ^k, the basis elements are monomials in u^2, v, G of degrees $2, 1, 3$, respectively, where the cubic identity is $G^2 + a_1 u^2 v G = u^6 + u^2 v^4$. By the ultrametric inequality $|G|_p \le H^3$. It follows that $|f(z)|_p / H(z)^k$ is bounded by the maximum absolute value of its coefficients in this basis.

Let y be torsion modulo Λ and put

$$A_y(z) = \exp\bigl(-\pi(z\bar{y} + |y|^2/2)/\operatorname{area}(\Lambda)\bigr), \qquad R_y(z) = \frac{|A_y(z)|_p H(z+y)}{H(z)}.$$

Raise the two twisted translates to an even multiple of the order of y. Their multipliers then agree with that of the same power of θ, so the preceding basis gives a finite bound for $R_y(z)$, uniformly in z. The Appell–Humbert multiplier changes R_y by the absolute value of a root of unity when y is shifted by a period, and this absolute value is one. In fact the extra phase is an exponential of πi times a rational Riemann pairing because y is torsion. Also $R_y(z)^4 = R_{2y}(2z)$ by the duplication of H and $A_y(z)^4 = A_{2y}(2z)$. Hence the positive finite suprema of $R_{2^j y}$ are successive fourth powers of one another on a finite orbit. Each equals one. At $z = 0$, $H(0) = 1$, and the actual translated value $A_y(0)\theta(y)$ is integral at every coefficient embedding.

The complex conjugate bound requires only the action on these torsion values. Let n be the order of y and choose k divisible by n. The algebraic degree-k theta line is $\mathcal{O}(k[t_0])$. Identify its canonical section with θ^k and normalize its evaluation at zero to be one. Translation by y preserves this line up to isomorphism. The analytic operator $f(z) \mapsto A_y(z)^k f(z+y)$ is a scalar multiple of an algebraic theta-group lift. Its n-th power is a root of unity times the identity, by the period multiplier, so that scalar is algebraic. After a complex automorphism, the transported operator and the corresponding analytic operator for the transported torsion point differ by a root of unity since they lift the same translation and both have finite order. Evaluating on the normalized section at zero makes $(A_y(0)\theta(y))^k$ algebraic. Taking k-th roots of the absolute values proves the equality of the moduli with those in the transported normalized model. The period lattice and marked half period up to lattice translation are specified by the cubic invariants and its invariant differential. There are only finitely many conjugate models, since a_1 is algebraic, and the lattice translation phases have absolute value one. On each model the norm $e^{-\pi|z|^2/(2\operatorname{area}(\Lambda))}|\theta(z)|$ descends to the compact elliptic torus. Its maximum, over these finitely many models, gives one constant $B_d \ge 1$ for all torsion values and all their complex conjugates.

The rank-one theta product and its finite projection. Write $\nu = N_1$ in this paragraph and let ζ have exact order ν. Put $\eta = \sqrt{2\pi/\sqrt{d}}$. The rank-one standard vacuum uses the lattice $L = r_\infty(\mathfrak{a}_0) \subset \mathbf{C}^\Phi$, where $r_\infty(\gamma) = (\eta\tau(\gamma))_{\tau\in\Phi}$ and $\mathfrak{a}_0 = E \cap \alpha\widehat{\mathcal{O}}_E$ with the prescribed finite idele α. The latter has component $P'(\zeta)^{-1}$ at places dividing ω, and component one elsewhere, with $P = \Phi_\nu$. The relative different is $(P'(\zeta))$. Indeed the cyclotomic and quadratic discriminants are relatively prime, so the compositum has $\mathcal{O}_E = \mathcal{O}_K[\zeta]$. This follows locally by tensoring with the unramified factor at each rational prime. At any place where $P'(\zeta)$ has positive valuation exactly one of $\omega, \bar{\omega}$ lies in the place, since $\omega\bar{\omega} = (d+1)/4$ is divisible by every prime dividing ν and $\omega + \bar{\omega} = 1$. Consequently $\mathfrak{a}_0\bar{\mathfrak{a}}_0 = \mathfrak{D}_{E/K}^{-1}$. Thus L is self-dual for the alternating form of the standard Fock Hermitian form. In the E coordinates this is the trace pairing $\operatorname{Tr}_{E/\mathbf{Q}}(\bar{x}y/\sqrt{-d})$, using $\mathfrak{D}_{E/\mathbf{Q}} = \sqrt{-d}\,\mathfrak{D}_{E/K}$.

Write $\Lambda = \Omega_0\mathcal{O}_K$ for the period lattice of the principal marked model just used. Define the complex linear forms and product

$$\ell_i(z) = \frac{\Omega_0}{\nu\eta}\sum_{\tau\in\Phi} z_\tau\,\overline{\tau(\zeta^i)}, \qquad S(z) = \prod_{i=0}^{\nu-1}\theta(\ell_i(z))^\nu.$$

For $\gamma \in \mathfrak{a}_0$, we have $\ell_i(\nu\, r_\infty(\gamma)) = \Omega_0 \operatorname{Tr}_{E/K}(\gamma\bar{\zeta}^i) \in \Lambda$, by the inverse different. The row orthogonality of the finite Fourier matrix at the distinct $\tau(\zeta)$, and $\operatorname{area}(\Lambda) = |\Omega_0|^2\sqrt{d}/2$, show that the sum over i of the pullbacks of the elliptic Hermitian form, each counted ν times, is the standard Fock Hermitian form.

The multipliers match too. On the elliptic lattice the theta multiplier is $\exp(\pi i E_\Lambda(\bar\omega l, l))$, where E_Λ is the Riemann form. To check this formula, on either idempotent line modulo 2, the sigma sign cancels the sign from translating by the mixed point t_0. Hence the multiplier is one on $\omega\Lambda$ and $\bar\omega\Lambda$. The Appell–Humbert law and $l = \omega l + \bar\omega l$ give the formula. The ℓ_i commute with the common scalar $\bar\omega$, so multiplying the elliptic phases gives the canonical multiplier of the standard vacuum on vL. In particular S is a section on this lattice and $S(0) = 1$.

Let Θ_{vac} be the entire theta function of the standard vacuum on L. Use the trace-one element ω in the adelic realization. For $\gamma \in E$ the rational translation formula takes the form

$$\Theta_{\varrho(\mathrm{finVec}(\alpha^{-1}\gamma))F} = U_{-r_\infty(\gamma)}\Theta_F, \qquad \chi_f(\alpha^{-1}\gamma) = \chi_\infty(-r_\infty(\gamma))^{-1},$$

where finVec places the vector in the finite adeles, ϱ is the Heisenberg translation, and χ_f, χ_∞ are the finite and canonical archimedean multipliers, and U_r is archimedean Heisenberg translation with its Gaussian factor. These identities use invariance under the rational Heisenberg translations and the product formula for the additive character, with the archimedean Gaussian divided out. If $\Theta_F = 0$, then F vanishes at the real-only arguments. Rational covariance moves these to arguments with dense finite components, so continuity gives $F = 0$.

The theta space on vL has dimension v^{2f}, since $(vL)^\vee = v^{-1}L$. The translates of Θ_{vac} indexed by $v^{-1}L/L$ are eigenvectors with distinct characters of L/vL, for the translations divided by their vacuum multiplier. The characters come from the perfect Riemann pairing of these two finite groups. The translates are therefore a basis. By the rational translation formula the expansion of S in this basis realizes it as the entire theta function of one vector F_s in the finite part of the actual standard model. To obtain the vacuum, choose a translate of S whose coefficient on Θ_{vac} is nonzero, and apply the sum of the normalized translations in L/vL. This kills the nontrivial characters. It gives a nonzero scalar multiple of the vacuum as a sum of v^{2f} finite Heisenberg translates of F_s with root-of-unity coefficients.

At any argument x of the finite adeles, $F_s(\mathrm{finVec}\, x) = \Theta_{\varrho(\mathrm{finVec}\, x)F_s}(0)$. For $x = \alpha^{-1}\gamma$, the rational formula gives the product of elliptic translated values at the arguments $-\ell_i(r_\infty(\gamma))$. They are torsion modulo Λ because their traces lie in K. The raw finite translation is locally constant on F_s, so density allows every finite adelic argument to be replaced by a rational one with the same value. Compositions of finite translations add only root-of-unity phases: for each argument of the finite additive character some integer multiple belongs to its open kernel. This gives the integrality and conjugate bounds on the finite adeles for every summand of the fixed vacuum projection.

The uniform bound for scaled vacua. The geometric input on this side concerns a single base K-CM elliptic theta model. For the standard rank-one space over E, it asks for one nonzero scalar c_s at each s, common to every finite adelic evaluation. All values of c_s times the vacuum are integral at every coefficient place, and every conjugate over K has modulus at most $\exp(C_{\mathrm{vac}}N_1^2)$, with C_{vac} independent of s. The standard rank-n formula uses the product of the same rank-one normalization in the n coordinates. A complex automorphism either fixes K or does so after composition with complex conjugation, which does not change the modulus.

The special theta product has N_1^2 of the preceding elliptic torsion factors in rank one. A nonzero multiple of the vacuum is a sum of N_1^{2f} Heisenberg translates of the product, with coefficients roots of unity. Thus the archimedean bound for the scaled vacuum is

$$N_1^{2f} B_d^{N_1^2} \le \exp\big((2 + \log B_d)N_1^2\big),$$

since $f \le N_1$ and $\log N_1 \le N_1$. The same finite sum proves integrality using the normalized elliptic torsion values at every coefficient embedding.

Conjugating the normalized Taylor series. Use the marked elliptic curve and the standard Fock coordinates of the vacuum construction. Write $\eta = (2\pi/\sqrt{d})^{1/2}$, $\nu = N_1$, and A_σ for the area of the marked conjugate elliptic lattice Λ_σ in its algebraic differential coordinate. For $p \in \mathbf{C}^\times$ write T_pF for the Taylor series of $\Theta_F(pZ)$. We use the scale

$$\mathrm{per} = \frac{s_0}{u_0}\frac{\eta}{\Omega_0}, \qquad s_0^f = \chi_{\mathrm{fin}}(\alpha)^{-1}, \qquad A_\sigma|\sigma(u_0)|^2 = \pi e^{\ell_d}\frac{2\pi}{\sqrt{d}}.$$

Here $\chi_{\mathrm{fin}} = v_{\mathrm{fin}}|\cdot|_f^{-1/2}$ is the finite multiplier character of the Weil model. The root s_0 and the area equalizer u_0 are chosen once. These are the normalizations of the weighted jet at every stage. We give the area and transport calculations explicitly.

The ideal class group of K has odd order h_K by genus theory. The main theorem of complex multiplication for an elliptic curve with full endomorphism ring $\mathcal{O}_K$ identifies the action of $\mathrm{Gal}(K^{\mathrm{Hil}}/K)$ on its j-invariant with multiplication by inverse ideal classes on the homothety class of its period lattice, where K^{Hil} is the Hilbert class field. It also identifies the finite Tate-module action in each chosen isogeny identification with the action of the corresponding finite idele. Apply the theorem to the principal lattice $\Lambda = \Omega_0\mathcal{O}_K$: its conjugate j-invariants over K are precisely the h_K maximal-order lattice classes. The archimedean CM type uses the chosen embedding of K, hence the reflex field is K and in this rank-one case the reflex norm is the identity. Milne, Introduction to Shimura Varieties, §11 and Theorem 11.2. Choose a marked conjugate curve for each class c, with lattice Λ_c and discriminant Δ_c in the differential coordinate of its normalized cubic. The good-reduction argument for the base cubic applies to its conjugates. In particular each Δ_c is a unit at all finite places. Write $\Lambda_c = u_c\mathfrak{d}_c\Lambda$, with $\mathfrak{d}_c = \mathfrak{a}_{c^2}\mathfrak{a}_c^{-1}$ for a system of ideal-class representatives $\mathfrak{a}_c$, and Λ the base lattice. These ideals represent the classes c and $\prod_c \mathfrak{d}_c = \mathcal{O}_K$. For $v \in \mathfrak{d}_c$ the scalar $u_c v$ is the differential of an isogeny to the c-curve. For $v' \in \mathfrak{d}_c^{-1}$ the scalar $u_c^{-1}v'$ is the differential of an isogeny in the reverse direction. These differentials are algebraic and integral on the specified good models at every finite place, by extension of isogenies to the elliptic schemes. Taking finite sums of products of the v's, and separately of the v''s, proves that $\prod_c u_c$ is a unit at every finite place. Applying σ to the same isogenies gives both lattice inclusions, hence $\Lambda_c^\sigma = \sigma(u_c)\mathfrak{d}_c\Lambda_\sigma$. The discriminant-normalized area of any lattice L' is $|\Delta(L')|^{1/6}\,\mathrm{area}(L')$. It depends only on the shape. The mean over the ideal classes is computed below. Consequently any root u of

$$u^{12h_K} = \prod_c u_c^{12}\Delta_c$$

has the required area property and is a unit at every finite place. The comparison uses these properties for the already fixed u_0 and the multiplier identity for the already fixed s_0 in per.

In detail, for a basis proportional to $1, \tau$, with $y = \Im\tau$, the discriminant-normalized area divided by π is $4\pi y|\eta_{\mathrm{D}}(\tau)|^4$, where η_{D} is Dedekind's eta. For the ideal lattice of this shape the partial ideal zeta function is

$$\frac{(2/\sqrt{d})^z}{|\mathcal{O}_K^\times|}\sum_{(j,k)\neq(0,0)}\frac{y^z}{|j\tau+k|^{2z}}.$$

For this normalization, Kronecker's first limit formula gives the ratio of constant term to residue at $z=1$ as $\log(2/\sqrt{d})+2\gamma-2\log 2-2\log(\sqrt{y}|\eta_{\mathrm{D}}(\tau)|^2)$. One can obtain the formula directly by Poisson summation in k for each $j \neq 0$. Before the prefactor $(2/\sqrt{d})^z/|\mathcal{O}_K^\times|$, the zero frequencies give

$$2\zeta(2z)y^z + \frac{2\sqrt{\pi}\,\Gamma(z-\frac{1}{2})}{\Gamma(z)}\zeta(2z-1)y^{1-z}.$$

For the nonzero frequencies the Mellin integral of the shifted Gaussian is uniformly convergent near $z=1$, with exponential decay in $|nj|y$ for frequency $n \neq 0$. At $z=1$, the Fourier transform of $y/(v^2+j^2y^2)$ at n is $\pi e^{-2\pi|nj|y}/|j|$, by integrating over the simple pole. Summing the terms with $j, n \neq 0$ gives $-4\pi\sum_{n\geq 1}\log|1-$

$e^{2\pi i n\tau}|$. The constant $2\zeta(2)y = \pi^2 y/3$ supplies the leading exponential in Dedekind's eta product. The other zero frequency has residue π and constant term $\pi(2\gamma - 2\log 2 - \log y)$, by the gamma duplication formula. Multiplying by the prefactor proves the stated ratio. The residues are the same in every class and the sum is $\zeta(z)L(z, \chi_d)$. The mean log normalized area divided by π is therefore $\ell_d + \log(2\pi/\sqrt{d})$. Also, for a finite idele a_f, choose $k > 0$ such that the fractional ideal of a_f^k is principal, say generated by a. The splitting character is unramified at the finite places and its archimedean component is $\prod_\tau \tau(a)/|\tau(a)|$. Triviality on $E^\times$ and the product formula give $\chi_{\mathrm{fin}}(a_f)^k = \prod_{\tau\in\Phi} \tau(\bar{a}) \in K$. Thus the modulus of $\chi_{\mathrm{fin}}(a_f)$ is preserved under all σ fixing K. Since $\mathfrak{a}_0\overline{\mathfrak{a}_0} = \mathfrak{D}_{E/K}^{-1}$, $|\alpha|_f^{1/f} = \mathrm{rd}_E/\sqrt{d}$, so $|\sigma(s_0)|^2 = \mathrm{rd}_E/\sqrt{d}$.

We next identify the full Taylor series after conjugation. The same normalization transports the torsion germs. For the odd reduced sigma section f_Λ with $f_\Lambda'(0) = 1$, the exponent subtracted from the logarithm of the usual sigma function is $s_2 z^2/2$, where $\eta_\Lambda(\lambda) = s_2\lambda + \pi\bar{\lambda}/\operatorname{area}(\Lambda)$ is the zeta quasiperiod. The de Rham class $[\wp(z)\,dz] + s_2[dz]$ has periods proportional to $\bar{\lambda}$. It is the eigenspace of a nonreal CM endomorphism for the eigenvalue on antiholomorphic differentials. In the algebraic de Rham basis $[\wp(z)\,dz], [dz]$ this determines s_2 algebraically and commutes with σ fixing K. The Taylor coefficients of the usual sigma function follow from the algebraic Weierstrass coefficients and $(\log\sigma_\Lambda)'' = -\wp$, with the odd, derivative-one normalization. Thus f_Λ transports coefficientwise.

For a torsion lift x choose $m_0 \geq 2$ such that $m_0 x/2$ is a period. Set $u = x/(2m_0)$, $M_0 = m_0^2$, and let U_x be reduced translation. The functions $D_j(z) = f_\Lambda(jz)/f_\Lambda(z)^{j^2}$ are algebraic elliptic functions by their prescribed torsion divisors and leading coefficient j. If u is a period, x is an even period and acts trivially. Otherwise the multiplier, oddness, and $M_0 u \in \Lambda$ give, with $B_u = (U_u f_\Lambda)^{4M_0}$,

$$\begin{aligned} U_x f_\Lambda(2m_0 z) &= B_u(z) D_{2m_0}(z+u), \\ B_u(z)^{M_0} &= f_\Lambda(2M_0 z)/D_{2M_0}(z+u), \\ B_u(0) &= -D_{M_0-1}(u)/D_{M_0+1}(u). \end{aligned}$$

The second quotient is a unit series. The indicated constant selects its unique M_0-th root, so translating and conjugating the series agrees once the deeper torsion lifts are tracked. The even section is $U_{t_0} f_\Lambda/(U_{t_0} f_\Lambda)(0)$ for the marked half-period t_0. Its conjugate uses the transported marked point. The composition scalar of two reduced translations is $\exp(\pm\pi i E_\Lambda(x, y))$ for their alternating Riemann pairing E_Λ, with the sign determined by the order of the two translations. On rational lifts it is a Weil pairing at a sufficiently divisible torsion order, with one argument possibly an integral multiple of the chosen smaller torsion label. For example if $mx, my \in \Lambda$, the pairing of $x/(2m)$ and $y/(2m)$ at order $2m^2$ gives this scalar or its inverse. These deeper lifts can be tracked together with the division-polynomial arguments. Choosing a different period lift of the marked point multiplies $U_{t_0} f_\Lambda$ by a constant which cancels in its normalization.

For each σ fixing K, choose $\Lambda_\sigma = \lambda_\sigma \mathfrak{b}_\sigma$ with $\mathfrak{b}_\sigma$ a fractional ideal of norm one. Such a representative is $\mathfrak{c}/\bar{\mathfrak{c}}$ since every ideal class is a square. Thus $|\lambda_\sigma|^2 = 2A_\sigma/\sqrt{d}$. The action on torsion commutes with $\mathcal{O}_K$, as the differentials of the CM endomorphisms are fixed by σ. The Tate modules of $K/\mathcal{O}_K$ and $K/\mathfrak{b}_\sigma$ are the respective completions of these two rank-one lattices. Their induced isomorphism is multiplication by an idele $\beta_\sigma \in \mathbb{A}_{K,f}^\times$ with $\beta_\sigma \widehat{\mathcal{O}}_K = \widehat{\mathfrak{b}}_\sigma$. Use representatives $v \in K$ as close to β_σ as needed for any finite list of deeper torsion lifts. This represents the image of $\Omega_0 x$ by $\lambda_\sigma v x$ simultaneously for that list. Compatibility with the Weil pairings gives $\beta_\sigma\bar{\beta}_\sigma = \chi_{\mathrm{cyc}}(\sigma)$ as an element of $\widehat{\mathbf{Z}}^\times$. Indeed, in the norm-one ideal frame the alternating pairing on rational lifts is the base pairing multiplied by $v\bar{v}$, and the equality of Weil pairings at all orders determines its residue units at every rational prime.

For standard rank r_0, use the product $S(Z)$ from the vacuum construction in all r_0 coordinates. Let F be its vector. The finite translation projection expresses the vacuum as a finite combination of translates of F. These translates therefore span the finite part. Put

$$S_\sigma(Z) = \prod_{k=1}^{r_0} \prod_{i=0}^{\nu-1} \theta_\sigma\left(\frac{\lambda_\sigma}{\nu\,\eta} \sum_{\tau\in\Phi} Z_{\tau,k}\,\overline{\tau(\zeta^i)}\right)^\nu .$$

Choose a positive integer e with $e \in \mathfrak{b}_\sigma$. On $v\,eL$, in each copy of the standard lattice L, the arguments are periods. The Hermitian form is the standard form by the same Fourier orthogonality since $A_\sigma = |\lambda_\sigma|^2\sqrt{d}/2$. The multiplier computation for the even marked section uses the same two idempotent lines modulo 2 and gives the canonical multiplier on $v\,eL$. As for S, the finite translation basis on this lattice realizes S_σ as Θ_{F_σ} on the same standard finite part.

At a rational finite argument $\alpha^{-1}\gamma$ the translation of S translates its i,k-th elliptic factor by $\Omega_0 x_{ik}$, with $x_{ik} = -\operatorname{Tr}_{E/K}(\gamma_k\bar{\zeta}^i)/v$. Applying per to the complex variables leaves the linear coefficients $(s_0/u_0)\overline{\tau(\zeta^i)}/v$. Their conjugates are the coefficients for S_σ at

$$t_\sigma = \sigma(s_0/u_0)\eta/\lambda_\sigma, \qquad |t_\sigma|^2 = \frac{\mathrm{rd}_E}{\sqrt{d}}\frac{\pi}{A_\sigma|\sigma(u_0)|^2} = \Omega^{-1},$$

after the permutation $\tau \mapsto \sigma\circ\tau$. Take v representing all the elliptic translates and their deeper torsion arguments and sufficiently close to β_σ that the finite translates of F_σ at $\alpha^{-1}v\gamma$ and at $\beta_\sigma\alpha^{-1}\gamma$ coincide. This is possible by local constancy. The conjugation of the elliptic germs then gives

$$\mathrm{rename}_\sigma\big(\sigma\,T_{\mathrm{per}}\varrho(\mathrm{finVec}\,x)F\big) = T_{t_\sigma}\varrho(\mathrm{finVec}(\beta_\sigma x))F_\sigma.$$

The equality holds at all finite arguments x by density of the rational arguments and local constancy on both sides. Taylor series at nonzero scale are injective on finite vectors, by entire continuation and rational translation covariance. The formula thus extends by semilinearity to a map L_σ on the finite part, with $\mathrm{rename}_\sigma(\sigma\,T_{\mathrm{per}}G) = T_{t_\sigma}L_\sigma G$. It respects the finite translations through $x \mapsto \beta_\sigma x$. Indeed the Heisenberg cocycle and the finite normalization phases are finite trace-character values, and $\sigma(\psi_f(z)) = \psi_f(\beta_\sigma\bar{\beta}_\sigma z)$. The covariance follows first on the translates of F, then on their span.

The moved vacuum and its entire majorant. Write $M_f = \widehat{\mathcal{O}}_E^{r_0}$ in the standard finite coordinates, v_0 for the unit vacuum and $T(x)$ for the vacuum-normalized finite translations. For g finite unitary the common fixed line of $T(x)$, $x \in gM_f$, is $\mathbf{C}\omega(g)v_0$. This follows from the corresponding statement for M_f, where averaging $T(x)$ is the vacuum projector, and from $T(gx)\omega(g) = \omega(g)T(x)$. Write $\beta_\sigma = z a_f$ with $z\bar{z} = 1$ and $a_f \in \widehat{\mathcal{O}}_E^\times$. At a nonsplit pair the valuation of β_σ is zero, and at a split pair its two valuations are opposite, so such a factorization is obtained by retaining just these valuations in z. With the same symbol for the central unitary scalar matrix, $L_\sigma v_0 = c_\sigma\omega(z)v_0$ with $c_\sigma \neq 0$. For every g, $L_\sigma\omega(g)v_0 = c_{\sigma,g}\omega(gz)v_0$. To compare the scalars, the pairing against v_0 on $T(x)v_0$ is $1_{M_f}(x)$. Pairing the image against $\omega(z)v_0$ gives $c_\sigma 1_{M_f}(x)$. By semilinearity this compares the pairings on the entire finite part. Apply it to $\omega(g)v_0$. The coefficient $\langle v_0, \omega(g)v_0\rangle = 1/\overline{\Delta_f(g)}$ is nonzero, where $\Delta_f(g) = \chi_{\mathrm{fin}}(\det u_f)$ and $u_f M_f = M_f \cap g^{-1}M_f$. For this normalized Weil formula, one averages over the translates $T(gx)v_0$ indexed by $M_f/(M_f \cap g^{-1}M_f)$. The averaging scalar is $\Delta_f(g)/|\Delta_f(g)|^2$, and only the identity coset survives in the pairing. Both $\Delta_f(g)$ and $\overline{\Delta_f(g)}$ have unchanged modulus under σ by the principal-ideal calculation above. Thus $|c_{\sigma,g}| = |c_\sigma|$.

Let $v_{\mathrm{int}} = c_s^{r_0}v_0$ be the scaled vacuum already chosen in the definite chain. The constant term of the transport, after finite translation, is the conjugate value of v_{int} at the corresponding finite point. Because β_σ is invertible, these values bound $L_\sigma v_{\mathrm{int}}$ at all finite arguments by $\exp(C_{\mathrm{vac}}v^2 r_0)$. The squared Hilbert norm of a theta vector is its normalized mean squared value on the adelic vector quotient by the rational vectors. In a Schrödinger polarization this follows by integrating the squared rational summation first in the dual translation variable, which kills the cross terms, then in the remaining translation variable. Absolute values descend to this quotient, and the finite-only arguments have dense image by archimedean approximation with rational vectors. Hence the Hilbert norm is bounded by the finite supremum. Since the unit vacuum and its unitary moves have Hilbert norm one, $|\sigma(c_s^{r_0})c_\sigma| \le \exp(C_{\mathrm{vac}}v^2 r_0)$.

For a unit moved vacuum one has the uniform pointwise bound

$$\sup_y |(\omega(g)v_0)(y)| \le \exp(O_d((f r_0 + 1)^2)).$$

Here is a proof retaining the measure and lattice of this model. Set $w = \varrho(y)\omega(g)v_0$. The squared value at y is $\sum_{\gamma\in E^{r_0}}\langle w, \varrho(\gamma)w\rangle$ with the prescribed rational embedding. In a rational Schrödinger polarization this identity is Poisson summation in the dual rational variable of the matrix coefficient, followed by rational summation in the other variable. The self-dual measures give quotient mass one. The translations by y affect the coefficients only by phases. Unitarity and the vacuum matrix coefficient bound their moduli by

$$1_{\{\alpha^{-1}g^{-1}\gamma_f\in M_f\}}\,\exp\left(-\frac{\pi}{\sqrt{d}}\sum_{\tau,k}|\tau(\gamma_k)|^2\right).$$

Two distinct vectors in the indicated lattice have distance at least one in the unscaled complex coordinates. For their difference γ_0, the finite conditions, unitarity, and $\alpha\bar{\alpha}\widehat{\mathcal{O}}_E = \mathfrak{D}_{E/K}^{-1}\widehat{\mathcal{O}}_E$ put $\sum_k \bar{\gamma}_{0,k}\gamma_{0,k}$ in the inverse relative different. Its trace to K is a positive rational integer equal to the squared distance. In real dimension $D_0 = 2fr_0$, a unit-separated set has at most $(2R+1)^{D_0}$ points in a ball of radius R. Bounding $e^{-cr^2} \le e^{(D_0+1)^2/(4c)}e^{-(D_0+1)r}$ and summing over integer shells proves the bound.

Consequently the entire theta of the transported vector satisfies

$$\left|\Theta_{L_\sigma\varrho(\mathrm{finVec}\,x)\omega(g)v_{\mathrm{int}}}(Z)\right| \le \exp\left(C_d(\nu r_0+1)^2 + \frac{\|Z\|^2}{2}\right). \tag{4.2}$$

One can take $C_d = C_{\mathrm{vac}} + 2d + 6$. In the squared moved-vacuum bound the count over integer shells is at most $3^{D_0}(1-e^{-1})^{-1}$ after the factor $\exp((D_0+1)^2\sqrt{d}/(4\pi))$ is removed. Taking the square root and using $f \le \nu$ gives the term $2d+6$ in this choice. Include C_d in the fixed constant of the scale condition $10^{12}C_0 \le \log\log b$. The factor $\exp(\|Z\|^2/2)$ removes the unit real Gaussian from the theta value. The finite translations move the evaluation point and introduce a phase. Each extra determinant character in a definite pair is algebraic of its prescribed weight, with modulus one after conjugation. To see this on a finite norm-one element, pass to a power whose ideal is principal, generated by a. The product $a\bar{a}$ is a unit. On doubling that power one can therefore use $a/\bar{a}$, a rational norm-one generator, and an integral compact factor. The compact character value is a root of unity by smoothness, and the rational norm-one point contributes only powers of conjugates of complex modulus one. Thus (4.2) applies to each moved vacuum on the standard Hom space of the definite pair. For a complex automorphism not fixing K, compose on its output with complex conjugation before taking the norm. The Weyl Gram matrix has real integral entries, so this preserves its norm as well as the scalar moduli.

4.3 The terminal restriction and the lifted forms

The integral terminal branching map. Use $n = n_m - 1$, and let $\beta_{\tau,h}$ be the number of kept columns of upper height h. Let $\gamma_{\tau,h}$ be the remaining upper columns, whose height drops by one. Thus the lower height-h count is $\mu_{\tau,h} = \beta_{\tau,h} + \gamma_{\tau,h+1}$. For the weighted shape $\beta_{\tau,h} = 0$ for $h < \kappa$ and for $h = n$. The second assertion follows since the positive kept heights are $\kappa \bmod t+1$, whereas $n = \kappa + m(t+1) - 1$. The determinant exponents on the two blocks therefore agree. Write J for the total number of lower boxes in heights $h \ge \kappa$. There is a nonzero integer branching matrix in the evaluation coordinates, with Weyl operator norm at most $\max(1,J)^{3J}$. We construct it at the outset so the same matrix enters the seesaw and the two length and index estimates.

We recall the evaluation-coordinate norm. For a shape ν' let $b_i(X)$ be the standard products of flag minors, and write $e_{\nu'}(X) = (b_i(X))_i$. Integral straightening expands every ordered product of flag minors of this shape with integer coefficients in the b_i. In the tensor product of exterior powers with its standard orthonormal exterior bases, the coordinate vector $w_{\nu'}(X) = \bigotimes_{\text{columns } h}(X_1\wedge\cdots\wedge X_h)$ is therefore $C_{\nu'}e_{\nu'}(X)$ for an integer matrix $C_{\nu'}$. It contains the identity rows indexed by the standard products. In particular its integral image in the tensor space is saturated. The kernel of the Weyl form is

$$w_{\nu'}(X)^{\mathsf{T}}w_{\nu'}(Y) = e_{\nu'}(X)^{\mathsf{T}}G_{\nu'}e_{\nu'}(Y), \qquad G_{\nu'} = C_{\nu'}^{\mathsf{T}}C_{\nu'},$$

as follows also by Cauchy–Binet. This identifies the tensor metric on the generated irreducible with the Weyl metric used on the definite values. The highest evaluation $X = 1$ has norm one. The associated polynomial of ξ is $e_{\nu'}(X)^{\mathsf{T}} G_{\nu'} \xi$. The normalized Fischer form on this polynomial model agrees with the Weyl form by unitary invariance and the highest vector normalization. If $H_{\nu'}$ is the squared Fischer norm of the highest product of minors, this shows that the integral matrix of Fischer pairings of the b_i, multiplied by $G_{\nu'}$, is $H_{\nu'}$ times the identity. Consequently $H_{\nu'}$ times any integral product of minors is the associated polynomial of an integral Weyl vector. For a diagram with Q boxes the monomial factorial is at most $Q!$. The sum of absolute highest-product coefficients is at most the product of the column-height factorials, also at most $Q!$. This gives $H_{\nu'} \le (Q!)^3$.

At each place insert the last unit vector e_* at the front of every shifted column and include the kept columns in the upper exterior space. Tensoring gives an isometry I of the lower tensor space into the upper tensor space. It is equivariant for $\mathrm{diag}(g, s)$ with the extra character $s^{\sum_h Y_{\tau,h}}$. Its bilinear kernel on evaluation vectors is

$$K(X, Y) = \langle I w_\mu(X), w_{\mathrm{upper}}(Y) \rangle_{\mathrm{bil}}.$$

It is the value at $\mathrm{diag}(X^{\mathsf{T}}, 1)Y$ of the product of upper minors with the added row first in the shifted columns and last in the kept columns. At $X = 1$ this product of flag minors is a nonzero polynomial.

For a matrix Y with last row $(y, 0, \dots, 0)$, expand each shifted minor in that row. Apart from the monomial $y^{\sum_h Y_{\tau,h}}$, the result at this place is the lower flag product at $Y_{\mathrm{old}}^{\mathsf{T}} X$, where the height-$h$ shifted rows are $2, \dots, h+1$ and the kept rows are $1, \dots, h$. Here Y_{old} is the n-by-n block in the old rows and the first n columns. Set $p = \kappa - 1$. After permuting the rows to put $2, \dots, p+1$ first, each column of height h contains its first $\min(p, h)$ standard rows. The remaining boxes of this lower diagram form a diagram $\mu^\flat$ on the last $n - p$ rows. Its total size over all places is at most J.

To verify that the denominator for this product is controlled by the remaining boxes, use the tensor realization. In each long column wedge the first p standard vectors onto the tail column, and in a short column insert its usual highest exterior vector. This is an integral isometry of the tail tensor space into the lower tensor space. The lower Weyl space intersects this inserted tensor space in exactly the inserted tail Weyl space. Indeed the prescribed tensor basis vectors, with these leading rows, are the simultaneous weight space of diagonal actions on the first p rows for the weights of the highest vector on these rows. The first row must appear in every column. Inductively the next row must appear in every column tall enough to contain it. This gives exactly the indicated leading rows. The projection to these weights sends $w_\mu(X)$ into the image of the tail Weyl space. For X with nonsingular first p-by-p block this follows by subtracting combinations of its first p columns from later columns to clear the upper p entries. The initial-column wedges are unchanged. Their projections are the leading minor scalars times the wedges of initial columns of the lower-right Schur complement. These X form a dense set. Conversely the image of the tail Weyl space is generated by the lower block orbit of the lower highest tensor. Thus orthogonal projection of an inserted exterior basis vector to the lower Weyl space equals the insertion of its projection to the tail Weyl space. The integral lattices in both Weyl spaces are saturated in their tensor spaces by the identity rows of the $C_{\nu'}$. The preceding Fischer identity for the tail diagram now shows that $c = H_{\mu^\flat}$ times the lower flag product is the associated polynomial of an integral lower Weyl vector. The row permutation and substitution by $Y_{\mathrm{old}}^{\mathsf{T}}$ act integrally on this polynomial model. In the tensor description the transpose substitution acts on the vector through the exterior powers of Y_{old}. It preserves the generated Weyl space and has integer polynomial coefficients in the ambient tensor basis. Saturation, with the coordinate left inverse given by the identity rows, gives the same integrality in the Weyl model. We have $c \le \max(1, J)^{3J}$.

For unrestricted generic Y, multiply the first column by its last entry y, and replace column $i > 1$ by $y Y_i - Y_{*,i} Y_1$. These are integral upper triangular column substitutions. The bottom row is now zero off the first column. Every flag minor of height h has been multiplied by y^h. It follows that a monomial in the y's times $cK(X, Y)$ has every coefficient in Y in the integral lower Weyl polynomial lattice in X. Comparing coefficients with their exponents shifted by this monomial removes the factor without introducing any

denominator. Finally, the coefficient vectors of the $e_{\mathrm{upper}}(Y)$ span the full upper coordinate lattice over $\mathbf{Z}$. For a diagonal term order the leading monomials of the standard products are distinct with coefficient one, giving a unitriangular coefficient submatrix. It follows that cI^{*} restricted to the upper Weyl space, with I restricted to the lower Weyl space in the adjoint, has integer matrix B in the evaluation coordinates. Its norm is at most c.

The last-coordinate character, including the common determinant shift $e_m(\tau)$, is $s_\tau^{b_\tau}$, $b_\tau = \sum_h \gamma_{\tau,h} + e_m(\tau)$. The ordinary multiplicity-one branching rule follows by removing the entries $n+1$ from a semistandard tableau of the upper partition, leaving a horizontal strip. This proves that the constructed nonzero map spans the branching line used in the terminal period. Use one such matrix in the restriction $\prod_\tau \tau(\bar\pi)^{-b_\tau} B$. Changing the generator once changes the whole restricted definite chain by one common scalar, absorbed in c_* in (4.3), retaining the tensor Ψ_ϑ for each twist. This also holds for the paired-chain expression, since both expressions apply the restriction to the same terminal form. At a coefficient embedding above p, the finite terminal gauge uses the inverse upper representation of the frame $\mathrm{diag}(g, s_\tau)$. In the resulting value B has matrix norm at most one over the valuation ring and contributes the last-coordinate factor $s_\tau^{-b_\tau}$. Off N the $\tau(\bar\pi)$'s are units and the ordinary frame uses $s_\tau = 1$. At N the adapted upper frame has $s_\tau = N$ when the place induced by τ contains $\bar\pi$, and $s_\tau = 1$ otherwise. Thus $|\tau(\bar\pi)|_p = |s_\tau|_p$, and the restriction's normalizing character agrees in absolute value with the factor in the upper gauge. The lower gauge is therefore bounded by the upper gauge at that adapted frame. This is valid for each embedding, including its composition with conjugation.

At a top, the a height-one columns all shift, and $e_m(\tau) \ge -w_m(\tau) \ge -2b$ at every place, so $b_\tau \ge a 1_{\tau\in\mathcal{T}} - 2b$. Since $\pi\bar\pi = N$, the restriction costs at most $-Ta\log N + 2fb\log N + 6J\log\max(1,J)$ in the logarithm of the squared norm at every embedding. Here $J = O(f n_m^2 r^8)$, since the large height-κ shifted part moves to height $\kappa - 1$ and only the kept part and the small higher columns remain at heights at least κ. Hence the final term is $O(x^{70})$ on the weighted range.

The following theta operators use the same Hermitian and oscillator frames throughout. For a lower edge $Y_j \to Y_{j+1}$, put $v_j = \mu_{j+1}^0 \vartheta^{j+1}$ and give its source and target ranks R_j, R_{j+1} the determinant characters $v_j v_{U(1)}^{\lfloor R_j/2\rfloor}$ and $v_j v_{U(1)}^{\lfloor R_{j+1}/2\rfloor}$, respectively. The first multiplies the target determinant and the second the inverse source determinant in the oscillator. At infinity use the joint highest polynomial for the signed parameters on the two groups, with its initial coefficient fixed once. Its construction in the joint harmonics is given with the signed type calculation below. The finite vector names the indicator of I_j modulo $\mathfrak{M}Q$ on the standard vacuum. At every lower stage let J_j be the kernel in K_j^0 of reduction at the chosen and twist places. Give K_j^0 mass one and multiply the ordinary theta integral on the source automorphic quotient by $[K_j^0 : J_j]$. At the bottom the input has level $J_g = \ker(K_0^0 \to \mathcal{G})$ on the stage-zero datum. At Y_m use the level with identity reduction at the chosen places and block upper triangular reduction at Q, for the blocks $r, t, \dots, t$. On the latter parabolic use the product of the stage tame characters of the added diagonal blocks, $\prod_{w\in Q}\prod_{l=1}^{m} \vartheta_w^{l}(\det k_{ll})$. The inclusion indicators transform by this character upon iteration: on an edge $k' I_j = I_j k$, the local ratio of determinants is precisely the determinant of its added block.

An m-fold twist ϑ is a weight-zero character of $U(1)(\mathbb{A}_{E^+})$ whose tame support Q is disjoint from its conjugate. Its norm $q_\vartheta = \prod_{v\in Q} \mathrm{N}(v)$ is prime to gN, and the j-th power of its tame character is nontrivial at every $v \in Q$ for $1 \le j \le m$. The trivial twist has $Q = \emptyset$. Let ι_ϑ be the chain of m theta lifts through Y_j with the corresponding twist powers and $\Theta_\vartheta = \iota_\vartheta|_W$ its restriction, realized at the final level.

For a tuple $t = (t_0, \dots, t_{m-1})$ of finite translation matrices in $\mathrm{Hom}(X_j, X_{j+1})(\mathbb{A}_{E,f})$, call it admissible at the integer u if $\pi_G t_j$ is integral at all finite places off u. Let $\mathscr{L}_{j,\vartheta}(\rho_f(t_j)v_{\mathrm{int}})$ denote the definite edge operator on the translated integral vacuum. It is the class sum of the rectangular Schur jet at $\bar\pi_G$ per, with the two rank-twisted characters as on the lower edge but in ranks n_j, n_{j+1}. The coefficient and its class measure are constructed below. For a finite linear combination Ψ of tuples, apply the definite compositions with coefficients $\Psi(t)$, and then the scaled branching restriction R_m. Denote this operator by $B_{\vartheta,\Psi}^{\mu}$. Joint level invariance means that acting on a tensor in consecutive Hom spaces by the same level element in

the two adjacent positions preserves the sum, with the last position fixed. The paired expression projects each translated vacuum to its integral vacuum part at the two halves above N. On each consecutive pair it then sums the opposed last-column and last-row translations, by $-y/N$ and its conjugate, over the residue vectors in the π-half of N.

The seesaw identity. There is one scalar $c_* \in \mathbf{C}^\times$ on the chosen datum D, and for each twist ϑ there is a good finite datum Ψ_ϑ, both independent of the ramified column. They give

$$\iota_\vartheta\big(\mathcal{P}(c_*\phi)\big) = B^{\mu}_{\vartheta,\Psi_\vartheta}(\phi) \qquad (\phi \in C_{\mathrm{ram}}). \tag{4.3}$$

The support is in the admissible translations for Nq_ϑ, the coefficients are integral at all coefficient places prime to q_ϑ, and the total absolute value of the conjugated coefficients is at most $e^{x^{60}}$ at the trivial twist, for every $\sigma \in \mathrm{Aut}(\mathbf{C})$. The tensor is jointly invariant at the column levels with last component the identity. Its chain also has the paired-chain expression using the projections onto the local vacua at the places over N. These last two properties allow the definite estimates in Theorem 2.

The functions to which the finite cuts apply. Let $F = E^+$, write $G_j = U(Y_j)$, and write $B_j \subset G_j$ for the Bessel subgroup of the fixed grading, with string length $\max(m-1-j, 0)$. Use the stage-zero level J_g from the finite chart. Put $J_j = J_j(\vartheta)$ for the level of the lower lift at stage j. Thus J_j is the kernel of the chosen and twist-place reductions inside the lattice stabilizer $K_j^0 \subset G_j(\mathbb{A}_{F,f})$. Write T_j for the unnormalized lower theta lift, with the harmonic at infinity and the finite inclusion datum of the lower edge. The actual arrow is $L_j = [K_j^0 : J_j]T_j$. Use Haar measure of mass one on K_j^0.

Let P_0 be the periodization of the seed of ϕ over $B_0(F)$. The seed in a class and Bessel coordinate is the prescribed finite coset indicator of J_g times the full real Bessel coordinate and the class cutoff β_c. The Poincaré series $\mathcal{P}\phi$ uses the same seed periodized over $G_0(F)$, with the class and Weyl-to-Bessel coordinates of the column right-translated by the fixed central lattice scalar. Choose representatives h_c of the double classes of the definite core at the column level K_0^X. Each seed coset is $h_c\eta^{-1}J_g$, a translate of J_g by a point of the embedded definite core. Its translates under $B_0(F)$ therefore have finite support in $B_0(\mathbb{A}_{F,f})J_g$. On the core the rational unipotent translates of the cutoff sum to one. Write y_0 in the stage-zero datum, so that $y_0 \in J_g$ amounts to $\eta^{-1}y_0\eta \in K(g)$ in the transported frame. The bottom normal form reads

$$P_0\big(\jmath_0(k)(\xi_0, y_0)\big) = \mathsf{A}\big(\rho_{X_0}(k_\infty^{-1})\phi(k_f)\big) \cdot \mathsf{W}(\xi_0).$$

Here $\jmath_0$ embeds the core, ρ_{X_0} is its conjugate Weyl action, A is the fixed change to Bessel coordinates, and the dot is bilinear with the product vector of real Bessel functions W. The central factor in the initial column coordinates cancels the inverse central factor in the seed coset dictionary. This supplies the same initial coefficient linear form for all twists.

5 The common real Bessel vector

Fix a real place of F, with chosen complex embedding Φ_v of E. All the constructions in this paragraph are made at that place. Write G_1 for the unitary group on the normalized frame, with ϱ hyperbolic pairs, a positive core X_0 of rank k and a last vector z of norm -1. Put $n = \varrho + 1$ and $r = k + 2\varrho + 1$. Give G_1 the Haar dy whose density at the identity is independent-entry Lebesgue measure on J_1X, for tangent X, divided by $\prod_{j=1}^{r} 2\pi^j/(j-1)!$. In the skew-Hermitian matrix J_1X use one complex coordinate for each off-diagonal pair and the imaginary part of its diagonal. This is the compact-dual normalization. Give $H = U(X_0)$ probability Haar. In the free coordinates for U, from row e_i take the complex entries in columns e_j, e_j^* with $j > i$ and in all columns of the core and z. Also take the imaginary entry in position (e_i, e_i^*). The unitary equations recover the other entries, with Haar du given by Lebesgue measure in these coordinates. The group law is triangular with unit absolute Jacobian, by multiplying successively up the

grades and using the Hermitian condition in the skew block. Use the Euclidean-unitary sign frame p_v, with $p_v^* J_1 p_v = J$, where J_1 is the normalized flag Gram matrix and $J = \operatorname{diag}(\varepsilon_i)$. Read the row coefficient of y at $p_v^{-1} y p_v$. Use $ds = du\, dh$ on $S = UH$. Write $d\dot{y}$ for the quotient Haar specified by $dy = du\, d\dot{y}$. Put $\operatorname{size}(y) = \inf_{s \in S}(1 + \|sy\|)$ in the operator norm. We use $\langle a, b \rangle = a^* b$. Let σ be the sum of the first superdiagonal entries on the frame $(e_0, \dots, e_{\varrho-1}, z)$, and put $\operatorname{wave}(c) = \exp(c - \bar{c})$. The basis scales at the frame entries are

$$s_{va} = \kappa_v^{\varrho - a} / \sqrt{N}, \qquad \kappa_v = g \Phi_v(\varpi)/(2\pi i) \in \mathbf{R}^\times.$$

Put $u = N_1/M$. Every prime dividing u divides M, so $\Phi_{N_1}(X) = \Phi_M(X^u)$. The different generator satisfies

$$|\Phi_v(\varpi)| = \sqrt{d}\, u\, |\Phi_M'(\Phi_v(\zeta)^u)|.$$

The derivatives on primitive M-th roots form a fixed finite set of nonzero values. With u polynomial in x and $b^{20} \le g \le b^{22}$, this gives $1 \le |\kappa_v| \le b^{23}$ uniformly for large x.

Conjugation by the inverse of the basis-scale matrix therefore divides the rational simple-root sum at v by κ_v. In these coordinates $\operatorname{wave}(\sigma)$ is exactly the real factor of the character with argument $\sigma/(g\varpi)$.

Let $m_\lambda(y)$ be the matrix coefficient of a unit highest vector in the lowest compact type of the row discrete series at v. Here λ is its Harish–Chandra parameter in the ordered unitary sign frame specified below. On H the column type is the symmetric power of exponent $a_v = \kappa(v, 1)$ in the column diagram. Its highest weight Λ' has entry a_v at the designated coordinate and zero elsewhere. The rank-k weight sums the columns of height strictly less than k. These have only the height-one contribution, and the determinant exponent at stage zero is zero.

The row coefficient on the actual sign frame. In that frame $G_1 = U(p, q)$, $p = k + \varrho$, $q = \varrho + 1$. Let $K = U(p) \times U(q)$ with probability Haar. Order the diagonal coordinates by decreasing λ_i, retaining the sign on each one. The first $k - 1$ entries, numbered from zero, are $k/2 - 1 - i$. The remaining entries are $-B_v - jh$, $1 \le j \le r - k + 1 = 2(\varrho + 1)$, where $B_v = w_0(v) - h_2$, $h_2 = r^8$ and $h = 2h_2$. Here k is even and r is odd. Since $x_1 \ge x$, the two rounded weights satisfy

$$\frac{xb}{x_1} \le B_v < \frac{a}{x} + N_1 < b + \frac{N_1}{x} + N_1 \le 2b$$

for all sufficiently large x, uniformly in v. Indeed $3^c \le 3^5 s$ and $2^s = m^{22}$ give $N_1 = O(x^{11} \log x)$, while $\log(b/x_1) = x^\varepsilon - x^{\varepsilon/6}$ makes the lower bound dominate $mh + 4r = O(x^{81/2})$. In particular $B_v \ge mh + 4r$. The first block is positive for the form and the signs on the second block alternate. In particular λ is regular, $\lambda_i \in (r - 1)/2 + \mathbf{Z}$, and opposite-sign entries differ by at least h.

Take the upper triangular Borel for this order. Write ρ_G for the half-sum of its roots and ρ_K for the compact half-sum. Set

$$\chi = \lambda - \rho_G, \qquad \mu = \lambda + \rho_G, \qquad \Lambda = \chi + 2(\rho_G - \rho_K).$$

The entries of χ are integers. They decrease weakly since successive entries of λ differ by at least one. Thus χ is a unitary character of the diagonal compact torus, dominant on the positive roots. Atiyah–Schmid's Theorem (8.1) gives a discrete series π_λ when the semisimple real group has finite center and a compact Cartan and λ is nonsingular in the lattice shifted by the half-sum of positive roots. Their Theorem (8.5) gives one copy of the lowest compact type, whose weight is $\lambda + \rho_G - 2\rho_K = \Lambda$. To apply the semisimple theorem, first restrict the parameter and roots to $SU(p, q)$. Its lowest type has the restriction of weight Λ. On $SU(p, q) \times U(1)$ give the circle the character $z^{\sum_i \lambda_i}$. The sum is the integer $\sum_i \Lambda_i$, so for $\zeta^r = 1$ the pair $(\zeta^{-1} I, \zeta)$ acts trivially, as one sees on the lowest type. The product representation descends to $U(p, q)$. On the derived group its infinitesimal character is the restriction of $\chi + \rho_G$. In Sun's convention the inducing line is tensored with the top exterior power of the upper nilradical before applying the derived functor in degree $\dim(\mathfrak{n} \cap \mathfrak{k})$. The parameter $\chi + \rho_G = \lambda$ pairs strictly positively with every root in that nilradical. The

good-range theorem for cohomological induction of a unitary Levi character gives an irreducible unitary module in the above degree and infinitesimal character $\chi + \rho_G$ [8, Theorems 1.2–1.3]. Here the Levi is the compact Cartan. Its bottom weight is $\chi + 2(\rho_G - \rho_K) = \Lambda$, as in Sun's one-dimensional inducing example. Within a sign block, the difference of the $\rho_G - \rho_K$ entries at indices $i < j$ counts the opposite-sign indices between them, so Λ is compact dominant. Vogan–Zuckerman's Proposition 6.1 identifies an irreducible unitary Harish–Chandra module with $A_{\mathfrak{q}}(\chi)$ when the admissible inducing character annihilates the split part of the Cartan, the indicated bottom type occurs, and the Casimir is that of $\chi + \rho_G$ [9, (5.1)–(5.2) and Proposition 6.1]. Here there is no split part and both modules have the required Casimir and compact type, so the modules agree on the derived group. The central circle acts by the inducing character throughout, since all roots kill the center. This identifies π_λ with the unitary realization of $A_{\mathfrak{b}}(\chi)$ on $U(p, q)$. [10, Theorems (8.1) and (8.5)]

Let τ_Λ denote this compact type, w its highest unit vector, $P_w = ww^*$ and $d_\Lambda = \dim \tau_\Lambda$. For an integral weight γ write

$$M_\gamma(Y) = \prod_{t \in \mathbf{Z}} \frac{\det Y_{\{i\,:\,\gamma_i \geq t\},\{i\,:\,\gamma_i \geq t\}}}{(\det Y)^{\mathbf{1}_{t \leq 0}}}.$$

The empty minor is one. Only finitely many factors differ from one on invertible matrices. On $K_{\mathbf{C}} = \mathrm{GL}_p(\mathbf{C}) \times \mathrm{GL}_q(\mathbf{C})$, $M_\Lambda(a) = w^* \tau_\Lambda(a) w$. Indeed each sign block contributes the coefficient of a highest wedge at every successive difference of its dominant weight, times the appropriate determinant power. Their tensor highest vector lies in the irreducible summand of weight Λ, and the product of these wedge coefficients is the stated minor product.

For a matrix Y with positive definite Hermitian part put

$$\mathcal{E}_\lambda(Y) = d_\Lambda \int_K \frac{\tau_\Lambda(u) P_w \tau_\Lambda(u)^*}{M_{-\mu}(u^* Y u)}\, du.$$

Every principal compression is accretive and invertible, so the integrand is holomorphic in Y, locally uniformly for $u \in K$. For $Y > 0$, identify Y with $Y^{1/2}U(r)$. On this whole cone Sun's Definition 3.2 and Theorem 6.3 give $d_\Lambda E_{\mathfrak{b}}(j_{\mathfrak{b}} \circ \psi_\chi)$. Indeed on the complexified torus modulo its compact part the inducing function is $\psi_\chi(tT) = \chi(tt^*)$, by the torus action and its value one at the identity. Its embedding $j_{\mathfrak{b}}$ sends the identity of that line to P_w. The factor $\delta_{\mathfrak{n}}^+$ from (9) multiplies this function by $\delta_{\mathfrak{n}}(tt^*)$. The torus coordinate tt^* of $u^{-1}Y^{1/2}U(r)$ is the positive diagonal D in $u^* Y u = NDN^*$, N upper unitriangular. Its trailing principal determinants are the trailing products of D_i. Thus $(\delta_{\mathfrak{n}}\chi)(D) = D^\mu = 1/M_{-\mu}(u^* Y u)$, including the full determinant factor when $\det Y \neq 1$. The torus bottom space is the line $\mathbf{C}w$. In Sun's formula (26) the factor is d_Λ for probability Haar. For a positive Hermitian $z \in G_1$, take $x = z^{1/2} \in \exp \mathfrak{p}$ there. The result is the compression of $\pi_\lambda(z)$ to τ_Λ. Compact left and right translation now give, for $p_1, q_1 \in K$,

$$m_\lambda(p_1 z q_1^{-1}) = \frac{\displaystyle\int_K M_\Lambda(p_1 u)\overline{M_\Lambda(q_1 u)}\, M_{-\mu}(u^* z u)^{-1}\, du}{\displaystyle\int_K |M_\Lambda(u)|^2\, du}. \tag{5.1}$$

The denominator equals d_Λ^{-1} by compact Schur orthogonality. This gives $m_\lambda(g) = w^* \pi_\lambda(g) w$ on all of G_1, including $m_\lambda(1) = 1$. It is of positive type, since the matrix with entries $m_\lambda(g_i^{-1} g_j)$ is the Gram matrix of the vectors $\pi_\lambda(g_j)w$. [11, Definition 3.2, Theorem 6.3 and (25)–(26)]

The same integral gives the local formula needed for expansion. For $Y > 0$, equation (5) and Theorem 3.3 of Sun give

$$\mathcal{E}_\lambda(aYa^*) = \tau_\Lambda(a)\mathcal{E}_\lambda(Y)\tau_\Lambda(a)^* \qquad (a \in K_{\mathbf{C}}).$$

Indeed a acts by congruence on the positive Hermitian model of $\mathrm{GL}_r(\mathbf{C})/U(r)$. On the compact complexification one has $\tau_\Lambda(a)^* = \tau_\Lambda(a^*)$. Both sides of the resulting identity with independent variables (a, Y, b) are holomorphic where Y, aYb are accretive. Near $(1, 1, 1)$, the conditions $Y = Y^*$, $b = a^*$ describe a

maximal totally real submanifold, and its positive part contains a relative neighbourhood of that point. Uniqueness of holomorphic continuation gives

$$\mathcal{E}_\lambda(aYb) = \tau_\Lambda(a)\mathcal{E}_\lambda(Y)\tau_\Lambda(b)$$

throughout the connected product of operator-norm balls $\|Y-1\| < 1/2$, $\|a-1\|, \|b-1\| < 1/16$ in the indicated spaces. There $\|aYb-1\| \le (17/16)^2(3/2) - 1 < 1$. In particular $w^*\mathcal{E}_\lambda(Y)w$ is unchanged by a lower unitriangular left factor and an upper unitriangular right factor within these bounds and within the two sign blocks. The upper factor fixes w, and the adjoint of the lower factor does so as well.

This local expression agrees with the splitting used to define the coefficient. To see the agreement on the whole real group, write $g = p_1 z q_1^{-1}$ as above and set $J = \mathrm{diag}(\varepsilon_i)$. The same-sign diagonal part of g is $a = p_1 S_0 q_1^{-1}$, where $S_0 = (z + JzJ)/2 = (z + z^{-1})/2$. Hence S_0 is positive block diagonal and commutes with z. Put $Z_0 = S_0^{-1} z > 0$. The unitary matrices S_0^{it} fix Z_0 by conjugation. Equivariance on the positive cone shows that $\mathcal{E}_\lambda(Z_0)$ commutes with their action, and hence with $\tau_\Lambda(S_0^{1/2})$ by differentiation. Equivariance under $S_0^{1/2}$ gives $\mathcal{E}_\lambda(z) = \tau_\Lambda(S_0)\mathcal{E}_\lambda(Z_0)$. It follows that

$$m_\lambda(g) = w^*\tau_\Lambda(a)\mathcal{E}_\lambda(a^{-1}g)w.$$

For $\|g-1\| < 1/16$, diagonal and off-diagonal block extraction are contractions in operator norm, by averaging $g-1$ and $J(g-1)J$ with either sign. Thus $\|a-1\| < 1/16$ and $\|a^{-1}g-1\| < 1/15$. The polarized identity then gives $m_\lambda(g) = w^*\mathcal{E}_\lambda(g)w$.

For the filling Haar dy put $\delta = \deg_v = (\int_{G_1} |m_\lambda(y)|^2\, dy)^{-1}$. With the convention $\langle u, v\rangle = u^*v$, Schur orthogonality in π_λ gives

$$\delta \int_{G_1} m_\lambda(xya)\overline{m_\lambda(y)}\, dy = m_\lambda(x)m_\lambda(a).$$

For the right action $R(g)F(y) = F(yg)$, use the kernel $e_\lambda(g) = \delta\,\overline{m_\lambda(g)}$. The decay estimate below makes e_λ integrable. Schur orthogonality gives

$$e_\lambda^* = e_\lambda, \qquad e_\lambda * e_\lambda = e_\lambda, \qquad e_\lambda * \delta_g * e_\lambda = m_\lambda(g)e_\lambda. \tag{5.2}$$

Indeed the last convolution at x has integrand $\delta^2\overline{m_\lambda(y)}m_\lambda(x^{-1}yg)$. Thus $Q_\lambda = R(e_\lambda)$ is an orthogonal projection and $Q_\lambda R(g)Q_\lambda = m_\lambda(g)Q_\lambda$. In π_λ the projection is P_w, again by Schur orthogonality. For any unitary representation R, a vector z fixed by $R(e_\lambda)$ has $\langle R(g)z, R(h)z\rangle = m_\lambda(g^{-1}h)\|z\|^2$. Thus $\pi_\lambda(g)w \mapsto R(g)z$ extends to an intertwiner, with operator norm $\|z\|$. Conversely evaluation at w sends each bounded intertwiner to the fixed space, since $\pi_\lambda(e_\lambda)w = w$. This correspondence commutes with operators commuting with R. The column coefficient is realized independently on the symmetric power. If j is its distinguished coordinate and $v_0 = e_j^{\otimes a_v}$ has norm one there, then $m_\eta(h) = (h_{jj})^{a_v}$ and $\langle v_0, \rho_H(h^{-1})v_0\rangle = \overline{m_\eta(h)}$. On the common block X of zero weights the row highest vector is fixed by $U(X)$. Indeed the raising operators there kill w, and the squared norms of the corresponding lowerings equal the weight differences, all zero. The diagonal torus acts trivially. The column vector is fixed there as well. These are the row and column coefficients, with their indicated conjugations, in the stationary average.

Global real coefficient estimates. In the sign coordinates write $D(g) = (g + JgJ)/2$ for the sign-block diagonal part, and $O(g) = g - D(g)$. We use the operator norm $\|\cdot\|$ and distinguish it from the Frobenius norm $\|\cdot\|_{\mathrm{F}}$. Fix a nonnegative integer c no greater than any opposite-sign Harish–Chandra gap $|\lambda_i - \lambda_j|$. The bound underlying both norms is

$$|m_\lambda(g)|^2 |\det D(g)|^c \le 1. \tag{5.3}$$

To prove it, let z be positive Hermitian in G_1. Then $z^{-1} = JzJ$ and $\det z = 1$. Denote its principal minor on B by $\Delta_B(z)$, with empty minor one. Jacobi's identity gives $\Delta_{B^c}(z) = \Delta_B(z)$. The opposite-sign

differences of $\mu = \lambda + \rho_G$ have absolute value at least c, since ρ_G follows the same order as λ. In the weight minors from (5.1), the determinant denominators are one for this z. Complementary minors give $M_{-\mu}(z) = M_\mu(z)$, and translating all thresholds gives $M_{\mu+c}(z) = M_\mu(z)$.

Suppose $c > 0$. Put $S_t = \{i : \mu_i \geq t\}$, $T_t = \{i : \mu_i + c \geq t\}$, and $Y_t = T_t \setminus S_t$. The entries in Y_t have weights in the integer interval $[t - c, t - 1]$, so each nonempty Y_t lies in one sign block. The determinant inequality for two principal sets of a positive Hermitian matrix gives

$$\Delta_{Y_t}(z) \leq \Delta_{T_t}(z)\Delta_{S_t^c}(z) = \Delta_{T_t}(z)\Delta_{S_t}(z).$$

Indeed, after taking the Schur complement of the intersection of the two sets, the determinant of the resulting two-block matrix is at most the product of its two block determinants. This follows by taking the Schur complement of one block and using monotonicity of the determinant on positive definite matrices. Here the union is the full set, whose minor is one. For each residue class of t modulo c, the sets Y_t in that class partition the indices, since index i occurs for $\mu_i < t \leq \mu_i + c$. On each sign block the product of the partition minors is at least the block determinant, by the same determinant inequality applied to disjoint sets. Consequently

$$\det D(z)^c \leq \prod_t \Delta_{Y_t}(z) \leq M_{\mu+c}(z)M_\mu(z) = M_{-\mu}(z)^2.$$

Write $g = pz$ by polar decomposition, with p in the maximal compact group $K = U(p_+) \times U(p_-)$. For the polar integral of the unit lowest-type coefficient, the numerator is $M_\Lambda(pu)\overline{M_\Lambda(u)}$, the denominator is $M_{-\mu}(u^* z u)$, and the outside factor is $\left(\int_K |M_\Lambda(u)|^2 \, du\right)^{-1}$ for probability Haar. The determinant of $D(u^* z u)$ equals $\det D(z) = |\det D(g)|$. Cauchy–Schwarz in u, with left invariance of probability Haar on K, proves (5.3). When $c = 0$ the estimate follows from unitarity of the row representation.

The block diagonal part of $g^* J g = J$ gives $D(g)^* D(g) = 1 + O(g)^* O(g)$. Taking determinants, the product of one plus the nonnegative eigenvalues is at least one plus their sum. Also $D(z) = (z + z^{-1})/2$, and the eigenvalues of z occur with their reciprocals. If $\|z\| > 1$, its largest eigenvalue and its reciprocal contribute at least $(\|z\|/2)^2$ to the determinant of $D(z)$, and all the other eigenvalues of $D(z)$ are at least one. If $\|z\| = 1$, then $z = 1$. Since $\|g\| = \|z\|$, the same bound holds for $|\det D(g)|$. Hence

$$|m_\lambda(g)|^4 \leq (1 + \|O(g)\|_{\mathrm{F}}^2)^{-c}, \qquad |m_\lambda(g)| \leq (2/\|g\|)^c. \tag{5.4}$$

These inequalities use $c = A$ in the off-sign comparison for a parameter of gap $500r^3 + A$. They also permit simultaneously all the operator-norm powers up to that gap used below.

Derived vectors and the local coefficient mass. Write π for the row discrete series, u_λ for its unit highest vector, and L for a bound on $|2\lambda_i|$. For a smooth vector let $E_{ab} = d\pi(e_{ab})$, using complexification of the real unitary Lie algebra. The adjoint on the smooth vectors is $E_{ab}^\dagger = \varepsilon_a \varepsilon_b E_{ba}$, as follows from $X \mapsto -JX^*J$, the real structure on the complexified Lie algebra. Let C_G be the sum of $E_{ab}E_{ba}$ over all pairs and C_K the sum over equal-sign pairs. Then

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

Let V_j be the space generated from the lowest K-type by products of at most j derived operators. It is a finite-dimensional K-module, as a quotient of the sum of tensor products of the lowest type with at most j adjoint factors. Each coordinate of its highest weights has absolute value at most $(L + r - 1)/2 + j$. Indeed, the lowest highest weight is $\lambda + \rho_G - 2\rho_K$, whose correction has coordinates of absolute value at most $(r - 1)/2$. The weights of that type lie in the convex hull of its compact Weyl orbit, and each adjoint weight shifts each coordinate by at most one.

On a highest compact weight ν, C_K has eigenvalue $\|\nu + \rho_K\|^2 - \|\rho_K\|^2$. On the row representation C_G has eigenvalue $\|\lambda\|^2 - \|\rho_G\|^2$. Put $N_j = L + r + 1 + j$. We have $|(\nu + \rho_K)_i| \leq N_j$ and $\|\rho_G\|^2 \leq rN_j^2$. Thus the last

sum of squares is bounded by $3rN_j^2\|u\|^2$ on V_j. Cauchy–Schwarz on the matrix entries and $\|X\|_{\mathrm{F}} \le \sqrt{r}\,\|X\|$ give

$$\|d\pi(X)u\| \le (r+1)^4(L+r+1+j)\|X\|\|u\|, \qquad u \in V_j, \tag{5.5}$$

with $d\pi(X)u \in V_{j+1}$. This also holds for the compact representation, with its own matrix size and doubled parameter bound.

If $\|g-1\| < 1/2$, the convergent logarithm satisfies $X = \log g \in \mathfrak{u}(J)$ and $\|X\| \le 2\|g-1\|$. The real Lie algebra condition follows from $g^{-1} = Jg^*J$ by taking logarithms locally. Integrate the derivative of $\pi(g_0\exp(tX))u_\lambda$ and apply (5.5) at $j=0$. For $\|g-1\| \ge 1/2$, the coefficient difference is at most two by unitarity. In both cases, with $N_0 = L+r+1$, we have

$$|m_\lambda(g_0 g) - m_\lambda(g_0)| \le 100N_0^5\|g-1\|. \tag{5.6}$$

In particular $|m_\lambda| \ge 1/2$ on $B_G((200N_0^5)^{-1})$, where $B_G(a) = \{g \in G_1 : \|g-1\| \le a\}$.

Volume and radial integration. Let $j : S_0 \to G_1$ be a proper homomorphism and ds_0 right Haar. Write V_D for the mass of $\{s : \|j(s)-1\| \le D^{-1}\}$. For $D, R_0 \ge 1$, any $g_0 \in G_1$ satisfies

$$\mathrm{vol}_{S_0}\{s : \|j(s)g_0\| \le R_0\} \le (3DR_0^2)^{2r^2} V_D.$$

Take a maximal $(DR_0)^{-1}$-separated subset of the indicated matrices. Disjoint ambient norm-balls of half this radius show that the number of centres is at most $(1+2DR_0^2)^{2r^2}$, since $M_r(\mathbf{C})$ has real dimension $2r^2$. The nearest centre $j(s_i)g_0$ has inverse of norm at most R_0, since $g^{-1} = Jg^*J$. Right division places ss_i^{-1} in the ball defining V_D. Right invariance then gives the inequality. For $a > 4r^2$, apply the identity $u^{-a} = \int_u^\infty aR_0^{-a-1}\,dR_0$ to $u = \|j(s)g_0\| \ge 1$, and integrate by Tonelli. The integral of $(2/u)^a$ is at most $2^a(3D)^{2r^2}\frac{a}{a-4r^2}V_D$. In particular, with $D=1$ and integral $a \ge 8r^2+1$,

$$\int_{S_0} (2/\|j(s)g_0\|)^a\,ds_0 \le 16^a V_1. \tag{5.7}$$

Properness makes the closed unit neighbourhood of the identity in this comparison have finite mass. For the closed subgroup S and any $x_0 \in G_1$, use the proper conjugation $j(s) = x_0^{-1}sx_0$ and $g_0 = x_0^{-1}y$. The multiplier for conjugation by H on du is a continuous character into $\mathbf{R}_{>0}$. Its compact image is trivial, so $ds = du\,dh$ is right invariant. Here V_1 is the mass of $B_x = \{s : \|x_0^{-1}sx_0 - 1\| \le 1\}$. The bound is uniform in y.

For $S_0 = G_1$ and $j = \mathrm{id}$, integrate in the compact-dual Haar dg. The formal degree here satisfies $\delta\int_{G_1}|m_\lambda(g)|^2\,dg = 1$. The local mass from (5.6) therefore implies $\delta\,\mathrm{vol}\,B_G((200N_0^5)^{-1}) \le 4$. With $D = 200N_0^5$ and $a \ge 8r^2+1$, the layer integral above gives the needed formal-degree weighted estimate

$$\delta\int_{G_1}(2/\|g\|)^a\,dg \le 8(2000N_0^6)^{2r^2}\,2^a. \tag{5.8}$$

The quotient measure used below is $d\dot{y}$ with $dg = du\,d\dot{y}$, so the subgroup and group applications involve the same real Haar as the Bessel seed and its Schur estimates. Left translation by H preserves this quotient measure by the same conjugation calculation in the U-fiber.

The Bessel average in the column model. Use the full Schur matrices $S_v(h)$ of the prescribed column coordinate space and its positive Gram $Q = B^*B$. In an orthonormal model the action is $\rho_H(h) = BS_v(h)B^{-1}$. The column basis and its dual are related by $QS_v(h) = S_v(h^{\mathsf{T}})^{\mathsf{T}}Q$, so $C = B^*$ satisfies

$$C\rho_H(h)C^{-1} = S_v(h^{\mathsf{T}})^{\mathsf{T}}.$$

Let $V = \mathbf{C}^{I_v}$ be this irreducible unitary column model of dimension d. Use its unit highest vector v_0 from the coefficient calculation. In the bilinear period the column uses the dual of the displayed transport. Write h_s for the compact component of $s \in S$. For a point $x_0 \in G_1$ to be specified, set

$$Y(y) = \int_S m_\lambda(x_0^{-1} s y)\ \mathrm{wave}(-\sigma(s))\rho_H(h_s^{-1})v_0\ ds. \tag{*}$$

Right reindexing gives $Y(uy) = \mathrm{wave}(\sigma(u))Y(y)$ and $Y(hy) = \rho_H(h)Y(y)$. The subgroup estimate (5.7) and the operator norm bound (5.4) give absolute convergence and boundedness for every y. The same Haar gives

$$\int_{U\backslash G_1} \int_S F(x_0^{-1} s y)\, ds\, d\dot{y} = \int_{G_1} F(g)\, dg \qquad (F \ge 0),$$

since H normalizes U, preserves du, and has mass one. This gives $Y \in L^1(d\dot{y})$ in norm. By the degree-weighted integrability (5.8), e_λ is in L^1. Apply the Schur identity with $a = 1$ inside the average to obtain $Q_\lambda Y = Y$. The interchange uses the uniform subgroup bound. Translation of e_λ is continuous in L^1, which gives continuity of the bounded average.

The local coefficient after removal of its phase. Write $\ell = 2\lambda$ in the sign coordinates. With the positive root order fixed above, $|(2\rho_K)_i| \le r$ and

$$2(\rho_G)_i = \#\{j : \ell_j < \ell_i\} - \#\{j : \ell_j > \ell_i\}.$$

Thus μ has the same order as ℓ. For an integral weight w, put $P_t = \mathrm{diag}(\mathbf{1}_{w_i \ge t})$. The minor product above has the equivalent finite expression

$$M_w(Y) = \det(Y)^c \prod_{c<t\le d} \det\bigl(1 + P_t(Y-1)\bigr), \qquad c \le \min_i w_i, \quad d \ge \max_i w_i.$$

Here c, d, t are integers. The expression is independent of the chosen endpoints. The local coefficient formula proved above uses the probability Haar dk on K, with normalization $D_K = \int_K |M_\Lambda(k)|^2\, dk$.

Let $g \in U(J)$, $\|g - 1\|_{\mathrm{F}} \le \varepsilon_{\mathrm{loc}}$, and suppose that

$$T \ge \max\{1, r^2, r^2 \log(r + L + 1)\}, \qquad |\ell_i| \le L, \qquad \varepsilon_{\mathrm{loc}} T \le 10^{-12}.$$

Write $Z = \log g$ using the series at 1. Then $\|Z\|_{\mathrm{F}} \le 2\varepsilon_{\mathrm{loc}}$ and $Z^* = -JZJ$, by $g^* = Jg^{-1}J$ and the same local logarithm for g^{-1}. In particular $\mathrm{tr}(\Lambda Z)$ is imaginary. Define

$$Q_w(A) = \sum_{i,j} |w_i - w_j|\,|A_{ij}|^2, \qquad q_g = Q_{-\mu}(Z)/4, \qquad \kappa_g = e^{-q_g}.$$

For the row coefficient with opposite-sign differences $|\lambda_i - \lambda_j| \ge 500r^3 + A$ as in the local integral formula, we will show

$$\bigl|m_\lambda(g) e^{-\mathrm{tr}(\Lambda Z)} - \kappa_g\bigr| \le \kappa_g/10 + e^{-100T}. \tag{5.9}$$

In each sign block, order coordinates by decreasing ℓ_i. Let Z_- and Z_+ be the strictly lower and strictly upper parts of these blocks, respectively. The triangular covariance in the local coefficient formula replaces g by

$$\widehat{g} = e^{-Z_-} e^{Z} e^{-Z_+}, \qquad m_\lambda(g) = D_K^{-1} \int_K \frac{|M_\Lambda(k)|^2}{M_{-\mu}(k^* \widehat{g} k)}\, dk.$$

The two unipotent factors fix the highest line and its dual on the appropriate sides with scalar 1. Each exponent has Frobenius norm at most $2\varepsilon_{\mathrm{loc}}$, so the matrices in the replacement are within $e^{6\varepsilon_{\mathrm{loc}}} - 1 < 1/16$ of 1 in Frobenius norm. This is inside the covariance neighborhood, and all denominator compressions are accretive.

We give the remainder estimate for a compressed determinant. For an orthogonal projection P and matrices A_1, A_2, A_3 with $\|A_i\|_{\mathrm{F}} \le d_0 \le 1/100$, set

$$T_P(A,B) = \operatorname{tr}\big(PA(1-P)B\big),$$
$$S_P = \tfrac{1}{2}\sum_i T_P(A_i,A_i) + \sum_{i<j} T_P(A_i,A_j), \quad \zeta = \max_i \|[P,A_i]\|_{\mathrm{F}}.$$

The series for the logarithm gives

$$\log\det\big(1 + P(e^{A_1}e^{A_2}e^{A_3} - 1)\big) = \operatorname{tr}\big(P\textstyle\sum_i A_i\big) + S_P + e_P, \qquad |e_P| \le 4000 d_0 \zeta^2. \tag{5.10}$$

The logarithm here is the trace of the matrix logarithm near 1. Indeed, expand the function

$$t \longmapsto \operatorname{tr}\log\big(1 + P(e^{tA_1}e^{tA_2}e^{tA_3} - 1)\big).$$

The terms linear in t give the trace displayed above. In degree two, the terms from the product of exponentials have coefficients $1/2$ on A_i^2 and 1 on A_iA_j for $i < j$. The quadratic term of the logarithm subtracts the same traces with P inserted between the two matrices, giving S_P. For the higher degrees split every A_i into its diagonal and off-diagonal blocks for P. The words using only diagonal blocks sum to the logarithm of the determinant of the product on $\operatorname{im} P$. This is linear in t. A trace with one off-diagonal factor vanishes. One with at least two such factors in degree j has modulus at most $\zeta^2 d_0^{j-2}$ by applying the Frobenius trace inequality at two of them and the operator norm at the others. Both block parts have operator norm at most d_0. Before the block splitting, the sum of absolute word coefficients in degree j is at most the coefficient of x^j in $-\log(2 - e^{3x})$. This series has nonnegative coefficients and value less than 2 at $x = 1/6$. Including at most 2^j terms per word, the terms of degrees $j \ge 3$ are bounded by $2\zeta^2 \sum_{j\ge3} 12^j d_0^{j-2} \le 4000 d_0 \zeta^2$.

Take $P = P_t$ for $w = -\mu$, $A_1 = -Z_-$, $A_2 = Z$, $A_3 = -Z_+$, and now put $P' = kPk^*$. Write $d_0 = 2\varepsilon_{\mathrm{loc}}$, $u = \|[P,Z]\|_{\mathrm{F}}$ and $v = \|[P,k]\|_{\mathrm{F}} = \|P' - P\|_{\mathrm{F}}$. The minor of $k^*\widehat{g}k$ for P has the compressed determinant in (5.10) for P'. Entry masking and $\|A_i\|_{\mathrm{op}} \le d_0$ give $\|[P,A_i]\|_{\mathrm{F}} \le u$ and $\|[P',A_i]\|_{\mathrm{F}} \le u + 2d_0v$. The diagonal of $P' - P$ has absolute sum v^2, by summing the squares of the entries of k across the cut for P. Since k preserves each sign block and $\sum_i A_i$ is diagonal within them, the linear trace for P' differs from $\operatorname{tr}(PZ)$ by at most d_0v^2. Replacing P by P' moves each block $PA_i(1-P)$ in Frobenius norm by at most $2d_0v$, by subtracting the two products. The same holds for the blocks in the reverse direction. The trace inequality then gives

$$|S_{P'} - S_P| \le 18 d_0 v(u + 2d_0v).$$

To evaluate S_P, let i be inside the cut and j outside it. The order of μ is the order of ℓ, so $\ell_i < \ell_j$. At equal signs the contribution through i, j is $-Z_{ij}Z_{ji}/2$, as the $(1,2), (1,3), (2,3)$ terms contribute $-1, +1, -1$ times $Z_{ij}Z_{ji}$, and the $(2,2)$ term contributes $1/2$ times it. At opposite signs only $(2,2)$ contributes. The relation $Z_{ji} = -\varepsilon_i\varepsilon_j\overline{Z_{ij}}$ makes both contributions $|Z_{ij}|^2/2$. Thus $S_P = u^2/4$. The three error bounds, with $2d_0 \le 1$, give an error at most $10^4 d_0(u^2 + v^2)$ away from $\operatorname{tr}(PZ) + u^2/4$.

For each pair i, j, exactly $|w_i - w_j|$ cuts separate the two indices, and $\sum_{c<t\le d} P_t = \operatorname{diag}(w_i - c)$. Also $\det(\widehat{g}) = e^{\operatorname{tr} Z}$. Multiplying the minor expansions therefore gives

$$M_{-\mu}(k^*\widehat{g}k) = \exp\big(q_g - \operatorname{tr}(\Lambda Z) - R_k\big), \qquad |R_k| \le 10^6 \varepsilon_{\mathrm{loc}}\big(q_g + Q_\Lambda(k) + r^2\big).$$

Here $Q_{-\mu}(k) \le Q_\Lambda(k) + 2r^2$ by $|\mu_i - \Lambda_i| \le r$ and $\|k\|_{\mathrm{F}}^2 = r$. The linear trace correction is at most $2\varepsilon_{\mathrm{loc}} r^2$.

For a cut P_t and unitary k, the square modulus of the compressed determinant on $\operatorname{im} P_t$ is $\det(1 - B^*B)$ there, where $B = (1 - P_t)kP_t$. Its eigenvalue factors are in $[0,1]$ and their product is at most $e^{-\|B\|_{\mathrm{F}}^2}$. The two off-diagonal blocks of k have equal squared Frobenius norm, as follows from the row and column norm identities on $\operatorname{im} P_t$. Consequently $|M_\Lambda(k)|^2 \le e^{-Q_\Lambda(k)/2}$. For a lower bound put $U_* = r + L + 1$ and

$h_* = (20U_*^3)^{-1}$. Since $|\Lambda_i| \le 2U_*$, and every compression of a matrix $\|k-1\|_{\mathrm{op}} \le h_*$ has least singular value at least $1-h_*$,

$$|M_\Lambda(k)| \ge (1-h_*)^{4rU_*} \ge 1 - 4rU_*h_* \ge 1/2.$$

An h_*-separated set in K, for the operator norm, has at most $(1+4/h_*)^{2r^2}$ elements by packing disjoint balls of radius $h_*/2$ in the ambient real matrix space. A maximal such set covers K by right translates of the h_*-ball about 1. Thus $D_K \ge \frac{1}{4}(1+4/h_*)^{-2r^2} \ge e^{-20T}$.

After removal of its phase the coefficient differs from e^{-q_g} by the average of $e^{-q_g}(e^{R_k}-1)$ for the probability measure $|M_\Lambda(k)|^2\, dk/D_K$. If $q_g + Q_\Lambda(k) + r^2 \le 50000T$, then $|R_k| \le 1/20$ and $|e^{R_k}-1| \le 1/10$. Otherwise, using $|R_k| \le (q_g + Q_\Lambda(k) + r^2)/4$, we have

$$|M_\Lambda(k)|^2 e^{-q_g}|e^{R_k}-1| \le e^{-q_g - Q_\Lambda(k)/2 + |R_k|} \le e^{-5100T}.$$

Integrating the two bounds and using the lower bound for D_K proves (5.9).

For the actual compact coefficient the column model gives $m_\eta(h) = h_{pp}^{a_v}$, with p the coordinate of the nonzero weight. If $Z_H = \log h$, take P to be the projection onto that coordinate. Apply (5.10) with only $A_2 = Z_H$ nonzero. Its quadratic term is $-\|[P, Z_H]\|_{\mathrm{F}}^2/4$ because $Z_H^* = -Z_H$. After multiplication by a_v the error is at most $32000\varepsilon_{\mathrm{loc}} q_H$, where $q_H = a_v\|[P, Z_H]\|_{\mathrm{F}}^2/4$. The same split at $q_H = 50000T$ gives

$$\left| m_\eta(h) e^{-\operatorname{tr}(\Lambda' Z_H)} - e^{-q_H} \right| \le e^{-q_H}/10 + e^{-100T} \qquad (\|h-1\|_{\mathrm{F}} \le \varepsilon_{\mathrm{loc}}).$$

This also covers $a_v = 0$, when the coefficient is 1.

A stationary graph on the prescribed frame. For the first $k-1$ entries of the row weight $\Lambda = \lambda + \rho_G - 2\rho_K$, the correction at index i is $(r-1)/2 - i - (p-1-2i) = 1 - k/2 + i$. It cancels the corresponding entry of λ, so Λ vanishes on their positive block X. On its complement, with $2n = m$, the entries are

$$u_i = -B_v - (i+1)h + (\rho_G - 2\rho_K)_{k-1+i}, \qquad 0 \le i < 2n.$$

They are integral and strictly decreasing, since the root correction has absolute value at most $2r$ and $h > 4r$. The bound $B_v \ge mh + 4r$ places them in $[-3B_v, -B_v]$, and their form signs are $(-1)^i$. The column weight is zero on X and equals a_v on the remaining coordinate p of the definite core. Set $t_i = u_i - a_v$. If c_*^2, c_i^2 are the absolute derivatives of the nodal polynomial at 0, t_i, respectively, Lagrange interpolation for $0 > t_0 > \cdots > t_{2n-1}$ gives

$$\sum_i \frac{(-1)^i t_i^j}{c_i^2} = \begin{cases} c_*^{-2} & j = 0, \\ 0 & 1 \le j < 2n, \\ -1 & j = 2n. \end{cases}$$

Indeed the leading coefficient is the sum of nodal values divided by the signed derivatives, which have signs $+1, (-1)^{i+1}$. Use $u_* = (c_*/c_i)_i$, $V_j = (t_i^{j+1}/c_i)_i$. The first vector has form norm 1 and is orthogonal to V_j for $j \le \varrho = n-1$. The V_j satisfy $h(V_i, V_j) = 0$ for $i + j < 2\varrho$ and -1 for equality. The Vandermonde matrix shows independence of $u_*, V_0, \dots, V_\varrho$. The orthogonal sum with X has the Gram of the core and the frame $f_0, \dots, f_\varrho = (e_0, \dots, e_{\varrho-1}, z)$ in the split coordinates. By Witt extension there is an isometry $q = x_0 p_v$ from the sign coordinates to these coordinates. Here p_v is the fixed unitary sign frame, $qV_j = f_j$, and q maps X and u_* to the two designated parts of the definite core.

Let $D = q\Lambda q^{-1}$. For $i < \varrho$, $Df_i = a_v f_i + f_{i+1}$, by multiplication by $u_j = a_v + t_j$ in the sign coordinates. On the definite core the compression of D is Λ', since $\Lambda u_* = a_v u_* + c_* V_0$. The group S is upper block triangular along the split flag, with identity diagonal blocks except for its compact action on the core. Only the middle diagonal block and the strictly lower entries of D contribute to $\operatorname{tr}(D(s-1))$. The shifts on f_i, paired with the adjacent entries of s, contribute $\sigma(s) = \sum_{i<\varrho} s_{f_i, f_{i+1}}$. The remaining lower entries are their reflected entries on the dual flag, by $D^* J_1 = J_1 D$, where J_1 is the split Gram matrix. To see their

contribution, write $J_1 f_i = \varepsilon_i' f_i^{\#}$, with $f_i^{\#} = e_i^*$, $\varepsilon_i' = 1$ for $i < \varrho$, and $f_\varrho^{\#} = z$, $\varepsilon_\varrho' = -1$. The reflected entry of D has the factor $\varepsilon_i'\varepsilon_{i+1}'$, and the paired entry of s has the same factor times $\overline{(s^{-1})_{f_i,f_{i+1}}}$, by $s^{-1} = J_1 s^* J_1$. On adjacent frame levels the latter conjugate equals $-\overline{s_{f_i,f_{i+1}}}$. The middle block $h_s \oplus 1$ acts trivially on z, so this includes the last adjacent pair. Since the other strict lower entries are zero, the reflected sum is $-\overline{\sigma(s)}$. For $g_s = p_v^{-1} x_0^{-1} s x_0 p_v$ we obtain

$$\begin{aligned} \operatorname{tr}(g_s - 1) &= \operatorname{tr}(h_s - 1), \\ \operatorname{tr}\big(\Lambda(g_s - 1)\big) - \operatorname{tr}\big(\Lambda'(h_s - 1)\big) &= \sigma(s) - \overline{\sigma(s)} =: \phi(s). \end{aligned} \tag{5.11}$$

The character ϕ is additive because each frame matrix is upper unitriangular. The common $U(X)$ inside S has the same compact action in both sign coordinates, by the choice of q. It fixes both highest vectors. For the row vector this also follows from the compact coefficient M_Λ, as all the weight entries on X are zero and each weight cut either contains X or is disjoint from it.

The same stationary point has the required quotient size. Since the distinct nodes are integers, $c_i \geq 1$, and the columns $p_v V_i$ through $i = 2\varrho$ have Euclidean lengths at most

$$M_0 = \sqrt{2n}(3B_v + a_v)^{2n-1}.$$

Gram–Schmidt on the first n columns is upper unitriangular without rescaling. Each resulting column $\widetilde{V}_j$ has length at most M_0, and pairing it with the old column $p_v V_{2\varrho-j}$ gives -1, so its length is at least M_0^{-1}. The frame change lifts to U. In fact, writing $E = (e_i)$, $E_{\mathrm{f}}^* = (e_i^*)$ as matrices of frame and dual columns, a frame action $E \mapsto EA,\ z \mapsto z + Ev$ with A upper unitriangular extends by

$$E_{\mathrm{f}}^* \longmapsto \big(E_{\mathrm{f}}^* + zv^* + Evv^*/2\big)(A^*)^{-1}$$

and by the identity on the core. The paired blocks preserve the split form and have the upper triangular levels of U. After this change the isotropic duals $J_1\widetilde{V}_j/\|\widetilde{V}_j\|^2$ for $j < \varrho$ have length at most M_0, pair to the identity with the $\widetilde{V}_i$, and are orthogonal to $\widetilde{V}_\varrho$. They can be installed by one further left multiplication in U fixing the frame. Indeed, in the new split coordinates the target dual columns have the form $E_{\mathrm{f}}^* + X_0 C + EB$, where X_0 is the matrix of core columns and $B + B^* + C^*C = 0$. This sends X_0 to $X_0 - EC^*$ and fixes E, z, giving an isometry in the subgroup of U fixing the frame.

For the resulting inverse matrix $G = (s_0 x_0)^{-1}$, $s_0 \in U$, $GG^* = J_1 + 2GP_-G^*$, with P_- the negative projection of J_1. Its negative basis is $(e_i - e_i^*)/\sqrt{2}, z$. The bounds on the frame and dual columns give $\|GP_-\|_{\mathrm{F}}^2 \leq (2\varrho+1)M_0^2$ and hence $\|G\|_{\mathrm{op}}^2 \leq 4nM_0^2$. Inversion preserves the operator norm on $U(J_1)$. It follows that

$$\operatorname{size}(x_0) \leq 8nM_0 \leq b^{5n}.$$

For the last inequality use $3B_v + a_v \leq 2b^2$, $b \geq 8n$, $8n\sqrt{2n} \leq b^2$, and $2^{2n-1} \leq b^n$. These left translations bound the quotient size of x_0. The average uses the stationary representative x_0 from (5.11).

Nonvanishing of the stationary average. Put $\beta = b^3$, $q_0 = (r+k)^2$, $T = q_0 \log\beta$, $\epsilon = 10^{-13}/T$, and write

$$\theta(s) = m_\lambda(g_s)\overline{m_\eta(h_s)}e^{-\phi(s)}, \qquad \omega_j(g) = (2/\|g\|_{\mathrm{op}})^j.$$

The actual parameters satisfy $\beta \geq 10^{12}$, $r + k + L + 1 \leq \beta$ with $L = 3a$, and the row has gap at least $500r^3 + A$ for $A = r^7 \geq 10^{30}T^3$. Indeed $B_v \leq 2b$, $a_v \leq a \leq exb$, $k \leq r \leq 3x^5$ and $b \geq x^{10^{12}}$ give $3B_v + a_v \leq b(6 + ex) \leq 2b^2$ and $r + k + 3a + 1 \leq b^3$ for sufficiently large x. The row gap is at least $2r^8$. Since $\log b \leq x^{1/100}$ and $r \geq x^5$, we have $T^3/r^7 \leq 1728x^{3/100-5}$.

Suppose first the off-sign blocks of g_s have combined Frobenius norm at most ϵ. With $A_1 = g_s - 1$, the row identity for $g_s J g_s^* = J$ gives $-2\operatorname{Re}(A_1)_{ii} = \sum_j \varepsilon_i\varepsilon_j|(A_1)_{ij}|^2$. The real weighted trace for Λ agrees by (5.11) with that for Λ' on $h_s - 1$, which is nonpositive. With weights $-\Lambda_i$ the sum of the unsigned squares

is therefore at most $6B_v\epsilon^2$. In particular the rows of A_1 outside X have total squared norm at most $6\epsilon^2$. Write $g_s = \left(\begin{smallmatrix} a & b' \\ c & d' \end{smallmatrix}\right)$ for $X \oplus X^\perp$, with the form $1 \oplus J'$. The equations $a^*a = 1 - c^*J'c$ and $a^*b' = -c^*J'd'$ give

$$\|b'\|_{\mathrm{F}} \le \frac{\sqrt{6}\epsilon(1+3\epsilon)}{\sqrt{1-6\epsilon^2}} \le 3\epsilon, \qquad \|(a^*a)^{1/2} - 1\|_{\mathrm{F}} \le 6\epsilon^2.$$

Remove the unitary factor of a by a common left element of $U(X)$. The new first factor is within 10ϵ of 1 in Frobenius norm by the three block bounds. The new second factor is within the same distance since its squared norm of the difference from 1 is $-2\,\mathrm{Re}\,\mathrm{tr}(h-1)$, equal to the sign-weighted sum of squares for the new first factor. The move preserves θ and $\|g_s\|_{\mathrm{op}}$.

At the moved point, let $X_s = (g_s - 1, h_s - 1)$, and denote by $\mathcal{L}$ the difference of the two weighted traces in (5.11). The homomorphism of the graph into the matrix pair and the additivity of ϕ imply $\mathcal{L}(X_s^j) = 0$ for $j \ge 2$. In detail, X_s^j is the j-th finite difference of $(g_s, h_s)^i$ at $i = 0$, and $\mathcal{L}((g_s, h_s)^i - 1) = i\phi(s)$. Thus the weighted difference of the two logarithms is $\phi(s)$. Their diagonal traces are imaginary, so the two phases of the coefficients cancel with $e^{-\phi(s)}$. Using (5.9) and its compact version with $\varepsilon_{\mathrm{loc}} = 10\epsilon$, write the product as $(\kappa_1 + e_1)(\kappa_2 + e_2)$ with $0 \le \kappa_i \le 1$, $|e_i| \le \kappa_i/10 + E_0$, $E_0 = e^{-100T}$. Its difference from $\kappa_1\kappa_2$ is at most $0.21\kappa_1\kappa_2 + 3.2E_0$. Thus $|\theta(s)| - 2\,\mathrm{Re}\,\theta(s) \le 10E_0$. The moved first factor has operator norm at most 2, so $\omega_{8q_0}(g_s) \ge 1$ also at the original point.

For the other points the off-sign bound gives $|\theta(s)|^4 \le (1+\epsilon^2)^{-A}$, so $|\theta(s)| \le e^{-1250T}$ using $\log(1+\epsilon^2) \ge \epsilon^2/2$. Since $16q_0 \le A$, the global coefficient bound also gives $|\theta(s)| \le \omega_{16q_0}(g_s)$. Their geometric mean bounds the modulus by $e^{-625T}\omega_{8q_0}(g_s)$. Combining this bound with the near estimate gives the following inequality on the whole stationary graph.

$$|\theta(s)| \le 2\,\mathrm{Re}\,\theta(s) + 32\beta^{-50q_0}\omega_{8q_0}(g_s). \tag{5.12}$$

Finally use B_x as above and $B_{\mathrm{sm}} = \{s : \max(\|g_s - 1\|_{\mathrm{op}}, \|h_s - 1\|_{\mathrm{op}}) \le h_0\}$, where $h_0 = (200\beta^5)^{-1}$. On B_{sm} the Frobenius differences are at most $\sqrt{r}\,h_0$. The inequality $\log\beta \le \beta$ gives $T\sqrt{r}\,h_0 \le (200\beta^{3/2})^{-1} \le 10^{-12}$. Both quadratic forms in the Gaussian estimates are at most $\beta^2 h_0^2$, since $\max_{i,j}|\mu_i - \mu_j| \le L + r$, $a_v \le L$, and the logarithms have Frobenius norms at most $2\sqrt{r}\,h_0$. Hence the modulus of each factor is at least $1/2$ there. The paired matrices of B_x have norms at most 2. An $h_0/2$-separated set for the maximum operator norm in the two matrix spaces has size at most $(1 + 1600\beta^5)^{2(r^2+k^2)} \le \beta^{11q_0}$ by disjoint ambient balls of radius $h_0/4$. Division by a covering centre on the right sends the covered points to B_{sm}, since inversion preserves each operator norm. The measure ds is right Haar, so $\mathrm{vol}(B_x) \le \beta^{11q_0}\,\mathrm{vol}(B_{\mathrm{sm}})$. This ball B_x has positive finite mass. Its interior contains the identity and it is compact, since the graph is closed and the second projection takes values in the compact group H.

Integrate (5.12) and use the subgroup tail estimate (5.7) with $8q_0 > 8r^2$. The bound on B_{sm} gives

$$\mathrm{vol}(B_x) \le 8\beta^{11q_0}\,\mathrm{Re}\int_S \theta(s)\,ds + 128\beta^{-39q_0}16^{8q_0}\,\mathrm{vol}(B_x).$$

The second term is at most half the left side. Thus

$$0 < \mathrm{vol}(B_x) \le \beta^{20q_0}\,\mathrm{Re}\int_S \theta(s)\,ds. \tag{5.13}$$

The scalar integral is $\langle v_0, Y(x_0)\rangle$ in the unitary column model, giving the required nonzero average in the same Haar. In particular the subgroup tail bound with $M = 8r^2+1$ gives $\int_S \omega_M(x_0^{-1}sy)\,ds \le b^{M+60q_0}\|Y(x_0)\|$ for every y, since $b \ge 16$.

Truncation in the same quotient measure. Keep the nonzero average Y at the stationary point x_0. Set $M = 8r^2 + 1$, $C_0 = 16^M$ and $R = b^{6n+2}$. The graph bound $\mathrm{vol}(B_x) \le b^{60(r+k)^2}\|Y(x_0)\|$, together with (5.7) and $b \ge 16$, gives

$$\int_S \omega_M(x_0^{-1}sy)\,ds \le P_0\alpha, \qquad P_0 = b^{M+60(r+k)^2}, \qquad \alpha = \|Y(x_0)\| > 0.$$

Let $B_1 = B_G(1)$ and put

$$\Phi(y) = \int_S \omega_{2M}(x_0^{-1} s y)\, ds, \qquad I_M = \int_{G_1} \omega_{2M}(g)\, dg,$$
$$d\mu_0(g) = \frac{C_0 P_0 \alpha}{\mathrm{vol}(B_1)} \omega_M(g)\, dg.$$

The positive function Φ is left S-invariant by Haar reindexing. Unfolding with the free-entry du and probability dh gives $\int_{U\backslash G_1} \Phi\, d\dot{y} = I_M$. Also $\|Y(y)\| \le \Phi(y)$, by the coefficient bound at exponent $2M$. This gives $Y \in L^1(d\dot{y})$, and its boundedness gives $Y \in L^2(d\dot{y})$.

For any $s \in S$ and $y \in G_1$, let $t = x_0^{-1} s y$. If $tg \in B_1$, then $\|tg\| \le 2$ and $\|g\| \le 2\|t\|$. The original orbit size satisfies $\mathrm{size}(yg) \le 2\,\mathrm{size}(x_0) < R$, using $\mathrm{size}(x_0) \le b^{5n}$. To see this, use left S-invariance and $\mathrm{size}(yh) \le \|h\|\,\mathrm{size}(y)$ for $h \in G_1$, which follows by taking the infimum and using $\|h\| \ge 1$. Moreover $\omega_M(g) \ge C_0^{-1} \omega_M(t)$. Left invariance of dg makes the mass of $t^{-1}B_1$ equal to that of B_1, so

$$\mu_0\{g : \mathrm{size}(yg) \le R\} \ge P_0 \alpha \omega_M(t).$$

Multiply by $\omega_M(t)$, integrate over s, and use the uniform estimate for its integral to conclude

$$\Phi(y) \le \mu_0\{g : \mathrm{size}(yg) \le R\}, \qquad \mu_0(G_1) \le C_0^2 P_0 \alpha.$$

For the second estimate use (5.7) with $S_0 = G_1$, $j = \mathrm{id}$, $g_0 = 1$, and $a = M$.

Apply the reproducing identity at x_0 and bound $|m_\lambda(g^{-1})| \le \omega_{2M}(g)$, since inversion preserves the operator norm. Cauchy–Schwarz with this weight gives

$$\alpha^2 \le \delta^2 I_M \int_{G_1} \|Y(x_0 g)\|^2 \omega_{2M}(g)\, dg = \delta^2 I_M \int_{U\backslash G_1} \|Y(y)\|^2 \Phi(y)\, d\dot{y}.$$

The equality unfolds the same group Haar, using the S-invariance of the squared norm. Since $\Phi \le C_0^2 P_0 \alpha$, this gives $\alpha \le \delta^2 I_M C_0^2 P_0 \|Y\|_2^2$. The group estimate (5.8) gives

$$\delta I_M C_0^2 \le 8(2000 N_0^6)^{2r^2} 2^{10M} \le b^{99r^2}.$$

Indeed $2000 \le b$ and $N_0 \le b^4$ give the factor b^{50r^2} for the power of $2000N_0^6$. The remaining factor is $2^{80r^2+13} \le b^{10r^2}$, since $b \ge 2000 > 2^{10}$. Thus the stated power has room to spare.

Set $c_0 = \sqrt{d}/\|Y\|_2$, which is finite by continuity and $Y(x_0) \ne 0$. Define the normalized full vector, its continuous scalar cut and their positive majorant by

$$W = c_0 Y, \qquad W^\flat(y) = \min(1, \max(0, R - \mathrm{size}(y)))W(y),$$
$$A(y) = c_0 \Phi(y), \qquad \mu_{\mathrm{cut}} = c_0 \mu_0.$$

The size is continuous by the right multiplication inequality and its application with the inverse. The cutoff preserves the left Bessel laws and vanishes for size at least R. Its support modulo U is compact. In fact, for each point with size at most R choose a translating element of $S = HU$ within one of the infimum, and discard the Euclidean-unitary left factor in H. This gives a representative in a fixed compact group norm-ball. We have $\|W^\flat\|_2^2 \le \|W\|_2^2 = d$. The matrix $\int W(y) W(y)^*\, d\dot{y}$ commutes with the irreducible compact column action, since H normalizes U and preserves du and $d\dot{y}$. Its trace is d, so it is the identity on the orthonormal column space. The full vector retains the row reproducer.

The majorant estimates are

$$A(y) \le \mu_{\mathrm{cut}}\{g : \mathrm{size}(yg) \le R\}, \qquad \mu_{\mathrm{cut}}(G_1) \int A\, d\dot{y} \le b^{2r^3} d. \tag{5.14}$$

For the product estimate use

$$c_0^2 C_0^2 P_0 \alpha I_M \le d(\delta I_M C_0^2 P_0)^2 \le d b^{2r^3},$$

where $M + 60(r+k)^2 + 99r^2 \le 347r^2 + 1 \le r^3$ at the actual ranks. If $W(y) \ne W^{\flat}(y)$, then $\operatorname{size}(y) + 1 > R$. For every s this implies $\|x_0^{-1} s y\| > 2b$, since otherwise $\operatorname{size}(y) + 1 \le 2b^{5n+1} + 1 \le R$. Apply the power bound at exponent $2M + 3r^3 f$. It applies because $2M \le 500r^3$ and $3r^3 f \le r^7$, with the row gap at least $500r^3 + r^7$. For both $W^{\natural} = W, W^{\flat}$ this yields

$$\|W^{\natural}(y)\| \le \begin{cases} A(y), & W(y) = W^{\flat}(y), \\ b^{-3r^3 f} A(y), & W(y) \ne W^{\flat}(y). \end{cases} \tag{5.15}$$

On the set where the vectors differ, the difference also has norm at most $\|W(y)\|$, since the cut multiplies by a scalar in $[0,1]$. Now $A \le \mu_{\mathrm{cut}}(G_1)$, so

$$\|W - W^{\flat}\|_2^2 \le b^{-6r^3 f} \int A^2 \, d\dot{y} \le b^{2r^3 - 6r^3 f} d \le b^{-10r^2} d.$$

All the magnitude bounds hold on the chosen large squares. There $x^5 \le r \le 3x^5$, $k \le r$, $n \le r$, $L = 3a \le 3exb$, and $f \le e^{C_0'} x^{11} (\log b)^{1/5}$. The scale $\log b = x^{\varepsilon}$ satisfies $100 \log\log b \le \log x$. In particular $f \le e^{C_0'} x^{11 + 1/500}$. Together these give $b \ge 2000$, $r + L + 1 \le b^4$, $3r^3 f \le r^7$ and the required row gap from the actual gap at least $2r^8$. They also give

$$(r+1)^4 (L + r + 1 + 2n^2) \le b^5,$$

which bounds every factor of (5.5) on the required normalized derivative words. The derivative Schur argument uses the full W in this same quotient. The absolute periodization uses (5.14) on the full seed, and the relative trace uses both estimates on the full and cut link tensors, before their separate arithmetic class cutoff.

Schur bounds for the right derivatives. On the right-regular Hilbert space of U-equivariant vectors $Q_\lambda W = W$, and $Q_\lambda R(g) Q_\lambda = m_\lambda(g) Q_\lambda$. Thus the cyclic representation of $W/\sqrt{d}$ is isometric to the row discrete series. Match the Gram matrices of the two sets of translates of the unit vectors and complete. The orthogonal projection onto the cyclic space commutes with right translation. For V in that space and square-integrable F in the ambient right-regular space, Schur orthogonality gives

$$\int_{G_1} |\langle F, R(g) V \rangle|^2 \, dg \le \|F\|_2^2 \|V\|_2^2 / \delta.$$

For matrix-unit words w with $|w| \le 2n^2$, (5.5) gives $\|W_w\|_2 \le P_1^{|w|} \sqrt{d}$ for $P_1 \ge b^5$. For each scalar test $\phi \in L^1(G_1) \cap L^2(G_1)$, duality and Cauchy–Schwarz in the Schur relation give

$$\int_{U \backslash G_1} \left\| \int_{G_1} \phi(g) W_w(yg) \, dg \right\|^2 d\dot{y} \le P_1^{2|w|} \frac{d}{\delta} \int_{G_1} |\phi(g)|^2 \, dg.$$

Here is an exact smoothing on the same cyclic space. Let K be the maximal compact of the row group G_1. Average a nonnegative smooth compactly supported approximate identity of integral one over conjugation by K, using probability Haar. The resulting functions η_j are again an approximate identity, since K is compact. On the cyclic representation generated by W, the operators $R(\eta_j)$ commute with K. The lowest K-type occurs there with multiplicity one and contains W under the Gram isometry. Consequently $R(\eta_j) W = c_j W$, with $c_j \to 1$ by strong continuity. Choose one $c_j \ne 0$ and put $\eta = \eta_j / c_j$. The equality

$$W(y) = \int_{G_1} \eta(a) W(ya) \, da$$

holds pointwise. Indeed both sides are continuous and the quotient Haar has full support.

For X in the real Lie algebra, define $\ell_X \eta(a) = \frac{d}{dt}\big|_0 \eta(\exp(-tX)a)$. Changing variable by left translation in the last integral gives the derivative of $W(y \exp(tX))$ with test $\ell_X \eta$. Iterating in the same order gives

the word test $\ell_{X_1} \cdots \ell_{X_j} \eta$. All these tests are smooth and compactly supported. Since W is bounded and continuous, the integrals give bounded continuous representatives of the Hilbert-space derivatives. This also follows in L^2 by differentiating the left translates of η in $L^1(G_1)$. Extend the derivatives complex-linearly to the matrix units E_{ab} in the normalized filling frame. The map $y \mapsto p_v^{-1} y p_v$ sends those units to $p_v^{-1} E_{ab} p_v$ in the row frame, each of operator norm one. Writing $W_{(a,b)w}$ for prepending such a derivative to w, one obtains along a smooth curve with $\gamma'(t) = \gamma(t)X$

$$\frac{d}{dt} W_w(\gamma(t)) = \sum_{a,b} X_{ab}\, W_{(a,b)w}(\gamma(t)).$$

Right differentiation preserves the two left Bessel transformation laws. Thus the word and convolution estimates apply to the full pointwise vector used on the positive Gram cell. The trace estimate for the radial cut $W^{\flat}$ uses its order-zero norm bound and its support.

5.1 Periodization and the initial Bessel coefficient

If D_v is the diagonal basis-scale matrix (with the inverse conjugate entries on the dual vectors), let $x_v = D_v p_v$ be the isometry of the signed real space onto the rational Gram space. In a different choice of real realization the two isometries differ by left multiplication in G_1. The real argument is

$$y_v = D_v^{-1} x_v g_v p_v^{-1} = (D_v^{-1} x_v p_v^{-1})(p_v g_v p_v^{-1}).$$

The sign-frame equivalence preserves the compact-dual Haar, and the remaining factor is a left translation. This map intertwines rational left multiplication with its action in the normalized real frame. At each place write μ_v for the finite positive measure μ_{cut}. The seed on its finite class coset is $\beta_c(y) W_i(y)$, with β_c the class cutoff and W_i the full Bessel coordinate. The orbital estimates use the radial cut, with its separate pointwise gain (5.15).

Absolute periodization at the weighted level. Write $\Gamma = G_0(F)$, U for the unipotent of B_0, and $U_c = U(F) \cap h_c K(g) h_c^{-1}$, using the filling level in the fixed rational coordinates. Take W_v, the integrable type kernel e_v and the finite positive measure μ_v from the real construction. The kernel fixes the Bessel coordinates by right convolution, and the change of basis intertwines the compact core action with the transpose Schur action. For the absolute folding argument use the finite positive translate majorant (5.14) for the full vector. Write $\mathrm{InCut}_v = \{y : \mathrm{size}(y) \le R\}$, with $R = b^{6(\varrho+1)+2}$. The estimate counts the right translates yh in this set against the finite μ_v. The pointwise gain (5.15) enters the full-to-cut trace comparison. Choose $\beta_c \ge 0$ with sum one on each left U_c-orbit in $G_{0,\infty}$. For existence one can take the indicator of a measurable fundamental set for the discrete left action. The quantitative trace calculation will use a cutoff supported in its frame box. The argument below applies to both choices. For $a = (c, i)$ let s_a be the seed with real factor $\beta_c W_i$ and finite support on $h_c \eta^{-1} J_g = h_c K(g) \eta^{-1}$ in the stage-zero variable x_f. Under $x_f \mapsto \eta^{-1} x_f \eta$ this coset becomes $\eta^{-1} h_c K(g)$, using the transported lattice frame. The original coset is right J_g-invariant and left U_c-invariant. The norm of a Bessel coordinate is invariant under U_∞, since the character has modulus one. Moreover $\sum_{u \in U_c} \beta_c(uy) = 1$, with nonnegative summands. Grouping $\sum_{\gamma \in \Gamma} |s_a(\gamma x)|$ by U_c thus removes the cutoff and gives a sum of absolute Bessel coordinates over the U_c-cosets satisfying the finite condition.

Here is a uniform bound for this sum. The real size cut has a compact set C_∞ of representatives modulo U_∞. Choose a Bessel translate within one of the infimum in the size cut. Its inverse norm equals its norm in the sign frame. The resulting closed matrix ball in the group is compact, and the core factor removed from the translating element ranges over a compact group. The group U_c is a cocompact lattice in U_∞. Indeed, in a central filtration of U its rational subgroup cut out by the fixed compact open level has full lattices in the additive successive quotients. Taking representatives successively gives a compact set C_U with $U_\infty = U_c C_U$. For fixed x and $h = (h_v)_v$, each coset whose real coordinate after right translation by

h is in the cut therefore has a representative γ with $\gamma x(h,1)$ in a fixed adelic compact C. In the fixed rational coordinates its real part lies in $C_U C_\infty$ and its finite part in $h_c K(g)\eta^{-1}$. The real frame changes carry compact sets to compact sets. In the rational coordinates the motions by Γ are left translations. For any two representatives in this compact, their quotient belongs to $\Gamma \cap CC^{-1}$, a finite set. Hence the number of contributing cosets is bounded by its cardinality, independently of x, h. Tonelli and the product measure $\otimes_v \mu_v$ now bound the absolute seed sum uniformly in x.

For the sum with signs, the rational U_c-character at infinity is one. At finite places it is trivial on $h_c K(g) h_c^{-1} \cap U$ by the level's simple-root conditions, and the global additive character is trivial on rational points. The same grouping leaves the finite-coset indicator times W_i in each coset. Consequently the sum is independent of the choice of the partition β_c. In particular the boxed cutoff used in the Gram integral gives exactly the same operator $\mathcal{P}$. The product kernel $e = \prod_v e_v$ fixes each such summand. The uniform absolute bound allows integration against $|e|$ before summing, so $R(e)$ fixes the bounded, measurable periodization. It is continuous at infinity because translates of e vary continuously in L^1, and right J_g-invariance gives continuity in the finite variables. It is left Γ-invariant by reindexing. These properties put the series in the indicated space of forms, with its actual Petersson norm. Its finite linear combination for any column defines $\mathcal{P}$. Since $[\Gamma\backslash G_0(\mathbb{A}_F)]$ has finite measure, Tonelli on this quotient gives $\int_{G_0(\mathbb{A}_F)} |s_a| < \infty$. For a bounded test z and either $\Delta = B_0(F)$ or Γ, unfolding therefore gives

$$\int_{\Delta\backslash G_0(\mathbb{A}_F)} z(g) \sum_{\gamma\in\Delta} s_a(\gamma g)\, dg = \int_{G_0(\mathbb{A}_F)} z(g) s_a(g)\, dg$$

when z is left Δ-invariant. This proves the two initial pairings are equal for the prescribed transposed test. In the Petersson convention linear in the second argument, the functional $z \mapsto \int_{G_0(\mathbb{A}_F)} \overline{s_a(g)} z(g)\, dg$ on the prescribed type has Riesz vector $\mathcal{P}_a$, the periodization just constructed. Its value is $\langle \mathcal{P}_a, z\rangle$ by unfolding. Multiplying the functional by the positive trace normalization multiplies this vector by the same factor.

To identify the initial linear form at $k_f y_0$, use $y_0 \in J_g$ to write $(k_f y_0)\eta = (k_f \eta)(\eta^{-1} y_0 \eta)$. The second factor belongs to $K(g)$. Decompose the first factor in its unique column double class. On the principal coset of that class a rational Bessel element contributing to the seed has core projection in its stabilizer. The stabilizer is trivial by the free level calculation, so the remaining rational translations are exactly U_c. Their cutoffs sum to one. For a finite left translate by $u_f \in U(\mathbb{A}_{F,f})$ at this class, choose $u_0 \in U(F)$ with $u_0^{-1} u_f \in h_c K(g) h_c^{-1}$. Such approximation follows from the additive quotients of the central unipotent filtration. The remaining finite translate preserves the coset, and its simple-root sum lies in $g\varpi$ times the additive annihilator. Put $\chi_f(u) = \psi_f(\sigma_U(u)/(g\varpi))$, where σ_U is the sum in the rational frame, and let χ_∞ be its real counterpart with the same sign. Rational left invariance moves u_0^{-1} to the real argument, where the Bessel factor is $\chi_\infty(u_0)^{-1} = \chi_f(u_0) = \chi_f(u_f)$. In the Riesz integral the conjugated period has phase $\psi_f(-\sigma_U(u_f)/(g\varpi))$. Thus the bottom formula above holds with the Weyl-to-Bessel map A equal to the tensor product of the given compact intertwining coordinate maps. Its contraction with $W(\xi_0)$ is bilinear. The value of the translated column at $k_f\eta$ is the original value at k_f, since the central translation defining its coordinates is by η^{-1}. In particular the resulting coefficient is independent of the twist. The positive rescaling $(\mu_f(K(g))v)^{-1/2}$ in the trace calculation applies to these same coordinate columns and this same $\mathcal{P}$.

6 The seesaw computation

Successive periods of the graded string. Inductively let P_{j+1} be the generic function obtained from P_j at the outer point $u_j = (\xi_j, y_j)$ by the constant term at the string map $Y_j \to Y_{j+1}$. In the mixed Weil model this means applying the Heisenberg translation by the string map and taking the zero coefficient in the nonzero graded directions. The one-step identity, in the measures of the lower chain, is

$$\int_{[B_j]} P_j(b\,u_j)(L_j^{\mathrm{t}} z)(b\,u_j)\, db = [K_j^0 \,:\, J_j] \int_{B_{j+1}(F)\backslash G_{j+1}(\mathbb{A}_F)} P_{j+1}(h) z(h)\, dh. \tag{6.1}$$

Here $[B_j] = B_j(F)\backslash B_j(\mathbb{A}_F)$ and L_j^{t} is the transposed lift on conjugate forms for the bilinear product integral. This formula includes absolute convergence for the tests of the given types and preservation of the generic condition for the next step. In the mixed model the nonmatching rational terms are killed by the unipotent character. The surviving string-map orbit has the next Bessel group as stabilizer. The factor $[K_j^0 : J_j]$ is the normalization of L_j^{t}. Here the kernel interchange uses the two absolutely convergent integrals of the same theta kernel against the two tests [12, §2]. In the bilinear pairing, Fubini interchanges the integrals over the two automorphic quotients. It carries the factor $[K_j^0 : J_j]$ into the transposed arrow. The rapidity and orbit estimates below verify absolute convergence for each pair of tests used here. The interior and boundary calculations below establish this identity for the normalized arrows, including genericity and absolute convergence at the indicated types.

The last three string lengths. For a string of length s, its ordered basis has grades $s-1, s-3, \dots, 1-s$. The core has grade zero. For $s = 1$ and $s = 0$ all grades are zero. The two final arrows have lengths $1 \to 0 \to 0$. Their string maps are zero, and the positive-degree spaces in the constant terms are zero. For the first arrow the Bessel group is the definite group fixing its one vector, and for each zero string the Bessel group is the whole definite unitary group. Accordingly the one-step formula on these arrows follows by Fubini: the kernel in P_{j+1} is exactly the theta kernel at $(b\,u_j, h)$, with both quotients compact. Its factor is $[K_j^0 : J_j]$ from L_j^{t}. The definition of P_{j+1} uses precisely the db on the left of the one-step formula. In particular the different choice of the grading constant at Y_m has no effect on this calculation.

When $m \geq 4$, the preceding arrow has lengths $2 \to 1$. In its mixed expansion the rational negative-degree entries form one column $v = w(e_+)$, where e_+ is the source vector of grade one. All target grades are zero. The central unipotent $e_- \mapsto e_- + te_+$, with $\bar{t} = -t$, acts on the term by the character of the Hermitian norm of v. The prescribed generic character supplies the inverse character at $v_* = cz'$, where z' is the sole target string vector and c is the common lower grading constant. Indeed the string map sends e_+ to cz'. Write T for this map and $\dagger$ for Hermitian adjoint. The off-diagonal entries of the source hyperbolic Gram are one. Thus the lower left entries of $T^\dagger T$ and $w^\dagger w$ are $-Nc\bar{c}$ and (v, v). Since $u - 1 = tE_{+,-}$, the character quotient $\beta(Tu, T)/\beta(wu, w)$ is $\psi((j_{\mathrm{Weil}}/2)t(-Nc\bar{c} - (v, v)))$, by $\beta(x, y) = \psi((j_{\mathrm{Weil}}/2)\,\mathrm{tr}(y^\dagger x))$. Integrating with its probability Haar measure on the compact additive quotient for t inside db kills exactly the terms with unequal norms. The norm $-Nc\bar{c}$ is nonzero. By rational Hermitian Witt extension the remaining vectors are the orbit of cz', with stabilizer $U(X') = B_{j+1}$. Unfolding the rational orbit gives the integral over $B_{j+1}(F)\backslash G_{j+1}(\mathbb{A}_F)$, with counting measure on the rational points. On compact sets of outer variables the same Schwartz seminorm bounds are uniform, giving continuity. Rational invariance follows by reindexing by the rational stabilizer. These are all the genericity conditions on the one-string target, whose unipotent group is trivial.

For the $2 \to 1$ arrow the mixed expansion has precisely the single column just described. The bound in the absolute orbit calculation applies with weight zero on the target. It permits the norm-character selection and the unfolding by counting measure on its rational orbit. The target has no other negative-degree entries.

The rational string orbit at the interior steps. Write $Y = Y_j$, $Y' = Y_{j+1}$, $s = m - 1 - j \geq 3$, and write λ, λ' for the two gradings. The string bases have grades $s - 1 - 2i$ and $s - 2 - 2i$. The core spaces have grade zero. Opposite string vectors pair to one, with middle norm $-N$ when there is a middle vector. The map $T : Y \to Y'$ has degree -1 and kills the core. On string vectors that have a successor in Y', its coefficients t_i are $1, -Nc, c$ in positive, zero and negative source grades if s is odd, and $c, 1$ in positive and negative source grades if s is even. The last source vector maps to zero. Thus $A = T^\dagger T$ is the string shift $e_i \mapsto a_i e_{i+1}$ with $a_i \neq 0$ for $i < s - 1$, and zero on the core. In the next grading $TT^\dagger$ is the same shift prescribed there, since the two constants c agree.

The mixed Fourier expansion in the two gradings. Keep the actual weighted datum, the outer point u_j, and one finite Schwartz vector and real Fock polynomial of the arrow. Write $M = \mathrm{Hom}_E(Y, Y')$ in the rational frames and give a matrix entry from source grade a to target weight b degree $b - a$. Use either target weight system $d = 0$ or $d = \lambda'$. Write M_-, M_0, M_+ for strictly negative, zero and strictly positive degree. The Hermitian trace pairing on M pairs degrees a and $-a$, since the Grams of the two strings pair opposite grades and the cores have grade zero. In particular $M = M_- \oplus M_0 \oplus M_+$ is a hyperbolic splitting with retained space M_0.

Let $K = M_+$ and write $[K] = K(E)\backslash K(\mathbb{A}_E)$ with probability Haar. In the fixed oscillator write Θ for the theta functional and ρ for Heisenberg translation, with $\rho(v)\rho(w) = \beta(v, w)\rho(v + w)$ and $\beta(v, w) = \psi((j_{\mathrm{Weil}}/2)\,\mathrm{tr}(w^\dagger v))$. The function $A_f(z) = \Theta(\rho(z)f)$ on $K(\mathbb{A}_E)$ is periodic modulo $K(E)$. Indeed β is trivial on $K \times K$, and the rational Heisenberg translations fix Θ. For each rational w, put

$$C_d(w; f) = \int_{[K]} \Theta(\rho(z)\rho(w)f)\, dz = \int_{[K]} \beta(z, w)^2 A_f(z)\, dz. \tag{6.2}$$

The second equality commutes $\rho(w)$ to the left and uses rational invariance. By adelic trace duality and the opposite-degree pairing, $\chi_w(z) = \beta(w, z)^2$ runs once through the characters of $[K]$ when w ranges over $M(E)/M_{\geq 0}(E)$. Thus the displayed integral is the Fourier coefficient in the χ_w convention. This checks both the sign of the phase and the counting on the rational labels.

For the fixed finite data the function A_f is invariant under an open compact subgroup of $K(\mathbb{A}_{E,f})$, since Heisenberg translation is smooth at the finite places. Strong approximation for the additive group identifies the resulting quotient of $[K]$ with a real torus: its lattice is the intersection of $K(E)$ with that open compact subgroup in the finite coordinates. In particular the characters that occur form a lattice of real frequencies. If its real dimension is D, application of $(1 - \Delta)^b$ for $2b > D$ bounds the absolute sum of coefficients by a constant times the maximum of finitely many real derivatives of A_f. Indeed the individual bound is the same maximum times $(1 + 4\pi^2\|\xi\|^2)^{-b}$ at frequency ξ. Fourier inversion at zero proves $\Theta(f) = \sum_w C_d(w; f)$ with absolute convergence.

For the agreement chart, let S consist of the entries with source grade $a > \max(0, \lambda'(b))$, where b now denotes the target basis index. The dual coordinates to S span a rational subspace K_S of K in both gradings. Requiring w to agree with T on S is exactly $\chi_w|_{[K_S]} = \chi_T|_{[K_S]}$. Summing the absolutely convergent series on this subtorus against χ_T^{-1} gives, in either model, the same probability mean

$$\sum_{w:\ w|_S = T|_S} C_d(w; f) = \int_{[K_S]} \beta(z, T)^2 \Theta(\rho(z)f)\, dz. \tag{6.3}$$

For $z = y^\dagger$, the phase on the right is $\psi(j_{\mathrm{Weil}}\,\mathrm{tr}(yT))^{-1}$, since j_{Weil} is anti-invariant under conjugation and ψ is given by the trace. This is the adjoint-coordinate reading of the common mean.

The normalized zero term. In a hyperbolic splitting $M = L \oplus U \oplus K$, choose the Schrödinger model retaining U and taking the mixed variable x along L. Use the same pair of rational dual frames at all places, so that K acts by $\rho(z)F(x) = \beta(x, z)^2 F(x)$ and L translates x. The retained vector action is the oscillator of U. In these coordinates the rational theta functional is a fixed nonzero scalar κ times the sum of the retained theta functional over the rational x. This follows from the rational Heisenberg invariance in the ordinary Schrödinger realization: invariance under rational modulations restricts the support of the functional to the rational lattice, and invariance under rational translations gives its counting measure. Use the ratio of the original theta functionals for κ. The average over $[K]$ kills every $x \neq 0$, so the normalized zero term is $Zf = \kappa F(0)$. For a retained Heisenberg translate y the same argument gives

$$\Theta_U(\rho_U(y)Zf) = \int_{[K]} \Theta(\rho(y)\rho(z)f)\, dz.$$

This fixes the normalization of Zf and gives

$$C_d(w; f) = \Theta_U(Z\rho(w)f).$$

At each place evaluate the variables dual to the same rational isotropic space. In this chart a finite Bruhat–Schwartz vector remains locally constant and compactly supported after evaluation. On a pure tensor the real and finite evaluations separate with the fixed scalar. For the paired arrows use the real adapted vacuum that defines the finite map and assign its fixed value to that map. Other real inputs give one retained real vector, independent of the finite input. The actual w is in rational coordinates, so the finite shift keeps the original frame and the factor α_W^{-1}.

For g stabilizing K, let $a = g|_K$ in the frame of K and g_0 be its induced action on $K^\perp/K$. We claim

$$Z\omega(g) = v(\det a)^{-1}|\det a|_E^{1/2}\,\omega_U(g_0)Z. \tag{6.4}$$

To calculate the scalar, first use the covariance of the Heisenberg representation. In the mixed x coordinate, invariance under all $K(\mathbb{A}_E)$-modulations characterizes the distributions that factor through $F(0)$. The characters separate every nonzero x; at a real place their derivatives also rule out derivatives of evaluation. At a finite place locally constant test functions on compact sets give the same support assertion. The covariance with the retained Heisenberg translations then determines the operator on $F(0)$ up to scalar, by irreducibility of the retained oscillator for the Heisenberg group. This argument also works for the continuous Schwartz operators, since their kernels commuting with modulation and translation are scalar multiples of the diagonal kernel.

The scalar on the unipotent radical of the stabilizer of K is a continuous character. It is trivial. At each place, scalar dilation t on K, the identity on U, and t^{-1} on L, for t in the base field, contracts the radical as t tends to zero. Conjugating leaves a character unchanged. Elements supported at finitely many places are dense in the adelic radical, giving the same conclusion globally. The Levi action on the retained factor is $\omega_U(g_0)$: the orthogonal sum has tensor vacuum and the normalized vacuum coefficients multiply. It remains to calculate the scalar for $g_a = \mathrm{diag}(a^{-*}, 1, a)$ on L, U, K. Covariance with translation along L sends $F(x)$ to $c(a)F(a^*x)$. Unitarity in the mixed L^2 measure gives $|c(a)| = |\det a|_E^{1/2}$.

For the phase use the adapted vacuum defined by the finite lattice action. In the adapted real frame the hyperbolic mixed vacuum is a positive Gaussian up to a constant, by the unitary Fourier transform of the Gaussian. In the adapted finite frame it is constant on the integral lattice and zero off it. Indeed the integral K-moves force self-dual lattice support by modulation, and the integral L-moves force constancy there by translation. The integral hyperbolic frame and the conductor in these coordinates are those of the standard vacuum, including at the places over two. Consequently the vacuum matrix coefficient of g_a is $c(a)$ times a strictly positive real number: it pairs this nonnegative profile with its a^*-dilate, whose support meets it in a neighborhood of zero.

Let k be the fixed adapted move from the standard vacuum frame. By unitarity, the coefficient of the adapted vacuum at g_a is the standard vacuum coefficient at $k^{-1}g_ak$, namely $\overline{\Delta(k^{-1}g_ak)}^{-1}$ in the Hermitian form linear in its second argument. In the following block calculation Δ is computed in this standard frame. At a real place the negative block of the transformed g_a is $(a^{-*} + a)/2 = a^{-*}(I + a^*a)/2$, and Δ_∞ is its conjugate determinant. Its phase is therefore $v_\infty(\det a)^{-1}$. At the finite places, write $D_f(A) = \chi_f(\det u_A)$, where the columns of u_A span $\widehat{\mathcal{O}}_E^p \cap A^{-1}\widehat{\mathcal{O}}_E^p$ and $\chi_f = v_f|\cdot|_E^{-1/2}$. This is independent of the integral basis since χ_f is unramified. The identity $A(\widehat{\mathcal{O}}_E^p \cap A^{-1}\widehat{\mathcal{O}}_E^p) = \widehat{\mathcal{O}}_E^p \cap A\widehat{\mathcal{O}}_E^p$ gives $D_f(A^{-1}) = \chi_f(\det A)D_f(A)$. Smith factors show $D_f(A^*) = \overline{D_f(A)}$, using $\chi_f(\bar{d}) = \overline{\chi_f(d)}$. The integral hyperbolic frame therefore gives

$$\Delta_f(g_a) = D_f(a^{-*})D_f(a) = \chi_f(\overline{\det a})\,|D_f(a)|^2.$$

Thus $\Delta(g_a)v(\det a) > 0$, and the same phase calculation and the modulus above prove (6.4). For the middle pair use $(uw)^r v^{-(2r-1)}$ on the positive-degree part and v^{-1} on the removed plane. These ideles have

modulus one, so their two contributions change the exponents by $(-r, 2r, -r)$ at (u, v, w). The terminal zero string uses no positive-degree factor.

For rational $\gamma \in B_j(E^+)$ and δ preserving the filtration defined by the target weights, $C_d(w; \omega(\gamma, \delta)f)$ is the coefficient at $\delta^{-1} w\gamma$. In the mean in (6.2) this follows by rational invariance of Θ and change of variables on $[K]$. For adelic moves raising the respective grades strictly, the action on M is the identity modulo a degree-raising operator. Its action on K has determinant one and its retained action is the identity. Hence the same substitution follows by (6.4); the determinant characters of the arrow are one on these unipotents. When a move changes w to $w + q$ with $q \in K(\mathbb{A}_E)$, the Heisenberg law and $\rho(q)$-invariance of Z give $C_d(w+q; f) = \beta(w, w+q)C_d(w; f)$. In particular the source degree-two test changes the integrand by $\beta(Tu, T)/\beta(wu, w)$. On the target the shift $v'^{-1}T - T$ is strictly positive when $v' - 1$ raises grades by two. The same rule gives $\beta(T, v'^{-1}T) = \psi((j_{\mathrm{Weil}}/2)\,\mathrm{tr}(TT^{\dagger}v'))$. Thus the retained T-coefficient has the required next genericity, using the original common constant c.

The real types of the fixed lower-chain tests. At a real place put $R_j = r + jt$ and $\epsilon_j = j \bmod 2$. In doubled Harish–Chandra parameters the entries on Y_j are the following entries minus ϵ_j:

$$\begin{aligned} &k - 2 - 2i && (0 \le i < k - 1),\\ &2h_2 - 2w_{a+1} && (0 \le a < m),\\ &t - 1 - 2w_\ell - 2b && (1 \le \ell \le j,\ 0 \le b < t). \end{aligned}$$

Only the middle list can have negative signs. Its sign at a is $(-1)^{a+\min(a,j)}$. Write $w = w_{j+1}$ and put $\gamma = \lfloor R_j/2 \rfloor - w$, $\gamma' = \lfloor R_{j+1}/2 \rfloor - w$. These are the two actual real weights of the rank-twisted character, since ϑ has real weight zero. The Fock rows are on Y_{j+1} and the columns are the conjugate side on Y_j. Their determinant exponents are γ and $R_{j+1} - \gamma'$. Consequently, for the actual doubled parameters l_j, l_{j+1}, the centered row and column entries are

$$l_{j+1} - 2\gamma + R_j, \qquad -l_j + 2\gamma' - R_{j+1}.$$

Since r, t are odd, $R_j \bmod 2 = \epsilon_{j+1}$ and $R_{j+1} \bmod 2 = \epsilon_j$. The centers of the common entries are therefore x and $-x$, with x the unshifted doubled entry plus $2w$. Every x in the first and the old added lists is $> t$ and keeps its sign. The stem centers are $2h_2 + 2(j - a)h$, with $h = 2h_2$. They are $> t$ for $a \le j$ and $< -t$ for $a \ge j + 1$. Exactly the latter signs reverse. The new block has positive sign and centers $t - 1 - 2b$, all of absolute value $< t$. This proves the signed correspondence in the paired Fock vacuum, including the final definite row.

The parameter is regular at both nodes, since the same-sign strings have distinct entries two apart and the centers of different lists are disjoint. For opposite signs the doubled stem-to-stem gap is at least $2h$. A negative stem index has $a \ge j + 1$, so its distance from any old added-block center is at least $3h$, before subtracting $t - 1$. The distance from the compact first list is larger still. On the target, the new positive string is at a distance of at least $3h - (t - 1) > 2h$ in doubled units from any negative stem entry, since the closest possible negative stem has $a = j + 2$. Thus the opposite-sign gap of the undoubled parameter is at least h. The datum has $100m \le k \le r$ and $t = 5r$, so $R_j + R_{j+1} < r^2$. With $h = 2r^8$ the required gap $500(R_j + R_{j+1})^3$ follows.

Here is the highest polynomial on this matching. For a common row with center x_i and sign a_i, its corresponding column has sign $a'_i = a_i \,\mathrm{sgn}(x_i)$. For any row i, subtract its vacuum weight from its lowest compact weight, and call the result d_i. The formula $\Lambda = \lambda + \rho_{\mathrm{n}} - \rho_{\mathrm{c}}$ gives

$$2d_i = x_i + \sum_u a_i a_u \,\mathrm{sgn}(x_u - x_i) - \sum_v a_i a'_v.$$

The row sum includes the t new entries; the column sum is over the old entries. For a new row numbered b in the added string, the matched terms cancel, and the remaining sum is $2b - t + 1$. Thus its relative

weight is zero. For a common row add its formula to the column formula for its mate. The two center terms cancel. A second matched pair with centers of the same sign gives zero in the two sums. For opposite signs, $\operatorname{sgn}(x_u - x_i) = \operatorname{sgn}(x_u)$, and the sign reversal gives zero again. The new rows have centers between the old positive and negative centers, so their contribution also cancels. Hence the two relative weights are negatives of each other. The sign of d_i on a common row is the sign of x_i, since the sum of the two ranks bounds the correction and $|x_i|$ exceeds that sum in the actual datum.

For each pair of row and column signs and each height $q \geq 0$ form the minor on the matched pairs of that signature with $|d_i| > q$. Multiply the nonempty minors. For either sign of d_i the rows chosen are an initial segment in the compact order appropriate to that variable sign, and so are the paired columns in the swapped order. Indeed the vacuum-subtracted weights are monotone among rows of the same sign in the Harish–Chandra order. For neighboring such rows this follows from the displayed formula: their highest-weight difference is at least half their doubled gap minus one. Rows with the other variable sign are outside the initial segment. The compact raising differential operators therefore give determinants with a repeated row or column. The diagonal weights count the occurrences with their variable signs and give the desired lowest compact weights. Each row and each column uses variables of only one sign, so all cross-sign Laplacians kill the product. It is nonzero on the matrix with the matching entries one and all other entries zero. Normalize it there by its value. This supplies the joint highest line of the two actual lowest compact types. Howe's joint harmonic decomposition for a real reductive dual pair in the polynomial oscillator pairs the two compact types occurring in the common harmonic subspace with multiplicity one for their product. The cross-sign Laplacians and the weights just checked put this polynomial in that subspace [13, Theorem 2.2(1)].

Rapidity and the real line in the two theta integrals. For the real correspondence divide the two determinant twists out of the covered oscillator. In the convention with the source acting by the right inverse, the source column vacuum after removing its determinant center is $-a_i'(p' - q')/2$, where (p', q') is the target signature. Passing to the source action by the inverse transpose negates the torus weights and preserves the Hermitian sign spaces. The source vacuum then has weights $a_i'(p' - q')/2$ and the target vacuum has weights $a_i(p - q)/2$ for source signature (p, q), precisely Paul's covers on the two signed groups. The determinant centers removed on those two sides are $\gamma' - R_{j+1}/2$ and $\gamma - R_j/2$. Thus the smaller genuine parameter is $(x_i/2)$ with the source signs, and the larger one has these entries with the target signs and the positive string of length t. Both are integral on the genuine covers, whose ranks have opposite parities. The source action by the inverse transpose also accounts for the dual in the bilinear integral. Equivalently one can use the determinant-twisted splitting after $\chi_V(z) = (z/|z|)^{2\gamma' - R_{j+1}}$ and $\chi_W(z) = (z/|z|)^{2\gamma - R_j}$ in the left-action convention. Their restrictions to $\mathbb{R}^\times$ have the required parities. The proof of Ichino's Lemma 7.10 uses this inverse transpose and (χ_W, χ_V^{-1}) in the right-action model. Its vacuum shifts are the ones just computed. The actual global kernel keeps its two original characters and the same finite vector.

Paul's Theorem 2.7(b) starts with a genuine discrete series whose positive and negative labels have magnitudes at least $(t + 1)/2$, where t is the increase in rank. Write a_+, a_- for the counts of positive and negative entries in its positive-sign block and b_+, b_- for those in its negative-sign block. For target signature $(a_+ + b_- + t, a_- + b_+)$, the theorem gives the discrete series with the common entries after reversing the labels of the negative entries, and with the string $(t - 1)/2, (t - 3)/2, \ldots, -(t - 1)/2$ in the positive-sign block. The common centers after halving meet the gap and the computed target parameter is precisely this one. For the reverse arrow Proposition 3.4(ii) deletes such a positive-labeled central string when the other entries have magnitudes at least $(t + 1)/2$, reversing the labels of the remaining negative entries. The proposition gives first occurrence in that tower for the maximal string. Here both adjacent integers are absent, so this string is maximal. For the conjugate convention, Proposition 2.9 takes the paired contragredients when the target signature is interchanged, which matches the sign change from swapping the oscillator grading. For the forward tower use Atobe's criterion in Ichino's Theorem 6.1(i).

With opposite rank parities this is the case $k_0 = -1$. The smaller centered parameter has no zero and is integral, so no positive odd central alternating string occurs and $k_\pi = -1$. In this case with no zero, Theorem 6.1(i) says that for a signature $(r_\pi + l + 2t' + 1, s_\pi + l)$, $t' \geq 1$, nonvanishing is equivalent to $l \geq 0$ and $|C_\pi^{\pm}(l + t')| \leq l$ for the two signed near sets. These are the surviving signed entries $(\nu, \pm)$ of the criterion with $0 \leq -1 \pm \nu < l + t'$. The counts r_π, s_π are the numbers of positive and negative labels after reversing the negative entries. The signature of Y_{j+1} is $(r_\pi + t, s_\pi)$, which uses $l = 0$ and $t_{\mathrm{Ichino}} = (t-1)/2$. The two near sets are empty by the strict gaps. Removing a hyperbolic plane makes $l < 0$ while keeping the offset t_{Ichino}, so every smaller member in that tower vanishes. In the zero negative-index case this condition is vacuous. For transposition the kernel is the same function with its arguments interchanged. The conjugate target input pairs with the larger representation on the oscillator side, and the output belongs to the conjugate of the smaller one. If a skew-Hermitian model is used for the other member of the pair, identify the two roles by the same trace-zero multiple of the forms as in the symmetric tower formula. [13, Theorem 2.7(b) and Propositions 2.9 and 3.4(ii)] [14, Theorem 6.1(i), Remarks 6.2 and 7.3 and Lemma 7.10]

The input at each arrow is a bounded automorphic vector fixed by the rank-one real discrete-series kernel. At fixed level it is in L^2. Formal-degree orthogonality puts it on the lowest compact line of copies of the prescribed discrete series, and its real translates generate those copies. A compactly supported smooth kernel acts as the identity on this line: average a sufficiently small smooth bump on the two sides against its rank-one compact idempotent and normalize the nonzero resulting coefficient. Thus the input is smooth. At each indefinite factor the restriction of its discrete series to $SU(p, q)$ is irreducible and square integrable: the compact central circle acts by a scalar, and the finite central cover transfers square integrability. Harish–Chandra's Corollary 1 to Lemma 65 puts every compact-finite matrix coefficient of such a discrete series in the Harish–Chandra Schwartz space [15, §33, pp. 77–78]. A discrete-series coefficient is an eigenfunction for the enveloping-algebra center. For a center-finite Harish–Chandra Schwartz function on the semisimple real group, its normalized constant term along any proper real parabolic vanishes, by the constant-term calculation of Wallach [16, §§7.2.1–7.2.2 and 7.7.1]. The statement applies to all left and right translates, since translations preserve these hypotheses. It gives zero integrals over the proper unipotent radicals, whose normalizing factor at the identity is one. The real radical is unchanged on passing from $SU(p, q)$ to $U(p, q)$. At a compact real factor there is no proper real parabolic.

For integrability of the kernel, Miličić's criterion states that for a discrete series of a connected semisimple group with finite center and parameter λ, its compact-finite coefficients belong to L^p, $1 \leq p < 2$, when the noncompact-wall distances satisfy

$$|(\lambda, \alpha)| > (2/p - 1)\tfrac{1}{4} \sum_{\beta} |(\alpha, \beta)| \quad \text{for every noncompact root } \alpha.$$

We use $p = 1$ on the same derived group. Type A_{R-1} with squared root length two gives the right side $R - 1$, less than the opposite-sign parameter gaps proved above [17, Part III, §3, Corollary to Theorem 1]. The compact central factor does not affect L^1. Take the automorphic constant term of the convolution reproducing the input. Boundedness, L^1 of the kernel, and compactness of the adelic unipotent quotient justify interchanging the integrations. At the chosen automorphic point the constant term is left invariant in the real convolution variable by the conjugated real radical. Disintegration along this radical makes the inner integral of the real kernel zero by the local result. Thus the input is cuspidal. On each finite-level component the Miller–Schmid decay theorem for smooth cuspidal automorphic forms with uniform moderate growth gives rapid decay, with bounds by smooth-vector seminorms [18, Theorems 2.23–2.24]. Here the input has the prescribed infinitesimal character and compact-finite vector, and the realization by an integrable reproducer gives smooth vectors with a common moderate exponent for their translates.

For the fixed finite vector and finite-level representatives, the Schwartz estimate for the ordinary theta kernel gives

$$|\Theta(\omega(g, h)(\phi \otimes f_f))| \leq C\, p(\phi)(\|g_\infty\|\|h_\infty\|)^A$$

on reduced real representatives. Here p is a fixed real Schwartz seminorm, A is fixed and the group norms dominate the matrix and inverse norms. This follows by the lattice sum and the polynomial seminorm bound for the Weil action given in the absolute-orbit calculation below. A right differential operator in h_∞ replaces ϕ by its corresponding oscillator derivative, leaving p and A fixed. Rapid decay on the input justifies differentiation under the theta integral and gives uniform moderate growth of the lifted function. Rational invariance comes from the theta sum, and an open compact stabilizer of the given finite vector on the output side gives its level. Write $\mathcal{T}(\phi, z)$ for the bilinear theta integral and R for right translation. Right Haar measure and the commuting actions give

$$\mathcal{T}(\omega(a,1)\phi, R(a)z) = \mathcal{T}(\phi, z), \qquad \mathcal{T}(\omega(1,b)\phi, z) = R(b)\mathcal{T}(\phi, z).$$

Fix a target point and the other local vectors in a cuspidal constituent of the input. Varying the oscillator vector and the source vector at one real place gives a continuous bilinear pairing invariant under simultaneous source translation: moving a right source translation off the oscillator gives its inverse on the form. The smooth-vector seminorm bound in the cuspidal decay theorem supplies continuity in the source vector, and the displayed bound supplies continuity in the Schwartz vector. In its polynomial and compact-finite vectors a nonzero pairing would give a surjection onto the contragredient source module, by irreducibility. The polynomial vectors are dense in the Schwartz realization, so the same vanishing criterion applies to the restricted Schwartz vector in the constant term. By Howe's correspondence all such pairings vanish against the smaller Witt-tower members at a real place. Wu's constant-term formula applies over a quadratic extension to the Weil kernel of a Hermitian and a skew-Hermitian unitary space, with a cuspidal source and compatible splitting characters. For a parabolic stabilizing an isotropic u-plane it expresses the constant term as the lift to the smaller Witt member on the restricted Schwartz vector, with the induced determinant character and modulus. The same calculation with the two roles exchanged applies to the transpose. For our oscillator multiplication by j_{Weil} makes the appropriate factor skew-Hermitian, and the source is cuspidal by the preceding calculation. When the target is lowered by $2u$ dimensions, the isotropic map block consists of the maps into its removed u-plane. The input element g acts there with determinant $(\det g)^{-u}$ by precomposition. The zero-term law (6.4) therefore changes the exponent in the inverse input character from γ' to $\gamma'-u$. The genuine center $\gamma'-R_{j+1}/2$ is preserved when R_{j+1} decreases by $2u$. In the swapped arrow lower the original source by $2u$. On the removed isotropic map block the input h has determinant $(\det h)^u$ by postcomposition. This changes γ to $\gamma - u$, preserving $\gamma - R_j/2$. All these determinants have adelic norm one, and the base character $\mu_{j+1}^0 \vartheta^{j+1}$ is unchanged. This checks the input splitting in the two tower applications, including the conjugate orientation. The lower vanishing now gives zero proper constant terms in each direction. Use the growth reduction in the proof of Miller–Schmid's Theorem 2.24, §5 after (5.11), directly on this smooth uniformly moderate function. On a Siegel set choose a maximal-parabolic radical for a largest split simple-root coordinate. The zero mean bounds the function on each compact unipotent quotient by a sum of its unipotent derivatives. Moving those derivatives to the right gives factors $e^{-\alpha(\log a)}$, with each root α in that radical bounded below by the chosen simple-root coordinate up to a constant. Right derivatives have the same zero means and the same growth exponent. Iteration gives any negative power of the reduction height. This proves rapid decay and puts the lift in the cuspidal L^2-space before identifying its infinitesimal character there. [19, Proposition 3.1, including its symmetric version]

On the cuspidal target take an irreducible constituent of the lift. Pair the theta integral against this constituent, fixing the vectors at all other places. The pairing is continuous in the local oscillator and the two local smooth vectors by the same Schwartz estimate and the input and output cuspidal seminorm bounds. On polynomial oscillator vectors the source and target compact equivariance put the pairing in the compact-finite parts of the two distribution duals. These are the usual Harish–Chandra contragredients by admissibility. With the conjugation on the output test this gives a nonzero quotient onto the tensor of the contragredient source and the output representation. Howe's smooth oscillator theorem says that for a real reductive dual pair, the irreducible admissible quotients of the two metaplectic preimages

occurring in the oscillator are paired with at most one partner, via the corresponding maximal isotypic quotient. Pulling the splitting characters back to these covers gives our actual oscillator and its source and target quotients. Thus the output is the paired discrete series calculated from the signed parameters [20, Theorem 1]. The actual joint highest polynomial restricts the output to the highest line of the multiplicity-one lowest compact type by compact equivariance. In the transpose it is the conjugate line of the smaller representation, with the reversed compact positive order. The formal-degree matrix coefficient of the specified unit vector in that line, in the original Haar measure, fixes the line exactly by Schur orthogonality. This argument applies to each finite-level cuspidal input by linearity on its fixed type, and to the actual finite inclusion vector without further normalization. It proves rapidity of the test and the required real line for its successive transposes. The sole index in the transpose remains $[K_j^0 : J_j]$.

Absolute sums in the orbit integrals. Let z be the rapid automorphic test of the conjugate target type. First take $d = 0$. The full absolute sum of mixed terms is invariant by reindexing under the rational target group, and under simultaneous rational source substitution with the labels. Choose a compact set of representatives for the source Bessel quotient, which is compact because its core is definite and its unipotent fibers have compact rational quotients. Keep u_j fixed. At the finite places use the finitely many adelic class representatives of the target at a common level of the finite datum and z. The finite translates on the compact source representatives take finitely many values. Hence one finite open compact subgroup and one torus lattice suffice in the bound following (6.2).

For the real bound choose reduced target representatives at those finite points. Each derivative on the fixed compact torus is a Heisenberg derivative applied to the real Schwartz vector. A lattice sum in a fixed Schrödinger polarization is bounded by a Schwartz seminorm of order greater than its real lattice dimension. The real metaplectic action on a fixed seminorm is bounded by a power of the norm of its matrix and inverse: the compact factors act boundedly, and the split diagonal action dilates the variables and their derivatives with its determinant factor. Heisenberg derivatives are polynomial differential operators in this chart. The source representatives and the torus representatives range over compact sets. Thus for some C, A the absolute sum on these representatives is at most $C\ \mathrm{ht}(h)^A$, where $\mathrm{ht} \geq 1$ is the reduction height of the target. The determinant phases have modulus one. The constants may depend on the fixed datum, test, and outer point, as allowed by the tested period. Since z is rapid and the target quotient has finite volume, the full absolute sum multiplied by $|P_j(ru_j)z(h)|$ is integrable in the two original quotient measures. The same seminorm bound on compact sets of h works in either grading and is locally uniform in the outer point.

Use Tonelli on this absolute sum, with counting measure on the rational orbit labels. If Γ is the source rational group and S_ξ the stabilizer of a label, summing $k_{\gamma\xi}(r) = k_\xi(\gamma^{-1}r)$ on Γ/S_ξ then integrating over $\Gamma\backslash B_j$ is integration of k_ξ on $S_\xi\backslash B_j$. This is the quotient integration formula with the same Haar and counting measures on the two discrete subgroups. For a nontrivial character test $k_\xi(ur) = c_u k_\xi(r)$, integrate next through the closed subgroup of left translations acting on the function by unit scalars. It contains S_ξ and u, and is closed by continuity and evaluation at a point where the function is nonzero. On each fiber right multiplication by u is defined and scales Haar by a positive factor. It also gives the unit scalar c_u, so the fiber integral vanishes if $c_u \neq 1$. This argument on the stabilizer quotient uses integration in stages and needs no invariance of that quotient under left multiplication by u.

For the target chart unfold the original $d = 0$ translates first. The absolute integral of their sum over the target rational quotient bounds the integral of the absolute chart sum over $B_{j+1}(E^+)\backslash G_{j+1}(\mathbb{A}_{E^+})$. At each fixed target point use (6.3) and the locally absolute actual grading to reduce the chart to T. This proves absolute integrability of the surviving period on the larger quotient, with its quotient measure. The factor $[K_j^0 : J_j]$ is still the normalization of the transposed lift. The case $2 \to 1$ is the same argument with the single column and trivial target grading, followed by the norm-character selection already computed for that arrow.

Here is the calculation of the surviving labels. For a vector v of strictly positive source grades and a vector x such that $(x, v) = 0$, use the Eichler matrix

$$u = 1 + vx^\dagger - xv^\dagger - \tfrac{1}{2}(x, x)vv^\dagger.$$

This is $1+D+D^2/2$ for $D = vx^\dagger - xv^\dagger$, with $D^\dagger = -D$ and $D^3 = 0$. In the tests below its nonconstant terms raise grades by at least two, and $w(u-1)$ is strictly positive in the indicated model. When $(x, x) \neq 0$ in these tests, Tv and wv are isotropic (the zero vector is allowed). Thus the last term makes zero contribution to the two traces. The phase comparison, with x multiplied by arbitrary adelic scalars, gives

$$(wx, wv) = (Tx, Tv).$$

Indeed the two remaining rank-one traces are conjugates with opposite signs, and multiplication by j_{Weil} and varying the scalar tests the difference by the nondegenerate adelic additive pairing.

For weights $d = 0$, first take v, x in the strictly positive source space $Y_{>0}$. Their adjoints are supported on the negative columns, so all the degree conditions hold. This gives equality of the two Gram forms on $Y_{>0}$. To prove injectivity there, if $wv = 0$ for a nonzero $v \in Y_{>0}$, choose its lowest nonzero grade a and take x to be dual to the successor of the string vector of grade a. Thus x has grade $2 - a$, and the terms of $u - 1$ raise grades by at least two. The term involving wx lies on negative columns and the terms involving wv vanish. If x has grade zero then Tv has strictly positive grades. The comparison gives $(x, Av) = 0$, whereas the coefficient of the successor of the lowest entry of v in Av is nonzero. Now rational Witt extension applies to the isometric embeddings of $Y_{>0}$ given by w and T, including its possibly degenerate Gram form. A rational target unitary move therefore makes the entire restriction to $Y_{>0}$ agree with T.

For the passage to the actual weights, keep a coarser chart of labels: require equality with T only on entries whose source grade $a > 0$ is strictly larger than the target grade. The chart is stable under the rational source $B_j(F)$. Its translates by $G_{j+1}(F)$ cover all labels not deleted in the weights-zero model. Two translates intersect only when their quotient belongs to $B_{j+1}(F)$. To see this, for each nonnegative target string vector f_i of grade b, take the source vector e_i of grade $b + 1$. Its image in the chart is $t_i f_i$ modulo grades $> b$. Descending through these b, a rational target move carrying one chart label to another must fix f_i modulo grades $> b$. These are the defining conditions of B_{j+1}, and they also imply that every element of B_{j+1} preserves the chart. Thus one unfolds the chart translates to $B_{j+1}(F)\backslash G_{j+1}(\mathbb{A}_F)$. The sum on this chart agrees in the two models by (6.3), applied to these prescribed entries. Use the original weights-zero absolute integral before the target chart unfolding and then use the actual-grading coefficient at each unfolded target point.

In the actual weights, the remaining entries to compare have target grade $b < 0$ and source grade $b < a \leq 0$. Work in increasing b. For a source basis vector x of grade a, wx has no grades below b by induction and the degree of T. Let $v = e_i$ be the source vector of grade $1 - b$ mapping to $t_i f_i$, where f_i is the target vector dual to the row of grade b. In the chart, wv is $t_i f_i$ plus terms of grades $> -b$. In the Eichler test the term $wvx^\dagger$ has degree at least $a - b$, and $wxv^\dagger$ has degree at least one. The correction term also has positive degree. Both Tv, wv have positive grades. The identity $(wx, wv) = (Tx, Tv)$ now reads off the coefficient in the specified row of grade b, because all higher grades in wv pair with grades below b. This proves $w - T$ has no negative degree entries, and hence only the class T contributes in the actual model. The source Bessel group preserves that rational class, since it fixes the half-string modulo higher grades and the nonconstant core block is killed by T.

The remaining coefficient is exactly the integral defining P_{j+1}. It is continuous by compactness of the rational source Bessel quotient and the same seminorm bound locally in the outer variables. The rational target Bessel group fixes the frequency class. If $v' - 1$ raises target grades by at least two, the shift from T to $v'^{-1}T$ is strictly positive. The Heisenberg law now gives the character $\psi((j_{\mathrm{Weil}}/2)\operatorname{tr}(TT^\dagger v'))$; the equality of c identifies it with the next prescribed shift. This proves the genericity and the period formula (6.1). Rational orbit unfolding uses the original quotient measures, and the only lift factor in that formula is $[K_j^0 : J_j]$.

The paired tensor at the outer points. Pair edges $2i, 2i+1$, $0 \le i < m/2$. For $y = (y_0, \dots, y_{m-1}, 1)$, translate the finite inclusion datum of arrow j by (y_j, y_{j+1}). Take the paired zero term in $\mathrm{Hom}(Y_{2i+1}, Y_{2i}^{-} \oplus Y_{2i+2})$ at the two translated string maps, with finite tensor $E_{i,\vartheta}(y)$. Put $E_{\vartheta}(y) = \bigotimes_{i<m/2} E_{i,\vartheta}(y)$. The next calculations realize the paired terms as coefficients of these definite lifts. The ramified support calculation then bounds the possible adapted profiles. The integral orbit comparison turns those bounds into finite cuts.

6.1 Primitive frames in the finite lattices

Let A be the integer ring of the unramified quadratic extension of a nonarchimedean local field, or the split quadratic algebra of integer rings, with conjugation. Write ϖ_v for a base uniformizer and q_v for the base residue cardinality. In a self-dual Hermitian lattice of rank r, a primitive l-frame has linearly independent columns in each residual factor. Two primitive frames with the same Gram matrix are in a single integral unitary orbit. At the residue level in the unramified case, apply Witt extension to the isometry of the two column spans. This applies to degenerate spans in the nondegenerate ambient Hermitian space. In the split case write the data as an injection $X : \mathbf{F}_{q_v}^{l} \to \mathbf{F}_{q_v}^{r}$ and a surjection $L : \mathbf{F}_{q_v}^{r} \to \mathbf{F}_{q_v}^{l}$, with LX equal to the Gram matrix in the chosen half. The match $Xv \mapsto X'v$ carries $\operatorname{im} X \cap \ker L$ onto $\operatorname{im} X' \cap \ker L'$. Extend it on the kernels and then on complements with matching images under L, L'. This is an invertible change of the first split factor, whose inverse transpose changes the second factor as prescribed.

For lifting, write J for the ambient Gram. A lift of the current unitary matrix from depth $b \ge 1$ has a Gram error $\varpi_v^b H$ with H Hermitian modulo ϖ_v. A change by $1 + \varpi_v^b A_0$ removes the error by solving $JA_0 + A_0^* J = -H$ on the residue field. One can take $JA_0 = -c_0 H$ since $c_0 + \bar{c}_0 = 1$. After this step the discrepancy D of the two frames at depth b, divided by ϖ_v^b, satisfies $X^* JD + D^* JX = 0$ in the residue algebra, for the current image X of the first frame. Extend the columns X to an invertible matrix P. Prescribe the first l columns of $P^* JA_1 P$ to be $P^* JD$. Their upper-left block is skew-Hermitian by the displayed identity. Complete the matrix to a skew-Hermitian matrix by the transpose rule, with zero lower-right block. Then $A_1 X = D$, so $1 + \varpi_v^b A_1$ matches the frames at the next depth and remains unitary there. Completeness gives the integral isometry. This proves the claim at every residue characteristic.

We also record the normalized number of primitive frames of any given integral Hermitian Gram S of size l. If the corank of its reduction is c and $l + c \le r$, the count at depth $b \ge 1$, divided by $q_v^{b(2rl-l^2)}$, is

$$\prod_{j=r-l-c+1}^{r} (1 - \xi^j q_v^{-j}), \qquad \xi = \begin{cases} 1 & \text{split,} \\ -1 & \text{unramified nonsplit.} \end{cases}$$

In the nonsplit case the first $l - c$ residual columns can be chosen as a nondegenerate frame, in $|U_r(\mathbf{F}_{q_v})|/|U_{r-l+c}(\mathbf{F}_{q_v})|$ ways. The size of U_j follows by choosing successive norm-one columns, using $q_v^{j-1}(q_v^j - (-1)^j)$ choices for the first. After choosing i of the isotropic columns, the residual orthogonal quotient has dimension $D = r - l + c - 2i$. There are $(q_v^D - (-1)^D)(q_v^{D-1} - (-1)^{D-1})$ nonzero isotropic vectors in it, each with q_v^{2i} lifts. Both vector counts follow by diagonalizing the form and using the $q_v + 1$ preimages of a nonzero residue norm under the norm map. More explicitly, if Z_D counts norm zero vectors and B_D counts vectors of any specified nonzero norm, then $Z_D = Z_{D-1} + (q_v^2 - 1)B_{D-1}$ and $Z_D + (q_v - 1)B_D = q_v^{2D}$. Multiplication gives the formula at depth one. At a split place choose X_1 injective in $X_1^{\mathsf{T}} X_2 = \bar{S}$. This gives the last l factors in the displayed product. The free block of X_2 must be injective on the c-dimensional kernel of $\bar{S}$, giving the remaining c factors. The derivative of the Gram map at a primitive frame is onto. In a field factor, $H \mapsto X^* JH + H^* JX$ is onto by primitivity and the trace-one element; in the split coordinates use $H_2 \mapsto X_1^{\mathsf{T}} H_2$. Each solution has exactly $q_v^{2rl-l^2}$ lifts at the next depth, as asserted.

The integral orbit of the initial half-string. Set $p_i = e_i$ for the string basis with $0 \le i < s$ and $i^\vee = s - 1 - i$. An adapted basis b is block upper triangular for the order which gives the core index $2s$

and e_i the index $4i+1$. Its columns span $x_F y\widehat{\mathcal{O}}_E^n$, where x_F is the finite frame and y the stage point. These bases exist at every finite place by triangular elimination over its valuation ring along this ordered flag, and the choices are integral frames at almost every place. Denote the string diagonal by β_i. For $i \le i^\vee$ the column b_i equals $\beta_i e_i$ modulo earlier string vectors, with no core entry. For i, j in this first half,

$$(b_i, b_j) = \overline{\beta_i}\,\beta_j(e_i, e_j).$$

Indeed $i + j \le s - 1$, and an earlier vector cannot pair with the other index or an earlier one, by the anti-diagonal string Gram matrix. Its only possibly nonzero entry on these columns is the norm at the middle vector in an odd string.

Suppose y, y' give equal valuations of β_i, β_i' at every finite place for $i \le i^\vee$. Multiplying the corresponding columns of b' by the units β_i/β_i' makes these diagonals equal. In the two integral orthonormal models the frames

$$B_i = (x_F y)^{-1} b_i, \qquad B_i' = (x_F y')^{-1} b_i' \quad (i \le i^\vee)$$

are primitive, being part of invertible integral matrices. They have equal Gram matrices by the displayed identity. We can therefore choose an integral unitary k carrying B_i to B_i' simultaneously, including when their Gram matrix is degenerate.

For completeness, at each finite place of F the quadratic algebra is split or unramified. Here E contains i and $\omega = (1 + \sqrt{-7})/2$, so $i\omega - \overline{i\omega} = i$ is a unit. The order generated locally by $i\omega$ has unit discriminant. Modulo the uniformizer the isometry between the primitive frames extends by Witt's theorem for Hermitian spaces. In the split algebra, the first frame specifies independent vectors and the conjugate frame independent covectors. Their pairings fix the induced map on the quotient by their common covector kernel. Extend the given map first on the intersection of the vector span with that kernel, then to the whole kernel and finally on complementary vectors using the prescribed quotient map. This gives the simultaneous change of basis for the two halves.

To lift an approximate integral isometry g matching the frames modulo ϖ^a, $a \ge 1$, write its correction as $g(1 + \varpi^a T)$. Over the residue algebra the two equations are

$$T + T^* = H, \qquad TB = D, \qquad H = (1 - g^* g)/\varpi^a, \quad D = g^{-1}(B' - gB)/\varpi^a.$$

The adjoint here is in the integral orthonormal frame. Equality of the original Gram matrices gives $B^*HB = B^*D + D^*B$ after reduction. Extend the prescribed map $B \mapsto D$ to any endomorphism T_0. The Hermitian error $H - T_0 - T_0^*$ vanishes on the product of the span of B with itself. Choose a vector space complement. This error is $T_1 + T_1^*$ with $T_1 B = 0$: in sesquilinear form coordinates assign the mixed block on the side with second argument in the complement, and split the complementary Hermitian block using an element of residue trace one. The same construction works over the split residue algebra and in characteristic two. Use $T = T_0 + T_1$ and pass to the inverse limit. Taking the product of these local solutions gives $k \in K_j^0$. This is the lifting calculation for primitive Hermitian frames, with the additional columns prescribed at each depth.

The rational-coordinate matrix $x_F y' k (x_F y)^{-1}$ now carries b_i to b_i'. Its action on e_i is the identity modulo the span of the earlier e_t, by induction up the half-string. These are the defining conditions of B_j; they include the middle vector. Thus $y' k y^{-1}$ belongs to $B_j(\mathbb{A}_{F,f})$, proving the orbit assertion used for the finite cut. Its separate support input is the bound on the valuations for a nonzero unfolded ramified column.

6.2 The definite coefficients and the paired edges

On the space of polynomials on $b \times a$ complex matrices, with $b \ge a$, fix the determinant twists in the oscillator action. Let V, V' be the matched irreducible source and target Schur modules for one partition

of length at most a, with their twisted actions ρ, ρ'. Suppose there is a nonzero bilinear family $B_{\ell,v}$ of coefficient polynomials with covariance

$$\Omega(k,h)B_{\ell,v} = B_{\ell\circ\rho'(h^{-1}),\rho(k)v}, \qquad v \in V, \quad \ell \in V'^{*}.$$

If the characters of the two determinants are fixed and the source is compact, then for any polynomial P and $t_0 \in V^*$ there is one $\ell \in V'^*$ such that

$$\int_{K_s} t_0(\rho(k^{-1})v)\,\Omega(k,1)P\,dk = B_{\ell,v},$$

where dk has mass one. Indeed translation in k makes the left side an intertwiner in v. The Cauchy decomposition gives one paired Schur space for each source shape. To see its multiplicity, in homogeneous degree d realize the rectangular polynomials as diagonal symmetric-group invariants in the tensor product of the d-fold tensor powers for the two coordinate spaces, with a dual as dictated by the action. The Schur–Weyl decompositions pair identical symmetric-group modules with multiplicity one by Schur orthogonality. Fixed determinant characters shift the weights and leave the source shape determined by the source highest weight. Thus the nonzero covariant family is the full V-intertwiner space. This argument tensors over the real places, and finite linear functionals in the polynomial coefficients may be applied after the average.

A compact edge in the weighted normalization. On the edge $X_j \to X_{j+1}$ let V, V' be its coefficient spaces, with real actions ρ, ρ' in the conjugate Weyl convention. For the two ranks $n < n'$ take the prescribed diagram, retaining the columns of heights $< n'$. There are no columns of heights strictly between n and n', since the heights above one have the form $k + i(t+1)$ and the consecutive ranks differ by $t+1$. At each embedding the incoming determinant exponent is the incoming real splitting weight $\lfloor n'/2 \rfloor - w_{j+1}$ plus the multiplicity $\kappa(n)$ of source-height columns. This is zero at $j = 0$ by the formula for $\kappa(k)$. At $j > 0$ it equals $\lfloor n_{j-1}/2 \rfloor - w_j$, by the constant weight increment h, the column count $h - (t+1)$ and $\lfloor n'/2 \rfloor - \lfloor n_{j-1}/2 \rfloor = t+1$. The outgoing exponent is $\lfloor n/2 \rfloor - w_{j+1}$. The twist has real weight zero, so both real actions are fixed for the entire family. These identities identify the action on the rectangular Schur coefficient polynomials $P_{\ell,v}(Z) = \ell(S_\kappa(Z)v)$:

$$\Omega(x,y)P_{\ell,v} = P_{\ell\circ\rho'(y^{-1}),\,\rho(x)v}.$$

Indeed the Fock substitution acts on the two sides of Z, and the maximal minors of height n contribute the $\kappa(n)$ extra powers of the source determinant. The remaining determinant powers are those of the real splitting characters with the given rank twists.

Let K_j^X be the actual source column level, and let $\mathcal{J}_\vartheta$ be the Schur jet matrix at scale $c_{\mathrm{jet}} = \overline{\pi_G}\,\mathrm{per} \neq 0$.

Write Θ for the theta functional and $\mathsf{v}(P) \otimes f$ for the insertion of the real Fock polynomial P in a finite datum f. In the matrix coordinates at all real places write $\alpha! = \prod_i \alpha_i!$. The generating series at scale c and its differential pairing are

$$T_c(f)(Z) = \sum_\alpha \frac{c^{|\alpha|}}{\alpha!}\,\Theta(\mathsf{v}(Z^\alpha) \otimes f)\,Z^\alpha, \qquad P(\partial)T_c(f)(0) = \Theta(\mathsf{v}(P(c\,\cdot)) \otimes f).$$

Indeed, for $P = \sum_\alpha P_\alpha Z^\alpha$, the derivative multiplies the α-coefficient by $P_\alpha \alpha!$. The pairing is complex bilinear. For the finite part of the positive definite model this series gives the ordinary theta jets as well. In the d_0 complex Fock coordinates put $d\mu(w) = \pi^{-d_0}\,dw$, with dw Lebesgue measure, and write $\mathsf{H}(w)$ for real Heisenberg translation in these coordinates. Set $G_f(w) = e^{-\|w\|^2/2}\Theta(\mathsf{H}(w)f)$. The integrated action gives $\Theta(\mathsf{v}(P) \otimes f) = \int P(\bar{w})G_f(w)\,d\mu(w)$. The real vacuum projection fixes f. The translated vacuum at w is $e^{-\|w\|^2/2}e^{-\|y\|^2/2}\exp(\sum_i y_i\overline{w_i})$. Dividing by the untranslated vacuum value $e^{-\|y\|^2/2}$ at y gives the entire theta function as $\int \exp(\sum_i y_i\overline{w_i})G_f(w)\,d\mu(w)$. On the line $y = tcz$, its m-th derivative at zero inserts

$(c\sum_i z_i\overline{w_i})^m$. The Heisenberg translates have bounded theta value, so Gaussian decay dominates these derivatives uniformly on compact t-disks. The multinomial formula gives $m![t^m]T_c(f)(tz)$.

The entries of $\mathcal{J}_\vartheta(f)$ are the jets of $T_{c_{\mathrm{jet}}}(f)$ paired with the entries of $S_\kappa(Z)$. Each entry is homogeneous of degree $D_j = \sum_\tau \sum_{h<n'} h\kappa_\tau(h)$, including all retained columns of source height n. The finite action changes f to $f' = \omega_f(g_f, h_f)f$: it commutes with the real Heisenberg operators used for the insertion. The real action on the polynomial is the covariance calculated above. Thus, for every $\ell \in (V')^*$ and $v \in V$,

$$c_{\mathrm{jet}}^{D_j}\,\Theta\big(\omega((x,g_f),(y,h_f))(\mathsf{v}(P_{\ell,v})\otimes f)\big) = \ell\big(\rho'(y^{-1})\,\mathcal{J}_\vartheta(f')\,\rho(x)v\big)\,.$$

This identity holds at each finite translate and for every retained degree, including degree zero.

The compact average applies to $B_{\ell,v} = P_{\ell,v}$, using the covariance just computed and the nonzero Schur minor on the first n rows. It gives one $\ell(P, t_0)$ for an arbitrary real Fock polynomial P and $t_0 \in V^*$. For $P = P_{\ell_0,v_0}$, the source average conjugates the rank-one endomorphism $v_0 t_0$ of V. It commutes with the source action and has trace $t_0(v_0)$, so $\ell(P_{\ell_0,v_0}, t_0) = t_0(v_0)\ell_0/\dim V$.

Combine the coefficient identity with the compact average. For all finite data f, smooth algebraic columns ϕ, twists and target points $h = (y, h_f)$, integrating the theta kernel with arbitrary Fock datum $P \otimes f$ against the coefficient $t_0(\rho(x^{-1})\phi(g_f))$ gives

$$\frac{c_{\mathrm{jet}}^{-D_j}}{[K^0_{X_j} : K^X_j]}\,\ell(P,t_0)\big(\rho'(y^{-1})(\mathscr{L}_{j,\vartheta}(f)\phi)(h_f)\big).$$

Here $\mathscr{L}_{j,\vartheta}$ is the prescribed algebraic class sum, normalized to give K^X_j volume one. In fact the finite integrand after the compact average is exactly $\mathcal{J}_\vartheta(\omega_f(g_f,h_f)f)\phi(g_f)$. It is left rational invariant and has an open stabilizer. At any finer level $K' \subset K^X_j$ fixing it, a class has mass $\mathrm{vol}(K')$ divided by the order of its rational stabilizer. The analytic measure gives $K^0_{X_j}$ volume one, so its class masses are $[K^0_{X_j} : K^X_j]^{-1}$ times the algebraic masses. The same ratio therefore holds at every refinement and is independent of the twist and finite data. The formula gives a single output functional uniform in all of them. This compact argument applies to both edges of the last pair as well as the edges of all preceding pairs.

A middle pair of graded arrows. Use Y, Y', Y'' for a pair of lengths $2r+1, 2r, 2r-1$ and respective full ranks d, d', d''. The lower ranks increase by t at each step, so $d'' - d = 2t$ and the outer ranks have the same parity. On $\mathrm{Hom}(Y', Y^- \oplus Y'')$ the first and second data combine by adjoint and flip in the first summand. The oscillator functional on the combined data is a fixed scalar times the product of the two theta functionals. In a product Schrödinger model the rational sums multiply, and the scaled adjoint identifies the first rational lattice sum with that in the flipped form. The intertwining scalar depends on the two chosen models. The transported Heisenberg frequency is the stacked map $S = (T^\dagger, T')$, with

$$S^\dagger S = -TT^\dagger + T'^\dagger T' = 0.$$

Each term on the right is the shift on Y' by the common constant c. By averaging the positive-degree maps, the product of the two constant terms thus becomes the zero term of the combined datum at S, with the reciprocal of the fixed theta scalar. The positive-degree variables agree under the adjoint because it preserves map degree when reversing source and target, as their Gram dualities negate the grades.

The zero-degree maps are from X' to $X^- \oplus X'' \oplus H$, where H is the plane of the two outer middle vectors z, z''. Indeed the outer nonzero grades are even and the middle string grades are odd. In $Y^- \oplus Y''$ the vector $z + z''$ is isotropic, and $(z - z'')/(2N)$ is its dual isotropic vector. To integrate over the source Bessel quotient, the source unipotent has only degrees at least two. Its character cancels the phase of the first constant term at T. Its fiber average therefore leaves only the source core variable. If u' is in the middle unipotent, the two frequency maps move to $u'^{-1}T$, $T'u'$, so the new stacked frequency is Su'.

The difference $S(u'-1)$ has degrees at least zero and its degree-zero part maps X' into the isotropic line $E(z+z'')$. More precisely its X' row along this line is

$$c\,(u'-1)_{\text{string grade 1, }X'}.$$

Here the grade-one vector maps under S to $c(z+z'')$, as follows by taking the adjoint across the middle entries of the string Grams. All other entries of $S(u'-1)$ have positive degree, because S has degree -1 and degree-zero maps have just the source and target indicated. The Heisenberg phase from composing with the shift by S is one by $S^{\dagger}S=0$. The positive-degree part disappears in the first zero term, and it pairs trivially with the retained degree-zero part.

The displayed grade-one row is additive on the unipotent: quadratic terms raise grades by at least two. It is onto, both rationally and adelically, by the Eichler matrix with the grade-one vector and any core vector. Thus the normalized unipotent quotient measure pushes to probability Haar on the additive rational quotient of this row. This is exactly the second mixed zero term, along $\mathrm{Hom}(X', E(z+z''))$. Both Bessel quotients are compact. Their fiber masses over the definite core quotients are fixed constants, so this pair of averages introduces a single scalar fixed by the graded spaces and measures.

For the action of the core triple let u, v, w now denote the three determinants. With v the fixed splitting character, the underlying Weil action before the zero terms has extra normalized v-exponents

$$d'-\lfloor d'/2\rfloor, \qquad \lfloor d/2\rfloor-\lfloor d''/2\rfloor-d, \qquad \lfloor d'/2\rfloor$$

at u, v, w, respectively. The $-d$ and d' come from the adjoint flip, and the two base characters are evaluated at the same determinants before and after restriction. On the positive-degree maps the determinant is $(uw)^r v^{-(2r-1)}$: each of the r negative middle grades admits both outer core blocks, and each of the $2r-1$ positive outer grades admits the contragredient middle core block. All remaining string blocks are fixed. On the removed plane the isotropic block has determinant v^{-1}. Use the normalized parabolic law (6.4) for both blocks. Their determinants are norm-one ideles, so the modulus is one. The exponents change by $(-r, +2r, -r)$. These are exactly the changes obtained by using the smaller ranks $(d-2r-1, d'-2r, d''-2r+1)$, by the equal parity of d, d''. This proves equivariance with the normalized triple Weil action of the three cores. The shift by S commutes with their simultaneous core action, since the string maps kill the cores and their images lie on the strings.

At any fixed tuple of real outer variables the moved data split into a real vector and the translated finite tensor. The two zero terms split in the same way: in adapted mixed models they evaluate the removed Fourier variables at zero, separately at the real and finite places. Their normalized finite map is precisely $E_{i,\vartheta}(y)$, including the finite $\alpha_{\mathrm{W}}^{-1}S$ shift. The fixed choice of polarization identifies the retained oscillator with the two definite oscillator spaces, and acts by the same finite intertwiner for all input data. For decomposable polynomial data the two compact integrals on the cores give the two successive weighted class sums by the compact-edge calculation above. A polynomial in the two sets of variables is a finite sum of such products. For arbitrary real vectors the same identity follows by Hermite approximation: the value of each test integral is continuous in the real Schwartz vector, and the family of evaluations furnished by the finite-dimensional last coefficient space is a closed subspace of the product of the scalar test spaces. Consequently one coefficient functional works for all finite tensors, twists, columns and evaluation points with those fixed real variables. The same argument for the final pair uses z/π instead of the plane. It has no positive-grade or plane term, and changing ranks to $(d-1, d', d''+1)$ preserves the displayed exponents by equal parity.

To state the induction, let $\Phi_{i,\vartheta,y}$ be the column after the first $i+1$ pairs with the finite tensors $E_{0,\vartheta}(y), \dots, E_{i,\vartheta}(y)$. Fix the real outer variables and a real point x at stage $2i+2$. For some functional $\ell_{\xi,x}$, all twists, columns, finite tuples with $y_0 \in J_g$, and core points k satisfy

$$P_{2i+2}\big(J_{2i+2}(k)(x, y_{2i+2})\big) = \ell_{\xi,x}\big(\rho_i((\mathrm{emb}_i\, k)_\infty^{-1})\Phi_{i,\vartheta,y}((\mathrm{emb}_i\, k)_f)\big).$$

Here ρ_i is the outgoing conjugate Weyl action. The bottom Bessel evaluation is the initial functional. The two compact integrals take its current value on the embedded core at the given outer point and produce the functional for the next pair. The finite component of that same outer point is used on both adjacent edge data. The level-independent scalar of the pair is absorbed in the functional. This proves the claimed coefficient identity on the whole family by induction. The zero-term calculation also gives equivariance of each finite evaluation under independent finite core moves at its three vertices. In the terminal pair the third group is $U(Y_m)$, acting on X_m by the block embedding fixing the extra line.

One valuation step on the ramified columns. Use additive valuations ord_w. For the finite Weil idele α_{W}, let $a = \alpha_{\mathrm{W}}/\bar{\gamma}$, where γ generates the chosen finite modulus for g, i.e. it has the valuations of $\overline{\pi_G}$. The modulus generator b for an edge with twist has $\mathrm{ord}_w\, b = \mathrm{ord}_w\, \gamma + \mathbf{1}_{w\in Q}$, where $Q = Q_\vartheta$ is the chosen half of the twist places. The weighted frames have $\mathrm{ord}_w\, c = \mathrm{ord}_w(a\bar{a})$. This uses the self-dual finite normalization of α_{W} for the different and $c = -(j_{\mathrm{Weil}} g\omega)^{-1}$. With $n = \lfloor (s+2)/2 \rfloor$ at string length s, the standard profile β_i^0 is that of

$$(a/\bar{a})^n \begin{cases} 1 & s \text{ odd}, \\ a^{-1} & s \text{ even},\ 2i < s, \\ \bar{a} & s \text{ even},\ 2i \geq s. \end{cases}$$

Writing $i^\vee = s - 1 - i$, these entries satisfy $\mathrm{ord}_w\, \beta_i^0 + \mathrm{ord}_{\bar{w}}\, \beta_{i^\vee}^0 = 0$. In particular the odd gluing column $e/\bar{\pi} + u_j z/N$ changes the finite frame but leaves the adapted middle line $u_j z$ in this profile.

The finite inclusion datum is fixed by the Heisenberg shifts of $b\widehat{\mathcal{O}}_E$-vectors. Indeed its finite-cell expansion has shifts in $\bar{b}^{-1}\widehat{\mathcal{O}}_E$; the translates by $b\widehat{\mathcal{O}}_E$ commute with these shifts by self-duality and fix the vacuum. Suppose an edge constant term at T is nonzero at finite points with rational-coordinate lattice bases R, R'. The mixed zero-term support argument gives, for an integral matrix L,

$$(\bar{b}/\alpha_{\mathrm{W}})T - R'LR^{-1} \quad \text{of degree at least zero.}$$

To see this, set g to be the finite unitary action on the Hom frame, so the moved datum is fixed by the lattice $b\,g\widehat{\mathcal{O}}_E$. Pair its intersection with the positive isotropic part against $\alpha_{\mathrm{W}}^{-1}T$ in finite coordinates. A nonzero zero term requires trivial commutator on that intersection. The intersection is a primitive submodule of the full lattice, so each integral functional there extends to the full lattice. By self-duality the frequency agrees on these pairings with an element of $\bar{b}^{-1}g\widehat{\mathcal{O}}_E$. The annihilator is exactly the maps of nonnegative degree, giving the formula after conversion to rational coordinates.

Take adapted bases B, B' with string diagonals β, β'. In the source use the dual basis in reverse string order, $\widetilde{B} = J^{-1}B^{-*}P_s$, with J the source Gram and P_s the permutation reversing its string and fixing the core. This is again a lattice basis, since $R^*JR = 1$. Both $\widetilde{B}$ and B'^{-1} preserve the grade filtrations. On a string vector of grade λ_i the degree-λ_i string entry of $\widetilde{B}$ is $(J_{i,i^\vee}\overline{\beta_{i^\vee}})^{-1}$. For the middle vector the degree-zero block may also have a core component, killed by T. In the target odd middle, the adapted order puts the middle vector before the core, so B'^{-1} sends this row component to itself times $(\beta_i')^{-1}$ modulo higher grades. Thus the degree -1 entry at matched indices in $B'^{-1}(\bar{b}/\alpha_{\mathrm{W}})T\widetilde{B}$ is

$$\frac{\bar{b}}{\alpha_{\mathrm{W}}}\,\frac{t_i}{J_{i,i^\vee}\overline{\beta_{i^\vee}}\beta_i'}.$$

It is integral by the support formula, since $B'^{-1}R'LR^{-1}\widetilde{B}$ is integral. The occurrence of $J_{i,i^\vee} = -N$ at the odd source middle cancels the factor $-N$ in t_i. Writing $A_w = \mathrm{ord}_w\, a$ and $C_w = \mathrm{ord}_w\, c = A_w + A_{\bar{w}}$, the coefficient $t_i/J_{i,i^\vee}$ has valuation 0 on positive odd source grades and on negative even source grades, and C_w on the remaining grades. Inserting the standard profile shows that its change from source to target is exactly that valuation minus A_w: n stays the same for odd to even and decreases by one for even to odd. Accordingly the deviations $e_i(w) = \mathrm{ord}_w(\beta_i/\beta_i^0)$ satisfy

$$e_i'(w) \le -e_{i^\vee}(\bar{w}) + \mathbf{1}_{\bar{w}\in Q}.$$

The integral Gram of the adapted target frame gives $e_i'(w) + e_{i^\vee}'(\bar{w}) \ge 0$ whenever the target Gram pairing is a unit. At a target odd middle and $w \mid N$ it yields the lower bound -1 instead. To exclude equality -1, name the place in the conjugate pair dividing π as w, and put $i = (s' - 1)/2$ for this target length s'. At w replace the target lattice by a diagonal lattice with core part standard, string diagonals $\beta_t'(w)$ for $t \le i$, and $1/\overline{\beta_{t^\vee}'}(w)$ for $t > i$. Its component at $\bar{w}$ is its Gram dual in the split factors. This defines a finite group point, since at a split place the factor at w may be changed freely and determines the factor at $\bar{w}$. The initial half valuations are unchanged at w and $\bar{w}$: for $t < i$ use the opposite diagonal entries, and at i use $\operatorname{ord}_w \beta_i' + \operatorname{ord}_{\bar{w}} \beta_i' = -1$. Outside $w, \bar{w}$ the lattice point is unchanged. The integral half-frame orbit calculation therefore permits a finite Bessel move of the target point x' to this split lattice, up to its integral lattice stabilizer. Some such move $x'' = d_B x'$ then satisfies

$$(x'')^{-1} J_j(\nu_w I) x'' \in J_j,$$

where j is the target stage and ν_w is the norm-one scalar of the ramification operator. It is integral, supported on $w, \bar{w}$, and preserves the split lattice. The level at these places is maximal because N is prime to the chosen and twist moduli.

In this odd-string target j is even. The nonzero stage is therefore read by the preceding paired definite chain. In reading it at x'', write the inverse Bessel move as a unipotent times a core move. The unipotent gives only the prescribed generic phase, since all its grades are even. Fix the real variables and use the even-stage coefficient formula at x'' and at the associated core point. Simultaneous translation by $\nu_w I$ on all definite vertices fixes each of the edge data, because it acts trivially on the Hom space and the two determinant characters are unramified on these units. Equivariance of the class sum therefore transfers the scalar translate of the input column to the same scalar translate of the final coefficient. By the finite triple equivariance, translating at the final vertex alone is the action induced by replacing x'' by $J_j(\nu_w I)x''$. The displayed level inclusion makes the finite tensor invariant under this change, since both edge data descend to the prescribed finite level cosets. The paired chain hence has the same final value on ψ and $R(\nu_w I)\psi$. Every ramified column is the difference of two such columns at this w, so the original stage value vanishes. This contradiction rules out the exceptional Gram bound -1.

Start with the standard profile at stage zero, where finite support lies in $B_0(\mathbb{A}_{F,f})J_g$. At each nonzero stage value the defining integral has a nonzero integrand at some preceding Bessel point, giving both a nonzero stage value and a nonzero edge constant term. Induction with the upper inequality therefore bounds every $e_i'(w)$ from above by j on $Q \cup \bar{Q}$ and zero elsewhere. Applying the Gram inequality and its ramified middle correction to the opposite index gives the corresponding lower bound. This proves the support condition for all finite places and outer variables of the cuts.

Cutting and integrating the outer variables. Take $A_0 = J_g$ and $A_j = K_j^0$ for $j > 0$, and set $S_j = B_j(\mathbb{A}_{F,f})$, $U_j = S_j \cap A_j$. The first support calculation gives a compact set $D_{j,\vartheta}$ whose left S_j-saturation covers all the finite support in stage $j > 0$. Indeed the valuation differences take only finitely many values. Choose a representative for each occurring profile, keeping its components at the twist places and their conjugates and setting the other components to one. The profile is unchanged by this replacement outside the twist places, by equality with the standard profile there. Take the union of the sets yK_j^0 for these representatives. For $j = 0$ use $D_{0,\vartheta} = J_g$.

For each S_j-orbit on $G_j(\mathbb{A}_{F,f})/J_j$ meeting $D_{j,\vartheta}/J_j$, choose one of its points there and form its U_j-orbit. Denote the union of these finite orbits by $\mathcal{F}_j$. The image of $D_{j,\vartheta}$ in the discrete quotient is finite by compactness. Each U_j-orbit is also finite. Writing $s_t = S_j \cap yJ_jy^{-1}$ for the stabilizer of $t = yJ_j$, the number of cosets of U_j whose translates of $\mathcal{F}_j$ contain t equals

$$m_j(t) = [s_t \,:\, s_t \cap U_j]. \tag{6.5}$$

In fact, in the orbit of $t \in \mathcal{F}_j$, the part of the cut is $U_j t$. The elements s with $s^{-1}t \in U_j t$ are precisely $s_t U_j$. This also shows that the multiplicity is constant under left U_j-translation.

The representatives y of the cut are integral off the twist primes. At a twist prime over ℓ, the component of J_j lies in the principal congruence subgroup, a pro-ℓ group. For any fixed y, a common q_ϑ-power of all elements in yJ_jy^{-1} belongs to K_j^0. To see this at a twist prime, take the normal core of the open intersection with the lattice stabilizer inside the conjugated pro-ℓ subgroup. Its quotient is a finite ℓ-group. Away from these primes the conjugation is by an integral matrix, so it preserves the lattice stabilizer. The same exponent works in the restrictions to s_t and to all conjugates there of $s_t \cap U_j$. Cauchy's theorem for the finite quotient by the normal core then shows that $m_j(t) \mid q_\vartheta^a$ for some $a \geq 0$. At stage zero $s_t \subset U_0$, since $y \in J_g$ and $J_0 \subset J_g$, so $m_0(t) = 1$. Also $\mathcal{F}_0 = J_g/J_0$, since two points of J_g/J_0 in the same S_0-orbit are in the same U_0-orbit. For the trivial twist all the multiplicities are one, and every $\mathcal{F}_j$ has at most

$$C_g = g^{f n_m^2}, \qquad n_m = \dim_E X_m, \tag{6.6}$$

points, by reduction at the chosen places from the integral level in rank at most n_m.

Normalize Haar measure on S_j by $\mathrm{vol}(U_j) = 1$, and take the complementary real Haar factor to give the fixed measure on $[B_j]$. The Bessel subgroup here is a unitary group acting on a unipotent group. The determinant of its action on the unipotent Lie algebra is an F-rational character of the unitary group, hence trivial. The Bessel subgroup and its finite and real factors are unimodular. For a function f_j left invariant under $B_j(F)$, integrate first over $[B_j]$ and then over $B_j(\mathbb{A}_F)\backslash G_j(\mathbb{A}_F)$. The latter quotient has separate real and finite factors. On the finite orbit $S_j y J_j$ its mass is $\mathrm{vol}(J_j)/\mathrm{vol}_{S_j}(s_t)$. Indeed the left S_j-average of the indicator of yJ_j is $\mathrm{vol}_{S_j}(s_t)$ on that orbit. Since the cut has $[U_j : U_j \cap s_t]$ points in this orbit, the mass is exactly the sum of $\mathrm{vol}(J_j)/m_j(t)$ over those points. It follows that when the inner period at (ξ_j, t) is $[K_j^0 : J_j]a(\xi_j)b(t)$, the outer integral is

$$\left(\int_{B_j(F_\infty)\backslash G_j(F_\infty)} a(\xi_j)\, d\xi_j \right) \sum_{t \in \mathcal{F}_j} \frac{b(t)}{m_j(t)}. \tag{6.7}$$

Here $F_\infty = F \otimes_{\mathbf{Q}} \mathbf{R}$ and the real measure is independent of the twist. The equality uses $[K_j^0 : J_j]\,\mathrm{vol}(J_j) = 1$. When a finite factor $b(t)$ is nonzero, the inner period and absolute integrability ensure that a is left invariant and integrable on the real quotient. Assign the real integration functional the value zero on functions outside this domain. This convention only affects terms with zero finite factor and gives a functional defined using the fixed real measure for all other outer variables.

The real coefficient of the last double lift. Assume the incoming coefficient on $U(X_{m-2})$ has been obtained. The two compact integrals for the last arrows now take place on $[U(X_{m-2})]$ and $[U(X_{m-1})]$. The only remaining string vector has norm $+N$ in the flipped source. Sending the extra unit vector of X_m to z/π is a rational isometry, since $\pi\bar{\pi} = N$. The retained group at the last evaluation is $U(Y_m)$ embedded by this isometry and the identity on Y_m. There is no positive-grade constant term or translation by a string map on these two arrows.

Fix the real outer variables ξ and hence the real Schwartz vector of this pair before varying its finite tensor, twist, or column. First take a real Fock polynomial which is a product of a polynomial on each edge, and a decomposable finite pair. In Schrödinger coordinates on the orthogonal sum, the Schwartz function is the product, and summing over rational vectors is the product of its two rational sums. The first variable is changed by adjoint and flip. Rational vectors correspond bijectively under that change. The intertwiner between the two choices of polarization contributes one scalar fixed by the choices. The determinant characters agree after reassigning the extra line. Indeed let d, d', d'' be the three original ranks. The additional splitting-character exponent at the middle determinant, including the adjoint flip, is $\lfloor d/2 \rfloor - \lfloor d''/2 \rfloor - d$. The new ranks are $d-1$, d', $d''+1$, and $d''-d = 2t$ from $d = n_{m-2}+1$, $d'' = n_m - 1$. The equal parity of d, d'' leaves this exponent unchanged. The other two splitting exponents are $d' - \lfloor d'/2 \rfloor$

and $\lfloor d'/2 \rfloor$. Block embedding preserves both determinants, and all factors of the two base characters use the same three determinants. This identifies the action used in the thick pair with the product of the two normalized edge actions. The two resulting compact integrals are successive definite lifts. Apply the compact edge calculation to their actual kernels. Evaluating the incoming functional on the first integral gives a functional on the intermediate type. Applying the second gives one on the last type, uniform over finite tensors, columns, twists and output arguments.

Finite sums of such products contain every polynomial in the two sets of variables. Polynomial vectors, with their Gaussian, are dense in the real oscillator Schwartz space by the Hermite expansion. For fixed finite tensor, twist, column and output point, both compact integrals are continuous in the real Schwartz vector. Indeed absolute theta summation on a compact set of representatives is bounded by a fixed Schwartz seminorm. Take the product of these spaces of values over *all* the tests. The subspace obtained from linear forms on the last finite-dimensional coefficient space is finite-dimensional and therefore closed in this Hausdorff product. Passage to the limit gives one linear form for the whole family. Thus the preceding compact calculation supplies the l_ξ on the last pair even for the real Schwartz vectors arising in the outer integration. Compact branching then gives the terminal Petersson pairing below with its stated factor.

The terminal pairing and the common scalar. Let $\mathscr{L}_\vartheta(v)\phi$ denote the definite chain on the tensor v. Denote by R_m the block restriction followed by the fixed branching map of the weighted datum and the factor $\prod_{\tau\in\Phi} \tau(\bar{\pi})^{-b_\tau}$, where b_τ is the prescribed branch exponent, then converted to an automorphic form by the inverse of the terminal algebraic realization. To pair the stage-m matrix coefficient against a conjugate form, integrate over the compact real group $U(Y_m)(F_\infty)$. If v_0 is the fixed highest vector of the test type, the average of $v_0 \otimes l_\xi$ under the two coefficient representations is an intertwiner from the restricted coefficient representation on X_m to the one on Y_m. The standard interlacing branching rule for $U(n_m) \downarrow U(n_m - 1)$ has multiplicity one, so the intertwiner is a scalar multiple of the fixed branching map. The Petersson pairing of the corresponding forms gives the same finite class integral, with a fixed constant by Schur orthogonality. The inverse of the displayed branching factor and this Petersson constant depend only on D. Thus for every final test form z the terminal period is

$$a_m(\xi)\,\langle z, R_m \mathscr{L}_\vartheta(E_\vartheta(y))\phi\rangle,$$

for a scalar function a_m independent of z, ϕ, ϑ, y. Here and below the Petersson pairing is linear in the second argument. The actual terminal type, branching map and scaling factor are those in the weighted column data.

Apply (6.1) and (6.7) successively, starting at $j = m-1$ and ending at $j = 0$. At each step transpose the last arrow in the pairing, holding the previous variables fixed. Its local period has the form in (6.7), with the real factor given by the preceding real integration and the finite factor given by the preceding finite sum. The finite support required for each such integration was uniform in the previous variables. At the bottom use the integral of the Poincaré series against the transposed test. The result, on testing against z, is the definite chain on the finite tensor

$$\mathscr{Z}_\vartheta = \sum_{t_j \in \mathcal{F}_j \ (0\le j<m)} \left(\prod_{j<m} m_j(t_j)^{-1} \right) E_\vartheta(t_0, \dots, t_{m-1}, 1) \tag{6.8}$$

followed by R_m, times one scalar a_* from the real integrals. It is independent of z, ϕ, ϑ. All summands are defined on the cosets. The chosen inclusion data are fixed by the two principal levels of each edge, and so changing the finite representatives leaves each translated datum unchanged. Taking all final tests identifies the forms, with the same conversion to algebraic forms on both sides.

For level elements b_j of the definite groups, the inclusions of the grade-zero cores carry b_0 into J_g and b_j into K_j^0 for $0 < j < m$, by the lattice calculation in Section 3.5. The stage-zero inclusion has the

same image in the transported filling frame since it commutes with η. For $b_m = 1$, the paired intertwining property makes simultaneous left translation of the t_j by these inclusions act on the kernel tensor by the joint action on its edges. The cuts and their multiplicities are invariant under these translations. Therefore $\mathscr{Z}_\vartheta$ has the required joint invariance.

Admissible finite tensors. Use the normalized integral vacuum on each definite edge. For y_j integral off q_ϑ, the calculation below writes $E_\vartheta(y)$ as a scalar C_*, depending only on D, times a combination of tensors V_t^{pair} of paired vacuum translates. The tuples t are admissible for q_ϑ before adding N to the modulus. The coefficients become integral at every coefficient place away from q_ϑ after multiplication by $d_* = C_g^{2m}$, and at the trivial twist the sum of their absolute values under each conjugate is at most one.

Finite zero terms on the two graded frames. Use the normalized finite oscillator operations for the given mixed polarization. The mixed zero-term covariance sends translations on the isotropic and retained frame to their retained component. The translate of a finite vacuum under a unitary move g is a sum over $gL/(L \cap gL)$ with the common multiplier $\Delta_f(g)/|\Delta_f(g)|^2$. Here $|\Delta_f(g)|^2 = [L : L \cap g^{-1}L]$. On the split side use the Schrödinger integral-box indicator and the adjoint flip.

Write L for the corresponding self-dual lattice, $\rho(t)$ for its vacuum-normalized Heisenberg translations, and $\mathbf{v}$ for the plain vacuum. We first compute the zero term Z for an isotropic frame k_0 with retained isometric embedding i_0. Suppose $B = i_0 g + k_0 C$ is integral, where g is unitary on the retained space. Then, for one scalar C_0 independent of t,

$$Z(\rho(t)\mathbf{v}) = \zeta_t C_0\, \omega(g)\rho(B^* t)\mathbf{v}_0, \qquad \zeta_t = 0 \text{ or a root of unity.}$$

Indeed translation by Bv, for v integral, fixes the vacuum. Under Z the isotropic component disappears and the retained component translates by gv. The jointly fixed line for gL_0 is $\mathbf{C}\omega(g)\mathbf{v}_0$, by conjugating the standard vacuum projector. If a translated zero term is nonzero, its commutators with translations by integral vectors of $K = \operatorname{im} k_0$ have trivial character. Hence $t \in L + K^\perp$. Locally this is $(L \cap K)^\perp_\psi = L + K^\perp$: split off the primitive submodule $L \cap K$ in L and extend its integral functional using self-duality. Write $t - l = i_0 w + k_0 b$ with $l \in L$. Then $w = gB^*(t - l)$ and $B^* l$ is integral. Commuting the translations and removing the integral parameters gives the displayed formula, since all the finite additive character values have finite order. This reasoning also applies when $C_0 = 0$.

Here are the matrices that make B integral in each actual pair. Put $\beta = \overline{\pi_G}$, $C =$ the chosen idele generating (β), $\delta = C/\bar{C}$, and $a = \alpha/\bar{C}$, for the Weil idele α. Let ϵ be the finite integral unit of norm -1 in the flip, and put

$$\lambda = \bar{\epsilon}\,\bar{\alpha}/\alpha, \qquad u_j = (a/\bar{a})^{\lfloor (m-j+1)/2 \rfloor}.$$

Thus $\lambda u_j = \bar{\epsilon}\delta u_{j+2}$ for the even source indices $j < m$, and $u_m = 1$. Let e, e'' be the last definite source and target coordinate in the pair. Put $D_\delta = 1$ for a middle pair. In the last pair it is diagonal with δ at e'' and one elsewhere. It is one above N.

Let $p, \bar{p}$ be the two idempotents of the chosen half of the places over N and its conjugate. With row vectors understood as columns when forming $\rho_0 \rho_0^*$, put

$$\rho_0 = e + \bar{\epsilon}\bar{p}\, e'', \quad A = p(\pi - 1) + \bar{p}(\bar{\pi}^{-1} - 1), \quad \gamma = 1 + A\rho_0\rho_0^*, \quad g = (D_\delta \gamma) \otimes 1.$$

Here $\rho_0^*\rho_0 = 1$ and $1 + A$ has norm one. For column $x = \gamma_\ell$ put $c = A\pi^{-1}(\bar{p}\, 1_{\ell=e} + \epsilon p\, 1_{\ell=e''})$. It is supported above N, so $\delta c = c$. The entries of x off e, e'' are integral and $\bar{\pi}x$ is integral. Multiplication on the three idempotent pieces $p, \bar{p}, 1 - p - \bar{p}$ gives

$$c - \pi^{-1}x_e = -\varpi_0\, 1_{\ell=e}, \qquad \bar{\epsilon}c - \pi^{-1}x_{e''} = -\varpi_0\, 1_{\ell=e''}, \qquad \varpi_0 = p + (1-p)\pi^{-1} \in \widehat{\mathcal{O}}_E.$$

In an odd string with middle vector z, the two pertinent lattice columns are $e/\bar{\pi} + u_j z/N$ and $u_j z$. Thus a core vector x with additional component $u_j H z$ is integral when $\bar{\pi}x_e$, its other core coordinates,

and $H - x_e/\pi$ are integral. In a middle pair the vectors z, z'' have norms $+N, -N$ in the flipped sum. The isotropic vector $z + z''$ is paired with $(z - z'')/(2N)$. Append $\lambda u_j c(z + z'')$ to the embedded column. Removing the source flip gives $H = c$ there, and on the target $H = \bar{\epsilon} c$. Both required differences are integral by the two congruences. In the last pair the extra vector embeds as z/π. Its coefficient is $\delta x_{e''}/\pi$, which after removing the flip and dividing by u_j is $-\epsilon x_{e''}/\pi$. The required difference $-\epsilon x_{e''}/\pi - x_e/\pi$ is integral, by adding the first congruence to ϵ times the second. The other target coordinates lie in its standard lattice. Tensor these columns with the dual core projection of the odd-index intermediate space. Its string has even length, so its core is an integral direct summand. This gives the desired integral B on the Hom space. For the specified stacked shift $t_0 = \alpha^{-1}(T^*, T')$ in oscillator coordinates, $B^* t_0 = 0$: both string maps in this column vanish on the intermediate core. Removing the shift therefore changes at most the phase and the zero indicator.

We compute the coefficient and support of the resulting finite expansion. At $\bar{p}$ use the Schrödinger coordinate v column by column. Put

$$n(y) = \epsilon \overline{y/N}\, e - \overline{y/N}\, e'',$$

where y/N denotes the selected residue lift supported on p. The e-entry of $\gamma v + n(y)$ on $\bar{p}$ is $\bar{\pi}^{-1} v_e + \epsilon \overline{y/N}$. For integral v_e exactly one residue class y in each column makes this entry integral. Existence of such a residue forces v_e integral, on multiplying by $\bar{\pi}$. Every other change in this column is $\rho_{0,i}$ times the change in the e-entry. Thus $\gamma v + n(y)$ belongs to the integral box for exactly one y when v belongs to the integral box, and for no y otherwise. This proves the finite indicator identity. All translations $n(y)$ are supported on one split isotropic side, on which the semicharacter and all mutual commutators are one. The first component after inverse adjoint and flip is $-y/N$, by $\epsilon\bar{\epsilon} = -1$; the second is $-\bar{y}/N$. The parameters already present are integral over N, since β is a unit there and the twist modulus q is coprime to N. Remove these integral components on the original vacuum before applying the box identity. The remaining parameters commute with every $n(y)$. Consequently the residue sum has equal coefficients in each pair; its overall scalar is independent of the residue.

Away from N the column correction c is zero and $\gamma = 1$. Give the source block of a stacked translation denominator β and the target block $\bar{\beta}$: multiplication by that denominator is required to give an integral parameter away from q. On the retained components the integral adjoint B^* preserves this condition with the new denominators $d_b(D_\delta)_{bb}$, where d_b is β or $\bar{\beta}$ according to that block. Indeed every nonzero block of the embedding retains its side and its valuation except for the terminal extra coordinate. That coordinate comes from the source and has $v(\bar{\beta}\delta) = v(\beta)$ at every finite place. More explicitly, a nonzero entry of B^* at a place w uses a conjugated embedding entry at $\bar{w}$. The respective denominators at w have equal valuations as just computed, so integrality follows termwise. Above N all denominators are units. Applying D_δ itself to the parameter then gives d_b as the denominator. Under the inverse adjoint and flip the source parameter is conjugated and multiplied by an integral unit, so both ordinary edge parameters have denominator $\bar{\beta} = \pi_G$.

Write $T(\xi) = \psi_0(c_0 \xi^* \xi/2) \rho_{\mathrm{Weyl}}(\xi)$ for the normalized translations used above. Here v_0 is the vacuum on L_0 and the underlying operators satisfy $\rho_{\mathrm{Weyl}}(\xi)\rho_{\mathrm{Weyl}}(\eta) = \psi_0(\eta^* \xi/2)\rho_{\mathrm{Weyl}}(\xi + \eta)$. The group-law discrepancy is $\psi_0(\overline{c_0}\eta^* \xi)$, using $c_0 + \overline{c_0} = 1$ and $\psi_0(z^*) = \psi_0(-z)$. It is one on L_0, since c_0 is integral and ψ_0 is trivial on the integral finite adeles. The unitary change D_δ preserves $\xi^* \xi$ and $\eta^* \xi$ exactly, including at the dyadic places. Thus the same normalized translations are a group action on $D_\delta L_0$, conjugate to the action on L_0. For $H = L_0 \cap D_\delta L_0$, every $T(\xi)$ with $\xi \in H$ fixes v_0. The probability-Haar average on $D_\delta L_0$ applied to v_0 is therefore the finite average over $D_\delta L_0/H$. Its image lies on the moved vacuum line. The Weyl commutator and self-duality of L_0 give zero vacuum pairing for $\xi \notin L_0$, so the identity coset gives a nonzero coefficient $\|v_0\|^2/h$, where $h = [D_\delta L_0 : H]$. Absorb the fixed scalar multiplying the average into C_0. The coefficients are $1/h$ up to roots of unity.

The quotient maps injectively into $\bar{\beta}^{-1}L_0/L_0$. Each diagonal factor satisfies $\bar{\beta}(D_\delta)_{bb} \in \widehat{\mathcal{O}}_E$ and also $d_b(D_\delta)_{bb} \in \widehat{\mathcal{O}}_E$. This shows that

$$h \mid |\mathcal{O}_E/(\beta)|^{(n_j+n_{j+2})n_{j+1}}$$

and that the averaged translations preserve the required denominators. The sum of the conjugated absolute coefficients of the average is at most one, including when the zero term vanishes.

Finally the finite inclusion datum on each lower edge is the characteristic function of its specified residue matrix. Finite Fourier inversion on its matrix module expresses it by the vacuum translations with coefficients roots of unity divided by the cardinality of the module. Its local pieces away from q are the ones over the chosen divisor of (β), with one matrix entry for each source and target index. Thus the part of the denominator prime to q divides $|\mathcal{O}_E/(\beta)|^{\dim Y_j \dim Y_{j+1}}$. The conjugated absolute coefficient sum is one. The moving level element is integral away from q. Its moved-vacuum formula has a finite-index denominator supported over q. Its Weil multiplier is $\chi_{\mathrm{fin}}(\det b)$ for the saturating lattice $bL = L \cap g^{-1}L$. At every $\ell \nmid q$, $\det b$ and the determinant for the additional splitting characters are units at every place over ℓ. The algebraic character values are then ℓ-adic units under every coefficient embedding: pass to a power with principal ideal, write its generator times a finite integral unit, and evaluate its infinity type on the generator, all of whose conjugates are ℓ-units. The finite norm factor uses only residue cardinalities over q. At the trivial twist the moving elements are integral everywhere. Their vacuum multipliers are one and the extra smooth character on the compact determinant group has finite image.

On each pair the two lower matrix exponents are $(\dim_E Y_j + \dim_E Y_{j+2}) \dim_E Y_{j+1}$ in sum. Each $\dim_E Y_k$ and each n_k, for $0 \le k \le m$, is at most n_m. Adding the retained exponent $(n_j + n_{j+2})n_{j+1}$ therefore gives at most $4n_m^2$. Thus the combined denominator divides C_g^4, where $C_g = |\mathcal{O}_E/(\beta)|^{n_m^2} = g^{f n_m^2}$. Tensoring $m/2$ pairs gives C_g^{2m}, uniformly in the twist and finite cut points, and conjugated coefficient sum at most one at the trivial twist. Replacing the plain vacua by the chosen nonzero integral vacua divides the common scalar by their fixed product of normalization factors. All the other scalars absorbed above are attached to the fixed pair, its frame or its polarization. This gives the single C_* needed before the outer cut sum.

By (6.5), the cut coefficients are integral away from q_ϑ. At the trivial twist the cut has total mass at most C_g^m by (6.6). Multiplying the finite expansions by the cut coefficients and summing consequently gives $\mathscr{Z}_\vartheta = C_* \sum_t r_t V_t^{\mathrm{pair}}$, with $d_* r_t$ integral away from q_ϑ and $\sum_t |\sigma(r_t)| \le C_g^m$ at the trivial twist. Choose $c_* = d_*/(a_* C_*)$ if $a_* C_* \neq 0$. Then the tensor for $\iota_\vartheta \mathcal{P}(c_* \phi)$ has coefficients $d_* r_t$ in these paired vacua. If $a_* C_* = 0$, the unfolded pairing vanishes for every column and twist, so one can take $c_* = 1$ and the zero finite datum.

In each pair the operator used in V_t^{pair} sums over residue vectors at the chosen places over N. It adds $-y/N$ on the last column of the first edge and $-\bar{y}/N$ on the last row of the second. Composition of Heisenberg translations contributes a root of unity and adds the translation parameters. Thus expanding the pairs gives a linear combination of ordinary edge vacuum tensors whose translations are admissible for Nq_ϑ, still with coefficients integral away from q_ϑ. Each pair has at most $N^{f n_m}$ summands. After collecting identical tuples the conjugated coefficient sum is at most

$$C_g^{3m} N^{f n_m m/2} \le e^{x^{60}}.$$

For the inequality, $f = O(x^{11} \log x)$, $n_m = O(x^{11/2})$, $m = O(\sqrt{x})$, $\log g = O(\log b)$, and $\log N = O(m \log b)$ by the level window. The expansion is an identity of the joint tensor, which retains its invariance.

Finally let P_N be the projector onto the vacuum at all the places above N, for each edge. The original paired translations have integral starting parameters there. On a translate of a vacuum, the projector multiplies the translate by the indicator of integrality of its parameter at those places. This follows by averaging the Heisenberg commutator character over the standard self-dual lattice. In a translated pair only the zero residue vector can survive both projections, since N is split and the chosen places contain one from each conjugate pair. Thus projecting each edge of a paired tensor and applying the same residue

sum to each pair returns that tensor. Apply this to the sum and use the linearity in the joint tensor of the definite chain. This gives exactly the required paired-chain expression at N for the ordinary coefficient function Ψ_ϑ obtained above. This proves (4.3) from the fixed-data inputs.

Projection at the chosen types. For $s_0 \in \mathcal{G}$ put

$$A_{s_0}\phi = |\mathcal{G}|\,\rho(e_0)\rho(s_0)\mathcal{P}(c_*\phi) \qquad (\phi \in C_{\mathrm{ram}}).$$

The following transfer applies to any principal-level form u on Y_0, including the auxiliary rows in the height comparison. Choose a lift of $s \in \mathcal{G}$ to integral matrices in the split charts. At every stage use the block diagonal lift with that same first block, the identity on later blocks and the identity at the twist places. The inclusion matrix $I_j : Y_j \hookrightarrow Y_{j+1}$ in the finite inclusion datum is fixed by $M \mapsto \widetilde{s}_{j+1}^{-1} M \widetilde{s}_j$. On the split lattice model this is the change of arguments in the finite Weil action. The modulus factors have absolute determinant one. The tame scalar at a selected place is the ratio of the two values of the edge residue character on the determinants. Both determinants are those of the same first block, so the ratio is one for $\mu_{j+1}^0 \vartheta^{j+1}$. The factors at the other places are unchanged. Thus the pair of lifts fixes the lower edge datum. Changing the source variable in the theta integral commutes with right translation by these simultaneous lifts. This proves $\iota_\vartheta(\rho(s)u)(h) = \iota_\vartheta(u)(h\widetilde{s}_m)$. The determinant ratios are one at every edge. The conversion to algebraic forms at the last stage commutes with finite right translation.

For each approved irreducible character χ of $\mathcal{G}$, Schur orthogonality gives its central projector. Consequently, writing $a_t = |\mathcal{G}|e_0(t)$,

$$a_t = \sum_{\chi \text{ approved}} \chi(1)\chi(t^{-1}).$$

Each character value is a sum of roots of unity, hence a_t is an algebraic integer. For any field automorphism σ of $\mathbf{C}$, applying σ to matrices of the irreducible representations permutes the irreducible characters and preserves their degrees. Therefore

$$\sum_t |\sigma(a_t)|^2 = |\mathcal{G}| \sum_{\chi \text{ approved}} \chi(1)^2 \le |\mathcal{G}|^2, \qquad \sum_t |\sigma(a_t)| \le |\mathcal{G}|^{3/2}. \tag{6.9}$$

The last bound is Cauchy–Schwarz. The convolution and the covariance prove the evaluated identity used in Theorem 2

$$(\Theta_\vartheta A_{s_0}\phi)(h) = \sum_{t\in\mathcal{G}} a_t\, B^{\mu}_{\vartheta,\Psi_\vartheta}(\phi)(h\widetilde{(ts_0)}_m). \tag{6.10}$$

It has the same scalar c_* before projection for every translate and twist. More generally, for any principal-level u,

$$\iota_\vartheta(|\mathcal{G}|\rho(e_0)\rho(s_0)u)(h) = \sum_{t\in\mathcal{G}} a_t\, \iota_\vartheta(u)(h\widetilde{(ts_0)}_m).$$

At any coefficient place the Weyl unit ball at h is a module over its valuation ring. The gauge uses the action of h^{-1} on the coefficient vector at the induced local places. Each $\widetilde{(ts_0)}_m$ and its inverse are integral at every finite place. The corresponding Weyl operators preserve the integral Weyl lattice, so the gauge of a given vector is the same at h and at $h\widetilde{(ts_0)}_m$. If all values of a scalar multiple of the unprojected lift have gauge at most one at their arguments, their translates and the combination with coefficients a_t have gauge at most one at h. This works with the same scalar and coefficient place. At the trivial twist, if all the conjugated unprojected values have squared Weyl norm at most e^B, the triangle inequality in the Weyl norm and (6.9) give an upper bound $|\mathcal{G}|^3 e^B$ for every squared projected value. In the class sum the weight of any class is the reciprocal of its finite stabilizer order. With H_m the number of classes of the final trivial-twist level, the logarithmic loss is at most

$$3\log|\mathcal{G}| + \log H_m \le x^{99}. \tag{6.11}$$

The definite class count below gives $\log H_m \le x^{36}$. Since $\log|\mathcal{G}| \le fr^2 \log g$, this proves the last inequality at all sufficiently large x.

7 Short integral forms

Throughout this step, D, the Poincaré maps A_s with their common seesaw scalar c_*, and the level prime N are fixed. Put $n_j = \kappa + j(t+1)$, the rank of X_j, and $C = \dim_{\mathbf{C}} \mathcal{C}$. We use $\|v\|_{\mathrm{W}}$ for the positive Weyl norm of a value and $\|\Theta_\vartheta y\|_\sigma$ for the class-sum norm of a lifted row after applying σ to its coefficients. At a coefficient embedding $j : \mathbf{C} \hookrightarrow \mathbf{C}_p$ we write $\mathcal{O}_j$ for the valuation ring pulled back to $\mathbf{C}$ and $v_j(p) = 1$. Integrality of a form at j means that its Weyl gauge is at most one at every adelic argument, with the lattice transported at that argument.

Theorem 2 (Short integral forms)**.** *For every sufficiently large square with the fixed d, ε, there are a positive rational q, a spanning family S of $\mathcal{C}_{\mathrm{ram}}$ of integral columns, and nonnegative allowances e_p of finite support with $\sum_p e_p \log p \le \mathcal{B}$. Put $B_s = qA_s$ for $s \in \mathcal{G}$. Then for every $s \in \mathcal{G}$, $\phi \in S$ and complex coefficient automorphism σ, one has*

$$\|\Theta_1 B_s \phi\|_\sigma^2 \le e^{h_{\mathrm{sh}}}.$$

For the local assertion let a finite family of rows be given, each with an assigned expression $B_{s_i}\phi_i$ for an integral ramified column. Let $\mathcal{S}$ be the finite set of translates used. For any prime p, any coefficient embedding $\iota : \mathbf{C} \hookrightarrow \mathbf{C}_p$, and any m-fold twist ϑ of level norm prime to p, at ι and $\iota \circ \mathrm{conj}$ there are integral matrices recombining the same lifted row family for ϑ into integral forms. The sum of the two negative logarithmic determinant norms is at most $F_p = 2|\mathcal{S}|Ce_p \log p$, with $\sum_p F_p \le 2|\mathcal{S}|C\mathcal{B}$. Here C is the dimension of the full column space and the maps and allowances depend only on the fixed datum.

A spanning family of short integral columns. At stage zero there are a columns of height one at each $\tau \in \mathcal{T}$ in the weight diagram. All other columns of the diagram have height at least κ. Thus the coefficient representation on X_0 has one symmetric-power factor of degree a for each top and a trivial factor at every other place, with determinant exponent zero. In the degree-a factor write $v(z) = (z^\alpha)_{|\alpha|=a}$ for the Weyl evaluation vector. The positive Weyl norm is invariant under the unitary group and gives $v(e_1)$ squared norm one. Unitary transitivity on the unit sphere and homogeneity give

$$\|v(z)\|_{\mathrm{W}}^2 = \Big(\sum_i |z_i|^2 \Big)^a.$$

It follows by comparing coefficients in $z, \bar{z}$ that the Gram matrix in these coordinates has diagonal entries $a!/\alpha!$. In products of factors the squared norms multiply.

The definite class lattices $L_g = E^\kappa \cap g\widehat{\mathcal{O}}_E^\kappa$ are self-dual for the standard Hermitian form. The free-level calculation applies to the chosen principal level within each definite stage. Evaluation at the class representatives therefore identifies the columns with a direct sum of copies of the coefficient representation.

Write $Q = dN_1$. Then Q belongs to the different of E. Indeed $E \subset \mathbf{Q}(\zeta_Q)$, and $\Phi_Q'(\zeta_Q)$ divides $Q\zeta_Q^{-1}$, by differentiating the factorization of $Z^Q - 1$. Since $\mathbf{Z}[\zeta_Q]$ is the ring of integers, Q kills its codifferent modulo that ring. The trace from $\mathbf{Q}(\zeta_Q)$ to E maps integers to integers, so the codifferent of E is contained in the cyclotomic codifferent. Intersecting with E proves the assertion.

Endow the real embedding of L_g into $\mathbf{C}^{f\kappa}$ with the Euclidean norm, using one embedding from each conjugate pair. Its Euclidean dual is $2\mathfrak{D}_E^{-1}L_g$ in the same coordinates because the Euclidean pairing is half the Hermitian trace pairing. Every nonzero $v \in L_g$ has $\|v\| \ge 1$. Indeed its squared length is the sum of the conjugates of the nonzero totally positive integer given by its Hermitian square. Thus every nonzero vector of the Euclidean dual has length at least $2/Q$.

Here is the weak transference estimate needed for this lattice. For a full Euclidean lattice L of rank D_0, Minkowski's first theorem on L and L^*, and the cube in the unit ball, give $\lambda_1(L)\lambda_1(L^*) \le D_0$. If $\lambda_1(L^*) \ge \rho > 0$, project along a primitive shortest vector $b \in L$. The projected lattice has dual $L^* \cap b^\perp$ and minimum in its dual at least ρ. Apply induction in dimension to the projected distance, then lift the approximate vector and round its remaining coordinate modulo b. The total distance is at most $(D_0 -$

$1)^2/\rho + D_0/(2\rho) \le D_0^2/\rho$, with the same argument in dimension one starting from the zero-dimensional lattice. For L_g the resulting radius is $R_{\mathrm{tr}} = (2f\kappa)^2 Q/2$. Choose successive E-independent elements of L_g by taking a target vector of Euclidean length $2R_{\mathrm{tr}}$ orthogonal to the archimedean image of the E-span already chosen and approximating it within R_{tr}. Each new length is at most $3R_{\mathrm{tr}}$ and the new vector is outside that span. The resulting frame $Y_g \in \mathrm{GL}_\kappa(E)$ satisfies

$$g^{-1} Y_g \in M_\kappa(\widehat{\mathcal{O}}_E), \qquad |\tau((Y_g)_{hi})| \le P_{\mathrm{frame}} = 6Q(f\kappa)^2$$

at all complex embeddings. The same proof works for n_j in place of κ.

Use a standard positive definite group of rank q with self-dual class lattices. At a free class represented by h, the level preserves the transported integral gauge. Let $v_\tau(X)$ be the vector of Weyl coordinate polynomials of the diagram at τ, with the determinant character treated separately. Choose Y_h by the same frame construction in rank q. For a nonzero determinant exponent a_τ, choose also the frame $Y_{h^{a_\tau}}$ for h^{a_τ}. Use the class values

$$\bigotimes_\tau v_\tau\big(\Phi_\tau(\overline{Y_h})X_\tau\big) \prod_{\tau\,:\,a_\tau \neq 0} \Phi_\tau(\overline{\det Y_{h^{a_\tau}}}).$$

Here the entries of X_τ run over $D+1$ distinct roots of unity, where D is any bound for the number of boxes of the diagram at each τ. For the symmetric-power evaluation one can take X_τ to be a single column. Extend each value by automorphy on its class and put zero elsewhere. Freeness makes this extension well-defined.

In the transported gauge at h, the conjugate on the frame matches the coefficient reading $\Phi_\tau \circ \mathrm{conj}$. The polynomial arguments use $h^{-1}Y_h$, and the extra factors use $\det(h^{-a_\tau} Y_{h^{a_\tau}})$ at their respective embeddings. Both are integral in these frames. The class transport preserves integrality. The values at h span the coefficient space since the Weyl coordinate polynomials are linearly independent, substitution by Y_h is invertible, and evaluation on the roots-of-unity grid is injective on polynomials of degree at most D in each entry. Indeed univariate interpolation in each variable proves this last assertion in succession. The determinant factors are nonzero scalars. Conjugate norm bounds on the representatives hold on their whole classes since rational unitary transport remains unitary at every complex embedding.

For the stage-zero space take the single-column evaluation at each top, of degree a and determinant exponent zero. The symmetric-power norm formula above bounds the squared norms uniformly by $(\kappa^3 P_{\mathrm{frame}}^2)^{aT}$. Apply the ramifying product $\mathcal{R}$ to this spanning family. Each $1 - z_v$ preserves integrality and multiplies a squared uniform norm bound by at most 4, by the triangle inequality and compactness of the finite translation in the adelic lattice. The number of factors is f, as in the central-orbit calculation. Consequently one set S spans C_{ram}, is integral at every coefficient embedding, and obeys

$$\|\sigma\phi(g)\|_{\mathrm{W}}^2 \le 4^f (\kappa^3 P_{\mathrm{frame}}^2)^{aT} \le e^{\mathcal{B}} \quad (\phi \in S,\ \sigma \in \mathrm{Aut}(\mathbf{C})). \tag{7.1}$$

For the last inequality use $T \le f$, $\log P_{\mathrm{frame}} = O(\log x)$ and $\log x = o((\log b)^3)$.

Class sums, conjugate bounds and the signed height. Use the finite datum Ψ_ϑ of the fixed seesaw. For the chain integral, the quotient Haar measure at each stage assigns one to the class of a free representative when the level has volume one. The combined integrand over the m stages, after summing over the support of Ψ_ϑ, is invariant under independent right translations by those levels. This follows by moving the input and output translations through each kernel; the intermediate actions cancel and joint invariance fixes the sum. Thus the chain is the sum over all tuples of classes of these combined values. In particular no denominators arise from passing from a joint integral to the joint class sum. If each individual term propagates an upper bound for the pointwise Weyl gauge, the ultrametric inequality gives the bound for the whole sum whenever the coefficients of Ψ_ϑ are integral. For complex norms the triangle inequality instead costs the number of class tuples and the sum of the absolute values of the coefficients.

The same reasoning applies after changes to the kernel data that commute with the indicated source and target levels.

For these standard positive definite groups with self-dual class lattices, let $q \le 2n_m$ and suppose the level contains the principal congruence subgroup modulo the rational integer M. The short integral frame in rank q gives a Gram matrix whose entries are algebraic integers of modulus at most $q(6Q(fq)^2)^2$ at all embeddings. Nonzero integral differences have maximum modulus at least one by their rational norm. Packing in the archimedean box therefore gives $\exp(O_d(fq^2 \log x))$ possible Gram matrices. If L' is spanned by the frame over $\mathcal{O}_E$, then $L' \subset L_h \subset (L')^\vee$. The index $[(L')^\vee : L']$ is the absolute norm of the Gram determinant and its logarithm is $O_d(fq \log x)$ by Hadamard. A finite abelian group of order J_0 has at most $J_0^{1+\log_2 J_0}$ subgroups, by choosing a sequence of generators each increasing the generated subgroup's order by a factor at least two. Frames with the same Gram matrix are identified by a rational unitary map, so these subgroups bound the possible class lattices for each Gram matrix. The classes at principal level modulo M surject onto the classes at the given level. Entrywise reduction modulo M bounds the principal level index by M^{2fq^2}. Thus the logarithmic class count is

$$O_d\big(fq^2 \log x + (fq \log x)^2 + fq^2 \log M\big).$$

For the definite stages of Theorem 2 take $M = gN$, including the final class-sum level on Y_m at the trivial twist. The bounds $f = O(x^{11} \log x)$, $q \le n_m = O(x^{11/2})$ and $\log(gN) = O(m \log b)$ give a logarithmic count at most x^{36} for large x.

7.1 Gaussian jet pairing and the archimedean height

Let Z_ρ be complex rectangular matrix variables, grouped into blocks Z_b of D_b complex entries. Write $s_{\rho,i}$ for the singular values of Z_ρ, in decreasing order. Let P be a polynomial of degree d_b, homogeneous in each block, and suppose

$$|P(Z)|^2 \le A \prod_{\rho,i} s_{\rho,i}^{2\lambda_{\rho,i}}, \qquad d_b = \sum_{\rho \in b,i} \lambda_{\rho,i}, \qquad \lambda_{\rho,i} \ge 0.$$

Let H be entire, with $|H(Z)| \le M e^{\|Z\|^2/2}$, where the norm is the Hilbert–Schmidt norm on all entries. With $P = \sum p_I Z^I$, $H = \sum h_I Z^I$, and $\phi(u) = u \log u - u$ for $u > 0$, $\phi(0) = 0$, one has

$$\left| \sum_I p_I I! h_I \right|^2 \le A M^2 e^{\sum_{\rho,i} \phi(\lambda_{\rho,i})} \prod_{b \,:\, d_b > 0} e(2d_b + 2)^{2D_b}. \tag{7.2}$$

Conjugating the coefficients of P preserves its bound, because Z and $\bar{Z}$ have the same singular values. Thus the same estimate holds with $\overline{p_I}$ in the pairing.

Evaluate H at zero in each block of degree zero. In every other block use $\epsilon_b = (2d_b + 1)^{-1}$ and the Gaussian measure $\pi^{-D_b} e^{-(1+\epsilon_b)\|Z_b\|^2}\, dZ_b$. The squared norm of a monomial of degree d_b in this measure is its factorial times $(1 + \epsilon_b)^{-D_b - d_b}$. Monomial orthogonality, Cauchy–Schwarz and the bound for H give $M^2 \prod_b (1 + \epsilon_b)^{D_b + d_b} \epsilon_b^{-D_b}$ times the squared norm of P in the Gaussian measure of parameter one. In its integral, the identity $\|Z_b\|^2 = \sum_{\rho \in b,i} s_{\rho,i}^2$ leaves the factor $e^{-\epsilon_b \|Z_b\|^2}$ after bounding $r^\lambda e^{-(1-\epsilon_b)r}$ by $(1-\epsilon_b)^{-\lambda} e^{\phi(\lambda)}$. The residual integral gives $\epsilon_b^{-D_b}$. Since $(1+\epsilon_b)/\epsilon_b^2 \le (2d_b+2)^2$ and $((1+\epsilon_b)/(1-\epsilon_b))^{d_b} \le e$, this proves (7.2).

For a rectangular Schur substitution in unitary norms with normalized highest vectors, the scalar matrix coefficient between unit vectors has the bound in this estimate with $A = 1$ and $\lambda_{\rho,i}$ the row lengths. In fact in the exterior-power column model, diagonal substitution contributes at most the product of the first h singular values in each column of height h. A coordinate inclusion is isometric since it is unitary-equivariant and preserves the normalized highest vector. Singular value decomposition reduces a rectangular matrix to this diagonal and inclusion. Scaling a block variable by a scalar c_ρ contributes $|c_\rho|^{2\sum_i \lambda_{\rho,i}}$ to A.

The archimedean estimate for a definite edge. Put $r_0 = n_j n_{j+1}$, $D_j = f r_0$, and $u = \log(g/\Omega)$. For one moved, normalized vacuum, Taylor transport at the actual coefficient scale $\overline{\pi_G}$per reduces the pairing with a Schur matrix coefficient to the Taylor pairing of an entire function $H(Z)$ with the conjugate polynomial, with the places permuted. In the polynomial the argument is scaled by $c'_\sigma = \sigma(\overline{\pi_G})t_\sigma$, with $|c'_\sigma|^2 = g/\Omega$. The function H has the majorant $M_j e^{\|Z\|^2/2}$, $M_j = \exp(C_d(\nu r_0 + 1)^2)$, by (4.2).

Choose unit input and dual test vectors in the conjugate Weyl spaces, and use their coefficientwise preimages, with the place permutation, for the scalar test on the original jet. The Hermitian dual is taken in the actual conjugate metric. The test polynomial has Schur row lengths $\lambda_{j,\tau,i}$ and includes the scale c'_σ above at each place. In (7.2) use one block of D_j entries and degree $d_j = \sum_{\tau,i} \lambda_{j,\tau,i}$. The squared scale factor is $A = e^{u d_j}$. Thus for $d_j > 0$ the squared operator norm is at most

$$M_j^2 (2d_j + 2)^{2D_j} \exp\left(1 + \sum_{\tau,i} [\phi(\lambda_{j,\tau,i}) + u\lambda_{j,\tau,i}]\right).$$

For $d_j = 0$ use $H(0)$. Since $\phi(\lambda) \le \log(\lambda!)$, the principal exponent is bounded by $\sum_{\tau,i} S(\lambda_{j,\tau,i})$. The sum over stages of the additional logarithmic costs $O_d((N_1 n_j n_{j+1} + 1)^2 + D_j \log(2d_j + 2) + 1)$ is $O(x^{70})$. Indeed $N_1, f = O(x^{11} \log x)$, $n_m = O(x^{11/2})$, $m = O(x^{1/2})$, and the stated row length bound gives $d_j \le 2afn_j$. The class-sum and finite coefficient factors are charged separately in the short-lift argument.

The weight at height κ is $w_1(\tau) - \lfloor(\kappa + t + 1)/2\rfloor$. The other heights $h > \kappa$, separated by $t + 1$, have multiplicity at most $2r^8$. There are at most m of them in each diagram. It follows that the row lengths are at most

$$a\, 1_{\{\tau\in\mathcal{T},\ i=0\}} + w_0(\tau) 1_{\{i<\kappa\}} + O(mr^8).$$

At a top $w_0(\tau) = a/x_1 + O(N_1 + r^8)$. Since $g/\Omega > 1$ eventually, the increment $S(v + 1) - S(v) = \log(v + 1) + \log(g/\Omega)$ is positive and $O(\log b)$ for $v \le 2a$. All the displayed lengths are at most $2a$. Consequently

$$\sum_{j,\tau,i} S(\lambda_{j,\tau,i}) \le m\big(TS(a) + \kappa \sum_{\tau} S(w_0(\tau))\big) + O\big(mfa \log b/x_1 + x^{70}\big).$$

The correction in the first row of a top costs $O((a/x_1 + mr^8)\log b)$, and all remaining corrections are at most $O(mr^8 \log b)$ per row. The same error bound dominates the sum of the remaining single-step Gaussian terms, the logarithmic class counts, and the squared-norm loss $2x^{60}$ from Ψ_1.

The terminal branching calculation gives the squared logarithmic norm cost

$$-Ta \log N + 2fb \log N + O(x^{70}).$$

Its positive terms and the class-sum and jet errors above are $o(\mathcal{B})$, using $\log N = O(m \log b)$, $a \ge xb$ and $x_1 \gg x^A$ for every fixed A.

Let γ_p be as in the prime-power decomposition of the height. More explicitly put

$$\mathrm{Ask}_p(v, \tau) = \frac{1 - \chi_d(p)}{2} \sum_{\substack{e\ge1 \\ e \text{ odd or } p=d}} (v/p^e - w([-p^e]\tau))_+,$$

$$\gamma_p = m\Big(\sum_{\tau\in\mathcal{T}} \mathrm{Ask}_p(a, \tau) + \kappa \sum_{\tau} \mathrm{Ask}_p(w(\tau), \tau)\Big).$$

Both asks and gains vanish for split primes and $p \ge x_1$. For $v \le a$, summation of $2\,\mathrm{Ask}_p(v, \tau)\log p$ gives $\mathrm{Ask}(v, \tau)$ by the weights of prime powers. Thus on all the values of the unprojected lift $\iota_1 \mathcal{P}(c_* \phi)$ the logarithmic squared-norm bound is

$$A + h_0 + 2 \sum_{p} \gamma_p \log p + \mathcal{B}. \tag{7.3}$$

Here A bounds the logarithmic squared norms of the input values, and the height with the same signed rem and levels is

$$h_0 = h_{\mathrm{main}} - m\,\mathrm{rem} + T\big(m(S(a) - \mathrm{MainAsk}) - a\log N\big).$$

Concavity of log on the intervals from 1 to a gives $\log(a!) \le a\log a - a + 1 + \frac{1}{2}\log a$. Hence $S(a) - \mathrm{MainAsk} \le a(\bar{w} + G) + 1 + \frac{1}{2}\log a$. The lower endpoint of the level window is $m(\bar{w} + G) - G - \frac{15}{2}\log b$. The terms $ma(\bar{w} + G)$ cancel against this choice of level. Since $G = O(\log b)$, $T \le f$, and $m/a \to 0$, this proves $h_0 \le h_{\mathrm{main}} - m\,\mathrm{rem} + \mathcal{B}$. For G one can use Q in the different to get $\mathrm{rd}_E \le Q$.

The projection formula preserves the integral gauge at each coefficient embedding. Its logarithmic squared-norm cost at the trivial twist is at most x^{99}.

Collision counts and the loss budget. Call p small non-split when $p < x_1$ and $\chi_d(p) \ne 1$. At such a prime write $c_p = \lceil \log_p m \rceil$ and

$$\ell_p = \frac{f a(\log b)^2}{p} \begin{cases} m & (p \equiv \pm 1 \bmod m), \\ 1 & (\text{otherwise}). \end{cases}$$

$$e_p = \begin{cases} 3\ell_p & (p \text{ small non-split}), \\ 4mfb & (p = N), \\ 0 & (\text{otherwise}). \end{cases}$$

These cases are disjoint. At the short residue classes write $p = mc \pm 1$ with $c \ge 1$, whence $m/p \le 2/c$. Bound the sum of $\log p/p$ at the other small primes by the harmonic sum over all integers up to x_1, multiplied by $\log x_1$. The two short residue classes have the same order of bound after using $c \le 2x_1$. Therefore

$$\sum_{p < x_1} e_p \log p = O\big(f a(\log b)^2(1 + \log x_1)^2\big).$$

Also $4mfb\log N = O(m^2 fb\log b)$. Since $(\log x_1)^2 = (\log b)^{1/3}$ and $m^2 b \le 4a$, these errors total $o(\mathcal{B})$, in particular

$$\sum_{p} e_p \log p \le \mathcal{B}. \tag{7.4}$$

For p small non-split, p is a unit modulo N_1. To see this, the primes dividing M are $2, 3, 5, 7, 11$. For each of the odd ones, $d \equiv -1 \pmod{p}$ and $\chi_d(p) = (-d/p) = 1$. For 2 use $d \equiv 7 \pmod 8$. At $p = d$ we have $d \equiv -1 + M/2 \pmod M$, so $d^2 \equiv 1 \pmod M$ and $-d^e$ is never 1 mod M for positive e.

For an inert p, write $u_p = [-p] \in (\mathbf{Z}/N_1\mathbf{Z})^\times$ and let o_p be the order of its reduction modulo M. When e is odd, a top σ has $[-p^e]\sigma$ a top exactly when $o_p \mid e$. We call an inert p served when $p \not\equiv \pm 1$ mod m and o_p is odd and at most c_p. For any set I of infinite places, define $\mathrm{coll}_p(I)$ as the sum of a/p^e over the live pairs $e \ge 1$ and σ for which both σ and $[-p^e]\sigma$ are tops outside I. Here live means that e is odd or $p = d$. If all such e are at least e_0, the sum is at most $2fa/p^{e_0}$ by the geometric series. At an unserved prime this gives

$$m\,\mathrm{coll}_p(\emptyset) \le \ell_p.$$

Indeed there is no collision for $p = d$ or for even o_p. For $p \equiv \pm 1$ mod m use the bound with $e_0 = 1$ and the extra m in ℓ_p. In the remaining unserved case use $e_0 = c_p + 1$ and $p^{c_p} \ge m$.

For a served p, consider the finite places induced by the tops at the given coefficient embedding j. Write $\lambda(\sigma)$ for the place induced by $j \circ \Phi_\sigma$. Decomposition in the compositum E gives

$$\lambda(\sigma') \in \{\lambda(\sigma), \overline{\lambda(\sigma)}\} \quad \Longleftrightarrow \quad \sigma' \in \langle u_p \rangle \sigma.$$

In fact the Frobenius generator, in the product Galois group of K and the cyclotomic factor, is $(\mathrm{conj}_K, [p])$, and complex conjugation is $(\mathrm{conj}_K, [-1])$. Their product acts as u_p on the cyclotomic label. Its odd powers

interchange the two conjugate places, and its even powers fix each. The places in a pair are distinct. Indeed the order of u_p is divisible by m, by reduction modulo 2^s. To see the divisibility, put $v = v_2(p^2-1)$. Since $p \not\equiv \pm 1 \bmod m$ and exactly one of $p-1, p+1$ has 2-valuation one, $v \le \log_2 m$. Successive squaring gives $v_2(p^{2^k}-1) = v+k-1$ for $k \ge 1$, so the order modulo 2^s is divisible by 2^{s-v+1}, hence by m. In particular the cyclotomic order of p is even and the Frobenius decomposition group cannot contain complex conjugation: if $[p]^k = [-1]$ for odd k, then $u_p^k = 1$, contradicting the even order.

Since o_p is odd, multiplication by $u_p^{o_p}$ pairs top places inducing conjugate primes. Choose one prime from each such pair, giving the marked set Λ. Let $I = \{\sigma : \lambda(\sigma) \in \Lambda\}$. For every live collision, one of the two places belongs to I, so $\mathrm{coll}_p(I) = 0$. An orbit through a top reduces modulo M to a cyclic group of odd order. It cannot contain -1, so it contains no thick place. Thus multiplication by u_p bijects the two halves of each marked pair, preserving every $w_j(\sigma)$. If $p \ge m$, then $c_p = 1$ and $o_p = 1$, so the entire marked orbit consists of tops. These facts apply for each coefficient embedding, including its composition with complex conjugation.

The lower exponent in one step. Write $B_{j,\sigma}$ for the number of boxes in columns of heights different from 1 and less than n_{j+1} at step j, and $r_{j,\sigma}$ for their number of columns. The height-κ term and the polynomial bounds on higher multiplicities give, uniformly in $j < m$ and σ,

$$\kappa w(\sigma) \le B_{j,\sigma} \le \kappa w(\sigma) + x^{60}, \qquad r_{j,\sigma} \le w(\sigma) + x^{60}.$$

For the lower bound, $w_1(\sigma)$ includes $3r^8$ in addition to the rounded-up nominal weight, exceeding the term $\lfloor(\kappa+t+1)/2\rfloor$ subtracted at height κ. For the upper bounds all added boxes come from rounding by at most N_1, the stage-one gap, and $O(m)$ polynomial-sized columns of height at most n_m.

For a prescribed I, put $t_\sigma = a1_{\sigma\in\mathcal{T}\setminus I}$ and write $\sigma' = [-p^e]\sigma$ for a live exponent. The coupling count in the Schur jet estimate at (σ, e) is

$$K_{\sigma,e} = \min(t_\sigma/p^e, r_{j,\sigma'} + t_{\sigma'}) + r_{j,\sigma}/p^e + \min(B_{j,\sigma}/p^e, B_{j,\sigma'}).$$

Let $B_j = Ta + \sum_\sigma B_{j,\sigma}$, $H_p = \lfloor \log_p B_j \rfloor$, and $\omega_p = (1-\chi_d(p))/2$. Put $\mathscr{E}_p = \{e \ge 0 : e \text{ odd or } p = d\}$. The Schur jet exponent before the charges is

$$J_{p,j}(I) = \omega_p B_j \sum_{\substack{e\in\mathscr{E}_p\\ e>0}} p^{-e} - \omega_p \sum_{\substack{e\in\mathscr{E}_p\\ e\le H_p}} \sum_\sigma K_{\sigma,e} - f n_{j+1}^3 (H_p+1).$$

Here $K_{\sigma,0}$ uses $\sigma' = [-1]\sigma$. The ramified normal form permits a quadratic term and equal Frobenius powers in the exponential, so the coupling set includes zero when $p = d$. The first term is the factorial gain less the CM contribution per box, using the geometric series over the positive live exponents.

7.2 Integral theta germs and jet divisibility

Use the marked model of θ above, over the valuation ring $\mathcal{O}$ of $\mathbf{C}_p$, with $v_p(p) = 1$. Let $X_0(z) = \wp(z+t_0) - \wp(t_0)$ and $Y_0(z) = (\wp'(z+t_0) - a_1 X_0(z))/2$. Here z is the invariant differential coordinate and the subscript only denotes the marked elliptic coordinates. Thus $Y_0 = z + O(z^2)$ and $X_0 = z^2 + O(z^3)$. The cubic and its derivative give

$$X_0 \in \mathcal{O}[[Y_0]], \qquad Y_0' = 1 + 3X_0^2 - a_1 Y_0.$$

For the first assertion solve the cubic for X_0 at the marked point, where its derivative in X_0 is a unit. Consequently $\mathcal{O}[[Y_0]] \subset \mathbf{C}_p[[z]]$ is closed under d/dz.

For $0 \ne \delta \in \mathcal{O}_K$ and $\delta y \in \Lambda$, put $n_\delta = \delta\bar\delta$. Then

$$(U_y\theta)(\bar\delta z) = \theta(z)^{n_\delta} G_{\delta,y}(z), \qquad G_{\delta,y} \in \mathcal{O}[[Y_0]]. \tag{7.5}$$

The two sides have the multiplier of a degree-n_δ section. The canonical multiplier $\exp(\pi i E_\Lambda(\bar\omega l, l))$ pulls back under $\bar\delta$ to its n_δ-th power. The translation phase is one since $E_\Lambda(\bar\delta l, y) = E_\Lambda(l, \delta y)$ is integral. The isogeny height ratio $H(\bar\delta z)/H(z)^{n_\delta}$ is bounded. This follows from the section basis of the marked cubic, applied to the pullbacks of θ and s_Λ^2, of degrees n_δ and $2n_\delta$. Under doubling the ratio becomes its fourth power. Since doubling on the complex uniformizing line is onto, its finite positive supremum is one. The torsion-translation height calculation above then gives $|(U_y\theta)(\bar\delta z)|_p \le H(z)^{n_\delta}$.

To read this bound in the formal ring, write the quotient by θ^{n_δ} as $P(X_0) + Y_0 Q(X_0)$ in the section basis. For any integral x whose reduction avoids the roots of $4x^3 + a_1^2 x^2 + 4x$, the two points above x have integral Y_0 values with unit difference. The discriminant polynomial has nonzero reduction, including at $p = 2$ where a_1 is a unit. At these two points $H/|\theta|_p = 1$, so the two quotient values give $P(x), Q(x) \in \mathcal{O}$. If a polynomial had a coefficient of negative valuation, divide by a coefficient of least valuation and choose a residue outside the finite excluded set and the roots of the resulting nonzero polynomial. This proves that P, Q have integral coefficients and gives (7.5). The same argument for an even section of degree n_δ gives a polynomial in X_0 alone of degree at most $\lfloor n_\delta/2 \rfloor$.

Put $L = -\log\theta$ as a formal series, with $L(0) = L'(0) = 0$. Logarithmic differentiation of the translated sigma function gives $L'' = L_2 + X_0$ for a constant L_2. The logarithmic derivative of $G_{\bar a,0}$ is integral in $\mathcal{O}[[Y_0]]$ for $0 \ne a \in \mathcal{O}_K$, since its constant term is one. Applying d/dz and evaluating at zero gives $(a\bar a - a^2)L_2 \in \mathcal{O}$. Subtract the readings for $a = \omega, \bar\omega$. As $\omega + \bar\omega = 1$, this proves $\sqrt{-d}\, L_2 \in \mathcal{O}$.

Now take p small non-split. Let $k_p = 2$ at inert p and $k_p = 1$ at $p = d$, set $q_0 = p^{k_p}$, $Q = q_0 - 1$, and use $\pi_0 = p$ or $\sqrt{-d}$, respectively. Then $\pi_0\bar\pi_0 = q_0$ and $v_p(\pi_0) = \omega_p$. All these p are at least 13. The Q-th roots of unity are in K_p with norm ζ^{p+1}. For a in the closure of $\mathcal{O}_K$, the formal action $z \mapsto az$ sends Y_0 into $\mathcal{O}[[Y_0]]$. For $a \in \mathcal{O}_K$ this follows by extending the endomorphism over the good elliptic model. General a follows by coefficientwise limits in the fixed invariant coordinate. Make the following rotation-equivariant choice of parameter

$$T_0(z) = \frac{1}{Q}\sum_{\zeta^Q=1} \zeta^{-1} Y_0(\zeta z) = z + O(z^{q_0}).$$

This is an integral change of formal coordinate, and it and its inverse $Z(T_0)$ have only degrees 1 mod Q. The reduction of $[\pi_0]$ is purely inseparable of degree q_0. At the ramified prime its differential is zero and its degree is p. At an inert prime, an ordinary special fiber would have a one-dimensional geometric étale p-torsion space on which the unital action of $\mathcal{O}_K/p = \mathbf{F}_{p^2}$ is impossible. Thus the reduction of $[p]$ has inseparable degree p^2. Here the CM endomorphism is taken on the good model at the origin corresponding to $z = 0$ under translation by the marked point. Its first nonzero coefficient in the reduced T_0 parameter is in degree q_0. If $Z_{q_0} = [T_0^{q_0}]Z$, this says that $(\pi_0 - \pi_0^{q_0})Z_{q_0}$ is a unit.

Choose $\varpi^Q = -1/Z_{q_0}$ and put $\xi = \varpi^Q$. We have $v_p(\xi) = \omega_p$. Also Z' is integral since dT_0/dz is an integral unit in the T_0 parameter. For $j \ge 2$, the coefficient in degree $1 + jQ$ of $\varpi^{-1}Z(\varpi X)$ has valuation at least $j\omega_p - v_p(1 + jQ) \ge \omega_p$. For $2 \le j \le 4$, $1 + jQ$ is prime to p. For $j \ge 5$, $1 + jQ < p^{(j-1)\omega_p}$ follows at $j = 5$ from $p \ge 13$ and then by induction. Thus this series is integral and is $X - X^{q_0}$ modulo ξ. Its compositional inverse $S(X) = T_0(\varpi X)/\varpi$ is integral and satisfies

$$S(0) = 0, \qquad S - S^{q_0} = X \pmod{\xi}.$$

Let $L^\flat$ be the part of L in degrees $p + 1$ mod Q. The same rotation average proves $W_0 = \exp(L^\flat - L) \in \mathcal{O}[[T_0]]$. In fact $\exp(a\bar a L(z) - L(az))$ is an integral unit of constant term one by (7.5), and this persists under the coefficientwise limit $a \to \zeta$. Taking the $-\zeta^{-(p+1)}$-th power preserves integrality by the integral binomial series, since the exponent is in $\mathbf{Z}_p$. Multiplying these powers over all rotations and taking the $1/Q$-th power proves the assertion.

For the exponent, first use (7.5) for π_0 and the even section $\theta(\pi_0 z)$. The quotient is $R(X_0)$, where $\deg R < q_0/2$, and it equals $\exp(q_0 L - L(\pi_0 z))$. In degrees $n < q_0$ all coefficients of $q_0 L - L(\pi_0 z)$ have

valuation at least ω_p. Indeed X_0 in these degrees has integral coefficients, since $T_0 = z \bmod z^{q_0}$. For $3 \le n < q_0$, $v_p(n(n-1)) \le \omega_p$, and the degree two uses the integral bound on $\sqrt{-d}L_2$. The factor $q_0 - \pi_0^n$ for $n \ge 2$ has valuation at least $2\omega_p$. The exponential in the same range is one modulo π_0, by $v_p(r!) \le (r-1)/(p-1)$. Because $X_0 = z^2 + \cdots$, successively reading the coefficients of degrees $2j < q_0$ shows that $R-1$ has all coefficients divisible by π_0. Its logarithm is then in $\pi_0\mathcal{O}[[T_0]]$, again by the factorial bound, since $R(0) = 1$. Taking degrees $p+1 \bmod Q$ gives

$$q_0 L^\flat - L^\flat(\pi_0 z) = \pi_0 T_0^g(\lambda + T_0^Q H_0(T_0)), \qquad g = \begin{cases} p+1 & (k_p = 2), \\ 2 & (k_p = 1), \end{cases}$$

where $\lambda \in \mathcal{O}$ and $H_0 \in \mathcal{O}[[T_0]]$. The projector reads the same congruence classes in T_0 because $T_0(\zeta z) = \zeta T_0(z)$.

Write $c_\varpi = \varpi^{p+1}/p$ and $\Xi(X) = -c_\varpi^{-1} L^\flat(\varpi X)$. Dividing the last identity by $-q_0 c_\varpi$ shows that $(1 - \pi_0^n/q_0)[X^n]\Xi$ is integral and congruent to $\lambda_p[X^n]S^g$ modulo ξ, with $\lambda_p = -\pi_0 \varpi^g \lambda/(q_0 c_\varpi) \in \mathcal{O}$. The factor multiplying λ has valuation zero. For $n \ge 3$, $1 - \pi_0^n/q_0$ is one modulo ξ.

Degrees zero and one vanish, and degree two vanishes in the inert case. At the ramified prime $1 - \pi_0^2/q_0 = 2$ and $[X^2]S^2 = 1$. Therefore Ξ is integral and

$$\Xi \equiv \begin{cases} \lambda_p S^{p+1} & (k_p = 2), \\ \lambda_p(S^2 - X^2/2) & (k_p = 1) \end{cases} \pmod{\xi}. \tag{7.6}$$

In the inert case λ_p is a unit. For this, set $h_j = [z^{p-1}]X_0^j$ and reduce modulo the maximal ideal of $\mathcal{O}$. The coefficients in degree $p-1$ of $(X_0^j X_0')'$ vanish, since $[z^p](X_0^j X_0')$ is integral. The cubic gives

$$(X_0^j X_0')' = (4j+6)X_0^{j+2} + (j+1)a_1^2 X_0^{j+1} + (4j+2)X_0^j.$$

If h_1 vanished on reduction, then $h_0 = 0$ and the recurrence would force $h_{(p-1)/2} = 0$, since $4j+6$ is prime to p for $0 \le j \le (p-5)/2$. But $h_{(p-1)/2} = 1$. Now $[X^{p+1}]\Xi = -h_1/(p+1)$.

Finally choose a positive integer n prime to p with $ny \in \Lambda$. Apply (7.5) at $z = \varpi X/n$, using $\delta = n$. This gives

$$(U_y\theta)(\varpi X) = V_y(X/n)\exp\bigl(n^2 c_\varpi \Xi(X/n)\bigr), \qquad V_y \in \mathcal{O}[[X]].$$

The factor V_y is the product of the rescaled $G_{n,y}$ and $W_0^{n^2}$; both are integral after scaling because $T_0(\varpi X) = \varpi S(X)$. The unit variable rescaling and the exponent scalar depend only on n. We have $v_p(n^2 c_\varpi) = 1/(p-1) - \omega_p$ in both cases. At a ramified prime $X = S - S^p$ modulo ξ, so $\lambda_p(S^2 - X^2/2) = \lambda_p(2S^{p+1} - S^{2p} + X^2/2)$ modulo ξ. This is the ramified shape used in the tensor calculation, with $2p = p + p^{k_p}$.

The normalized products and the source gauge. The same one-variable model applies to the standard-data module in the placed Weyl coordinates. Write $\nu = N_1$, $s_1 = s_0/u_0$, $\mathfrak{a}_0 = E \cap \alpha\widehat{\mathcal{O}}_E$, and use $L = r_\infty(\mathfrak{a}_0)$ and η from the vacuum construction. We first construct a tight frame with frame constant $M_p'^2$ in $\mathfrak{a}_0$, where M_p' is an integer prime to p. List each vector with its multiplicity. For $p \nmid \nu$ use $M_p' = \nu$ and the powers of ζ, each repeated ν times, as in the original vacuum calculation.

For $p \mid \nu$ write $A_0 = (d+1)/4$, $\nu' = \nu/p^{v_p(\nu)}$, $r_0 = \nu/p$, $\zeta_p = \zeta^{r_0}$. All primes dividing ν divide A_0. There is an integer z_0 with $A_0 z_0^2 + z_0 + 1 \equiv 0 \bmod \nu$. Indeed $u \mapsto A_0 u^2 + u$ is injective modulo ν, since the difference at u, v is $(u-v)(1 + A_0(u+v))$ with the second factor a unit. Put $t_1 = \bar{\omega} + A_0 z_0$ and $c_1 = t_1 \nu'/\nu$. Use the list

$$w_{jkh} = \zeta^j\bigl(1_{k=0}\zeta_p^h + c_1(\zeta_p^h - \zeta_p^k)\bigr), \qquad 0 \le j < r_0, \quad 0 \le k, h < p.$$

These vectors lie in $\mathfrak{a}_0$. On the denominator side of α at a prime over p, this uses

$$v_p\bigl(\Phi_\nu'(\zeta)\bigr) + v_p(\zeta_p - 1) = v_p(\nu).$$

In fact the prime-to-p cyclotomic factor is unramified and the different of the p^u-part is generated by $\Phi'_{p^u}(\zeta_{p^u})$, of valuation $u-1/(p-1)$, by differentiating $(X^{p^u}-1)/(X^{p^{u-1}}-1)$. On the opposite side $\bar{\omega}$ vanishes modulo the prime and $\bar{t}_1$ is a unit. Since $t_1\bar{t}_1 = A_0(A_0z_0^2+z_0+1)$, t_1 is divisible by $p^{v_p(v)}$ there. At every other prime v'/v is a unit.

For two embeddings σ,τ, the sum over h,k of $\Phi_\sigma(w_{jkh})\overline{\Phi_\tau(w_{jkh})}$ vanishes if their ζ_p values differ, by geometric sums. If the values agree, that sum is

$$p(1+c_1+\bar{c}_1+2pc_1\bar{c}_1)\Phi_\sigma(\zeta)^j\overline{\Phi_\tau(\zeta)^j}.$$

The sum over j then vanishes unless $\sigma=\tau$. The common diagonal value is

$$M'_p = v(1+c_1+\bar{c}_1+2pc_1\bar{c}_1) = v+v'(1+2A_0z_0)+2p(v')^2t_1\bar{t}_1/v.$$

It is a positive integer congruent to v' modulo p. Repeating this list M'_p times gives the desired identity $\sum_i \Phi_\sigma(w_i)\overline{\Phi_\tau(w_i)} = (M'_p)^2\delta_{\sigma\tau}$.

Use the forms

$$L_{i,c}(Z) = \frac{\Omega_0}{M'_p\eta}\sum_\tau \overline{\Phi_\tau(w_i)}Z_{\tau,c}$$

in each matrix coordinate c. With $\prod_{i,c}\theta(L_{i,c})$, the trace-dual identity $\mathfrak{a}_0\bar{\mathfrak{a}}_0 = \mathfrak{D}_{E/K}^{-1}$, the frame identity and the canonical multiplier give the same construction on M'_pL as the original vacuum construction on vL. Let F be its finite Weil vector. Its evaluation at zero is one. This product recovers the previously normalized vacuum with integral coefficients, as follows. Expand F in the translates of the normalized vacuum indexed by $(M'_p)^{-1}L/L$ in each coordinate, and for each translate apply the inverse translation followed by the projection to the vacuum. This projection averages the multiplier-normalized translations from L/M'_pL. Multiply its output by the value at zero of the corresponding translated vacuum and sum. The result is $F(0)$ times the normalized vacuum. Each value used as coefficient is integral by the established vacuum bound, and the averaging denominator is a power of M'_p. Every phase is a root of unity, and each shift in the formula is integral after multiplication by M'_p in the raw finite coordinates. This also proves the expansion with the density of rational finite translations used in the original construction.

A datum in Std_δ is an integral linear combination of finite translates of this normalized vacuum whose vectors become integral above p after multiplication by δ. For each term of the vacuum recovery choose a rational finite translation $\alpha^{-1}\gamma$ with the same value as the actual translating operator on F. This is possible by smoothness and weak approximation, while retaining integrality above p after multiplication by $M'_p\delta$. After multiplication by some integer n' prime to p, $n'M'_p\delta\gamma$ belongs to $\mathfrak{a}_0$. Under the rational Heisenberg formula the elliptic torsion argument is $-\Omega_0\,\mathrm{Tr}_{E/K}(\gamma\bar{w}_i)/M'_p$. Thus $n\delta y_{i,c}\in\Lambda$ with $n=n'(M'_p)^2$ prime to p. This works for a common n through each finite sum by taking a common multiple of the finitely many denominators prime to p. In the normalized jet of scale c_0 per, the linear coefficients after reading by j_p are

$$z_{\tau i} = j_p\left(\frac{c_0s_1}{M'_p}\overline{\Phi_\tau(w_i)}\right). \tag{7.7}$$

Here c_0 is the scalar in the kernel scale, so it is $\bar{\pi}_G$ in the chain of short lifts. In particular for the small non-split p prime to the twist we use $\delta=1$. Such p is prime to v and off g_0,g_1,N. Consequently s_1,c_0,M'_p are all p-adic units and (7.6) applies with a common exponent to the elliptic factors of each standard-data summand. This supplies the placed products and the coefficient $\beta = j_p(c_0s_1/M'_p)$ in the tensor calculation below; the common $1/n$ in the germ may be included in β.

For every p, $j_p(s_1\overline{\Phi_\tau(w_i)})$ is integral. To see this, the finite splitting-character valuation formula on α reads the product of the local absolute values of α through $j_p\Phi_\tau$ conj. This formula follows by taking a power of the ideal of α which is principal and using the infinity type on its generator. The absolute values of α in the product are all equal. They are one for $p\nmid v$, and for $p\mid v$ every weight reading lies on the

same chosen quadratic side and the relative different valuation is constant there. Since $s_0^f = \chi_{\mathrm{fin}}(\alpha)^{-1}$ and u_0 is a unit, $|j_p(s_1)|_p$ is the inverse of that common value. Now $w_i \in \mathfrak{a}_0$ proves the claim.

For the ordinary source-gauge bound, $d(d/dz)$ preserves the algebra $\mathcal{O}[[Y_0]][L', \theta]$. It sends L' to $d(L_2 + X_0)$ and θ to $-dL'\theta$, with $dL_2 \in \mathcal{O}$ by the logarithmic derivative calculation. Every member of the algebra has integral constant term. It follows from (7.5) that all derivatives of $(U_y\theta)(\bar{\delta}dz)$ at zero are integral when $\delta y \in \Lambda$. Replace δ by $n\delta$ for each product above, which changes no absolute value. The multivariable chain rule and (7.7) show that a homogeneous differential polynomial of degree B with integral coefficients has value bounded by

$$\left| j_p\left(\frac{c_0}{\bar{\delta}\, d} \right) \right|_p^B .$$

The integral Schur jet polynomials have this degree with B equal to the number of boxes below the target rank. This gives the bound on coefficient vectors at the identity.

For arbitrary source and target arguments, use rational unitary changes to make both arguments integral at all places above p. Weak approximation follows from the Cayley chart and approximation in its vector space. Rational equivariance transports the two Weyl gauges exactly. Integral unitary changes preserve Std_δ. Indeed their action on a vacuum is a sum of finite Heisenberg translates indexed by the quotient of the standard lattice by its intersection with the inverse image of that lattice. The coefficient is the finite Weil multiplier divided by its complex squared norm. The translating vectors are integral above p, and the coefficient is a p-adic unit since the local determinant factors of the integral automorphisms are units. The two additional character values in the paired action are units for the same reason. Heisenberg covariance carries the integral condition on δx through these changes. With a marked diagonal inserted on the target, the same calculation leaves the diagonal action on the target tableau and determinant character. Each tableau entry then contributes the absolute value of its diagonal coordinate, which is the charge calculation in the marked chain.

Away from N, the good datum Ψ_ϑ for a twist of level norm prime to p gives standard data with $\delta = \pi_G$ and $c_0 = \bar{\pi}_G$. The ratio per box is $|j_p(1/d)|_p$. At N, the last-column cut uses the same source-gauge calculation after clearing the last-column translation at the places over the quadratic place induced by j_N. Take $\delta = \omega - c$, where $c \in \mathbf{Z}$ has the same residue as ω at that place. This clears the simple poles there, and $\bar{\delta}$ is a unit at j_N because $\omega - \bar{\omega} = \sqrt{-d}$ and $N \nmid d$. Also $\bar{\pi}_G$ is a unit at j_N, so the derivative gauge has unit bound per box. Before that derivative bound, the residue-modulus averaging is cleared by $N^{2fn_jn_{j+1}}$. It contributes at most N^{fb} per lift. The subsequent paired Fourier move and its character cost have their separate bounds in the level-prime calculation, yielding the displayed N^{3mfb} bound under the starting gauge condition.

The tensor and Schur calculations for the charged jets. Use $v_p(p) = 1$ and the valuation ring $\mathcal{O}$ of $\mathbf{C}_p$. Every small non-split prime in this application is at least 13, since $2, 3, 5, 7, 11$ split. The normalized elliptic factors have the common exponent and residue shape of (7.6), with $v_p(\varpi) = \rho_p$, $v_p(\xi) = \omega_p$, and $v_p(\alpha_p) = 1/(p-1) - \omega_p$. The placed product gives the coefficients (7.7). Include the common unit change of variable from the torsion denominator in β. The frame constant and the torsion denominator are prime to p in this case.

For one product let $L_{i,c} = \sum_\sigma z_{\sigma i} X_{\sigma,c}$ be its linear arguments, with c a matrix coordinate. The frame identity gives $\sum_i \Phi_\tau(w_i)\overline{\Phi_{\sigma'}(w_i)} = l\, 1_{\tau=\sigma'}$, where l is the tight-frame constant. In the residue calculation the $z_{\sigma i}$ are a common integral scalar β times $j_p(\overline{\Phi_\sigma(w_i)})$. Any cyclotomic denominator of w_i divides a power of N_1 and is a unit here. Raising $j_p(\overline{\Phi_\sigma(w_i)})$ to the p^e-th power modulo ξ gives $j_p(\Phi_{[-p^e]\sigma}(w_i))$ for $e \in \mathscr{O}_p^\circ$. For inert p this uses the odd Frobenius power to cancel conjugation on K. At $p = d$, the two embeddings of K have equal reductions modulo ξ because $v_p(\sqrt{-d}) = v_p(\xi) = 1/2$. On the cyclotomic

factor the statement is the Frobenius action on roots of unity of order prime to p. These readings give, modulo ξ,

$$\sum_{i,c} L_{i,c}^{p^e} L_{i,c} = l\beta^{p^e+1} \sum_{\sigma,c} X_{\sigma,c}^{p^e} X_{[-p^e]\sigma,c}.$$

In particular the matrix coordinate c agrees at both ends of a coupling. Substitutions with zero constant term that kill the couplings on the right through degree B_j also kill the exponent of the product modulo ξ in these degrees. To see this, iterate $S = X + S^{p^k}$ in the truncated ring to write $S(L) = \sum_{r\geq 0} L^{p^{kr}}$. Each term in the sums of $S(L)^{p+1}$ and $S(L)^{p+p^k}$ is a Frobenius power of $\sum_{i,c} L_{i,c}^{p^e} L_{i,c}$ with a live e. When $k = 2$ the two indices of its Frobenius powers differ by an odd integer. When $k = 1$ their difference can be zero, and the term $u_0 L^2$ also uses $e = 0$. Terms with $e > H_p$ have degree greater than B_j. Because $v_p(\alpha_p \xi) = 1/(p-1)$ and $v_p(r!) \leq r/(p-1)$, the exponential after any such substitution has integral coefficients through degree B_j. The integral factors and their common unit rescaling preserve this property. It passes to the expansion of a standard datum by the ultrametric inequality.

For the Schur calculation use $n = n_j$, $n' = n_{j+1}$, and let $\kappa_{\sigma,h}$ be the number of columns of height $h < n'$ in the jet. Fix a target tableau η and integral n-row matrices Y_σ. The flag test is the product $P_Y(Z)$ of the minors of $Z_\sigma Y_\sigma$ on the rows of the target columns and the first h columns. The coefficient of a source tableau in the expansion of this product in the source flag basis is the corresponding Schur polynomial. This expansion also includes the common determinant power at the source when full source columns occur. Its coefficient estimates follow from the estimates of $P_Y(\partial)\Theta(0)$ for all integral Y. Indeed, the flag basis polynomials reduce to linearly independent polynomials over the residue field. With a diagonal leading-term order their leading monomials record the multiset in each row of the standard tableau, which determines its filling. Multiplying by a common determinant power preserves independence. If a coefficient had smaller valuation than the claimed bound, division by a coefficient of minimal valuation would give a nonzero polynomial over the infinite residue field, with a nonzero value at the reduction of some integral Y. This would contradict the bound for that flag test.

Write $y_{\sigma,a} \in \mathcal{O}^n$ for column a of Y_σ. We can change the first h directions at (σ, h) by an integral h-by-h matrix $M_{\sigma,h}$ at controlled determinant cost. A congruence on each changed row of a matrix, given by a linear form on that row with modulus $q_0 \in \mathcal{O}$, costs at most $v_p(q_0)$. For one form pick a row with a nonzero value of minimal valuation, subtract multiples of it from the other rows, then multiply the chosen row by q_0. If all values vanish the identity suffices. Previously imposed congruences survive integral row operations.

For a live e, the desired pairing of directions u, v at $\sigma, \sigma' = [-p^e]\sigma$ is $\sum_a u_a^{p^e} v_a \equiv 0 \pmod{\xi}$. When the restrictions are put on u, use linear forms with p^e-th roots of the coefficients and modulus ξ^{1/p^e}. Taking the p^e-th power modulo ξ recovers the pairing. When the restrictions are put on v, use modulus ξ. The following choices cover every pair of heights except those with first height one and either $\sigma \in I$ or both second height one and $\sigma' \in I$.

- For first height one outside I, either require its sole change coefficient to be a multiple of ξ^{1/p^e} or impose one pairing congruence on each relevant partner height against the initial height-one source direction. The two costs are t_σ/p^e and $r_{j,\sigma'} + t_{\sigma'}$.

- For first height different from one and second height one, put one congruence on the first height against the initial height-one direction of the partner. This costs $r_{j,\sigma}/p^e$.

- For both heights different from one, clear all the coefficients in each row of the change matrix on one side. Take the smaller of the two costs $B_{j,\sigma}/p^e$ and $B_{j,\sigma'}$.

The costs in this list are in units of $v_p(\xi) = \omega_p$ and count each change matrix once for every column of its height. In the last case clearing the coefficients means requiring each entry of the matrix row to be divisible by the relevant modulus. In the other cases where a pairing is imposed, the opposite height-one

change is a scalar times its initial direction. All these congruences can be imposed together using the same matrices, so that

$$\sum_{\sigma,h} \kappa_{\sigma,h}\, v_p(\det M_{\sigma,h}) \le \omega_p \sum_{\substack{e\in\mathscr{E}_p \\ e\le H_p}} \sum_{\sigma} K_{\sigma,e}.$$

For an exempt pair, use the prescribed charges $g_{\sigma,h}(b)$ with a positive charge at one end for each common matrix row b. In that row take independent variables for the changed directions at all heights, multiplying the direction from height h by $p^{g_{\sigma,h}(b)}$. The substitution made from these rows kills the live couplings modulo ξ. For a nonexempt pair the direction pairing vanishes by construction. For an exempt pair it is multiplied by at least one factor of p. Apply the preceding product-integrality argument to this substitution and the common factor ϖ. Applying the product of determinant derivatives in the new variables multiplies the original flag test by $\varpi^{B_j} p^{c(\eta)} \prod_{\sigma,h} \det(M_{\sigma,h})^{\kappa_{\sigma,h}}$, where $c(\eta)$ is the sum of charges in the target tableau. This follows by the chain rule and multilinearity of each determinant in the h direction slots. The number of variables is at most fn'^3, since it is the sum over $\sigma, h < n'$ of $n'h$. Every derivative monomial has degree B_j. Its factorial multiplier has valuation at least $B_j/(p-1) - fn'^3(H_p+1)$, by the base-p digit formula for each factorial. Subtracting the valuations of the scaling factors proves that every jet entry has valuation at least

$$(1/(p-1) - \rho_p)B_j - \omega_p \sum_{\substack{e\in\mathscr{E}_p \\ e\le H_p}} \sum_{\sigma} K_{\sigma,e} - fn'^3(H_p+1) - c(\eta) = J_{p,j}(I) - c(\eta).$$

For the final equality, if p is inert then $1/(p-1) - 1/(p^2-1) = \sum_{e>0,\ e \text{ odd}} p^{-e}$. If $p = d$, then $1/(p-1) - \rho_p = \frac{1}{2}\sum_{e>0} p^{-e}$.

For the nonheight-one part of the exponent, the difference $B_{j,\sigma}/p^e - \min(B_{j,\sigma}/p^e, B_{j,\sigma'})$ is bounded below by $\kappa(w(\sigma)/p^e - w(\sigma'))_+ - x^{60}$. If $\sigma \in \mathcal{T} \setminus I$ and $\sigma' \notin \mathcal{T} \setminus I$, the first minimum in $K_{\sigma,e}$ similarly costs at most x^{60} beyond the ask for these height-one boxes. If both are tops outside I, charging a/p^e covers the whole shortfall from this ask. If $\sigma \in \mathcal{T} \cap I$, the first minimum is zero and the height-one part contributes all of a/p^e. For $e = 1$, its excess over $(a/p - w(\sigma'))_+$ is at least a/x_1. Outside the tops the height-one part is zero.

The remaining term $r_{j,\sigma}/p^e$ in the coupling count is at most $(a/x + x^{60})/p$. The asks vanish beyond H_p since $B_j \ge a > x_1$. The tail of the first sum in $J_{p,j}(I)$ is nonnegative. Finally $H_p + 1 = O(\log b)$ uniformly in p, since $\log B_j = O(\log b)$ and $p \ge 2$. The per-place error per exponent, including the factorial rounding penalty, is bounded by

$$2x^{60} + (a/x + x^{60})/p + n_{j+1}^3 = O(a/(pm))$$

uniformly for $p < x_1$. Here $x_1 = b^{o(1)}$ and $a \ge xb$. At $p = d$ the zeroth coupling comes from the quadratic term in the ramified germ of (7.6). The positive ask begins at exponent one. For $\sigma' = [-1]\sigma$, at most one place is thick and at most one is a top, since their respective classes are -1 and 1 modulo M. The first minimum is at most $r_{j,\sigma'}$ if $t_\sigma \ne 0$, and the last minimum is bounded by the number of boxes at the thin place. It follows that

$$K_{\sigma,0} \le 2(a/x + x^{60}) + \kappa a/x_1 + x^{60} = o(a/(dm)).$$

Here d is fixed, $m \le 2\sqrt{x}$, and x_1 dominates every fixed power of x. Summing over $e \le H_p$ absorbs this error in $fa(\log b)^2/(pm)$ for large x. We obtain

$$J_{p,j}(I) \ \ge\ \frac{\gamma_p}{m} - \mathrm{coll}_p(I) + \omega_p \#(\mathcal{T} \cap I)\frac{a}{x_1} - \frac{fa(\log b)^2}{pm}. \tag{7.8}$$

7.3 The marked operators and relative index

Fix a coefficient embedding $j_p : \mathbf{C} \hookrightarrow \mathbf{C}_p$, a twist of level norm prime to p, and the good datum of the same seesaw. Use M_p for the module of columns integral at j_p. For $p \ne N$, use the source gauge calculation

on the translations allowed in Ψ_ϑ. It gives the factor $|j_p(1/d)|_p$ per box, which is one when $p \neq d$. The final restriction preserves the integral gauge off N. the branching matrix is integral and its extra scalar is a unit. This proves the required bound at primes outside $\{N\} \cup \{p < x_1 \,:\, \chi_d(p) \neq 1\}$, since the fixed d lies below x_1 for large x.

At a small non-split prime, using $I = \varnothing$ in (7.8) and the plain class sum gives

$$v_{j_p}\big(B^{\mu}_{\vartheta,\Psi}(\phi)\big) \;\geq\; \gamma_p - \ell_p - m\,\mathrm{coll}_p(\varnothing) \qquad (\phi \in M_p \cap C_{\mathrm{ram}}).$$

Here the valuation of a form denotes the lower bound for all its integral-gauge coordinates at every argument. At an unserved prime this is at least $\gamma_p - 2\ell_p$, with no index cost. At any small non-split prime the trivial scalar correction p^{s_p}, where $s_p = \lceil 2mfa/p \rceil$, gives valuation at least $\gamma_p - \ell_p$ on the image of $p^{s_p} M_p$, by the geometric-series bound for unrestricted collisions. The estimates before restricting to Y_m apply to every integral column, whether or not it is ramified. This lets us add the two images below before restricting.

For p served, use the marked set chosen above and put

$$A_\lambda = a\#\{\tau \in \mathcal{T} \,:\, \lambda(\tau) = \lambda\}, \qquad \Delta_p = c_p f(n_m + c_p), \qquad Y_\lambda = p^{A_\lambda} T_{\lambda, c_p}.$$

Here $T_{\lambda,i}$ is the spherical operator for the double coset of $\mathrm{diag}(pI_i, I)$ at λ. It acts at a split place of E over p, away from the chosen and level places. The marked place is spherical for each stage kernel because its vacuum and characters are unramified, and it lies away from the twist primes.

The split Hecke polynomial on the spherical kernel. The kernel identity here holds on the whole spherical vector space. We compute it by counting minuscule modifications over the residue field. Let $w = \lambda$ be the marked split place with residue order $q = q_\lambda$, and let $n < n + k$ be the source and target ranks. Give each maximal compact subgroup volume one. For rank v, write

$$P_{v,w}(X) = \sum_{i=0}^{v} (-X)^i q^{\binom{i}{2}} T_{w,i}.$$

The double coset for $T_{w,i}$ consists of the minuscule modifications of index q^i. A representative g_U can be chosen integral with valuation of determinant i and reduction kernel U, as U runs through the i-planes of the residue vector space. To see the coset parametrization, lift a basis adapted to U and multiply the U-coordinates by a uniformizer. The lattice $g_U^{-1}\mathcal{O}_w^v$ is the inverse image of U in $\varpi_w^{-1}\mathcal{O}_w^v/\mathcal{O}_w^v$. It determines the left compact coset, and g_U^{t} gives the corresponding right compact coset.

Use the split matrix realization of the same local Weil kernel, with matrix variable A of size $(n+k)$-by-n. The vacuum has integral support in this matrix variable. The action on the target replaces A^{t} by $A^{\mathrm{t}} g_U^{-1}$ and multiplies the integral-box indicator by $\alpha^{\dim U}$. On the source at the conjugate place it replaces A by $A g_U^{-1}$ and multiplies by $\zeta^{\dim U}$. Here $\alpha = q^{n/2}\chi_S(\varpi_w)$ and $\zeta = q^{(n+k)/2}\chi_T(\varpi_w)$. These scalars are the splitting characters times $|\det g_U|^{-n/2}$ and $|\det g_U|^{-(n+k)/2}$, respectively. In each case the translated indicator is supported on integral A with U in the kernel of the corresponding residue matrix. This follows by testing the linear map on $g_U^{-1}\mathcal{O}_w^v$.

For a residue matrix of rank r, summing the two Hecke polynomials therefore gives $\prod_{i<n+k-r}(1 - q^i\alpha X)$ on the target and $\prod_{i<n-r}(1 - q^i\zeta X)$ on the conjugate source. Indeed the coefficient of $(-Y)^i$ in $\prod_{a<D}(1 - q^a Y)$ is $q^{\binom{i}{2}}$ times the number of i-planes in a D-space. The identity follows by fixing a hyperplane and counting i-planes lying in it and those meeting it in an $(i-1)$-plane, with q^{D-i} extensions for each such $(i-1)$-plane. Factoring the first k terms separates the two polynomials at each residue matrix. In the theta pairing, change variables by each source representative to move its inverse action on the kernel to the source vector. The source double coset at the conjugate place is the inverse double coset at w, since the two actions on the split unitary group are inverse transpose. In the resulting paired terms g_U can be replaced by g_U^{t}: they lie in the same double coset, the source vector is fixed by the compact subgroup on

the right, and the lift is invariant under its action on the left by the source vacuum. The resulting equality for the actual stage lift L and every spherical vector u is

$$P_{n+k,w}(X)Lu = \prod_{i<k}(1 - q^i\alpha X)\, L\big(P_{n,w}(q^k\alpha\zeta^{-1}X)u\big).$$

The datum and the lifted vector are fixed by the relevant compact subgroups because the characters are unramified there.

In the marked chain the rank twists give the same base-change character χ_{e+1} on the two sides of arrow e. Its value at a uniformizer has p-adic valuation zero. For this assertion, raise the prime to a power principal in the ray class group for the character conductor. The character identity on its generator computes the valuation, divided by that power, as the signed sum of its weight over the two halves of the marked pair. The bijection by u_p preserves the weights and gives zero. Finite-order factors also have valuation zero. It follows that $q^k\alpha/\zeta = q^{k/2}$. For the chain C_{j+1} of the first $j+1$ lifts use the variable $X = -Uq^{(n_{j+1}-n_e)/2}$ at stage e. The coefficient of U^i at that stage is

$$y_e(i)T_{\lambda,i}, \qquad y_e(i) = q^{(n_{j+1}-n_e)i/2+\binom{i}{2}}.$$

Iterating the polynomial equality gives $P_{n_{j+1},w}(-U)C_{j+1} = G(U)C_{j+1}P_{n_0,w}(-Uq^{(n_{j+1}-n_0)/2})$. The product G contains the factors

$$1 + Uq^{i+(n_e+n_{j+1}-n_{e+1})/2}\chi_{e+1}(\varpi_w), \qquad 0 \le e \le j, \quad 0 \le i < n_{e+1} - n_e.$$

It has integral coefficients and constant term one. The recurrence for its reciprocal gives integral coefficients. Extracting degree c_p from $G(U)^{-1}P_{n_{j+1},w}(-U)$ transfers the operator $y_0(c_p)T_{\lambda,c_p}$ across the chain. With $q = p^{2h_\lambda}$, one has $h_\lambda \le f$ by $[E : \mathbf{Q}] = 2f$, and

$$v_p(y_0(c_p)) = h_\lambda\big((n_{j+1} - n_0)c_p + c_p(c_p - 1)\big) \le \Delta_p.$$

Multiplying the extracted identity by $p^{\Delta_p}/y_0(c_p)$ keeps the coefficients integral. They depend on the ranks and the chain characters at λ, which agree for every term of Ψ_ϑ.

Before moving the operators across the chain, consider a Hecke translate of a stage datum at each $\lambda \in \Lambda$, of index at most c_p. Its diagonal representative has p in a set J_λ of at most c_p coordinates and 1 in the others. Right and left compact factors preserve the gauges. In the transported gauge the two corresponding powers are reciprocal at paired embeddings. For $\sigma \in I$, charge each height-one box whose index is outside $J_{\lambda(\sigma)}$. At the conjugate half charge every box with index in the same $J_{\lambda(\sigma)}$. These charges cover each exempted coupling in the jet estimate since live partners induce conjugate primes. At the marked half, the diagonal change of gauge pays one power of p for each box in J_λ. For height one this fills the complement of the charged boxes, costing a total of A_λ at that prime. For the other columns it costs at most c_p per column because a tableau column has distinct entries. At the conjugate half the reciprocal gauge cancels the charges in J_λ. The determinant exponents cancel between the two halves by the same balancing of stage weights. Thus a single step increases the gauge by at most

$$p^{\sum_{\lambda\in\Lambda} A_\lambda + c_p \sum_{\sigma\in I} r_{j,\sigma} - J_{p,j}(I)}.$$

This uses the local covariance of the normalized Schur kernel in the transported Weyl gauge.

The coefficient extraction permits the operator Y_λ to be moved from the initial column across the first $j+1$ lifts to a modification of the j-th datum. In that modification replace the λ-translate by a finite linear combination of translates of index at most c_p and multiply by $p^{A_\lambda - \Delta_p}$. Its coefficients are integral. Thus the gauge cost A_λ in the last display is replaced by Δ_p. For distinct marked primes, the operators commute with each other and the modifications off their own places. Their target Hecke operators commute with the level actions on the vacuum data, so the modified joint sum still has the plain class-sum estimate.

Equivariance of the lift moves the modified target action from the kernel into the target form. With $Y = \prod_{\lambda\in\Lambda} Y_\lambda$, induction on the chain length now shows that applying all the modifications, one per stage, gives the original chain applied to $Y^m\phi$. The resulting total exponent of the gauge is at most

$$\sum_{j<m}\left(\#\Lambda\,\Delta_p + c_p\sum_{\sigma\in I} r_{j,\sigma} - J_{p,j}(I)\right).$$

To check its sign, all $\sigma \in I$ are thin. If $p < m$, then $c_p(a/x_1 + x^{60}) = O(a/(pm))$ uniformly because x_1 dominates every fixed power of x. If $p \ge m$, then $c_p = 1$ and these σ are tops. In both cases $\#\Lambda\,\Delta_p = O(fa/(pm))$ uniformly, by the polynomial bounds on f, n_m, c_p and $x_1 = b^{o(1)}$. Therefore

$$\#\Lambda\,\Delta_p + c_p\sum_{\sigma\in I} r_{j,\sigma} \;\le\; \frac{\ell_p}{m} + \#(\mathcal{T}\cap I)\frac{a}{x_1}$$

for all large x. Here $\omega_p = 1$ and $\mathrm{coll}_p(I) = 0$. Substituting (7.8) proves that the output on $Y^m M_p$ has valuation at least $\gamma_p - 2\ell_p$.

Recombination and the served-prime index. We use the following elementary determinant rules, valid over any valuation ring. A recombination of a tuple is a matrix with entries in that ring acting on the tuple, with cost $v_p(\det)$. If M is a finite free module and A is an integral square matrix invertible over the fraction field, any tuple in M can be recombined into AM at cost at most $v_p(\det A)$. To see this, put the tuple in the columns of X. Perform elementary row and column operations over the valuation ring to diagonalize $A^{-1}X$, after initially multiplying by a common denominator. Such diagonalization follows by moving an entry of least finite valuation to the first position, using it to clear its column and row, and inducting. Scale the columns at each negative-valuation diagonal entry by its inverse. The sum of the negative valuations in absolute value equals the negative valuation of the corresponding maximal minor in the diagonal form. All minors of the original matrix have valuation at least $-v_p(\det A)$, by complementary minors for A^{-1} and Cauchy–Binet. The same holds after integral row and column operations. This proves the cost bound. The argument uses only finite matrices, including when $\mathcal{O}_{j_p}$ is nondiscrete.

If X and Y preserve M, costs for recombining into $XM + p^sM$ and $YM + p^sM$ add for $XYM + p^sM$. Indeed first recombine into $x_i = Xu_i + p^sv_i$. Next perform the second recombination on the tuple u_i and apply that same matrix to the tuple x_i. The new error is in p^sM, and the determinants multiply. Similarly one can pull back a recombination by evaluating a linear map on a tuple and using its recombination matrix on the original tuple.

The marked operator in the actual column coordinates. Fix a served p, a coefficient embedding j_p, and its marked set Λ. Use $k = \kappa$, $r_p = c_p$, and $C = \dim_{\mathbf{C}} \mathcal{C}$. For $\lambda \in \Lambda$, let L_λ be the length of the top cycle in $\{\lambda, \bar\lambda\}$. This length is the order of $u_p^{o_p}$, where $u_p = [-p]$, so $\mathrm{ord}(u_p) = o_pL_\lambda$. It is even, since m is a power of two, $m \mid \mathrm{ord}(u_p)$ and o_p is odd. Write $\tau_i = u_p^{o_p i}\tau_0$, with τ_0 inducing λ. Read $W_\lambda = \mathcal{O}_{E,\lambda}$ through

$$g_i = j_p\Phi_{\tau_i}\quad (i\ \text{even}), \qquad g_i = j_p\Phi_{\tau_i}\,\mathrm{conj}\quad (i\ \text{odd}).$$

The full arithmetic Frobenius at this inert prime is $\Phi_{u_p}\circ\mathrm{conj}$ in these coordinates. It acts by $[p]$ on the cyclotomic factor and by conjugation on K. Since o_p is odd, the step from g_i to g_{i+1} reads its o_p-th power. Put $P = p^{o_p}$ and $q_\lambda = P^{L_\lambda}$. The embeddings satisfy $g_{i+1}(a) \equiv g_i(a)^P \bmod p$ for $a \in W_\lambda$. This is congruence modulo p in $\mathcal{O}_{j_p}$, since p is unramified in E and the congruence before the embedding is in pW_λ. Also $\#(W_\lambda/pW_\lambda) = q_\lambda$. Indeed $4 \mid \mathrm{ord}(u_p)$ here, so the order of $\Phi_{u_p}\circ\mathrm{conj}$ is $\mathrm{ord}(u_p)$. The two marked halves have $L_\lambda/2$ letters each, so $A_\lambda = aL_\lambda/2$.

Let X be the definite class set at the stage-zero level and $K_\lambda = \mathrm{GL}_k(W_\lambda)$. The group K_λ may be used to choose representatives of the finite parahoric cosets in the completed local group. Put

$$d_\lambda = \mathrm{diag}(pI_{r_p}, I), \qquad \mathscr{P} = K_\lambda \cap d_\lambda K_\lambda d_\lambda^{-1}.$$

Choose the class representatives ρ_x integral above p, using rational unitary changes as in the source-gauge calculation. Their components at $\lambda, \bar{\lambda}$ can be made the identity by a right level element supported on this split pair. They give M_p by arbitrary integral Weyl values on X, since the class stabilizers are trivial and the level is integral above p. The level contains the supported maximal compact at $\lambda, \bar{\lambda}$, as p is away from the chosen and level primes. Haar mass one gives the Hecke action as a sum over the right compact cosets of $K_\lambda d_\lambda K_\lambda$, indexed by $b \in K_\lambda/\mathscr{P}$.

An edge $e = (x, b)$ replaces the λ-component of ρ_x by $b\,d_\lambda$, with conjugate inverse transpose at $\bar{\lambda}$ and no change elsewhere. Let this new adele be $\gamma_e \rho_{o(e)} u_e$, with γ_e rational unitary and u_e in the stage-zero level. At λ this gives $\gamma_e^{-1} = h_e d_\lambda^{-1} b^{-1}$ for $h_e \in K_\lambda$. At the other places over p outside the paired places, the rational part is integral invertible. Both edge parametrizations used below are bijective. One is the original (x, b). For the other, suppose two edges have the same $o(e)$ and the same coset $(h_e^{\mathrm{t}})^{-1}\mathscr{P}$. Then $d_\lambda h_e^{-1} h_{e'} d_\lambda^{-1}$ is integral invertible, by the definition of $\mathscr{P}$ and transposition. Hence $\gamma_e \gamma_{e'}^{-1}$ is integral invertible at the split pair. Away from the pair its comparison between ρ_x and $\rho_{x'}$ is in the level because the origins are equal. The supported maximal compact puts the comparison in the level everywhere. Thus $x = x'$, and freeness gives $\gamma_e \gamma_{e'}^{-1} = 1$. The equation at the split pair then gives $b^{-1} b' \in \mathscr{P}$. This proves injectivity and the two finite sets have equal size.

Let J index the degree-a monomials in k variables in each letter, and let H index the Weyl tableaux outside the cycle. The integral stage-zero Weyl basis consists of these $J \times H$ indices. For the letters define the polynomial substitution action ρ by the symmetric powers of $g_i(h)$ at even i and of $g_i(h)^{-\mathrm{t}}$ at odd i. In the monomial substitution matrix the entry at (α, β) is the coefficient of z^α in the substitution of z^β by the columns of the base matrix. Thus the coefficient action on the Weyl evaluation vector is the transpose of the substitution matrix for the transposed base matrix. This is a bijection of integral tableau and monomial indices, using the height-one columns, and requires no change by reciprocal multinomial coefficients.

Let $\mathrm{dep}(\beta)$ be the degree in the variables outside the first r_p at even letters and in the first r_p at odd letters. Set $D_p = \mathrm{diag}(p^{\mathrm{dep}(\beta)})$. The summand of $Y_\lambda = p^{A_\lambda} T_{\lambda, c_p}$ along e, on the cycle and the other tableaux respectively, is

$$\big(\rho(h_e) D_p \rho(b^{-1})\big)^{\mathrm{t}} \otimes U_e, \qquad U_e \in \mathrm{GL}_{|H|}(\mathcal{O}_{j_p}).$$

Indeed at an even letter, transpose followed by conjugation of γ_e gives γ_e^{-1}. The p per box changes d_λ^{-1} to $p d_\lambda^{-1}$. At an odd letter the transpose reads γ_e^{t} through the redirected embedding and keeps d_λ. The off-cycle factors are integral invertible because their top readings see only places outside the pair, and the other stage-zero factors are trivial. The formula proves that $Y_\lambda M_p \subseteq M_p$ and identifies the marked operator on all the columns.

The truncated colength of the marked operator. The colength calculation uses the residue depth sum and the adjoint pairing below. Put

$$\mathfrak{c} = L_\lambda a \left(P^{-2r_p} + \frac{r_p^2}{kP} \right), \qquad c_0 = k^2 o_p L_\lambda, \qquad A_0' = L_\lambda c_0.$$

Write $\mathcal{O} = \mathcal{O}_{j_p}$ in this paragraph, choose $w \in \mathcal{O}$ with $w^2 = p$, and put $\Delta = \mathrm{diag}(w^{\mathrm{dep}(\beta)})$. Let M_* be the integral space of edge coefficients v for which $\Delta \rho(g) \Delta^{-1} v$ is integral for every $g \in K_\lambda$. Over the residue field, let $\mathcal{D}_j$ be the space of polynomials in the specified degrees all of whose $\rho(g)$-translates have zero coefficients in depths below j.

The depth sum for the marked residue cycle. For the served cycle $0 < 2r_p \le k$. Use the residue field $\mathbb{F}$ of $\mathcal{O}_{j_p}$ and the marked residue field $\mathbf{F}_{q_\lambda}$, with embeddings $\sigma_i = \overline{g_i}$ and $\sigma_{i+1} = \sigma_i^P$. Here and below the latter notation means raising each output to the P-th power. Write $R = r_p$ within this paragraph, and

$\ell = L_\lambda$, so $q_\lambda = P^\ell$ and ℓ is even. Use variable indices $0 \le b < k$ in each letter i mod ℓ. Call $v = (b, i)$ deep if $b \ge R$ at even i or $b < R$ at odd i. Give a deep variable the weight $1/\kappa_v$, where

$$\kappa_v = \begin{cases} P & (b < 2R), \\ P^{2R} & (b \ge 2R), \end{cases}$$

and give every other variable weight zero. Denote the set of degree-a-in-each-letter monomials by J. We prove that if $0 \ne F \in \mathcal{D}_j$, then F contains a monomial of weight at least j.

Replace all shallow variables by independent indeterminates $u_{b,i}$. The resulting polynomial f in the deep variables is nonzero over $L = \mathbb{F}(u_{b,i})$. Let a matrix $A(y)$ over $\mathbf{F}_{q_\lambda}$ have arbitrary entries y_{cb} in the positions

$$(c, c+R), \quad (c, b) \text{ with } 2R \le b < k, \qquad 0 \le c < R,$$

and zero entries elsewhere. Choose lifts with this same support over W_λ, so their squares vanish and $g = I + A(y)$ is invertible there. The corresponding action of g^{-1} on a deep variable, after specializing the shallow variables, is $Z_v \mapsto Z_v - z_v(y)$ with

$$z_{b,i}(y) = \sum_{c<R} u_{c,i}\sigma_i(y_{cb}) \qquad (i \text{ even},\ b \ge R),$$
$$z_{c,i}(y) = -\sum_{b\ge R} u_{b,i}\sigma_i(y_{cb}) \quad (i \text{ odd},\ c < R).$$

Unspecified y_{cb} in the sums are zero. Indeed the base matrix of g^{-1} in the substitution is $I - A(y)$ at even letters and $I + A(y)^{\mathrm{t}}$ at odd letters. All the shifted coefficients involve only shallow variables. Since gF has depth at least j and $F = g^{-1}(gF)$, the polynomial f belongs to the j-th power of the point ideal of $z(y)$ in $L[Z_v]$ for every such y.

For each letter i and distant column $b \ge 2R$ use the R integers $0 \le h < 2R$ for which $i - h$ is even. Frobenius gives the simultaneous equations

$$z_{b,i-h}(y)^{P^h} = \sum_{c<R} u_{c,i-h}^{P^h}\sigma_i(y_{cb}).$$

Their coefficient matrix is invertible over L. A monomial in its determinant has for each c exactly one factor $u_{c,i-h}^{P^h}$ and therefore determines h from its exponent, even if some $i - h$ coincide modulo ℓ. Distinct permutations give distinct monomials. For the paired position $(c, c+R)$ at even i, divide $z_{c+R,i}(y)$ by $u_{c,i}$. At odd i, subtract the already recovered distant entries from the equation for $z_{c,i}(y)$ and divide by $-u_{c+R,i}$. In both cases the coefficient is nonzero. It follows that every $\sigma_i(y_{cb})$ is the evaluation at $z(y)$ of a polynomial with coefficients in L and all monomial weights below one. The powers used for distant columns have weight $P^h/P^{2R} < 1$, and for the paired positions the degree-one terms have weight $1/P$. In particular the points $z(y)$ are distinct, since the entries of y can be read at any fixed i.

For each deep variable $v = (b, i)$, raising $z_v(y)$ to the power $\kappa_v = P^h$ gives a linear combination of the entries at letter $i + h$, with coefficients in L. Therefore for each v there is a polynomial H_v all of whose monomials have weight below one such that

$$G_v = Z_v^{\kappa_v} - H_v(Z)$$

vanishes at all these points. There are $q_\lambda^{R(k-2R+1)}$ points. The product of the κ_v is exactly this number. In each even letter there are R near and $k - 2R$ distant deep variables, and in each odd letter there are R near and no distant deep variables. Their product thus has exponent $\ell R(k - 2R + 1)$ in P.

For $j > 0$, impose the jets of order below j at all these points. The Chinese remainder theorem gives a surjective jet map, because the j-th powers of distinct point ideals are comaximal. Reduce products $Z^\alpha G^\gamma$ by substituting $Z_v^{\kappa_v} = G_v + H_v$ whenever $\alpha_v \ge \kappa_v$. In each expanded term the weight of the part Z^α strictly

decreases, and $|\gamma| + \sum_v \alpha_v/\kappa_v$ never increases. The reduction terminates because the weights are rational with fixed denominators. Modulo the common jet kernel the products with $0 \le \alpha_v < \kappa_v$ and $|\gamma| < j$ therefore span. Their number is $\prod_v \kappa_v$ times the number of monomials of degree below j in the deep variables, exactly the dimension of the jet target. Thus they are independent there. If every monomial of f had weight below j, reduction would express f itself as a linear combination of those products. Since its jets vanish, every coefficient in that expression would be zero, which contradicts $f \neq 0$. The case $j = 0$ is automatic. Specialization preserved the weights of the deep parts, so the costly monomial in f comes from one in F.

For each j, projection to the coefficients of monomials of weight at least j is therefore injective on $\mathcal{D}_j$. After summing these dimension bounds for $1 \le j \le M$, each monomial is counted at most its weight. The mean exponent in any variable is a/k by symmetry within its letter. Thus

$$\sum_{j=1}^{M} \dim_{\mathbb{F}} \mathcal{D}_j \le |J| \frac{a}{k} \left(\frac{\ell R}{P} + \frac{\ell(k-2R)}{2P^{2R}} \right) \le |J|\, \ell a \left(P^{-2R} + \frac{R^2}{kP} \right).$$

This is the required $\mathfrak{c}|J|$, on the residue representation of the same integral columns.

For the adjoint pairing take the diagonal matrix $B_{\beta\beta} = p^{A_0'} / \prod_i \binom{a}{\beta_i}$ on $\mathcal{O}^J$, where $\binom{a}{\beta_i} = a!/\prod_b \beta_{bi}!$. It has integral complementary diagonal B' with $BB' = p^{A_0'} I$. Indeed $a < q_\lambda$ and Legendre's formula give $v_p(\binom{a}{\beta_i}) \le k \log_p q_\lambda \le c_0$, by bounding the carry at each power of p by k. The tensor symmetric pairing gives $B\rho(g^{\mathrm{t}}) = \rho(g)^{\mathrm{t}} B$, by realizing polynomials in each letter as symmetric tensors with the induced bilinear form. The monomials in that form have squared weights $\binom{a}{\beta_i}^{-1}$.

The edge average at the served prime. Use the marked Frobenius cycle, source gauge and monomial pairing of the marked-colength calculation. For this paragraph put

$$L = L_\lambda, \quad n = a, \quad r = r_p, \quad \ell_0 = L/2, \quad F = o_p L, \quad P = p^{o_p}, \quad q = P^L = p^F,$$
$$A_0' = Lc_0, \quad c_0 = k^2 F, \quad c = 5A_0', \quad T_0 = 3kc.$$

Here ℓ_0 is a number of letters. The residue field of the marked ring W_λ is $\mathbf{F}_q$. The maps $\iota_s : W_\lambda \to \mathcal{O} = \mathcal{O}_{j_p}$ reduce to embeddings of $\mathbf{F}_q$ in the residue field $\mathbf{k}$ of $\mathcal{O}$, and $\iota_{s+1} \equiv \iota_s^P \bmod p\mathcal{O}$. The action ρ on the degree-n monomials in each letter uses $\iota_s(g)$ at even letters and $\iota_s(g)^{-\mathrm{t}}$ at odd letters. Let G be the block diagonal subgroup of K_λ with blocks of sizes $r, k-r$. We use the integral pairing $\langle v', v\rangle = v'^{\mathrm{t}} B v$ of that calculation, so $\rho(g^{\mathrm{t}})$ is adjoint to $\rho(g)$. The depth dep, $\Delta = \mathrm{diag}(w^{\mathrm{dep}(\beta)})$, $w^2 = p$, and $M_* = \{v : \Delta\rho(g)\Delta^{-1} v \in \mathcal{O}^J \text{ for all } g \in K_\lambda\}$ use the same marked reading.

For an adjacent root $(b, b+1)$ and a letter i define the linear defect on real exponent vectors by

$$d_{i,b}(\beta) = \beta_{u_i,i} + \frac{\sum_s P^{[s-i]_L} \epsilon_s (\beta_{b+1,s} - \beta_{b,s})}{q-1}, \qquad \epsilon_s = (-1)^s, \qquad u_i = \begin{cases} b+1 & i \text{ even}, \\ b & i \text{ odd}. \end{cases}$$

The notation $[s-i]_L$ uses the representative in $[0, L-1]$. Put $d_i = d_{i,r-1}$. Define $\mathscr{C}$ by requiring all coefficients of $\rho(g)v$ at $d_i(\beta) > 3T_0$, for all $g \in G$ and i, to lie in (w^c). Define $\mathscr{D}$ over $\mathbf{k}$ by the same requirements with zero as the coefficient ideal. We will construct an integral matrix V and prove

$$M_* + w^c \mathcal{O}^J \subseteq \mathscr{C}, \qquad V\mathcal{O}^J \subseteq \mathscr{C},$$
$$\langle v', Vv\rangle \equiv p^{A_0'} \langle v', v\rangle \pmod{w^c} \quad (v' \in M_*,\ v \in \mathcal{O}^J),$$
$$w^{c-1} v \in \mathscr{C} \implies \bar{v} \in \mathscr{D}, \qquad [K_\lambda : \mathscr{P}] \dim_{\mathbf{k}} \mathscr{D} \le (1+\theta_0)|J|.$$

These are the simultaneous inputs to the marked-edge block argument at the selected s_p. The size bounds used below are

$$Lkn(\ell_0 + 2) + n + k < P^{\ell_0}, \qquad 2(k+1) \le T_0, \qquad 8LkT_0P^{2r} \le \theta_0 n.$$

For the served primes, $P^{\ell_0} = q^{1/2} \geq 13^{m/2}$ dominates a times the polynomial factors in x. Also T_0 is polynomial in x, $P^{2r} = b^{o(1)}$, $\theta_0^{-1} = m^2 f$ is polynomial in x, and $a \geq xb$. These give the displayed bounds uniformly using the same scale as the marked-colength calculation.

The root filtration on the marked columns. Use the same marked ring W_λ, coefficient ring $\mathcal{O}$, P, $q = P^L$, and embeddings ι_s. We use that $P \geq 5$ and that these embeddings satisfy the Frobenius congruences modulo $p\mathcal{O}$, where $p = w^2$. For coefficients in $\mathcal{O}$ write $\mathrm{ord}_w(c)$ for the greatest integer e for which $c \in (w^e)$, with order $+\infty$ at zero. Equivalently $\mathrm{ord}_w(c) = \lfloor 2v_p(c) \rfloor$ for $c \neq 0$. For the residue-field argument take the same substitutions reduced to $\mathbf{k}$, set $w = 0$, and give each nonzero coefficient order zero. For variable weights $h(z)$, a polynomial has level at least B if each coefficient at z^α has order at least $B - h \cdot \alpha$. Sums preserve a common lower level and products add the levels. In particular a substitution sending each variable to a polynomial of at least its weight preserves levels. If a second such substitution differs on each variable by terms of strictly greater weight, its difference on a polynomial of level B has level strictly greater than B. For this last assertion, expand the difference of each monomial as a sum with at least one changed factor. All sums are finite.

Fix a root $(b, b+1)$ and letter i. Denote the two variables at s by X_s, Y_s, in that order by increasing row at even letters and by decreasing row at odd letters. Put $P_s = P^{\lceil s-i \rceil_L}$. For an exponent array α write

$$A_\infty(\alpha) = q \sum_s P_s \alpha_{X_s}, \qquad A_0(\alpha) = q\alpha_{Y_i} + \sum_{s \neq i} P_s \alpha_{Y_s}.$$

The defect $d_{i,b}(\alpha)$ of the marked edge is $(A_0(\alpha) - A_\infty(\alpha)/q)/(q-1)$. Suppose the increments $\delta_s = h(Y_s) - h(X_s)$ are in $[1/2, 1]$. Suppose also that every translate of F by the $\mathrm{SL}_2(W_\lambda)$ on this root has level at least $e + h \cdot \alpha$, where e is a nonnegative integer and $F_\alpha \notin (w^{e+1})$. We show $d_{i,b}(\alpha) \leq 0$.

Put $\zeta = \delta_i / q > 0$ and $\eta_s = -1$ when s has the parity of i, and $+1$ otherwise. Make the substitution $\mathscr{L}$ over $\mathcal{O}[t]$, as follows:

$$X_i \mapsto t^q X_i, \quad Y_i \mapsto X_i,$$
$$X_s \mapsto t^{qP_s} X_s - \eta_s t^{P_s} Y_s, \quad Y_s \mapsto Y_s - \eta_s X_s \quad (s \neq i).$$

The other variables are fixed. In $\mathscr{L}F$ take the coefficient of the monomial obtained from z^α by merging Y_i into X_i; reduce this polynomial in t modulo w^{e+1} and denote it by $\Phi(t)$.

For the expansion at infinity give t weight $-\zeta$ and X_s weight $h(X_s) + qP_s\zeta$, keeping the other weights. The terms $t^{qP_s} X_s$ and the unaltered variables, with $Y_i \mapsto X_i$, preserve the original weights. The two cross terms at $s \neq i$ have excesses respectively $\delta_s - P_s\zeta \geq 1/2 - 1/P > 0$ and $qP_s\zeta - \delta_s \geq P/2 - 1 > 0$. At the merged monomial, the weight of the term with t-degree j is $h \cdot \alpha + \zeta(A_\infty(\alpha) - j)$, because the new weights of X_i, Y_i agree. The level lower bound $e + h \cdot \alpha$ kills coefficients with $j > A_\infty(\alpha)$ modulo w^{e+1}. At $j = A_\infty(\alpha)$ the cross terms also vanish modulo w^{e+1}. The merged exponent and this t-degree uniquely recover α in the substitution keeping only the weight-preserving terms. Indeed the only forgotten exponents were the division between X_i and Y_i, and the exponent of X_i is recovered from the t-degree after subtracting the already fixed contributions from $s \neq i$. The leading coefficient is therefore $F_\alpha \bmod w^{e+1}$ and $\deg \Phi = A_\infty(\alpha)$.

For the expansion at a finite point take $x \in W_\lambda$, $u = \iota_i(x)$ and $u_s = \iota_s(x)$. Give t weight ζ, keep the weights of X_s, and subtract $q\zeta$ from $h(Y_i)$ and $P_s\zeta$ from $h(Y_s)$ at $s \neq i$. In the root rows take $\left(\begin{smallmatrix} 0 & 1 \\ -1 & x' \end{smallmatrix}\right)$, with $x' = -x$ for even i and $x' = x$ for odd i. It sends X_s to Y_s and Y_s to $X_s + \eta_s u_s Y_s$, up to the sign -1 on each input variable with row index b. This follows by taking the columns of the matrix in the even letters and of the inverse transpose in the odd letters. Apply $\mathscr{L}$ after this root action and shift t to $t + u$. For X_s the result is the signed $\mathscr{L}(Y_s)$. For Y_s at $s \neq i$ it is the signed $\mathscr{L}(X_s)$ plus the two signed errors

$$\big((t+u)^{qP_s} - t^{qP_s} - u_s\big) X_s, \qquad -\eta_s \big((t+u)^{P_s} - t^{P_s} - u_s\big) Y_s.$$

At i only the first error occurs, with $P_i = 1$ and $u_i = u$. The two errors follow by expanding $\mathscr{L}(X_s)(t + u) + \eta_s u_s \mathscr{L}(Y_s)(t+u)$; the constant terms cancel with these signs. Every parenthesis has coefficients in (p) by the Frobenius congruences, including for the exponent qP_s since the cycle returns after L steps.

In these finite-point weights $\mathscr{L}(Y_s)$ has weight at least $h(X_s)$, since $\delta_s - P_s\zeta > 0$ for $s \neq i$, and $\mathscr{L}(X_s)$ has weight at least $h(Y_s)$ by $qP_s\zeta \geq \delta_s$ for $s \neq i$. At i the two weights are exact. Each error has coefficient order at least two, while the finite-point t-weight is nonnegative. For the X_s-error use $h(X_s) + 2 \geq h(Y_s)$, and for the Y_s-error use $2 \geq P_s\zeta$. Thus acting, substituting and shifting is level-preserving. Apply this to the inverse root translate of F, whose level is at least $e+h\cdot\alpha$ by hypothesis. It gives that level for $(\mathscr{L}F)(t+u)$. At the merged monomial the term of t-degree j weighs $h \cdot \alpha + \zeta(j - A_0(\alpha))$, since the adjusted weights of X_i, Y_i again agree. The coefficients with $j < A_0(\alpha)$ vanish modulo w^{e+1}. Hence $(t-u)^{A_0(\alpha)}$ divides Φ.

Choose one x from each of the q residue classes of W_λ. Differences of distinct u's are units and remain units in the coefficient quotient. These powers of linear polynomials are pairwise comaximal, so their product divides Φ. The divisor is monic, so $qA_0(\alpha) \leq \deg\Phi = A_\infty(\alpha)$, including over the possibly nonreduced quotient. This proves the required extreme-term inequality for the original exponent and either coefficient ring.

For the between-block root, the depth weight rises by one from X_s to Y_s. For each i replace it by $\mathrm{dep} - d_{i,r-1}/3$, with $d_{i,r-1}$ here viewed as a weight on individual variables. The rise of the defect weight on this pair is $1+2/(q-1)$ at i and $2P_s/(q-1)$ at $s \neq i$. Both lie in $[0, 3/2]$, so the new increments lie in $[1/2, 1]$. If all root translates have original depth level at least u, they also have new level at least u. Otherwise choose a term of least new level below u over the root orbit. The minimum exists because support degrees are bounded and finite coefficient orders are nonnegative integers. Its defect is nonpositive by the extreme-term inequality, so its original depth level is at most its new level, a contradiction. In particular, $d_{i,r-1}(\beta) \geq 3j$ forces the depth-scaled coefficient into (w^{u+j}) for integers $u, j \geq 0$, with the same integer exponent when $\mathcal{O}$ is nondiscrete. For a within-block root in the residue field, the signed height rises by one from X_s to Y_s. The same extreme-term inequality with $e = 0$ applies to a term of least signed height over the root orbit. These are the two root inputs used in halfway divisibility and the Levi flag selection below.

Halfway divisibility. For the root $(r-1, r)$, the deep indicator rises by one in the root direction at each letter. Use the root filtration on this marked action.

For $v' \in M_*$ choose an integer u large enough that $H = w^u\Delta^{-1}v'$ is integral. For every $g \in K_\lambda$, the coefficients of $\Delta\rho(g)H$ are divisible by w^u. Apply the estimate to $\rho(g)H$ and all its root-group translates. Since Δ commutes with $\rho(g)$ for $g \in G$, cancelling w^u gives

$$d_i(\beta) \geq 3c \quad \Longrightarrow \quad (\rho(g)v')_\beta \in (w^c) \qquad (g \in G).$$

In particular $M_* \subseteq \mathscr{C}$. For $g \in G$, the adjoint relation and $g^{\mathrm{t}} \in G$ give

$$d_i(\beta) > T_0 \quad \Longrightarrow \quad \langle v', \rho(g)e_\beta\rangle = B_{\beta\beta}(\rho(g^{\mathrm{t}})v')_\beta \in (w^c).$$

The first containment for $w^c\mathcal{O}^J$ follows from the definition of $\mathscr{C}$. If $w^{c-1}v \in \mathscr{C}$, then cancelling w^{c-1} puts each tested coefficient of v in (w), which reduces to zero in $\mathbf{k}$. This gives the claimed implication to $\mathscr{D}$.

An integral functional for one matrix. Let $\xi = e_r - e_{r-1}$ on the row indices and put

$$D = Lkn(\ell_0 + 2)P^{\ell_0}, \qquad H_0 = 2T_0P^{\ell_0}.$$

Then $D + k < q$. In the j-th copy of a generic matrix X, let $g \in G$ act by $X \mapsto X^g = \iota_j(g^{-1})X\iota_j(g)$. The row and column weights of X^α are $\sum_{b,d}\alpha_{bd}\xi_b$ and $\sum_{b,d}\alpha_{bd}\xi_d$. We construct I_j on polynomials of degree at most D with the following properties. It is invariant, vanishes on monomials of row weight $> H_0$, and agrees with $E_j = p^{c_0}\,\mathrm{ev}_1$ modulo (w^c) on a polynomial whose entire G-orbit has those coefficients in (w^c).

Fix a homogeneous polynomial ψ of degree $d \leq D$. The diagonal monomial pairing on the entries of X has coefficient $\eta_\alpha = p^{c_0}/\binom{d}{\alpha}$ at X^α. It is integral. For $d < p^F$, Legendre's formula bounds each floor-function contribution to the valuation of the multinomial by $k^2 - 1$, and the contributions at powers at least p^F vanish. Units in the denominator are invertible in $\mathcal{O}$. After restricting X to the two diagonal

blocks, form an operator $D_d(\psi)$ on $U_d = (\mathcal{O}^k)^{\otimes d}$ by pairing ψ with the entries of $(X^{\mathsf{t}})^{\otimes d}$. Polynomials with an off-block variable contribute zero. The diagonal pairing in degree d is proportional over the fraction field to the factorial pairing $(X^{\alpha}, X^{\alpha}) = \prod_{b,e} \alpha_{be}!$. Applying a linear substitution to the first argument therefore applies its transpose to the second, as follows by the constant-coefficient differential operator pairing at zero. For $X \mapsto \iota_j(g^{-1}) X \iota_j(g)$ the transpose substitution is $X \mapsto \iota_j(g)^{-\mathsf{t}} X \iota_j(g)^{\mathsf{t}}$. Consequently, writing $g_j = \iota_j(g)$, we have

$$D_d(\psi(X^g)) = g_j^{\otimes d} D_d(\psi)(g_j^{-1})^{\otimes d}, \qquad \operatorname{tr}_{U_d} D_d(\psi) = p^{c_0}\psi(1).$$

For the trace, a diagonal monomial in X of degree d occurs on $\binom{d}{\alpha}$ diagonal tensor entries, giving the stated scalar.

The coefficient matrices of the block-restricted tensor polynomial $X^{\otimes d}$ lie in the $\mathcal{O}$-span of $g_j^{\otimes d}$ for $g \in G$. Indeed, evaluation vectors of the degree-d block monomials at the invertible block matrices over $\mathbf{F}_q$ span over $\mathbf{k}$. An annihilator would give a polynomial vanishing on those invertible matrices. Multiplying by the product of the two determinants gives a polynomial vanishing on the entire $\mathbf{F}_q$ grid, of degree at most $d + k < q$. Interpolation one variable at a time makes that polynomial zero, and cancellation gives a zero annihilator. Choosing a residue evaluation basis and lifting its points to G yields a matrix with unit determinant over $\mathcal{O}$. Inverting this matrix expresses the tensor coefficient matrices in the indicated span. The same span contains the coefficients of $(X^{\mathsf{t}})^{\otimes d}$, by transposing the matrices used for evaluation.

Let $U_{d,+}$ be generated by the G-translates of those basis tensors $e_{\mathfrak{b}}$ with $\sum_{\nu=1}^{d} \xi_{b_\nu} > H_0$. The span result shows finite generation, since all translates are linear combinations of the finitely many applications of coefficient matrices to these tensors. It also shows that $D_d(\psi)$ preserves $U_{d,+}$. A finitely generated torsion-free module over a valuation domain is free. In this case one can prove it by eliminating from the generator matrix a nonzero entry dividing all its entries, then continuing on the remaining block. Use the trace difference

$$I_j(\psi) = \sum_{0 \le d \le D} \left(\operatorname{tr}_{U_d} D_d(\psi_d) - \operatorname{tr}_{U_{d,+}} D_d(\psi_d) \right),$$

where ψ_d is the homogeneous part of ψ. Conjugation invariance follows from the operator conjugation identity and stability of $U_{d,+}$. For a monomial of row weight $> H_0$ the operator is zero on every low-weight basis tensor. The transpose makes each matrix input index of $(X^{\mathsf{t}})^{\otimes d}$ a row index of the corresponding monomial in X. The nonzero columns map to $U_{d,+}$ since their inputs are among its generators and the operator preserves it. The two traces are equal by trace cyclicity for the maps through $U_{d,+}$.

Suppose next that all conjugates of ψ have high-row coefficients in (w^c). On a high-weight basis tensor $D_d(\psi_d(X^g))$ uses only the high-row part of $\psi_d(X^g)$. Dividing that part by w^c gives a polynomial whose operator preserves $U_{d,+}$. The conjugation identity, with g or g^{-1}, now gives $D_d(\psi_d) U_{d,+} \subseteq w^c U_{d,+}$ on the generating translates. This proves divisibility of the restricted trace even if the lattice $U_{d,+}$ is unsaturated. It proves $I_j(\psi) \equiv E_j(\psi) \pmod{w^c}$. Finally both I_j and E_j kill monomials whose row and column weights differ. Conjugation by the diagonal $\left((1+p)^{\xi_b}\right)_b \in G$ acts on the monomial by $(1+p)$ to the power of the column minus the row weight. No nonzero power of $1+p$ equals 1 in characteristic zero.

Contraction at the marked precision. Take L matrices X_s with independent entries. The substitution matrix $U(X)$ on the monomial indices is the tensor product of the symmetric powers of the following letter matrices. In letter s form the products of $X_{s-d}^{[p^d]}$, $0 \le d < m_s$, in descending order of d, where brackets denote entrywise powers. Take the row at b from the product of length $m_s = \ell_0 + \operatorname{deep}(b, s)$. For odd s, take this row from the transposed adjugate of the product of that length. Thus $U(1, \dots, 1) = 1$. Each entry has degree at most D in each copy. A factor has degree at most P^{ℓ_0}, there are at most $\ell_0 + 1$ factors in a row product and an adjugate entry multiplies at most $k - 1$ entries. The symmetric powers and tensor product then give the degree bound.

For $g \in G$ write $X^g = (X_j^g)_j$. The matrix

$$U_g(X) = \rho(g)\, U(X^g)\, \rho(g^{-1})$$

has the same description by generic factors, with each factor replaced by $\iota_s(g)(X_{s-d}^g)^{[P^d]}\iota_s(g^{-1})$. For the verification, conjugate the full row-length products and use that the two row-length projectors commute with the block diagonal matrices. At odd letters the transposed adjugate changes the conjugating matrix to its inverse transpose. This follows from the polynomial adjugate identity, so it holds for generic factors.

We prove the coefficient estimate for this U_g. Fix a copy j and use the row vector $\xi = e_r - e_{r-1}$. For an exponent ν in the copies X_s, write $R(\nu_s)$ and $C(\nu_s)$ for its row and column weights against ξ. Assign the entry (b, b') of X_u the weight

$$\omega_{u,b,b'} = \frac{\epsilon}{P^{\ell_0}} \begin{cases} \xi_b, & u = j, \\ P^{[u-j]_L}(\xi_b - \xi_{b'})/q, & u \neq j, \end{cases} \qquad \epsilon \in \left\{ \tfrac{1}{2}, -\tfrac{1}{2} \right\}.$$

Use the coefficient order ord_w above. The weight of X^ν is $\epsilon W_j(\nu)$, where

$$W_j(\nu) = P^{-\ell_0}\left(R(\nu_j) + \sum_{u \neq j} \frac{P^{[u-j]_L}}{q}(R(\nu_u) - C(\nu_u)) \right).$$

Consider the generic factor $\iota_s(g)(X_{s-d}^g)^{[P^d]}\iota_s(g)^{-1}$ at letter s, with $0 \le d \le \ell_0$. Entrywise P^d-th power is a ring map modulo p, and it carries the readings of g and g^{-1} at $s-d$ to those at s. The two conjugations therefore cancel modulo p. The entry (b, b') is congruent to $X_{s-d,b,b'}^{P^d}$ modulo p as a polynomial, and all its monomials have degree P^d. Since every $|\omega_{u,b,b'}| \le P^{-\ell_0}$, their weights lie in $[-1, 1]$. The error coefficients have order at least two, so the entry has level at least $P^d\omega_{s-d,b,b'}$, the level of the pure power.

Put $u_0 = 1/(2P^{\ell_0})$, and define the real numbers

$$\lambda_s(d, b) = \frac{\epsilon\xi_b}{P^{\ell_0}} \frac{P^{d+[s-d-j]_L}}{q}.$$

For any s, d these sum to zero over b, and $|\lambda_s(0, b)| \le u_0$. The difference $\lambda_s(d+1, b) - \lambda_s(d, b')$ agrees with $P^d\omega_{s-d,b,b'}$ unless $s - d = j$. In fact if $[s - d - j]_L > 0$, decrementing that representative by one leaves the exponent $d + [s - d - j]_L$ unchanged. If $s - d = j$, advancing d replaces the representative zero by $L - 1$. The difference then equals $P^d\omega_{j,b,b'}$ minus $\epsilon\xi_{b'}P^d/(P^{\ell_0}q)$, whose absolute value as an error is at most u_0. Thus the factor entry has level at least $\lambda_s(d+1, b) - \lambda_s(d, b') - u_0$.

Use any fixed allowable length $m_s \le \ell_0 + 1$ for a letter product. Along each matrix product summand these potential differences telescope. Every entry in row b has level at least $\lambda_s(m_s, b) - (m_s + 1)u_0$. A transposed adjugate entry in row b is a signed sum of products taking one entry from every row other than b. The sum of their potentials is $-\lambda_s(m_s, b)$. The lower level for a row chosen from either the ordinary product or the transposed adjugate is consequently

$$(-1)^s\lambda_s(m_s, b) - k(\ell_0 + 2)u_0.$$

For odd letters this calculation uses the full product at that fixed length before selecting the adjugate row.

For the prescribed length $m_s = \ell_0 + \mathrm{deep}(b, s)$, the first term of this level is the variable weight in $\epsilon(1 - 1/q)d_{j+\ell_0}$. To see this, both are zero off the two root rows. On the shallow row the exponent contributes $P^{[s-j-\ell_0]_L}/q$ with a negative sign. On the deep row the power in the numerator is $P^{1+[s-j-\ell_0-1]_L}$: it equals $P^{[s-j-\ell_0]_L}$ off letter $j + \ell_0$, and q at that letter. This is exactly the extra term in the defect formula.

In the degree-n symmetric power, each coefficient at a row exponent uses the row levels with the multiplicities of that exponent. This follows by expanding the substitution of each input monomial as a product of linear forms in the output variables. Multiplying over the letters gives

$$\mathrm{ord}_w\left([X^\nu](U_g)_{\beta\gamma}\right) + \epsilon W_j(\nu) \;\ge\; \epsilon(1 - 1/q)d_{j+\ell_0}(\beta) - \frac{Lnk(\ell_0 + 2)}{2P^{\ell_0}}.$$

The scale condition bounds the last quotient by $1/2$. If the coefficient survives modulo w^c, its order is less than c. Use both signs of ϵ to get $|(1 - 1/q)d_{j+\ell_0}(\beta) - W_j(\nu)| \le 2c + 1$. There are n variables counted with

multiplicity in each letter, so $|d_{j+\ell_0}(\beta)| \le 2n$ by the geometric sum in the defect formula. Since $2n \le q$, this proves

$$[X^{\nu}](U_g)_{\beta\gamma} \notin (w^c) \quad \Longrightarrow \quad \left|d_{j+\ell_0}(\beta) - W_j(\nu)\right| \le 2c + 2 \le T_0 = 3kc.$$

All the coefficients here use the original monomials and the alternating action in the marked edge calculation.

The conjugated factors and the same signed bound are recorded in the fixed proof.

Contract each copy of $U(X)$ by I_j to obtain V. For the range and pairing, choose a copy j. Contract the other copies by $L_u \in \{I_u, E_u\}$ and denote by $F_{j,g}(X_j)$ the vector obtained from $U_g(X)\rho(g)v$. The invariance of the other L_u gives

$$F_{j,1}(X_j^g) = \rho(g^{-1})F_{j,g}(X_j).$$

If $[X_j^{\alpha}]F_{j,g,\beta}$ survives modulo w^c, at least one contributing monomial of U_g survives and has $R(\nu_u) = C(\nu_u)$ for every contracted copy $u \ne j$. The coefficient estimate therefore reads

$$\left|d_{j+\ell_0}(\beta) - P^{-\ell_0}R(\alpha)\right| \le T_0.$$

With all other $L_u = I_u$, the coefficient $(\rho(g)Vv)_{\beta}$ is $I_j(F_{j,g,\beta})$. If $d_{j+\ell_0}(\beta) > 3T_0$, I_j kills all surviving monomials, since their row weights exceed $2T_0P^{\ell_0} = H_0$. This proves $V\mathcal{O}^J \subseteq \mathscr{C}$.

For the pairing put $\psi = \langle v', F_{j,1}\rangle$. In $\psi(X_j^g)$, the coefficient at a monomial of row weight $> H_0$ is the sum over β of $[X_j^{\alpha}]F_{j,g,\beta}\langle v', \rho(g^{-1})e_{\beta}\rangle$. Each summand whose first factor survives has $d_{j+\ell_0}(\beta) > T_0$. Its second factor is divisible by w^c by halfway divisibility. Thus $I_j(\psi) \equiv E_j(\psi) \pmod{w^c}$. Replace the E_j's by I_j's one copy at a time. The initial all-E_j contraction is $p^{Lc_0}I$, which proves the pairing congruence with the stated $p^{A_0'}$.

Packing the small-defect residue space. Work over $\mathbf{k}$. Let $M = q^j$ for a positive integer j. For each array u of residues in $[0, M-1]^{kL}$, take a form $f_u \in \mathscr{D}$. Choose a nonnegative $\eta(u) \equiv u \bmod M$ in each coordinate by adding $M(k - \lfloor \sum_b u_{b,s}/M \rfloor)$ at row $r-1$ in letter s. Its degree in that letter lies in $[Mk, M(k+1))$. Form the polynomial

$$f(z) = \sum_u z^{\eta(u)} f_u(z^M).$$

Packing is injective by the disjoint exponent classes modulo M. The reduced substitutions commute with $z \mapsto z^M$, since the entries of the substitution matrices lie in the embedded $\mathbf{F}_q$. A nonzero term of any Levi translate of f therefore has exponent $\alpha = M\beta + \gamma$, where β occurs in a Levi translate of some f_u and γ is a nonnegative exponent in the translate of a tag $z^{\eta(u)}$. Those tags retain their degrees in each letter. The defect formula gives $|d_i(\gamma)| \le 2M(k+1)$, by summing the geometric weights in the second term. Thus

$$d_i(\alpha/M) \le 3T_0 + 2(k+1) \le 4T_0, \qquad n + k \le \sum_b \alpha_{b,s}/M \le n + k + 1.$$

For a nonzero f, choose a nonzero term of least signed height among all Levi translates. The variable $z_{b,s}$ has height b at even letters and $-b$ at odd letters. Let $\mathscr{I}$ be the inverse image of the lower triangular subgroup modulo p. We can find this coefficient in a representative of $(G \cap \mathscr{I})\backslash G$. Indeed $y \in G \cap \mathscr{I}$ substitutes a diagonal unit times each variable plus variables of strictly higher height. Writing the chosen translate as yx with x a coset representative, the coefficient of least height can only come from the same term in the translate by x, with a nonzero diagonal multiplier. In the translate by x, for any adjacent root inside a block, all translates by its root-SL_2 have no coefficient at smaller height. The signed height rises by one along the root at each letter. The residue root estimate for a term of least height gives $d_{i,b}(\alpha) \le 0$ for every within-block root and all i. The between-block root already has $d_i(\alpha/M) \le 4T_0$.

Define the cyclic linear map

$$(A_1v)_{b,s} = \epsilon_s(Pv_{b,s+1} - v_{b,s})$$

on real arrays. Substituting in the defect formula gives $d_{i,b}(A_1 v) = (A_1 v)_{u_i,i} + v_{b+1,i} - v_{b,i}$. Indeed the geometric sum of successive differences telescopes with the wraparound factor $q - 1$. Let $\mu_b = P^{2(r-1-b)}$ for $b < r$, and $\mu_b = 0$ for $b \geq r$. Apply A_1 to the array which equals μ at even letters and $P\mu$ at odd letters. The resulting vector σ is $(P^2 - 1)\mu$ at even letters and zero at odd letters, so is nonnegative integral with degrees at most $P^{2r} - 1$ in each letter. The displayed defect identity gives $d_{i,b}(\sigma) \leq 0$ for every root. For $b = r - 1$ it gives -1 at even letters and $-P$ at odd letters. For $b < r - 1$, the even defect is zero since $\mu_b = P^2 \mu_{b+1}$, and the odd defect is nonpositive. Both defects vanish when $b \geq r$.

Put

$$D_s = n + k + 1 + 4T_0 \sum_b \sigma_{b,s}, \qquad E_0 = 4T_0 P^{2r},$$
$$\Omega = \{\beta \geq 0 \,:\, d_{i,b}(\beta) \leq 0 \text{ for all } i, b\},$$
$$\mathcal{B} = \{\beta \geq 0 \,:\, D_s - 1 \leq \sum_b \beta_{b,s} \leq D_s \text{ for all } s\}.$$

All $D_s \leq n + E_0$, and $\alpha/M + 4T_0\sigma$ belongs to $\Omega \cap \mathcal{B}$. For each of the chosen Levi flag representatives and each point of this intersection on $M^{-1}\mathbf{Z}^{kL}$, record the coefficient of the translated packed polynomial at the exponent $M(\beta - 4T_0\sigma)$. Use zero if this array has a negative coordinate. The least-height argument proves injectivity of this linear coefficient map. The domain has dimension $M^{kL} \dim_{\mathbf{k}} \mathscr{D}$. As $M = q^j \to \infty$, the grid count divided by M^{kL} converges to the volume, since the region is a bounded polytope and its boundary has measure zero. Consequently

$$\dim_{\mathbf{k}} \mathscr{D} \leq [G \,:\, G \cap \mathscr{I}]\, \mathrm{vol}(\Omega \cap \mathcal{B}).$$

The flag factor in the volume. Let $\mathcal{F}_{k,q} = \prod_{j=1}^{k} \frac{q^j - 1}{q - 1}$. We show $\mathcal{F}_{k,q}\, \mathrm{vol}(\Omega \cap \mathcal{B}) = \mathrm{vol}(\mathcal{B})$. From the defect identity, $A_1 v \in \Omega$ is equivalent to interlacing at every link. At even s, the arrays $x_b = P v_{b,s+1}$ and $y_b = v_{b,s}$ satisfy $x_b \geq y_b$ and $y_b \geq x_{b+1}$ for consecutive indices. At odd s use $x_b = v_{b,s}$ and $y_b = P v_{b,s+1}$ with the same inequalities. Either link forces its two arrays to be in decreasing order.

Away from the finitely many tie hyperplanes, define $M_s(b, d) = \mathbf{1}_{\epsilon_s(P v_{b,s+1} - v_{d,s}) > 0}$. For decreasing arrays x, y, the matrix with entries $\mathbf{1}_{x_b > y_d}$ has determinant one exactly when $x_b > y_b > x_{b+1}$, with only $x_b > y_b$ at the last row, and determinant zero otherwise. To see this, each row is a suffix of ones and the suffixes are nested. A nonzero determinant forces their k lengths to be distinct and nonzero, hence $k, k - 1, \ldots, 1$, and then the matrix is upper triangular with diagonal one. This applies by transposition to an odd link. When the rows of v are independently reordered in each letter, the sign from that reordering occurs once as a row sign and once as a column sign in $\prod_s \det M_s$, so cancels. After sorting, the product is the indicator that all interlacings hold. There is precisely one ordering giving interlacing when that indicator is one. Therefore $\prod_s \det M_s$ equals the sum of the indicators that $A_1 v$ satisfies the interlacing after the independent letter reorderings.

For a tuple of permutations $\tau = (\tau_s)_s$, set

$$(A_\tau v)_{b,s} = \epsilon_s (P v_{b,s+1} - v_{\tau_s(b),s}).$$

Expanding the determinant product expresses it as the signed sum of $\mathbf{1}_{A_\tau v > 0}$ over all τ, with sign $\prod_s \mathrm{sgn}(\tau_s)$. The bounds on the sums of the letters are the same for every τ and for every reordering of the rows within letters. Multiply by the indicator of these bounds and integrate. The interlacing interpretation gives $(k!)^L\, \mathrm{vol}(\Omega \cap \mathcal{B})/|\det A_1|$, whereas the expanded determinants give $\sum_\tau \mathrm{sgn}(\tau)\, \mathrm{vol}(\mathcal{B})/|\det A_\tau|$.

To compute these Jacobians, let D_ϵ multiply by ϵ_s in letter s, let $(Sv)_{b,s} = v_{b,s+1}$, and let $(R_\tau v)_{b,s} = v_{\tau_s(b),s}$. Then

$$A_\tau = P D_\epsilon S (I - P^{-1} H), \qquad H = S^{-1} R_\tau.$$

The unsigned permutation H shifts around the cycle and permutes the rows. Its return permutation on one letter is a product π of the τ_s, up to the convention for representing permutation matrices. A cycle of length j of π corresponds to a cycle of length Lj of H. Thus

$$|\det A_\tau| = P^{kL} |\det(I - q^{-1}\pi)|.$$

The signs of the tuples equal the signs of their return permutations. Each π occurs $(k!)^{L-1}$ times. The two integrals show that the ratio $\operatorname{vol}(\Omega \cap \mathcal{B})/\operatorname{vol}(\mathcal{B})$ depends only on k, q and is the same as for one even letter with step q.

In that one-letter calculation the cone condition on $\beta = A_1 v$ is $\beta_b \geq q\beta_{b+1}$. The triangular linear map

$$y_b = (1 + q + \cdots + q^b)(\beta_b - q\beta_{b+1}), \qquad \beta_k = 0$$

maps the cone bijectively to the nonnegative orthant and preserves the sum of the coordinates by telescoping. Its determinant is $\mathcal{F}_{k,q}$. It therefore gives the asserted volume ratio on a slab of unit width.

Modulo p, $\mathscr{P}$ is the lower two-block parabolic and $\mathscr{I}$ the lower Borel. The quotient of the two-block parabolic by its full unipotent radical is the product of the two general linear groups. Hence choices of a parabolic flag and of the full flags within the two blocks give a full flag, so $[K_\lambda : \mathscr{P}][G : G \cap \mathscr{I}] = \mathcal{F}_{k,q}$. The slab volume is at most $\prod_s \left(D_s^{k-1}/(k-1)!\right)$, by the volumes of the positive simplices with degree bounds D_s and $D_s - 1$. Since $2LkE_0 \leq \theta_0 n$, it follows that

$$\begin{aligned}
[K_\lambda : \mathscr{P}] \dim_{\mathbf{k}} \mathscr{D} &\leq \operatorname{vol}(\mathcal{B}) \\
&\leq \exp(L(k-1)E_0/n) \left(\frac{n^{k-1}}{(k-1)!} \right)^L \\
&\leq (1 + \theta_0) \binom{n+k-1}{k-1}^L = (1 + \theta_0)|J|.
\end{aligned}$$

Here $e^u \leq 1 + 2u$ for $0 \leq u \leq 1/2$. This completes the construction and the dimension bound used in the integral edge block at s_p. The separate sum of the depth-space dimensions enters the two half-restriction determinants.

We give the passage from these estimates to the truncated determinant. On $\mathcal{O}^{H \times X \times J}$, form the two restriction matrices J_1, J_2. They send the value at the origin to $\Delta\rho(h_e^{\mathrm{t}})$ times that value on the forward edge, and the value at the terminus to $\Delta\rho(b^{-1})$ times that value on the reverse edge. They have full column rank. Write $J_i = Q_i d_i$ by Smith reduction, with Q_i having an integral left inverse and d_i square. Smith reduction here uses a least-valuation entry at each finite step. This works over the nondiscrete coefficient valuation ring.

For either end, every source and parahoric coset occurs. If $u \in \mathscr{P}$, then $\Delta\rho(u)\Delta^{-1}$ is integral invertible. Indeed the entries from a deep variable to a shallow variable are divisible by $p = w^2$ in each letter, balancing the one power of w in the half-depth scale. This is true for u^{-1} as well. Consequently a vector in the saturation of J_i has values in M_* in every edge coordinate. For the same reason, if t_v is a Smith divisor of J_i, the corresponding primitive source column has all normalized translated coefficients divisible by t_v. If $v_p(t_v) > (j-1)/2$, its reduction lies in $\mathcal{D}_j$ at each source. These columns are independent on reduction. Thus for each j their number is at most $|H||X| \dim \mathcal{D}_j$. Their valuations are bounded by half the maximum depth, by testing the identity translate on the primitive source column. Summing over j gives

$$v_p(\det d_i) \leq \frac{1}{2}\mathfrak{c}C, \qquad C = |H||X||J|.$$

On the edge coordinates let D_U act by U_e on H and by the identity on J. Extend B, B', V diagonally over the edges and the H-coordinates, and let R_2 be a left inverse of Q_2. The actual paired middle matrix and its near inverse on the edge module are obtained from

$$A = Q_2^{\mathrm{t}} D_U B, \qquad K_1 = V B' D_U^{-1} R_2^{\mathrm{t}}.$$

The image of K_1 has every value in $\mathscr{C}$. Since values of Q_2 lie in M_*, the averaging congruence gives $AK_1 \equiv p^{2A_0'} I \bmod w^c$. Here the U_e cancel because they act in H and commute with the operators on J.

Let $R_1 Q_1 = I$, put $\epsilon = w^{c-1}$, and apply Smith reduction to $(I - Q_1 R_1)K_1$ in the free complement of Q_1. Keep the primitive column for each divisor of valuation at most $v_p(\epsilon)$. Each such column multiplied by ϵ belongs to $\mathscr{C}$ edge by edge, by dividing the corresponding column image and using that the original image is in $\mathscr{C}$. The columns of Q_1 have the same property. All these primitive columns are independent modulo the maximal ideal. The residue-dimension bound shows that the number h of kept columns satisfies

$$C + h \le |H||X|[K_\lambda \,:\, \mathscr{P}] \dim \mathscr{D} \le (1 + \theta_0)C.$$

The other terms in the complementary image are divisible by ϵ. Apply A to the decomposition of K_1 into its projection on Q_1, the kept terms and the discarded terms. Since $v_p(\epsilon) > 2A_0'$, the result is

$$p^{2A_0'} \mathcal{O}^C \subset (AQ_1)\mathcal{O}^C + \operatorname{span}_{\mathcal{O}}\{h \text{ vectors}\}.$$

The inclusion follows because the matrix on the left after moving the error is $p^{2A_0'}$ times a matrix congruent to the identity modulo the maximal ideal.

If $p^u \mathcal{O}^C \subset B_0 \mathcal{O}^C + \operatorname{span}(h \text{ vectors})$, at most h Smith divisors of B_0 have valuation greater than u. To see this, diagonalize B_0 and project to those coordinates. In this block the inclusion, applied to each column of $p^u I$, gives $p^u I = D_0 L_0 + U_0 S_0$, with D_0 the diagonal block and U_0 the h-column generator matrix. The equality $p^{-u} U_0 S_0 = I - p^{-u} D_0 L_0$ makes the left side an integral invertible matrix, because $p^{-u} D_0$ is zero modulo the maximal ideal. Therefore the rank of U_0 is at least the size of the block. It follows that the truncated determinant cost of AQ_1 at any cutoff s is at most $2A_0' C + \theta_0 s C$. The costs of d_1 and d_2 add to it under multiplication, with the same s. The truncated valuation of each diagonal divisor t_ν is $\min(v_p(t_\nu), s)$, because the lattice is enlarged by $p^s \mathcal{O}^C$. The tuple recombination rule gives subadditivity.

Finally $d_2^{\mathrm{t}} A Q_1 d_1 = J_2^{\mathrm{t}} D_U B J_1$ is the actual edge matrix of Y_λ multiplied at the source by B. This follows edge by edge from the adjoint identity and $\Delta B \Delta = D_p B$. Its image is therefore contained in $Y_\lambda M_p$. We conclude that any integral tuple can be recombined into $Y_\lambda M_p + p^s M_p$ at determinant valuation at most

$$C(\mathfrak{c} + 2A_0' + \theta_0 s).$$

This recovers the truncated statement from the two prescribed estimates on the actual edge module.

Uniform sizes and the served-prime budget. Take $\theta_0 = 1/(m^2 f)$. The two size inequalities in the averaging estimate hold uniformly for served p. We have $L_\lambda \le f = O(x^{11} \log x)$, $k \asymp x^5$, $1 \le o_p \le r_p = O(\log x)$, and $2r_p \le k$. Also $p < x_1$, so

$$\log(P^{3r_p}) = O\big((\log x)^2 \log x_1\big) = o(\log b).$$

All the other factors of the degree lower bound are at most polynomial in x, including $\log q_\lambda$. Thus the required degree is $b^{o(1)}$, and $a \ge xb$. On the other side, $m \mid o_p L_\lambda$ gives $q_\lambda \ge 13^m$, which exceeds a^4 eventually. In particular A_0', which is bounded by a polynomial in x, is at most $L_\lambda a P^{-2r_p} \le \mathfrak{c}$ for large x. Since $p^{r_p} \ge \max(m, p)$, $P \ge p$ and $m r_p^2 / k = o(1)$, we get

$$3\mathfrak{c} \le \frac{6 L_\lambda a}{mp} = \frac{12 A_\lambda}{mp}.$$

The one-operator cost is bounded by $C(12A_\lambda/(mp) + s/(m^2 f))$. This is an absolute determinant cost for each tuple, with C the dimension of the complete stage-zero column space.

At $s = s_p = \lceil 2mfa/p \rceil$, recombination with products of the integral operators adds the costs. For $Y = \prod_{\lambda \in \Lambda} Y_\lambda$ there are m occurrences of every operator in Y^m. Since $\sum_\lambda A_\lambda \le fa$ and $\#\Lambda \le f$, the resulting cost at the original cutoff is no more than

$$C\left(\frac{12fa}{p} + \frac{s_p}{m}\right) \le C\left(\frac{14fa}{p} + \frac{1}{m}\right) \le C e_p.$$

Here $e_p = 3fa(\log b)^2/p$ at a served prime. The marked chain bounds the output on $Y^m M_p$ and the scalar correction bounds the output on $p^{s_p} M_p$. As the bound holds for every tuple, we may apply it to each assigned translate block of integral ramified columns and use linearity to retain ramification of the recombined vectors. The same q from $\beta_p = \lceil \gamma_p - e_p \rceil$ works on the two conjugate readings, using the respective marked sets. Thus the paired height cost is $F_p = 2|S|Ce_p \log p$ for the fixed selected rows. Its sum stays within the allowance for (7.10).

The index at the level prime. Use $p = N$ and a fixed coefficient embedding j_N. Write $\lambda(\tau)$ for the place induced by $j_N \Phi_\tau$ and let $P_N = \{\overline{\lambda(\tau)} : \tau \in \Phi\}$. Since N splits completely, P_N contains exactly one place of each conjugate pair. These are the places used to measure the coefficient action in the Weyl gauge, which is taken through $j_N \Phi_\tau$ conj. Embed the local general linear groups at P_N in the unitary group using inverse transpose at their conjugates. For $S \subseteq P_N$, let d_S be the resulting adelic diagonal with last entry N at S, all other entries 1. Put $d_0 = d_{\{w \in P_N \,:\, w|\pi\}}$ and $u_0 = d_{P_N}^{-1}$.

For an adelic diagonal d', let $M(d')$ consist of the columns whose Weyl gauge in the lattice at $g d'$ is at most one for every g. Thus $M(1) = M_N$. After conjugation by d_0, the local column level at P_N becomes the group of integral invertible matrices with last column congruent to the last standard basis vector. At $w \mid \pi$ the original condition is on the last row; conjugating by the last-coordinate diagonal with entry N converts it to the last-column condition. At $\bar{w} \mid \pi$ the condition is on the last row of the inverse transpose, which already says that the original matrix fixes this last column modulo N. The same block calculation shows that $(d_0 d_S^{-1})^{-1} K_0 (d_0 d_S^{-1})$ is integral above N for $S = P_N$ and for $S = \{w \in P_N \,:\, w \mid \pi\}$, where K_0 is the stage-zero column level.

For each fixed S use weak approximation to move every class representative c by a rational unitary matrix so that $l = c d_0 d_S^{-1}$ and $d_S^{-1} l d_S$ are both integral invertible above N. This is an open condition containing the identity on l. The required weak approximation follows from the Cayley map $A \mapsto (1 + A)(1-A)^{-1}$ on the space of skew-Hermitian matrices over E^+. Its image contains the dense open set of local unitary matrices with $1 + g$ invertible. Weak approximation in this vector space meets the finitely many prescribed local opens at the places over N. Now $M(d_0 d_S^{-1})$ is determined by ordinary integral values at these representatives, because the shifted arguments there are the integral points l. The gauge for $M(d_0)$ is the gauge at $c d_0 = l d_S = d_S (d_S^{-1} l d_S)$, hence agrees with the gauge at d_S on each representative. Testing only the representatives suffices since the rational actions cancel in transported gauges and $d'^{-1} k d'$ is integral invertible for $k \in K_0$. Thus in the same representative coordinates the two lattices are related by the rational diagonal action of d_S at the given coefficient embedding j_N. In the symmetric-power factors its exponents are the powers of the last coordinate in the monomials at tops measured in S. Permuting the κ coordinates shows that their mean is a/κ in each factor. The inclusion $M(d_0) \subset M(d_0 d_S^{-1})$ has determinant valuation at most

$$CTa/\kappa.$$

The finite-matrix argument above supplies the same bound for recombining any tuple from the larger lattice into $M(d_0)$.

In the conjugated local representation let $R_0(k)$ denote translation by $d_0 k d_0^{-1}$, using the indicated embedding of the local product group. The actual finite Fourier operator is

$$F = R_0(u_0) \sum_{\substack{y \in ((\mathbf{Z}/N)^{p_N})^\kappa \\ y_{\text{last}} \text{ a unit}}} \psi_N(y_{\text{last}})\, R_0(k_y),$$

where k_y is integral invertible and has last column and last row congruent to y. The coefficients ψ_N are roots of unity. To see that F preserves the column level, take k in the conjugated last-column stabilizer. Then $u_0^{-1} k u_0$ is integral invertible with last row congruent to the last standard basis row. It acts bijectively on the indexing vectors y, preserving their last coordinate. The matrix formed by multiplying k_y by $u_0^{-1} k u_0$ on the left has the same last column modulo N as the representative of the new index. Their

quotient on the right fixes the input column. Thus left multiplication by k permutes the sum defining F. Away from N the column level commutes with the sum.

If $\phi \in M(d_0)$, evaluate the y-summand at g. Its input argument is $g d_0 u_0 k_y d_0^{-1}$, and the lattice after the d_0 shift is the one at $g d_0 u_0 k_y$. This agrees with the transported lattice at $g d_0 u_0$ because k_y is integral invertible. The ultrametric inequality gives $FM(d_0) \subset M(d_0 u_0)$. Start with a tuple in $M(1)$ and recombine into $M(d_0)$. Each of the subsequent $m/2$ applications of F lands in $M(d_0 u_0)$ and allows another recombination back into $M(d_0)$. Pull the matrices back to the input tuple by linearity. We obtain an integral recombination whose columns have $F^{m/2}$-images in $M(d_0)$ and whose determinant valuation is at most

$$(m/2+1)CTa/\kappa \;\le\; 4mfbC = e_N C$$

for large x, since $T \le f$ and $a/(\kappa b) = O(x^{-4})$. This proves the index estimate for the test used at N.

7.4 Divisibility at the level prime

Let k be a finite field of cardinality q, with a nontrivial additive character $\mathbf{e}$. The pairing on matrices of opposite rectangular shapes is $(Y, X) \mapsto \operatorname{tr}(YX)$. It is perfect. For X with u entries, every residue function has Fourier coefficients $q^{-u} \sum_X f(X)\mathbf{e}(-\operatorname{tr}(YX))$. Indeed summing the characters between two points is q^u for equal points and zero otherwise. For integral values of f, each coefficient is integral after multiplication by q^u. The dual-half translations of the split oscillator restrict to these characters on its integral matrix box.

For the determinant variant, take $X = \binom{A}{0}$ with $A \in \mathrm{GL}_l(k)$, give it value $\xi(\det A)^{-1}$ for a nontrivial multiplicative character ξ, and use zero off this support. If B is the leading $l \times l$ block of the opposing matrix, the unnormalized transform is

$$S_\xi(B) = \sum_{A \in \mathrm{GL}_l(k)} \xi(\det A)^{-1} \mathbf{e}(\operatorname{tr}(BA)).$$

For invertible B, the substitution $A' = BA$ gives $S_\xi(B) = \xi(\det B) S_\xi(I)$. For singular B, choose an invertible D fixing $\operatorname{im} B$ pointwise with $\xi(\det D) \ne 1$. The substitution $A \mapsto AD$ preserves the trace and multiplies the character by a nontrivial scalar, so $S_\xi(B) = 0$. Applying the unnormalized square-matrix transform twice gives $S_\xi(I) S_{\xi^{-1}}(I) = q^{l^2} \xi((-1)^l)^{-1}$, by character orthogonality. Hence both the normalization $q^{-u} S_\xi(I)$ and its inverse are integral away from the residue characteristic. In the rectangular lift only the block B pairs with A, so the denominator is still q^u for the full rectangle. The local calculation at N will use additive inversion on its actual paired matrix boxes.

The paired residue calculation. Use $R_N = (\mathbf{F}_N)^{p_N}$ and e for the last basis vector in each rank. In rank n_j, take d_j with last entry N at $P_N \cap \{w : w \mid \pi\}$ and last entry 1 elsewhere, as for d_0. Write b_y for the indicator that the last column of an integral matrix is y modulo N, a_y for the corresponding last-row indicator, and $c_y(S) = \psi_N(e^{\mathrm{t}} S y)$ on integral matrices. All are extended by zero outside the integral matrices. In the split Schrödinger model the N-part of the glued pair is a sum of translates by $-y/N$ in the last column on the π-half and by its conjugate in the last row. At $w \in P_N$ over π, the translated first box tests integrality of the other columns and $T_{\mathrm{last}} - y/N$. Equivalently it tests $b_y(Td_j)$. The other translate multiplies the second box by ψ_N of its last row against y. At $w \in P_N$ over $\bar\pi$, the second box instead tests $S_{\mathrm{last}} - y^{\mathrm{t}}/N$, using conjugation to identify the two $\mathbf{F}_N$ coordinates, and the first box gives the character of its last column against y. On this support the phase is $\psi_N((NS_{\mathrm{last}} \bmod N) \cdot (T_{\mathrm{last}} \bmod N))$. In each integral first box on this half, z is its last column and the phase is the same dot product. The product over all places thus gives

$$\sum_z b_z(Td_j) c_z(d'_{j+2} S),$$

where d'_{j+2} has last entry N on $P_N \setminus \{w : w \mid \pi\}$. The other matrix coordinates have the same integral tests in both expressions. In the Schrödinger formula the self-pairing in the phase vanishes for a translate

supported on one split half. The remaining linear term is the character ψ_N of the original translation. It is unchanged by integral changes of the lift of y/N.

Write $L_j^{\mathrm{raw}}[f]$ and $R_j^{\mathrm{raw}}(g)$ for the lift and right translation in the unshifted coordinates at P_N. With $C_s(g) = \chi_{j,n_{j+1}}(g)$ and $C_t(k) = \chi_{j,n_j}(k)$, the local Weil action of (g,k) on the residue function is $C_s(g)^{-1}C_t(k)f(k^{-1}Sg)$. Here $\chi_{j,r}$ is the base-change character for arrow j at the determinant times $|\det|^{-r/2}$ at P_N. Its character is unramified there because the chosen places and the twist conductor are prime to N. Thus equivariance gives

$$L_j^{\mathrm{raw}}[f(k^{-1}\cdot g)] = C_s(g)C_t(k)^{-1}R_{j+1}^{\mathrm{raw}}(k)L_j^{\mathrm{raw}}[f]R_j^{\mathrm{raw}}(g^{-1}).$$

Applying the lift to the Weil translate by (d_j, d_{j+1}) defines the placed map

$$L_j[f] = R_{j+1}^{\mathrm{raw}}(d_{j+1})L_j^{\mathrm{raw}}[f]R_j^{\mathrm{raw}}(d_j^{-1}).$$

Write $R_j(g) = R_j^{\mathrm{raw}}(d_j g d_j^{-1})$. Put $B_j' = L_j[b_e]$ and $u_j = d_{P_N}^{-1}$ in rank n_j. For integral invertible A, B, split Weil covariance says

$$R_{j+1}(B)L_j[f(B\cdot A)]R_j(A) = L_j[f].$$

Both determinant characters on the compact groups are one. For $u_{j+1}^{-1}Su_j$, write S in blocks with its last row and column separated. Conjugation divides the top part of the last column by N, multiplies the left part of the last row by N, and fixes the bottom right entry. Hence $a_e(u_{j+1}^{-1}Su_j) = b_e(S)$, including their support conditions. The full covariance gives

$$R_{j+1}(u_{j+1})L_j[a_e] = \gamma_j B_j' R_j(u_j), \qquad \gamma_j = \prod_{w\in P_N} N^{(n_{j+1}-n_j)/2}.$$

The same unramified base-change character occurs in numerator and denominator at each place. Its values at the two last-coordinate determinants cancel. The remaining absolute-value powers are exactly the displayed rank difference. In particular $|j_N(\gamma_j)|_N \le 1$.

When a component of $y_{\mathrm{last}} \in R_N$ vanishes, the sum testing that last entry is invariant under a source central unit at that place. Its lift kills a vector in the image of 1 minus this unit's translation. The ramifying differences give this condition separately at every $w \in P_N$. Expanding c_e along its last column and last row, respectively, therefore gives for ramified ϕ

$$R_{j+1}(u_{j+1})L_j[c_e]\phi = F_{j+1}B_j'\phi = \gamma_j B_j' F_j \phi,$$

where $F_j = R_j(u_j)\sum_{y_{\mathrm{last}}\in R_N^\times}\psi_N(y_{\mathrm{last}})R_j(k_y)$. In the last-column expansion use $k_y e = y$, and in the last-row expansion use $e^{\mathrm{t}}k_y = y^{\mathrm{t}}$. If $y = (v,t)$ with t a unit, one such residue matrix is

$$k_y = \begin{pmatrix} I + t^{-1}vv^{\mathrm{t}} & v \\ v^{\mathrm{t}} & t \end{pmatrix}.$$

Its determinant is t, by the Schur complement. The operators F_j commute with the central translations. The maps B_j' intertwine them, since simultaneous multiplication by a central unit on both sides fixes the residue argument. These two maps thus preserve the ramifying-image conditions.

In the first arrow of a pair use $(g,k) = (d_{2i}, 1)$ in the unshifted covariance formula. In the second use $(g,k) = (1, (d_{2i+2}')^{-1})$. This removes the two complementary diagonals from the matrix arguments and gives the scalar

$$\chi_{2i,n_{2i+1}}(d_{2i})\ \chi_{2i+1,n_{2i+1}}(d_{2i+2}').$$

Indeed the source determinant factor in the first arrow is read in rank n_{2i+1}, and the target factor in the second arrow in rank n_{2i+1}. The intermediate translations by d_{2i+1} cancel when the maps are written in placed coordinates. At the exterior stages there are the d_{2i} and d_{2i+2} shifts of the placed maps, with

the last factor $R_{2i+2}(u_{2i+2})$ since $d_{2i+2}d'_{2i+2} = u_{2i+2}^{-1}$. Each $w \in P_N$ contributes to exactly one of the two determinants. The N-valuation of its character value there is the signed infinity weight at the matching place, by the same ray-class argument as in the marked transfer. Its absolute value as an exponent is at most $2b$ by the stage weight bound. The absolute-value power in $\chi_{j,r}$ has N-adic size at most one for these diagonals. Thus the scalar for this pair has N-adic size at most N^{2fb}.

In the remaining composition $R_{2i+2}(u_{2i+2})\sum_y L_{2i+1}[c_y]L_{2i}[b_y]$, the terms with $y_w = 0$ at any w kill the ramified input. Every other y is ke for some integral invertible k. Covariance gives $L_{2i}[b_y] = R_{2i+1}(k)B'_{2i}$ and $L_{2i+1}[c_y]R_{2i+1}(k) = L_{2i+1}[c_e]$. The sum is therefore the number of these vectors times $F_{2i+2}B'_{2i+1}B'_{2i}$ on the input. Move the F through all preceding B'_j using the same commutation identity. Induction over the pairs gives a single scalar times

$$B'_{m-1}\cdots B'_0F_0^{m/2}\phi.$$

The counts of residue vectors are integers and every γ_j is N-adically bounded by one. The scalar is shared by all terms of the good datum and has N-adic size at most N^{mfb}, which suffices for the allowance N^{2mfb}.

The gauge bound for B'_j follows from the placed source gauge at N and the finite Fourier calculation for residue functions. The last-column indicator on the integral box has a Fourier expansion on the residue matrices at P_N, with coefficients integral after multiplication by $N^{2fn_jn_{j+1}}$. Indeed the inversion denominator is the order $N^{fn_jn_{j+1}}$ of the matrix group under addition, and each coefficient numerator is an integral sum of roots of unity. Naming the residue function first projects the datum at N onto its vacuum part. For a Heisenberg translate, this projection is zero if its translation vector fails the integral-box test there. Each additive character on the box is then a Heisenberg translation of the projected vacuum, with possible simple poles only at $\overline{P_N}$. Choose $\delta \in \mathcal{O}_K$ zero modulo the prime of K induced by j_N and one modulo its conjugate. All places of $\overline{P_N}$ are over the first prime, and hence multiplication by δ clears the poles. The conjugate $\bar{\delta}d$ is a unit at j_N. In the normalized standard-data estimate the per-box ratio $|j_N(\bar{\pi}_G)/j_N(\bar{\delta}d)|_N$ is one. After removing the Fourier denominator the gauge cost is at most $N^{2fn_jn_{j+1}} \le N^{fb}$. The last-column indicator is stable under the two conjugated stabilizers. For A, B fixing the respective last vectors modulo N, $B^{-1}SAe = e$ in the residue field exactly when $Se = e$, and integral invertibility preserves the integral box. Thus the shifted lattices and the class-sum calculation use the same levels and gauges.

For the final restriction the d_m-shift contributes N^{-b_τ} in the last-coordinate action at embeddings with $\overline{\lambda(\tau)} \mid \pi$. By the branching identity these factors match, in the valuation, the scalar $\prod_\tau \Phi_\tau(\bar{\pi})^{-b_\tau}$ in the restriction. The integer branching map is integral in the remaining rank. Consequently the usual output Weyl gauge is at most N^{3mfb} on the columns passing the test, and the index bound just proved pays for the required local submodule.

7.5 A common scalar and the paired index

Take the one rational scalar

$$\beta_p = \lceil \gamma_p - e_p \rceil, \qquad q = \prod_p p^{-\beta_p} > 0, \qquad B_s = qA_s.$$

The product has finite support. At an embedding above p the factor $q/p^{-\beta_p}$ is a unit. At a small non-split prime, $\beta_p \le \gamma_p - 2\ell_p$ for large x because $\ell_p \ge 1$ uniformly. At N we have $\beta_N = \lceil -4mfb \rceil \le -3mfb$. At the remaining primes $\beta_p = 0$. Thus the local estimates give, for every finite tuple of integral ramified columns and every twist of level norm prime to p, an integral coefficient matrix of determinant valuation at most Ce_p for which the recombined definite lifts multiplied by q are integral at the specified embedding. By the common seesaw and the integral projection this also holds for the corresponding $\Theta_\vartheta B_s$-images at any fixed s.

At the archimedean embeddings the same q satisfies

$$2\log q + 2\sum_p \gamma_p \log p \;\le\; 2\sum_p e_p \log p \;\le\; 2\mathcal{B}.$$

Use the spanning set S from (7.1) in (7.3) and add the projection cost. Since $h_0 \le h_{\mathrm{main}} - m\,\mathrm{rem} + \mathcal{B}$ and $5\mathcal{B} + x^{99} \le L$ for large x, the result is

$$\|\Theta_1 B_s \phi\|_\sigma^2 \le e^{h_{\mathrm{sh}}}, \qquad h_{\mathrm{sh}} = h_{\mathrm{main}} - m\,\mathrm{rem} + L \quad (\phi \in S,\ s \in \mathcal{G},\ \sigma \in \mathrm{Aut}(\mathbf{C})). \tag{7.9}$$

Take a finite set of translates $\mathcal{S}$ and assign each row of a family y in their images of integral ramified columns one representation $y_i = B_{s_i}\phi_i$, with $s_i \in \mathcal{S}$. For a prescribed $\iota = (p, j_\iota)$, take a twist θ with level norm prime to p. At each of j_ι and $j_\iota \circ \mathrm{conj}$, group the rows by s_i, recombine each block of columns, and put its matrix on the diagonal in the given row order. The resulting matrices U_+ and U_- over the respective valuation rings make the lifts for the same θ integral at the two places. The determinant valuation of each matrix is at most $|\mathcal{S}|Ce_p$. Put

$$E_p = |\mathcal{S}|Ce_p \log p, \qquad \delta = |\mathcal{S}|C\mathcal{B}.$$

The paired index bound is

$$-\log|j_\iota(\det U_+)|_p - \log|(j_\iota \circ \mathrm{conj})(\det U_-)|_p \le F_p, \qquad F_p = 2E_p. \tag{7.10}$$

All finite sums of F_p are bounded by 2δ by the summed allowance for e_p. The maps, row preimages and costs precede the choice of the coefficient place. The matrices are chosen for the embedding and twist used there. This proves Theorem 2. The rank argument uses the ranges of the same maps B_s, the dimension C, and the bound $\dim\ker \mathcal{R} \le fC/(N-1)$.

8 Relative trace and selection of short rows

Let $W = \mathrm{im}\,\rho(e_0)$ be the good row space and $R = \dim_{\mathbf{C}} W$. Use the scaled maps B_s, the spanning set of integral ramified columns, the squared norm bound $e^{h_{\mathrm{sh}}}$, and the finite allowances e_p from Section 7.

Theorem 3 (Short rows of nearly full rank)**.** *On every sufficiently large selected square, $R > 0$. There is a set of translates $\mathcal{S}$ and a fixed n-tuple of independent rows $y_i = B_{s_i}\phi_i$, with $s_i \in \mathcal{S}$ and ϕ_i in the short spanning set, such that, for $\eta_0 = b^{-3/2}$,*

$$n \ge (1 - 4\eta_0)R, \qquad |\mathcal{S}|C \le 5R\log(1 + b^2), \qquad \mathrm{covol}_\sigma^2(\Theta_1 y) \le e^{nh_{\mathrm{sh}}} \quad (\sigma \in \mathrm{Aut}(\mathbf{C})).$$

At any coefficient prime p, for each embedding j over p and each twist θ of level norm prime to p, the same tuple of rows has the two local integral recombinations in (7.10), at j and $j \circ \mathrm{conj}$ for this θ. Their combined logarithmic determinant cost is at most $F_p = 2|\mathcal{S}|Ce_p \log p$, with $\sum_p F_p \le 2|\mathcal{S}|C\mathcal{B}$. The translates, preimages, untwisted norm bounds and allowances are fixed before p, j, θ.

8.1 Normalization of the relative trace

Use the same datum. Write $G = \mathcal{G} = G_0 \times G_1$, where G_i is the product at the chosen places over g_i. Let $K(g)$ be the row principal level in the common finite adelic unitary group of rank r and K_r its full lattice stabilizer. Thus $K_r/K(g) = G$. Let $S = U \rtimes H_0$ be the Bessel subgroup on Y_0, with U unipotent and $H_0 = U(X_0)$ fixing the frame. Use the class representatives h_c in $H_0(\mathbf{A}_f)$ and the Bessel coordinates of the filling datum. The column coordinate isomorphism first translates by the central lattice-scale element and then takes the Weyl-to-Bessel coordinates of the class values. We give the resulting coordinate space its Euclidean inner product. Its basis consists of $e_{c,i}$ for a class c and a tuple i of Bessel-coordinate indices. Each of the index sets is nonempty, so $C > 0$.

8.1.1 Unfolding the normalized pairings

For a discrete subgroup Γ of a second-countable unimodular group D, let Φ_i be left Δ_i-invariant seeds, $\Delta_i \subseteq \Gamma$. Use Haar on D and counting measure on the discrete subgroups. If the product of the two periodizations of $|\Phi_i|$ is integrable on $\Gamma\backslash D$, then

$$\begin{aligned}\langle \mathrm{Per}_1 \Phi_1, \mathrm{Per}_2 \Phi_2\rangle &= \int_{\Delta_1\backslash D} \overline{\Phi_1(y)}\, \mathrm{Per}_2 \Phi_2(y)\, dy \\ &= \sum_{\gamma\in\Delta_1\backslash\Gamma/\Delta_2} \int_{\Delta_\gamma\backslash D} \overline{\Phi_1(y)}\Phi_2(\gamma^{-1}y)\, dy,\end{aligned} \tag{8.1}$$

where $\mathrm{Per}_i \Phi_i(y) = \sum_{\alpha\in\Delta_i\backslash\Gamma} \Phi_i(\alpha y)$ and $\Delta_\gamma = \Delta_1 \cap \gamma\Delta_2\gamma^{-1}$. Unfold the first periodization. In the second sum the left Δ_2-cosets in the double coset of γ^{-1} have representatives $\gamma^{-1}\alpha$ for $\alpha \in \Delta_\gamma\backslash\Delta_1$. Joining the translates of the first fundamental domain gives the last integral. Tonelli for the absolute-value seeds justifies all regroupings. The inner product is linear in its second entry.

Use the product of the filling Haar measures at infinity and the common finite Haar measure $d\mu_f$ on the row group. The quotient by rational points uses counting measure. At infinity write dy as the disintegration along U_∞, with the filling Haar measure du on U_∞ and quotient measure $d\dot{y}$ on $U_\infty\backslash G_\infty$. The product of full Bessel vectors is $W(y)$, with $\int W_i(y)\overline{W_j(y)}\, d\dot{y} = \delta_{ij}$. It transforms on U_∞ by its Bessel character and on $H_{0,\infty}$ by the unitary Bessel-coordinate representation. The unipotent cutoff β_c satisfies $\sum_{u\in U_c} \beta_c(uy) = 1$ for $U_c = U(F) \cap h_c K(g) h_c^{-1}$, where $F = E^+$. Its integral along each U_∞ orbit is $\mathrm{vol}(U_c\backslash U_\infty) = \nu$, the same number for each c. Use the quantitative partition of Section 8.6. Replacing the partition leaves the periodized map unchanged by the cutoff independence in the periodization argument. Take the signed frame scales s_{va} from the real Bessel frame. Put $\varrho = m/2 - 1$, $k = \kappa$, $d_U = \varrho(\varrho + k)$ and

$$\begin{aligned} D_U &= \prod_{v\in\Phi} \prod_{a=0}^{\varrho} (s_{va}^2)^{k+2(\varrho-a)}, \\ \nu &= \frac{(g^{2d_U+\varrho}/N^k)^f\, |\operatorname{disc} E|^{\varrho/4}}{\mathrm{vol}(\widehat{\mathcal{O}}_E)^{d_U} D_U}, \quad \mathrm{vol}(\widehat{\mathcal{O}}_E) = 2^f |\operatorname{disc} E|^{-1/2}. \end{aligned}$$

Here $g = g_0 g_1$ and ϖ is the chosen generator of the different with $\bar{\varpi} = -\varpi$. Write $\sigma_U(u)$ for the sum of the entries just above the diagonal of the frame block. For the standard additive character $\psi = \psi_\infty \psi_f$ of $\mathbb{A}_E/E$, put

$$\chi_\infty(u) = \psi_\infty\big(\sigma_U(u)/(g\varpi)\big), \qquad \chi_f(u) = \psi_f\big(\sigma_U(u)/(g\varpi)\big).$$

The real Bessel law is $W(uy) = \chi_\infty(u) W(y)$ in rational coordinates. The character χ_f is trivial on the principal congruence subgroup of U, and $\chi_f(u)\chi_\infty(u) = 1$ for $u \in U(F)$.

Covolume of each unipotent fibre. The computation of ν uses successive outer rows of the unipotent group. Use rational coordinates with frame $e_0, \dots, e_{\varrho-1}, z$ and dual isotropic vectors $e_0^*, \dots, e_{\varrho-1}^*$. In the finite class lattice, the basis matrix B_c is the product of the Bessel lattice change and the action of the class representative in H_0. The class lattice is the orthogonal sum of its middle block and the hyperbolic planes. The matrix is the identity on the hyperbolic pairs. In the middle block its determinant has finite idele modulus N^f. Indeed the Bessel change replaces one vector x of X_0 by $x/\bar{\pi} + z/N$, where $\pi\bar{\pi} = N$. The determinant of the class representative has norm one. The finite modulus is unchanged by conjugation, so $|\bar{\pi}|_f^2 = |N|_f = N^{-2f}$. The finite class lattice is self-dual. In the two coordinates changed by the Bessel basis the Gram matrix is $\left(\begin{smallmatrix} 0 & -1 \\ -1 & -N \end{smallmatrix}\right)$, which has determinant -1, and the class representative is an isometry.

At row a remove the outer hyperbolic plane from the subspace remaining after rows $0, \dots, a-1$. Write J_a for the form on the space Y_a after this removal, of dimension $d_a = k + 1 + 2(\varrho - 1 - a)$ over E. The row subgroup is formed by the matrices

$$n_a(l,t) = \begin{pmatrix} 1 & l & t \\ 0 & I & -J_a^{-1} l^* \\ 0 & 0 & 1 \end{pmatrix}, \qquad t + \bar{t} = -l J_a^{-1} l^*.$$

Here l is a row on Y_a. Projection to Y_a gives the next unipotent group. This projection splits both over the reals and on the class lattice at principal level since its basis respects the decomposition into the outer pair and Y_a. Let $B_{c,a}$ be the restriction of the class basis matrix to Y_a. In this basis the conditions on the row and its conjugate column are equivalent to $(l/g)B_{c,a} \in \widehat{\mathcal{O}}_E^{d_a}$, by self-duality. For such an l, the scalar $l J_a^{-1} l^*/g^2$ is in $\mathcal{O}_F$. Take $\omega = -\zeta_3 \in \mathcal{O}_E$, which has trace one since the field contains the 5040-th roots of unity. Thus $t = -\omega l J_a^{-1} l^*$ satisfies the equation and is integral modulo the prescribed level in the sense $t \in g\mathcal{O}_E$. All other admissible t differ from it by $gi\mathcal{O}_F$, since the datum contains i and $\mathcal{O}_E \cap F = \mathcal{O}_F$. Consequently every row in the displayed lattice occurs, and its central fibre is a translate of the same skew lattice.

In a row subgroup the Haar measure is Lebesgue in the free complex row and in $\operatorname{Im} t$. Multiplication shifts the row additively and shifts the central coordinate by a function of the row. In the product $n_a(l,t)\operatorname{diag}(1,h,1)$ the free row is (lh, t), giving Jacobian one at fixed h since $|\det h| = 1$. Successive disintegration therefore gives the product measure in the free entries. This is also the Haar measure used for the Bessel quotient. On the group fixing the frame, the free complex entries in X_0 and the independent real entries of the skew matrix are the defining coordinates of that Haar measure. The quadratic correction to the skew matrix gives a translation on each fibre over the rectangular block. Adding the frame entries by the upper unitriangular lift preserves the product measure. At fixed frame entries, right multiplication by that lift is triangular in the remaining free entries of each row, with diagonal one.

The standard $\mathcal{O}_E$-lattice in the complex embeddings has covolume $2^{-f}\sqrt{|\operatorname{disc} E|}$. For the row lattice cut out by $lB_{c,a} \in \widehat{\mathcal{O}}_E^{d_a}$ the covolume is this quantity to the power d_a times $|\det B_{c,a}|_f = N^f$. The determinant factor follows by elementary divisors at the finite places. The imaginary lattice obtained from $i\mathcal{O}_F$ has covolume $\sqrt{|\operatorname{disc} F|}$. Both i and ζ_3 generate E over F. Their quadratic polynomial discriminants -4 and -3 show that the relative discriminant is the unit ideal. Thus $\sqrt{|\operatorname{disc} F|} = |\operatorname{disc} E|^{1/4}$. In rescaled real coordinates the complex entry in column j is multiplied by the scale of j divided by that of e_a, and the central entry by s_{va}^{-2}. The two scales at an inner hyperbolic pair are inverses. The complex entries in X_0 have column scale one and the entry in column z has squared column scale $s_{v\varrho}^2 = 1/N$. Hence the real Jacobian at v on row a, including the level factor g, is

$$g^{2d_a+1} s_{v\varrho}^2 (s_{va}^2)^{-k-2(\varrho-a)}.$$

In multiplying over the rows, $N^{f\varrho}$ from the finite determinants cancels the factor $\prod_v s_{v\varrho}^{2\varrho} = N^{-f\varrho}$. The extra factor for $a = \varrho$ in D_U contributes $\prod_v s_{v\varrho}^{2k} = N^{-kf}$ when the remaining product is written over D_U. As $\sum_a d_a = d_U$, this gives the displayed ν on each class. Integration of the partition of unity on a fundamental domain gives the same number on each unipotent fibre.

For a set T of rational terms define the blocks

$$A_T(c,s,c')_{ij} = \sum_{\substack{\gamma \in T \\ \overline{h_c^{-1}\gamma_f h_{c'}} = s}} \int_{h_c^{-1}\gamma_f h_{c'} \in K_r} \beta_c(y)\beta_{c'}(\gamma_\infty^{-1} y) W_i(y) \overline{W_j(\gamma_\infty^{-1} y)}\, dy.$$

In this formula the bar on a finite group element denotes reduction in the lattice frame of the filling datum. Use the Petersson unfolding (8.1) with both $\Delta_i = \{1\}$. The first seed is $s_{(c,i)}$. Represent s by $k \in K_r$ in the filling frame. In the ambient coordinates the right level is $\eta K(g) \eta^{-1}$ and the second seed $s_{(c',j)}$ is right-translated by $\eta k \eta^{-1}$, as in the stage-zero lattice chart. Absolute integrability follows from the uniform

absolute sum and L^1 seed bound in the weighted periodization. Taking the conjugate gives the order of factors in A_T. Write a point of the first seed's finite support as $x_f = h_c k_1 \eta^{-1}$, with $k_1 \in K(g)$. The shifted second seed has finite support there precisely when

$$\gamma_f^{-1} x_f(\eta k \eta^{-1}) = \gamma_f^{-1} h_c k_1 k \eta^{-1} \in h_{c'} K(g) \eta^{-1}.$$

Since $K(g)$ is normal in K_r, this is equivalent to $h_c^{-1} \gamma_f h_{c'} \in kK(g)$, independently of k_1. The finite integral then has the full coset volume $\mu_f(K(g))$ by unimodularity. The real integrand is the product of the two cutoffs and two full Bessel coordinates displayed above. Write $\mathcal{Q}$ for the unscaled Poincaré map $\mathcal{P}$ in the Euclidean column coordinates, and $P = (\mu_f(K(g))\nu)^{-1/2}\mathcal{Q}$. Then

$$\langle \rho(s) P e_{c',j}, P e_{c,i} \rangle = A_{\mathrm{all}}(c, s, c')_{ij}/\nu.$$

We use an inner product linear in its second entry. Thus the matrix entries of $K(s) = P^* \rho(s) P$ are the conjugates of the displayed entries. If i_{c_*} takes the coordinates of $c_* \phi$ on $\mathcal{C}_{\mathrm{ram}}$, the same projection as in (6.10) gives

$$B_s = q|\mathcal{G}|\, \rho(e_0)\rho(s)\mathcal{Q} i_{c_*}.$$

The coordinate isomorphism and the nonzero scalars preserve preimage dimensions for the rank calculation.

At $s = 1$, a term in $S(F)$ forces $c = c'$ and $\gamma \in U_c$. Indeed its H_0 component identifies the two column classes at column level, and the class stabilizers there are trivial. The rational unipotent phase at infinity is trivial on U_c by the finite congruence and the product formula for the additive character. The sum of the second cutoff over U_c is one. Integration of the first cutoff along U_∞ then gives νI_C for this block matrix. More generally, write J for the image of the integral points of S in G, and $T_S(s)$ for the normalized matrix obtained from the $S(F)$ terms. For $t \in J$, each class c has exactly one class c' with an $S(F)$ term over (c, t, c') modulo the unipotent actions. The existence follows by decomposing the H_0 part using the column classes and approximating the unipotent part by a rational unipotent in its prescribed finite open coset. Multiplying one term by the inverse of another reduces uniqueness to the case $s = 1$. It follows that $T_S(t)$ permutes the class fibres by unitary operators. It vanishes outside J and $T_S(1) = I_C$. Translating a rational term by one such $S(F)$ element permutes the rational terms over the corresponding class and residue labels. To see that this transport preserves the cutoff normalization, group the absolutely summable terms into U_c–$U_{c'}$ double orbits. In an orbit represented by γ, summing the second cutoff over $U_{c'}$ gives one and summing the first cutoff over U_c modulo $U_c \cap \gamma U_{c'} \gamma^{-1}$ unfolds the first integral to the quotient by that intersection. The two Bessel factors have trivial phase under the respective class unipotents. Thus the orbit integral is independent of the choices of the two partitions. Applying the Bessel transformation law to the transported orbit gives

$$K(ts) = T_S(t) K(s), \qquad T_S(ts) = T_S(t) T_S(s) \quad (t \in J).$$

For s supported in G_i, only the fixed class fibres contribute to $\operatorname{tr} T_S(s)$. A term in such a fibre has its H_0 part integral and congruent to one above g_{1-i}, so freeness at that single level forces this part to be one. Conversely an integral unipotent reduction acts on all fibres by its phase, by rational unipotent approximation. A lift u gives phase $\chi_\infty(u)$ in the unfolded block $A_{S(F)}$, from conjugating $W_j(u^{-1}y)$. The matrix $T_S(s)$ conjugates that block, so its phase is $\chi_f(u)$. Thus $\operatorname{tr} T_S(s) = C\chi_U(s)$ for these s, where χ_U is the reduction of χ_f on integral U, extended by zero.

8.2 Finite types, transfer and the row-space dimension

8.2.1 The share of allowed finite types

The good idempotent e_0 is the product of two central orthogonal idempotents $e^{(i)} \in \mathbf{C}[G_i]$. At a chosen place w above g_i, an allowed label is a monic squarefree polynomial of degree r with nonzero constant

term and no root among $x \in k_w$ with $x^{M_\omega} = 1$, where $M_\omega = (Md)^{s+6}$. A good tuple uses allowed labels everywhere and at least one irreducible label whose root generates $\mathbf{F}_{q^r}^\times$, where $q = g_i$. Use the Green irreducible representations for these labels, with the choice of characters in the datum. Use the squarefree case of Green's formula (4.1) for this label with irreducible-factor degrees $d_1, \dots, d_l$ and residue size $q = g_i$. It gives

$$\frac{\dim(\tau_c)^2}{|\mathrm{GL}_r(k_w)|} = q^{-r} \frac{\prod_{j=1}^{r}(1-q^{-j})}{\prod_{h=1}^{l}(1-q^{-d_h})^2} \geq q^{-r}(1-r/q),$$

since $\sum_h d_h = r$ and each numerator factor is at least $1-1/q$. The corresponding idempotent has coefficient $\dim(\tau_c)\chi_{\tau_c}(u^{-1})/|\mathrm{GL}_r(k_w)|$ at u. Its identity coefficient is its share in the regular representation.

There are at most $(1+M_\omega)q^{r-1}$ monic degree-r polynomials with a forbidden linear factor, including X. Those divisible by a nonconstant monic square number at most $\sum_{1\leq j\leq r/2} q^{r-j}$, which is at most $2q^{r-1}$. The regular share of all allowed labels is therefore at least $1-(M_\omega+3+r)/q$. The number of primitive labels is $\varphi(q^r-1)/r$, since a primitive root lies in no proper subfield. If n has a distinct prime factors, its j-th prime factor is at least $j+1$, so $\varphi(n)/n \geq 1/(a+1) \geq 1/(1+\log_2 n)$. Consequently the primitive share is at least $c_0/(r^2 \log q)$ for an absolute $c_0 > 0$ at these scales. There are f factors in G_i, with $f \asymp x^{11}\log x$, $r \asymp x^5$, $M_\omega = x^{O_d(1)}$ and $b^{10} \leq q \leq b^{11}$. For x sufficiently large the union bound for forbidden factors and the product bound for missing all primitive factors give

$$\beta_i := (1-e^{(i)})(1) \leq \frac{f(M_\omega+3+r)}{q} + \exp\left(-\frac{c_0 f}{r^2 \log q}\right) \leq b^{-8}/2. \tag{8.2}$$

The first term is $b^{-10+o(1)}$, and $f/(r^2\log q)$ is at least a fixed positive multiple of $x^{1-\varepsilon}\log x$.

8.2.2 Transfer of traces on the good types

Work over $F = E^+$ with base field $\mathbf{Q}$ for restriction of scalars. The Hermitian space W_j has diagonal entries consisting of $\kappa + j - 1$ copies of 1 and $m - j$ copies of -1. The finite adelic frames on Y_0 and on each W_j identify their finite groups, compact subgroups, and reduction maps with the same $G_f, K_r, K(g)$ and G. The spaces of interest each use a single archimedean discrete series at each real place, with the common regular Harish–Chandra parameter up to permutation and the signature of the group in question. The archimedean convolution kernel is $e_\pi = d_\pi \overline{m_\pi}$, where $m_\pi(u) = \langle v_\pi, \pi(u)v_\pi\rangle$ for its specified unit vector and d_π is the formal degree. The parameter gaps ensure $e_\pi \in L^1$. By the formal-degree projection (5.2), evaluation at v_π identifies the multiplicity space with the e_π-fixed space and preserves the commuting finite Hecke action. In particular the two spaces give the same finite Hecke trace. Let p be a common level-preserving Hecke operator away from the chosen places and their conjugates. Its realizing element is a finite linear combination of translations which are trivial there. Sandwich each translate by the averaging idempotent of the fixed compact subgroup away from the chosen places. This gives a finite-away convolution operator with the same action on principal-level forms. Use Haar probability on this compact subgroup and the same finite test function on every inner form.

For each good tuple, insert the tensor product of its finite type idempotents. At a chosen split place v over each g_i, the full group is $\mathrm{GL}_r(F_v)$ and the primitive Green label gives an irreducible cuspidal representation τ of $\mathrm{GL}_r(k_v)$. The depth-zero compact-induction theorem in this case says that the inflation of τ to $K_v = \mathrm{GL}_r(\mathcal{O}_v)$, extended to its normalizer $F_v^\times K_v$ by any compatible unitary central character, compactly induces to an irreducible supercuspidal representation. [21, Proposition 6.6] The extension exists by choosing a unitary value on a uniformizer. In the induced model a vector supported on the inducing subgroup has its diagonal coefficient supported there, with restriction to K_v equal to the corresponding coefficient of τ. Choose a unit basis and take the conjugate coefficients to give $\chi_\tau(u^{-1})$. The truncator $1_{\mathcal{O}_v^\times}(\det u)$ on $F_v^\times K_v$ cuts out exactly K_v. With probability Haar on K_v, the inflated idempotent is $\dim\tau\, 1_{K_v}(u)\chi_\tau(u^{-1})$. It is therefore a sum of truncated diagonal coefficients of the contragredient supercuspidal representation, with common scalar $\dim\tau$.

A truncated coefficient here is a normalized diagonal coefficient times a smooth compactly supported function on the quotient by the derived group, which is determinant for GL_r. For a supercuspidal representation π of $\mathrm{GL}_r(F_v)$ such a function f_v is strongly cuspidal. Its orbital integrals on nonelliptic elements vanish and any unitary irreducible representation contributing a nonzero trace is $\pi^\vee \otimes (\xi \circ \det)$ for some character ξ of $F_v^\times$. [22, §1.9 and Lemma 1.9.1, pp. 667–668] The two split places over g_0 and g_1 are distinct, so these tests force supercuspidality at both.

Every good form is cuspidal by the local constant-term test. The discrete-series vectors used in the real projection are smooth. For each proper rational parabolic use a chosen split place with a primitive label. A lattice basis adapted to the local flag makes the compact part of the radical reduce to that of a proper finite parabolic. The cuspidal Green type at this place has zero invariants under this radical. The test gives a zero constant term on each good tuple and hence on the good space. Conversely Miller and Schmid's rapid decay theorem gives, on each generalized Siegel set S, the bound $|f_v(g)| \le C_{S,A}\, \nu_j(v) \|g\|^{-A}$ for every $A > 0$, for the function associated to a smooth vector v in a cuspidal automorphic representation, with ν_j a continuous seminorm depending on the representation and A. Apply this to the finitely many arithmetic quotients at the fixed principal level, for the restriction of scalars to $\mathbf{Q}$. Its center is anisotropic and the chosen discrete-series lines are smooth. The bound gives bounded functions on all the quotients by reduction theory, so every such cuspidal principal-level vector belongs to the bounded model. [18, Theorem 2.24, p. 12] Thus the cuspidal trace computes the trace on the full good space.

At a real place use the following normalized pseudo-coefficient theorem. For a real reductive group with compact center and a discrete-series representation π_v, there is a compactly supported smooth K_v-finite function f_{π_v} with $\operatorname{tr} \pi_v(f_{\pi_v}) = 1$, trace zero on the other discrete series, and trace zero on the basic representations induced from proper parabolics. The traces are computed using the chosen group Haar. The expansion of an irreducible character by basic characters preserves its infinitesimal character. [23, §5.2, Lemma 6 and the Corollary following Proposition 4, pp. 212–213] Our $U(p, q)$ has compact center, and the row type is discrete series. Thus a nonzero trace requires the designated infinitesimal character. For the decreasing regular Harish–Chandra parameter λ_v, the differences and parity of its coordinates make $\lambda_v - \rho_{\mathrm{roots}}$ dominant integral. This is the highest weight of the finite-dimensional algebraic coefficient system with infinitesimal character λ_v. For $p, q > 0$, restrict to $SU(p, q)$, using the finite central covering by $SU(p, q) \times U(1)$. The restriction of an irreducible unitary representation is irreducible on $SU(p, q)$ since the center acts by a scalar. Its infinitesimal character is regular integral on this derived group. An irreducible unitary Harish–Chandra module of $SU(p, q)$ with this character is of the form $A_{\mathfrak{q}}(\lambda)$, cohomologically induced from an admissible one-dimensional character of the Levi of a θ-stable parabolic $\mathfrak{q} = \mathfrak{l} + \mathfrak{u}$. The inducing weight vanishes on the Levi roots and is nonnegative on the positive nilradical roots in a compatible order. [24, Theorem 1.2, p. 253, and §4, pp. 257–259] Its shift by half the positive roots and the shifted algebraic coefficient weight are dominant representatives of the same regular infinitesimal character. They coincide. For $A_{\mathfrak{q}}(\lambda)$ with an admissible inducing character, the Vogan–Zuckerman formula gives cohomology with the finite-dimensional module whose lowest weight is $-\lambda$ in that order. In degree $\dim(\mathfrak{u} \cap \mathfrak{p})$ the cohomology is one-dimensional, where $\mathfrak{p}$ is the noncompact part of the Cartan decomposition. [9, Theorem 5.5, p. 74] Here that module is the dual of the algebraic coefficient system. Its central character cancels that of the original representation. The relative cohomology on the finite central cover therefore descends to $U(p, q)$. At a compact signature the matching cohomology is in degree zero.

Harris and Labesse's similitude lifting states that a cuspidal automorphic representation of the restriction of scalars of the unitary group occurs in the restriction of a cuspidal automorphic representation of its rational similitude group. Their discrete-series theorem states that if the unitary representation has such a lift, is cohomological at all real places and has supercuspidal local Jacquet–Langlands transfers at two distinct split finite places, then its real factors belong to discrete series. [22, Proposition 1.8.1, p. 667, and Theorem 3.1.5(B), pp. 679–680] We have established the real cohomology and the two supercuspidal components. The local transfers are the identity on the split GL_r groups. Thus on the contributing spectrum

each pseudo-coefficient counts exactly the specified member, once per copy, as does the formal-degree projector.

For the real comparison, the normalized pseudo-coefficients of matching discrete-series parameters on real inner forms have matching stable orbital integrals on corresponding elliptic classes, with relative sign $(-1)^{pq-p'q'}$ for signatures (p, q) and (p', q') of rank r. On nonelliptic regular classes the cuspidal integrals vanish. Use Haar from transported invariant densities on the groups and matching Haar on corresponding compact Cartans. [25, §4, Lemma 5.3 and Theorem 6.3] [22, Remark after Proposition 1.10.6 and proof of Theorem 3.1.6] Since r is odd, pq and $p'q'$ are both even. Thus the sign is $+1$ at every real place in this comparison.

The inner-form trace identity used here is the following case of Harris and Labesse's Proposition 1.10.6. For two unitary inner forms over F, suppose compactly supported smooth test functions have a common transfer to the quasi-split form. At one split finite place use a truncated coefficient of a supercuspidal type whose local Jacquet–Langlands transfer is supercuspidal. At a distinct place use a cuspidal test, and at every place where the groups are nonisomorphic use a cuspidal test. Then the traces of these functions on the cuspidal spectra of the two groups are equal. [22, Proposition 1.10.6, case (a.1), (b.1), (c), pp. 671–672] The two primitive split factors supply the first and second tests. The finite frames match all finite factors, and compactly supported real discrete-series pseudo-coefficients supply the real cuspidal tests with the matching stable orbital integrals just verified. The common finite measures and the normalized real tests give no scalar in the equality. Summing over the diagonal coefficients and good tuples gives equality of the finite Hecke traces. The argument also applies to every positive power of the common Hecke operator. Equality of dimensions and Newton identities therefore give equality of its characteristic polynomial on the good spaces.

8.2.3 The dimension bound from the definite inner form

Take $p = 1$ to carry R to the good space of the definite W_m. It is bounded by the full principal-level dimension there. Let $\mathrm{mass}_r = \sum_a |\Gamma_a|^{-1}$ be the mass of the genus of the standard rank-r lattice at K_r. Refining each class at K_r to $K(g)$ acts on its $|G|$ level choices by Γ_a. The action is free by the principal-level stabilizer bound. Hence the number of new classes is $|G|\,\mathrm{mass}_r$. At each class the value space is the full algebraic representation of the compact real group with the common parameter. For the local parameters $\lambda_{v,1} > \cdots > \lambda_{v,r}$ without the factor two, its dimension is

$$\frac{\prod_{a<b}(\lambda_{v,a} - \lambda_{v,b})}{\prod_{a=0}^{r-1} a!}.$$

Here is a derivation in the $U(r)$ normalization. For a highest weight μ put $l_i = \mu_i + r - i$, so $l_i - l_j$ is the difference of the corresponding infinitesimal-character coordinates. The multiplicity-one branching rule restricts to the highest weights $\mu_i \geq \nu_i \geq \mu_{i+1}$ of $U(r-1)$. Set $t_i = \nu_i + r - 1 - i$, which ranges independently between l_{i+1} and $l_i - 1$. By induction the numerator of each dimension is the Vandermonde in the t_i. Use its determinant in monic falling powers $(t_i)_{r-1-j}$ with descending columns. Summing each row uses $\sum_{t=l_{i+1}}^{l_i-1}(t)_e = ((l_i)_{e+1} - (l_{i+1})_{e+1})/(e+1)$. The product of the denominators is $(r-1)!$. The difference matrix is obtained by subtracting adjacent rows in the rank-r Vandermonde in the l_i and expanding in its constant column. This proves the formula, starting from dimension one at rank one.

8.2.4 Compact-dual Haar and formal degree

On $U(j)$ use the invariant density given at the identity by real Lebesgue in the independent imaginary diagonal entries and complex strict lower entries of a skew Hermitian matrix. The first-column map to the

unit sphere in $\mathbf{C}^j$ takes the independent first-column coordinates to the sphere density, with the remaining coordinates giving this same invariant density on $U(j-1)$. Thus the compact volume satisfies

$$\operatorname{vol} U(j) = \frac{2\pi^j}{(j-1)!} \operatorname{vol} U(j-1), \qquad \operatorname{vol} U(0) = 1.$$

The sphere area follows by integrating the Euclidean Gaussian in polar coordinates. At $U(p,q)$ the analogous coordinates are those of JX, where $J = \operatorname{diag}(I_p, -I_q)$ and X is the tangent matrix. The compact-dual identification keeps the diagonal blocks of X and multiplies the off-diagonal blocks by i. In the independent entries of JX this changes a sign on one diagonal block and rotates the off-diagonal complex coordinates, preserving the absolute density.

Here the real filling Haar is the compact-dual normalization in the formal-degree formula. At a real place write the group as $U(J)$ with $J = \operatorname{diag}(I_p, -I_q)$, $p + q = r$. The differential of the Cayley chart $Y \mapsto (1 + J^{-1}Y/2)(1 - J^{-1}Y/2)^{-1}$ identifies the filling density with independent-entry Lebesgue measure on JX, divided by $\operatorname{vol} U(r)$ in those coordinates. This density corresponds to probability Haar on the compact dual by the preceding calculation. If $p = 0$ or $q = 0$ the group is compact and Schur orthogonality gives the degree equal to the representation dimension.

For $p, q > 0$, use the semisimple group $S = SU(p,q)$ with $K_S = S(U(p) \times U(q))$. On S start with Haar induced by the metric volume of S/K_S and probability Haar on K_S. Use the corresponding dual metric on $X^c = SU(r)/K_S$ with the same probability Haar on K_S. For a connected linear semisimple group without compact factors, Olbrich's formal-degree formula in these measures gives $\dim \tau / \operatorname{vol}(X^c)$ for a discrete series with the infinitesimal character of the finite-dimensional irreducible representation τ. The discrete Plancherel mass is the Schur formal degree for this Haar measure. Here $SU(p,q)$ is connected, semisimple and noncompact, with a compact Cartan. The finite-dimensional representation is obtained from the algebraic extension of the compact type of highest weight $\lambda_v - \rho_{\mathrm{roots}}$. On restriction it has the indicated infinitesimal character. The same quotient disintegration gives total mass $\operatorname{vol}(X^c)$ on the compact group. Divide both Haar measures by this mass. The compact Haar is now a probability measure and the formal degree on S is $\dim \tau$. [26, §2, pp. 7–8, and Corollary 5.2, p. 15]

The maps $SU(p,q) \times U(1) \to U(p,q)$ and $SU(r) \times U(1) \to U(r)$ by scalar multiplication have the same kernel $\{(\zeta^{-1} I_r, \zeta) : \zeta^r = 1\}$. Push forward the measures just compared, using probability Haar on each central circle. Both maps have the same local determinant under the tangent identification and have r sheets, so the comparison is unchanged. The measure on $U(p,q)$ is the filling Haar by the Cayley calculation. An irreducible unitary representation stays irreducible on $SU(p,q)$ since the central circle acts by scalars. Its matrix coefficient norm squared is unchanged in the product integral, so the normalized formal degree remains $\dim \tau$. In this application τ is the restriction of the compact $U(r)$ representation of highest weight $\lambda_v - \rho$, extended algebraically to $\mathrm{GL}_r(\mathbf{C})$ before restriction. It has the required infinitesimal character on $SU(p,q)$. The weighted Bessel system takes $\deg_v$ to be the formal degree of this real row type. We consequently obtain exactly the same Weyl dimension in the orbital energy and the definite row count. The finite factor $|G| \operatorname{mass}_r$ still comes from the free refinement of classes, with the residues in counting measure.

Thus, with $\deg_{\mathrm{Bes}} = \prod_v \deg_v$,

$$R \le |G| \deg_{\mathrm{Bes}} \operatorname{mass}_r . \tag{8.3}$$

8.3 Orbital energies and local coordinate bounds

8.3.1 Orbital energy and the fixing sets

For a set $\Omega \subseteq G$ and a set T of rational terms stable under rational unipotents on both sides, write $\mathscr{E}_T(\Omega) = \nu^{-2} \sum_{s \in \Omega, c, c'} \|A_T(c, s, c')\|_{\mathrm{HS}}^2$. This is a sum in counting measure. In particular $\mathscr{E}_{\mathrm{all}}(\Omega) = \sum_{s \in \Omega} \|K(s)\|_{\mathrm{HS}}^2$. Put $\mathscr{M} = C^2/(\deg_{\mathrm{Bes}} \operatorname{mass}_r)$, $\mathscr{D} = |J|C$ and $\tau = b^{kf/4}$. Write $A_T^{\mathrm{cut}} = A_T^{\flat}$ for the same blocks with the cut Bessel vectors in place of the full vectors. The link sum uses these cut vectors, and $\mathscr{E}$ uses the full vectors.

The field contains the 5040-th roots of unity since $5040 \mid N_1$. The admissibility size ratios below are those of the filling datum attached to D.

For a subset A of the chosen places above g_0, prescribe at each $w \in A$ a nonzero vector of the Bessel frame span, in either the standard module at w or the contragredient module at $\bar{w}$. Write $F_A \subseteq G$ for the subgroup fixing the prescribed vectors and Ω_A for the portion of F_A having all leading frame minors of sizes $1, \dots, \varrho + 1$ nonzero at one and the same chosen place above g_0. The last minor is the determinant of the augmented frame pairing. The open-orbit parametrization uses it to give a nonzero terminal pivot. When A is the full set take $\Omega_A = F_A$. Define $\mathscr{M}_A = b^{-6k|A|}\mathscr{M}$. For X the span of the isotropic frame, the link of γ is the image of $X^{\perp} \cap (\gamma X)^{\perp}$ in $(X^{\perp}/X)^2$ under $y \mapsto ([y], [\gamma^{-1}y])$. Let Z be the cross-pairing matrix of the two augmented frames and Z_{ϱ} its leading block. Their joint Gram matrix has diagonal blocks $-Nee^*$, where e is the last coordinate vector, and off-diagonal blocks Z, Z^*. Swapping the block columns and taking the Schur complement of Z gives the frame discriminant

$$(-1)^{\varrho+1} \det \begin{pmatrix} -Nee^* & Z \\ Z^* & -Nee^* \end{pmatrix} = |\det Z|^2 - N^2 |\det Z_{\varrho}|^2. \tag{8.4}$$

For invertible Z the rank-one determinant formula in the Schur complement gives the second term, using the last cofactor of Z. The polynomial identity then holds for every Z. The link energy squares the sum within one link at fixed (c, s, c') before summing over these labels and links.

Summing the finite denominator shells. Enlarge the finite set of places to a conjugation-stable set containing the support of g and the denominators of the finite packet under consideration. Let A be the product of these completions and R its ring of integers. Use the restriction $\psi_T(x) = \psi_f(x/\varpi)$ to these places. The annihilator of R for ψ_T is R. Thus the normalization $\mathrm{vol}(R) = 1$ is self-dual, and the finite Plancherel constant is one. Write $\mathrm{md}(c)$ for the modulus of multiplication by c. Put $n = \varrho + 1$. Write M for the group Haar measure of the unitary level, K for its integral group, $K(g)$ for its principal congruence group and Z for the intermediate subgroup imposing the prescribed residue conditions. Put

$$v_U = \mathrm{vol}_U(U \cap K(g)), \qquad M_Z = M(Z)M(K(g))/v_U^2.$$

Here the measures on the successive unipotent coordinates are those of the finite unfolding. The finite vector has one coordinate for each right coset $K(g)a$ in Z. At y that coordinate is $\chi_g(ya^{-1})$, where $\chi_g(uk) = \psi_T(\sigma_U(u)/g)$ for $u \in U$, $k \in K(g)$, extended by zero elsewhere. In particular there is no factor $1/[Z : K(g)]$ in its inner product. Indeed on $uy \in Z$ there is a unique coset containing uy. Integrating first over the u's giving that coset, each nonempty fibre has measure v_U, and the two characters give

$$v_U \langle W_Z(y), W_Z(yx) \rangle = \int_U 1_Z(uy) \chi_g\big((uy)x(uy)^{-1}\big)\, du.$$

The sign here uses the inner product linear in its second argument.

Let dM_0^+ be the upper measure at floor zero. The ordered factorization is $dM = du * dM_0^+$, with du the unipotent Haar. At a point y put $B(y) = \int_U 1_Z(uy)\, du$. Whenever this fibre is nonempty, its translate in U is contained in $U \cap K$, whose volume is one in the integral unipotent coordinates. Thus $0 \le B(y) \le 1$ and $\int B(y)\, dM_0^+(y) = M(Z)$. The coset formula bounds the pairing by $B(y)/v_U$. If it pairs nontrivially with $W_Z(y')$, then $y' \in UzK(g)$ for any chosen $z \in Uy \cap Z$. Indeed the common nonzero coset coordinate gives two representatives differing on the right by $K(g)$, which is normal in K. Unfolding $M(zK(g))$ makes the measure of $UzK(g)$ in M_0^+ equal to $M(K(g))/v_U$. Since $B(y)^2 \le B(y)$, the pairing norm is

$$\mathcal{N}(W_Z, W_Z) = \iint |\langle W_Z(y), W_Z(y') \rangle|^2\, dM_0^+(y')\, dM_0^+(y) \le \frac{M(Z)M(K(g))}{v_U^3}.$$

Consequently the common restricted integral uses $v_U \mathcal{N}(W_Z, W_Z) \le M_Z$.

Use three estimates on the same finite or complex local floor and the same pair of head gauges. The restricted positive-weight integral is proved after the scalar-volume bound. The capped estimate and the reconstruction from the dual coordinates are proved here. The numerical applications to the two local measures follow in their respective passages.

The absolute orbital majorant with a capped coordinate volume. Use the local level at the finite product or at $\mathbf{C}$, with level scalar g, box K, vector ball $\mathcal{B}$, gauge constant $\Lambda \geq 2$, and volume exponent s. Write $n = \varrho + 1$. Number the hyperbolic floors from 0 to $\varrho - 1$. For $l \leq \varrho$, H_l fixes the first l hyperbolic pairs. After the marked floor the remaining group is the core group $H_{\varrho+1}$, and $H_{\varrho+2} = \{1\}$. The group Q_l acts upper unitriangularly on the first l frame vectors, with $Q_{\varrho+2} = U$. The earlier unipotent radical is $\mathcal{N}_{<l}$, with $Q_l = H_l\mathcal{N}_{<l}$. Nearness at l means $Q_l y \cap K \neq \emptyset$, and is left Q_l-invariant. The groups Q_l and H_l decrease with l, and $U, H_l \subseteq Q_l$. The Weyl elements w_p, w_q of two heads of length $j \leq \varrho$ commute with H_j. The ratio $d = w_q^{-1}w_p$ is diagonal on the earlier hyperbolic pairs and normalizes U.

Fix an allowed position x. For $a < \varrho$ let $L_a(y, x)$ be the scalar values of the dual coordinate $h(e_a, w)$ on the vectors w with $h(e_b, w) = 0$ for $b < a$ and

$$y^{-1}w \in \mathcal{B}, \qquad (y^{-1}w - (yx)^{-1}w)/g \in \mathcal{B}.$$

Use the local additive Haar on these scalar sets. Put $r_a(y, x) = \operatorname{vol} L_a(y, x)$ and $P_l(y) = \prod_{a<l} r_a(y, x)^s$. Each scalar set contains a neighbourhood of zero and lies in the image of a compact set. Indeed the two ball conditions are open, contain zero, and the first bounds w in $y\overline{\mathcal{B}}$. Projection onto the dual coordinate is open on the subspace with the earlier coordinates zero. The same projection shows that the condition $c \in L_a(y, x)$ is open in (y, x, c), so r_a is measurable. The first a dual-zero conditions and the a-th dual coordinate are preserved by Q_{a+1} on that subspace. Changing w by this group element preserves the two pullbacks after changing y by the same left multiplication. Thus r_a is left Q_{a+1}-invariant, $0 < P_l < \infty$, and

$$P_l(\tau y) = P_l(y) \quad (\tau \in Q_l), \qquad P_{l+1}(y) = P_l(y) r_l(y, x)^s \quad (l < \varrho).$$

Write θ_l for the actual floor measure, supported on H_l, and $M_j = (\eta_{j-1} * \theta_{j-1}) * \cdots * (\eta_0 * \theta_0)$ for the preceding stack, where η_l is the Heisenberg measure. Convolutions are read with the highest floor on the left. Let W be a bounded measurable cut Bessel vector of frequency $1/g$, so $W(y) \neq 0$ implies $Uy \cap K \neq \emptyset$. Let W' be bounded and measurable with the same character on U, supported where y is near at j. Use the two kernels

$$k(y, z) = \langle W(y), W(z) \rangle, \qquad k'(y, z) = \langle W(y), W'(dz) \rangle.$$

For either kernel b put

$$b_l(y, z) = \int b(\tau y, \tau z)\, d(\theta_{\varrho+1} * \cdots * \theta_l)(\tau),$$
$$m_l^b(y) = \int |b_{l+1}(\tau y, \tau yx)|\, d\theta_l(\tau).$$

The empty convolution has unit mass at 1. In the weighted square sums defined in the floor calculation, write $A_l^b(y) = P_l(y)|b_{l+1}(y, yx)|^2$, $B_l^b(y) = \int (A_l^b + B_{l+1}^b)(\tau y)\, d\theta_l(\tau)$ for $l \leq \varrho$, and $B_n^b = 0$. Let Z_l^b integrate B_l^b against $\theta_{l-1} * \cdots * \theta_0$. With $C_\Gamma = \Lambda^{10^6 n^3}$ and $J_j = \prod_{a<j} \mathrm{fl}_a(p_a^{-1})$, the square product is $\mathcal{D}_j(x) = C_\Gamma^3 J_j^2 Z_0^k Z_j^{k'}$. The positive majorant and the active set at x are

$$\mathcal{F}_j^+(x) = \int m_j^k(y) m_j^{k'}(w_p^{-1}y)\, dM_j(y),$$
$$\mathcal{U}_j(x) = \{y : k(y, yx) \neq 0,\ w_p^{-1}y \text{ near at } j,\ w_q^{-1}(yx) \text{ near at } j\}.$$

For the weighted floor form with multiplier $\omega(\sigma\tau^{-1})$ of modulus at most one, use $W' = W$. At fixed y, x the two integrations against θ_j have upper arguments $\sigma y, \sigma yx$ and $\tau w_p^{-1}y, \tau w_p^{-1}yx$, because the

heads commute with H_j. The triangle inequality and Tonelli bound the modulus by $m_j^k(y)m_j^{k'}(w_p^{-1}y)$. The absolute integral over $dM_j\,dM$, times the head density κ_j, therefore majorizes the form whenever the exterior integral is finite. The real mixed application uses the displayed positive majorant directly with its choice of W'.

Suppose that for each y the θ_j-mass of $\{\tau : \tau y \text{ is near at } j+1\}$ is at most $V \geq 1$. Take V finite in the estimates below. Either local level has a finite bound from the size-seven argument, so an infinite mass bound can be replaced by that finite bound when needed. For $0 \leq \sigma < 1$, assume every $y \in \mathcal{U}_j(x)$ has an a with $j \leq a < \varrho$ and $r_a(y,x) \leq \sigma$. For $\sigma \geq 1$ there is no extra volume condition. Then

$$\mathcal{F}_j^+(x)^2 \leq \sigma^s V^2 \mathcal{D}_j(x). \tag{8.5}$$

The majorant vanishes if $\mathcal{U}_j(x)$ is empty. This covers the case $\sigma = 0$ under the cap hypothesis, and the case $j = \varrho$ when that hypothesis is imposed.

The measured one-floor bounds. The map $\tau \mapsto \tau^{-1}e_l$ takes θ_l to the quadric measure at floor $l \leq \varrho$ used in the measured floor coordinates. Here e_ϱ denotes the last marked vector. On a lower floor the Weyl parameter c gives the last hyperbolic coordinate and scales the d_l horizontal coordinates by c up to conjugation and signs. Its density $\mathrm{md}(c)^{d_l}\,dc$ is the Jacobian of this change. The skew coordinate has the translated skew Haar of the quadric chart. The terminal floor measure has this pushforward by its definition. The H_l-action preserves both the ambient additive Haar and the Hermitian value. It therefore takes the continuous family of quadric measures to another disintegration for the same norm. Uniqueness almost everywhere, continuity against compactly supported tests and full support of the skew measure give equality at the fixed Gram value, including zero. The stabilizer of e_l in H_l is $H_l \cap Q_{l+1}$. The left quotient by this closed stabilizer is a standard Borel space. The orbit map induces an injective Borel map on the quotient with Borel inverse on its image. Thus invariant measurable functions of τ descend measurably to the quadric orbit point. In particular, if a measurable function is invariant on the left by Q_{l+1}, its integral at τy against θ_l is invariant when y is replaced by ζy with $\zeta \in Q_l$. To see this, factor ζ in $H_l\mathcal{N}_{<l}$, conjugate the unipotent factor to the left, and use that $\mathcal{N}_{<l} \subset Q_{l+1}$ is normalized by H_l. The H_l factor on the right changes the orbit point by its inverse action and preserves its measure. The same argument works with invariance under $Q_{l+1} \cap H_b$ and $\zeta \in Q_l \cap H_b$ for $b \leq l$, using the factorization within H_b. At $l = \varrho + 1$ it is the invariance of the core-group measure. These identities apply to integrable complex functions by taking their real and imaginary parts.

For $l < \varrho$ the inverse moment needed at one floor is

$$\int_{\{\tau\,:\,\tau y \text{ near at } l+1\}} r_l(\tau y, x)^{-s}\, d\theta_l(\tau) \leq \Lambda^{10^4 n}. \tag{8.6}$$

If y is not near at l, the integration set has zero measure. Otherwise take $\zeta y = y_0 \in K$ with $\zeta \in Q_l$ and use the preceding invariance, since the near indicator and the reciprocal volume are left Q_{l+1}-invariant. If $\xi\tau y_0 \in K$ with $\xi \in Q_{l+1}$, the projection of $y_0(\xi\tau y_0)^{-1}e_l$ to the remaining floor space is $\tau^{-1}e_l$. Indeed $\xi^{-1}e_l - e_l$ lies in the earlier frame span and $\tau \in H_l$ fixes that span. Both box matrices have size one. Inversion adds four and multiplication adds the two sizes and one, putting the orbit point in size seven.

Use the inverse-volume moment at y_0 with no extra row equations. Its set D consists of orthonormal coordinates v for which the pullbacks of $E_0 v$ at y_0, $y_0 x$ are a close pair, where E_0 embeds the orthonormal part. If J is the split Gram matrix, set $u = E_0^* J \tau^{-1} e_l$, the orthonormal projection of the orbit point. The vector $\tau E_0 v$ has the required earlier dual coordinates zero and its two pullbacks are unchanged. Its l-th dual coordinate is $h(\tau^{-1}e_l, E_0 v) = u^* v$. Therefore $\{u^* v : v \in D\} \subset L_l(\tau y_0, x)$. The inverse volume to the power s is bounded by $\phi(u)^s$, with a measurable inverse-volume majorant ϕ supplied by that moment. On the size-seven quadric the paired-coordinate fibre integral costs Λ^{1287n} and the orthonormal coordinate has size at most eleven. Integrating $\phi(u)^s$ costs at most Λ^{1188n}. Their product is at most $\Lambda^{10^4 n}$, proving (8.6). The same size-seven argument with integrand one, including the terminal floor, bounds the θ_l-mass of a translated near set by $\Lambda^{1287n}\,\mathrm{vol}(\text{size } 11)^{d_{\mathrm{orth}}}$. Here d_{orth} is the number of orthonormal coordinates.

Supports, squares and the downward induction. The last mass bound and the mass at most one on the core factor over the cut support justify the upper averages. More explicitly, for a product of upper-floor factors to put the first argument of b in UK, the partial product through each floor l must put it in the near set at $l+1$, since the remaining factors are in $H_{l+1} \subset Q_{l+1}$. Successive integration bounds the measure of this set by the product of the floor mass bounds. Boundedness of W, W' now gives absolute convergence. A nonzero upper average or pivot therefore yields some point of the corresponding H_l-orbit where the unaveraged kernel is nonzero. In particular $b_{l+1}(\tau y, \tau yx)$ vanishes off the near-$(l+1)$ set. Since $P_l(\tau y) = P_l(y)$,

$$\int |b_{l+1}(\tau y, \tau yx)|^2 \, d\theta_l(\tau) \le B_l^b(y)/P_l(y).$$

Cauchy–Schwarz on the near set at $j+1$ thus gives

$$(m_j^b(y))^2 \le V B_j^b(y)/P_j(y).$$

For $l \ge j$ the simultaneous left action of $Q_l \cap H_j$ preserves b_l. At the unaveraged top this follows from the two Bessel characters, which cancel because the ratio d commutes with H_j. Descending induction uses the right-invariance identity for the floor measures. For $\nu \in \mathcal{N}_{<l}$ the simultaneous left action on b_{l+1} only introduces a unit character. To check this, move ν through each $\tau \in H_{l+1}$ in the upper average. The conjugate stays in $\mathcal{N}_{<l}$ and has the same simple-root entries. Indeed H_{l+1} fixes the frame through the entry indexed by l, and the later simple-root entries of this radical are zero. For the translated kernel the second conjugate is also conjugated by the diagonal ratio d, which multiplies each simple-root entry by its fixed ratio. Both characters are consequently independent of τ. Factoring $Q_{l+1} \subset Q_l = H_l \mathcal{N}_{<l}$ now gives left Q_{l+1}-invariance of A_l^b. Its support lies in the near set at $l+1$. The definition of B_l^b and the right-invariance identity for θ_l show by downward induction that B_l^b is invariant on the left by Q_l and zero away from the near set at l.

Put $C_1 = \Lambda^{10^4 n}$. Suppose some $l \ge j$ and some base y have the following property. For each $\tau \in H_l$ with $k(\tau y, \tau yx) \ne 0$, a coordinate at an index a with $l \le a < \varrho$ has volume at most σ. For $0 < \sigma < 1$ we claim

$$m_l^k(y)^2 \le \sigma^s C_1^{\varrho - l} B_l^k(y)/P_l(y). \tag{8.7}$$

At $l = \varrho$ the hypothesis makes the raw kernel zero on the whole H_ϱ-orbit. Thus its pivot is zero. Suppose $l < \varrho$ and write $z = \tau y$ for $\tau \in H_l$. When $r_l(z, x) \le \sigma$,

$$|k_{l+1}(z, zx)|^2 \le \sigma^s A_l^k(z)/P_{l+1}(z).$$

When $r_l(z, x) > \sigma$, that volume is unchanged along the H_{l+1}-orbit of z. The cap hypothesis on the H_l-orbit therefore forces a later small coordinate whenever the raw kernel is nonzero on this smaller orbit. The triangle inequality $|k_{l+1}(z, zx)| \le m_{l+1}^k(z)$ and the induction hypothesis bound the square by $\sigma^s C_1^{\varrho - l - 1} B_{l+1}^k(z)/P_{l+1}(z)$. In both cases it is bounded by

$$r_l(z, x)^{-s} \, \frac{\sigma^s C_1^{\varrho - l - 1}}{P_l(y)} \big(A_l^k(z) + B_{l+1}^k(z) \big).$$

Cauchy–Schwarz on the near-$(l+1)$ set uses (8.6) in its first factor and the defining integral of $B_l^k(y)$ in its second. This proves (8.7).

For the product of the two pivots, take a y where it is nonzero. The second pivot gives $\tau \in H_j$ with $W(\tau w_p^{-1} y)$ and $W'(\tau w_q^{-1} yx)$ both nonzero. The cut condition for W and the near-j support of W' imply both translated near conditions. Since the Weyl elements commute with H_j, they hold also with any $\rho_0 \in H_j$ in place of τ, by multiplying on the left by $\tau \rho_0^{-1} \in Q_j$. A nonzero first pivot then gives an active point by its raw kernel witness. This proves the vanishing assertion and treats $\sigma = 0$ directly, even for a vacuous mass bound $V = \infty$. If $0 < \sigma < 1$, the first pivot satisfies (8.7), with $C_1^{\varrho - j} \le C_\Gamma$. The second

pivot costs V. For $\sigma \geq 1$ use the bound V on each pivot. In both branches, using $V \geq 1$, the square of the product is bounded by

$$\sigma^s V^2 C_\Gamma \frac{B_j^k(y)}{P_j(y)} \frac{B_j^{k'}(w_p^{-1} y)}{P_j(w_p^{-1} y)}.$$

Cauchy–Schwarz in the two Weyl coordinates. Write $I_j(y) = 1/P_j(y)$ on the near set and zero elsewhere. On the supports of the two weighted square sums the last product is unchanged on replacing its reciprocals by I_j. Cross the reciprocal weights in Cauchy–Schwarz with dM_j. The resulting two integrals are

$$\int B_j^k(y) I_j(w_p^{-1} y)\, dM_j(y), \qquad \int B_j^{k'}(w_p^{-1} y) I_j(y)\, dM_j(y).$$

Their estimates use $M_j = \eta_{<j} * (\theta_{j-1} * \cdots * \theta_0)$, by moving each new floor measure past the earlier Heisenberg convolution, which is invariant under conjugation by that floor group. The first integral uses the written preceding inverse-weight bound with the head $\bar{p}$, since $w_p^{-1} = w_{\bar{p}}$ and conjugation preserves all local moduli. The cost is $C_\Gamma J_j Z_j^k$.

In the second use $y = w_p z$ in the integral of the left H_j-invariant function $B_j^{k'}(z) I_j(w_p z)$. For this change of variables, write the full Haar as $M = S_j * M_j$ with the upper stack S_j on H_j. That stack is right H_j-invariant. Induct from the core group. At a lower floor the higher stack convolved with the current Heisenberg measure is right invariant under the stabilizer $Q_{l+1} \cap H_l = H_{l+1} \mathcal{N}_l$, by conjugation invariance of the Heisenberg Haar under H_{l+1}. The right-invariance identity for θ_l now gives the assertion on H_l. At the marked floor the stabilizer is just $H_{\varrho+1}$. For any positive integrable function a on the group, put $A(z) = \int a(hz)\, dS_j(h)$. For a nonnegative left H_j-invariant function Φ, unfolding gives $\int \Phi\, dM_j = \int \Phi\, a/A\, dM$, as A is positive and finite M_j-almost everywhere. In this identity a left translation by w_p on Haar replaces a by $a(w_p \cdot)$ and its denominator by $A(w_p \cdot)$, since w_p commutes with H_j. This proves invariance of the integral of such functions under the Weyl change. The inverse-weight estimate on the second integral now costs $C_\Gamma J_j Z_j^{k'}$ with head p. Finally $Z_j^k \leq Z_0^k$: integrating the recursion for B_l^k against the preceding θ-stack gives Z_l^k as Z_{l+1}^k plus a nonnegative piece. The two costs and the pointwise factor give (8.5).

Coincident pivots at the bottom. At $j = 0$ with $W' = W$, both heads are empty and M_0 is the point mass at 1. Thus $\mathcal{F}_0^+(x) = m_0^k(1)^2$, $P_0(1) = 1$, $Z_0^k = B_0^k(1)$ and $\mathcal{D}_0 = C_\Gamma^3 B_0^k(1)^2$. Here all points are near at 0, since $H_0 = Q_0$ is the whole group and the box is nonempty. If every raw nonzero pair at position x has a coordinate of volume at most $e < 1$, apply (8.7) once to their common pivot and square the bound. This gives

$$\mathcal{F}_0^+(x)^2 \leq e^{2s} \Lambda^{2 \cdot 10^4 n \varrho} B_0^k(1)^2 \leq e^{2s} \mathcal{D}_0(x). \tag{8.8}$$

At $e = 0$ the raw support is empty and the same conclusion follows by vanishing. This is the input on the exterior of the principal coset in the finite bottom integral.

Reconstruction from the dual coordinates. Work over the local coefficient ring A with its involution. Order the split basis as $e_0, \dots, e_{\varrho-1}$, the middle block, and $e_{\varrho-1}^*, \dots, e_0^*$. The Gram matrix J pairs each e_a with e_a^* with value one and has no other entries on these rows. Write $n = \varrho + 1$ and $r = \dim_A V$. The Bessel group U is block upper unitriangular in this order. The group Q_j sends each e_a, $a < j$, to itself plus a combination of earlier e_b's. It contains U. For $h \in Q_j$ the row at e_a^*, $a < j$, of h differs from the identity row only in the columns e_b^*, $b < a$. Indeed $Jh = (h^{-1})^* J$, and $h^{-1} e_a$ has the stated upper-triangular form. Consequently h preserves the a-th dual coordinate of a vector whose earlier dual coordinates vanish.

In a gauge, products add sizes, a sum of at most $r + 1$ entries of size i has size $i + 1$, and negation and conjugation preserve size. The elements $0, 1$ belong to every scalar size set. Multiplication by 1 gives monotonicity in the size. All matrix and vector sizes below are entrywise. Thus a matrix product or matrix–vector product costs one additional size beyond the sum of the two sizes. Since J, J^{-1} have size one, $K^{-1} = J^{-1} K^* J$ has size at most $i + 4$ when K is a unitary matrix of size i.

First consider a head p of length $j \le \varrho$ in a gauge S. This means $p_a \in A^\times$ for $a < j$ and the inverse of $P_l = \prod_{a<l} p_a$ has size one for each $l \le j$. The Weyl element w_p sends e_a to $p_a e_a^*$ and e_a^* to $\bar{p}_a^{-1} e_a$ on these pairs. Suppose that sz and $hw_p^{-1}z$ lie in the box, with $s, h \in Q_j$. Their quotient sw_ph^{-1} has size seven in S. In its first $a+1$ primal columns, for $a < j$, right multiplication by h^{-1} changes the columns $p_b se_b^*$, $0 \le b \le a$, by an upper unitriangular matrix. Take the minor on the rows $e_0^*, \dots, e_{a-1}^*, i$ and these columns. Its determinant is

$$P_{a+1}(se_a^*)_i.$$

In the unchanged columns the first a rows form a lower-triangular block with diagonal $p_0, \dots, p_{a-1}$ and have zero last entries. Expansion in the last column gives the identity, also when a row is repeated. Expansion by one row at a time puts a $t \times t$ determinant with size-seven entries in size $8t$, starting with the empty determinant 1. At $i = e_a^*$ the minor is P_{a+1}, so all P_l, $l \le j$, have size at most $9n$. Using the inverse partial products gives size $10n$ for se_a^*, $p_a = P_{a+1}/P_a$ and $p_a^{-1} = P_a/P_{a+1}$.

For the matrix estimate let $0 \le \ell \le \varrho$, and let ν be unitary and block upper unitriangular for the first ℓ split pairs, with the remaining coordinates treated as one middle block. Let $g = \bar{g}$ be a unit of size one in the estimating gauge Γ, and set

$$D_\nu = (I - \nu)/g, \qquad \widetilde{D}_\nu = (I - \nu^{-1})/g.$$

Suppose for each $a < \ell$ that $f_a(e_b^*) = \delta_{ab}$ for $b \le a$, and both f_a and $D_\nu f_a$ have size c. Consider any matrix row for which the entries on the nondual columns of the ℓ pairs (the primal columns and the entire middle block) are known to have size d, and whose evaluations at all f_a also have size d. It suffices to recover the dual columns in decreasing a. To find the column at e_a^*, subtract the evaluation at $f_a - e_a^*$ from the evaluation at f_a. The vector $f_a - e_a^*$ has size $c+1$ and vanishes in the as yet unknown columns e_b^* for $b \le a$. If the current size bound is d', the product costs at most $d' + c + 2$ and the subtraction gives $d' + c + 3$. All entries of the row have size $d + (c+3)\ell$.

Apply this rule to the rows of D_ν in the middle block and at the dual vectors of the first ℓ pairs. The upper-triangular pattern gives zero in all their nondual columns. These rows have size $c + (c+3)\ell$. Since $J\widetilde{D}_\nu = D_\nu^* J$, the nondual columns of $\widetilde{D}_\nu$ have size $c + 4 + (c+3)\ell$. For such a column of J, only the rows of D_ν just bounded can contribute to $D_\nu^* J$. The two multiplications by Gram matrices cost four. Now use $\nu f_a = f_a - gD_\nu f_a$ as test vectors for $\widetilde{D}_\nu$. Their initial dual coordinates are still δ_{ab}, their size is at most $c+2$, and $\widetilde{D}_\nu(\nu f_a) = -D_\nu f_a$ has size c. A second application of the rule gives

$$\widetilde{D}_\nu \quad \text{of size} \quad c + 4 + (2c+8)\ell. \tag{8.9}$$

For the relative position x, use the pair (z, zx). Let p, q be heads of length $j \le \varrho$ in S. Assume $w_p^{-1}z$ and $w_q^{-1}zx$ are near at j, so their left Q_j-cosets meet the box. Choose $u, v \in U$ so that $K_1 = uz$, $K_2 = vzx$ lie in the box and $\Delta = (K_2 - K_1)/g$ has size two in Γ. When $j > 0$, assume every scalar of size $10n$ in S, divided by g, has size one in Γ. For every $j \le a < \varrho$, assume that $L_a(z, x)$ contains a unit c_a of A with c_a^{-1} of size one in Γ.

Take $\nu = vu^{-1}$ and $\ell = \varrho$ in (8.9). For $a < j$, set $f_a = ue_a^*$. The head calculation at the two box points gives size $10n$ in S for ue_a^* and ve_a^*. Hence f_a has size at most two in Γ and $D_\nu f_a = (u - v)e_a^*/g$ has size at most two there. For $a \ge j$, take a vector w realizing c_a in $L_a(z, x)$. Write $b = z^{-1}w$, $b' = (zx)^{-1}w$ and $e = (b - b')/g$. Both b and e are in the level ball and have size one. Set $f_a = c_a^{-1}uw = c_a^{-1}K_1 b$. It has the required initial dual coordinates and has size at most four. Moreover

$$D_\nu f_a = c_a^{-1}(K_2 e - \Delta b).$$

The two terms inside the parentheses have sizes at most three and four. Subtraction and multiplication by c_a^{-1} give size at most six. All the test vectors therefore work with $c = 6$. Equation (8.9) gives size $10 + 20\varrho$ for $\widetilde{D}_\nu$. Since $K_1^{-1}\nu^{-1}K_2 = x$,

$$(x - I)/g = K_1^{-1}(\Delta - \widetilde{D}_\nu K_2). \tag{8.10}$$

The right side has size at most $19 + 20\varrho \le 99n$. For $j = 0$ the head scaling condition and the first set of tests are empty. For $j = \varrho$ the condition on remaining L_a's is empty. The proof includes $\varrho = 0$, when U is the identity.

For the finite test use $W = W' = W_Z$, the preceding unipotent factor v_U, and Plancherel constant $c = 1$. Fix $0 \le j \le \varrho$ and let p, q be heads for the gauge α, with $\alpha/g \in R$ when $j > 0$. Put $\Lambda = 2/\operatorname{md}(\alpha)$. The orthonormal part has $d_{\mathrm{orth}} = k - 6\varrho - 13$ coordinates and the level has $3\varrho + 6$ additional hyperbolic pairs, so the volume exponent is $s_{\mathrm{f}} = k - 12\varrho - 23$. Put $\tau = -\zeta_3^{\zeta} \in R$, so $\tau + \bar\tau = 1$. In the marked integral block with Gram matrix $\left(\begin{smallmatrix} 0 & -1 \\ -1 & -N \end{smallmatrix}\right)$, take $u = -e_o$, $v = z - N\tau e_o$. They have norm zero and pairing one, and $z = v - N\tau u$. For two remaining norm-one slots a, b, take a unit c with $c\bar c = -1$, using the split or unramified norm on units. Then $a + cb$ and $\tau a - \bar\tau cb$ have norm zero and pairing one. Their change-of-basis determinant is $-c$. The marked pair has determinant -1. Applying the construction on $3\varrho + 6$ disjoint pairs leaves the stated number of orthonormal coordinates and preserves the integer-box measures. Put $m_0 = 3\varrho + 5$. There are at most $2f$ local factors above any rational prime, so the elementary ideal product gives $\zeta_R(m_0)^n \le \zeta(m_0)^{2fn} \le 1 + 2^{-n}$ at the selected large scales. This also bounds the ideal sums with larger exponents needed in the finite shell calculation. For floor mass use the integral gauge $\alpha = 1$, whose constant is two. The size-seven quadric fibre is bounded by $2^{99n(7+6)} = 2^{1287n}$. Size eleven is still R, of additive volume one. Hence the floor mass bound is $V_{\mathrm{f}} = 2^{1287n} \le \Lambda^{9999n}$ in any head gauge. On the shell $I_x = cR$, with $I_x = \{c \in R : c(x - I)/g \text{ is integral}\}$ and $h = [R : I_x]$, the gauge of depth $\beta = \bar c c$ makes $(x - I)/g$ of size one. Its constant is $2h^2$.

In the self-dual integral coordinates each unipotent fibre at the principal level has volume $v_U = \prod_{a<\varrho} \mathrm{fl}_a(g)$. Indeed the dilation by g has modulus $\mathrm{fl}_a(g)$ on the horizontal and skew coordinates at floor a. Write $\boldsymbol{D} = \boldsymbol{D}_j(x)$ for the square product in the capped majorant, on this same pair of finite coset vectors.

On the finite product $A = \prod_w E_w$ with $R = \prod_w \mathcal{O}_w$, the box is K and the level ball is R^r. A nonzero inner product of $W_Z(z)$ and $W_Z(zx)$ gives a common nonzero coset coordinate $K(g)b$. By the definition of the phase there are $u, v \in U$ with $K_1 b^{-1}, K_2 b^{-1} \in K(g)$, where $K_1 = uz$ and $K_2 = vzx$. Thus $K_1, K_2 \in K$ and $K_1^{-1}K_2 \in K(g)$, since $b \in Z$ normalizes $K(g)$. Consequently $\Delta = K_1(K_1^{-1}K_2 - I)/g$ has integral entries.

For $L_a(z, x)$, put $w' = uw$. The earlier zero dual coordinates and the a-th value are preserved in both directions by the triangular action of U. The condition $z^{-1}w \in R^r$ is equivalent to $w' \in R^r$, since $K_1 \in \mathrm{GL}_r(R)$. The second ball condition is the integrality of $g^{-1}(z^{-1} - (zx)^{-1})u^{-1}w'$. The permissible w''s form an R-submodule of R^r. Taking w' to be a sufficiently divisible multiple of e_a^* clears all the coefficients in the second condition at each place and gives an a-th coordinate in $A^\times$. The image under the a-th coordinate is therefore an ideal of R with a nonzero component at every place. Write it $d_a R$ with $d_a \in A^\times \cap R$ by choosing a uniformizer power in each component. In the normalization $\mathrm{vol}(R) = 1$, $\mathrm{vol}\, L_a(z, x) = \operatorname{md}(d_a)$.

Let $d = \prod_{j \le a < \varrho} d_a$. For the head gauge of depth α, size i means $\alpha^i c \in R$ and $\Lambda = 2/\operatorname{md}(\alpha)$. Use the estimating gauge of depth $\alpha' = \alpha^{10n+1}\bar d d$. This is integral, invertible in A and self-conjugate. For each remaining coordinate $\alpha'/d_a \in R$. When $j > 0$, $\alpha g^{-1} \in R$ in the head gauge. If c has size $10n$ there, $\alpha' c/g = (\bar d d)(\alpha g^{-1})(\alpha^{10n} c)$ is integral. These facts and the integral Δ give all the hypotheses of the reconstruction at the same active pair. Hence $(\alpha')^{99n} \in I_x$, the denominator ideal of $(x - I)/g$. If $h = [R : I_x]$, taking moduli gives

$$\operatorname{md}(d)^{198n} \le h^{-1} \operatorname{md}(\alpha)^{-99n(10n+1)} \le h^{-1}\Lambda^{99n(10n+1)}.$$

For $\sigma^{400n^2} = h^{-1}\Lambda^{99n(10n+1)} < 1$, all the remaining volumes being greater than σ^2 would contradict this inequality. Indeed there are at most $\varrho < n$ factors, and their product to power $198n$ would exceed σ^{400n^2}. For the empty product the left side is one.

Take $e \ge 0$ and a nonnegative integer t. Let $0 \le f$ be bounded by the absolute majorant, and suppose that when $e < 1$ it satisfies $f^2 \le e^{2s_{\mathrm{f}}}\boldsymbol{D}$ on the positions. The majorant is supported there by the common-coset test. Put $e_0 = \min(1, e)$, $B = 99n(10n + 1)$, and use the threshold $\sigma^{400n^2} = h^{-1}\Lambda^B$ of the finite reconstruction. Its small-coordinate conclusion gives cap σ^2 if $\sigma < 1$, including an empty active locus

at $j = \varrho$. With cap one for $\sigma \geq 1$, (8.5) now supplies the pointwise factor $E = \Lambda^{9999n} \min(e_0, \sigma)^{s_{\mathrm{f}}}$. For $t + 10^8 n^4 \leq k$, we have $s_{\mathrm{f}} \geq t + 400n^2(200000n^2)$. Putting t powers on e_0 and $400n^2(200000n^2)$ powers on σ gives

$$\min(e_0, \sigma)^{s_{\mathrm{f}}} \leq e_0^t h^{-200000n^2} \Lambda^{200000n^2 B}.$$

This includes $\sigma \geq 1$. If $e_0 = 0$, the hypothesis gives a zero integral on each shell since the weighted self-pairings are finite there almost everywhere. The restricted integral uses M_Z and $(2h^2)^{50000n^2}$. This spends $100000n^2$ powers of h, leaving $h^{-100000n^2}$. The actual level has $\sum_I [R : I]^{-\mu} \leq 2$ for some $\mu < 9n$, so the residual sum is at most two. Using $\Lambda \geq 2$, all factors beyond $e_0^t M_Z$ are bounded by $\Lambda^{B_{\mathrm{f}}}$, where

$$B_{\mathrm{f}} = 200000n^2 B + 1500000n^3 + 9999n + 50000n^2 + 2 \leq 10^9 n^4.$$

Consequently

$$\kappa_j \int f \, dM \leq \min(1, e)^t M_Z \Lambda^{10^9 n^4}.$$

Taking $e = 1$ proves absolute integrability of the finite majorant.

Plancherel on one floor. Let A be $\mathbf{C}$ at a real place of the base field, or $\prod_{w \in T} E_w$ at the finite places, with the involution of E. Let ψ_A denote $z \mapsto \exp(z - \bar{z})$ on $\mathbf{C}$ and ψ_T on the finite product. Write d for the number of orthonormal coordinates of the local floor basis, and $h(v, v') = v^* J v'$. In that basis write a vector as (p, w, z_0), where p consists of one coordinate from each hyperbolic plane, z_0 of its partners and $w \in A^d$. Thus

$$h(v', v) = (z_0')^* p + (w')^* w + (p')^* z_0.$$

The order in p pairs the second coordinate of the distinguished hyperbolic plane with its first. Write B_i for the vectors all of whose ambient entries have gauge size i, and $n = \varrho + 1$.

Use the following measure condition at size i. In floor coordinates the measure q is the image of $dp\, dw\, d\zeta(z)$ by a map $(p, w, z) \mapsto (p, w, z_0(p, w, z))$. The ζ-mass of the points in the coordinate box of size $i + 4$ is at most $\mu_i(p)$ for every p, w, and $\int \mu_i(p)\, dp \leq M_i(q)$. The increase by four covers passage from ambient to floor coordinates. For the floor quadric measure and for the additive volume of the floor one may use

$$M_i(q) = \Lambda^{99n(i+6)}.$$

Indeed in the distinguished plane the quadric chart has coordinates $(\bar{t}^{-1} \tau(\gamma - h(v, v)) + t\beta, t)$, where $\tau + \bar{\tau} = 1$, $\bar{\beta} = -\beta$. Its measure is $dt\, dv\, d\beta$, with the normalized skew Haar $d\beta$. Fix p, w and subtract one admissible set of residual coordinates from all the others. The terms γ and $w^* w$ cancel. The resulting residual parameters belong to the homogeneous hyperbolic cone chart at the same p, with coordinates of size $i + 5$. Translation preserves their measure. The integral in p is therefore bounded by the mass of this cone box. For additive volume the ordinary product box in the paired coordinates gives the same upper bound.

Here are the cone masses for the gauges in use. Let e be the number of hyperbolic planes including the distinguished plane. Over $\mathbf{C}$ the box radius is $M = L^{i+5}$, $\tau = 1/2$ and β is purely imaginary. Since $h(v, v)$ is real, the bound on the first coordinate forces $|\beta||t| \leq M$. Its skew length is at most $2M/|t|$. The integral of $|t|^{-1}$ over $|t| \leq M$ is $2\pi M$. The other $2(e - 1)$ coordinates contribute at most $(\pi M^2)^{2(e-1)}$. This fits M_i with $\Lambda = L^2$, $e \leq n$, $L \geq 4$.

At a finite product put $R = \prod_{w \in T} \mathcal{O}_w$, with $\mathrm{vol}(R) = \nu(R^-) = 1$. In the unit box write the hyperbolic form as $\mathrm{Tr}(x^* y)$, $x, y \in R^e$, where $\mathrm{Tr}(z) = z + \bar{z}$. For almost every x the ideal generated by the entries of x^* is tR for a unit of A. At fixed x, the image of dy under $\delta \,\mathrm{Tr}(x^* y)$ is normalized Haar on $\delta \,\mathrm{Tr}(tR)$, with δ the skew unit. This image contains $t\bar{t} R^-$, since $\tau + \bar{\tau} = 1$, $\tau \in R$ and $\delta \in R^\times$. Its density relative to the skew Haar is at most $\mathrm{md}(t)^{-1} = N(tR)$. Indeed multiplication by $t\bar{t}$ has modulus $\mathrm{md}(t)$ on skew elements. Averaging the quadratic level over a neighborhood of zero gives this same upper bound at fixed

x. On integrating in x, bound the content index by $\sum_{I\subset R} N(I)\mathbf{1}_{x^*\in I^e}$, over open integral ideals. This gives $\zeta_R(e-1)\le 2$ by the rank of the finite chart. Shrinking the neighborhood yields the unit-box mass on the zero quadric by continuity on compact open tests. For the box $\alpha^{-(i+5)}R$ in each coordinate, α self-conjugate, dilation multiplies this bound by $\mathrm{md}(\alpha)^{-(i+5)(2e-1)}$. The rescaling back to the unit box acts on the $2e$ ambient variables by α^{i+5} and on the skew level variable by $\alpha^{2(i+5)}$. Here $e\le 1+9n+\varrho$ and $\Lambda = 2/\mathrm{md}(\alpha)$, which gives M_i.

Suppose q_1, q_2 satisfy this condition at sizes i_1, i_2. For f integrable and square integrable on q_1, supported in B_{i_1}, a unit $\lambda\in A^\times$, and a fixed shift y_0, one has

$$\int_{B_{i_2}} \left| \int \psi_A(\lambda h(v'+y_0,v)) f(v)\, dq_1(v)\right|^2 dq_2(v') \le M_{i_1}(q_1)M_{i_2}(q_2)\left(\frac{c}{\mathrm{md}(\lambda)}\right)^d \|f\|^2_{L^2(q_1)}. \tag{F}$$

To see this, fix both bases p, p'. In the input residual variables integrate f against the character of $\lambda((p')^* z_0 + h(y_0, v))$, calling the result $a_{p,p'}(w)$. By Cauchy–Schwarz,

$$\int |a_{p,p'}(w)|^2 dw \le \mu_{i_1}(p) \iint |f(p,w,z_0(p,w,z))|^2 dw\, d\zeta(z).$$

The output residual variables contribute the character of $\lambda(z_0')^* p$, of modulus one. At fixed p', the absolute output is bounded by the integral in p of the absolute Fourier transforms of $a_{p,p'}$ in the w variables. Minkowski's inequality, ordinary Plancherel in w, and Cauchy–Schwarz in p bound its squared $L^2(dw')$ norm by $M_{i_1}(q_1)(c/\mathrm{md}(\lambda))^d\|f\|_2^2$, uniformly in p'. Integrating the output residual bound and then $\mu_{i_2}(p')$ proves (F).

Over $\mathbf{C}$, $d = k-1$ and the character is $\exp(2\mathrm{i}\,\mathrm{Im}\, z)$. Relative to the ordinary 2π Fourier transform on the real plane, the dual-variable Jacobian in one coordinate is $|\lambda|^2/\pi^2$. Thus $c=\pi^2$, $\mathrm{md}(\lambda) = |\lambda|^2$. Over the finite product $d = k - 6\varrho - 13$. The character has the integral product ring as its own annihilator, that ring has mass one, and conjugation preserves this Haar measure. Hence $c = 1$ and $\mathrm{md}(\lambda) = \prod_{w\in T} |\lambda_w|_w$.

Number the frame $(f_1,\dots,f_n) = (e_0,\dots,e_{\varrho-1}, z)$ in order. For the inserted character take both measures in (F) to be the quadric measure on floor $j\le\varrho$. At a given position and preceding translate, the two upper kernels K_{j+1}, K'_{j+1} occur with one factor on each of the two orbits. The orbit map $\sigma\mapsto t\sigma^{-1}f_{j+1}$ takes $d\theta_j$ to the quadric measure for every $t\in H_j$. The upper kernel descends to this orbit, by invariance under its stabilizer. When both kernels contribute, y and $w_p^{-1}y$ are near at j. Choose $tu\in H_j\mathcal{N}_{<j}$ with $tuy\in K$, and put $y_0 = ty$, $v = tut^{-1}$. Then $v\,y_0\in K$, and $w_p^{-1}y_0$ is still near. The head minors give size $10n$ for the $v\,e_a^*$ and $p_a^{\pm1}$. The triangular inverse estimate with level scalar one and $c = 10n+1$ bounds $I - v^{-1}$ by size $c+4+(2c+8)j$. Adding the identity and multiplying by $v\,y_0$ gives size $60n^2$ for y_0. Multiplication by w_p^{-1} gives size $99n^2$. The empty head uses the box representative itself. For active $\sigma\in H_j$, choose $\xi'\in Q_{j+1}$ with $\xi'\sigma y\in K$. Project $(ty)(\xi'\sigma y)^{-1}f_{j+1} = t\sigma^{-1}(\xi')^{-1}f_{j+1}$ onto the remaining floor. This gives $t\sigma^{-1}f_{j+1}$ because $(\xi')^{-1}f_{j+1} - f_{j+1}$ lies in the earlier frame fixed by t,σ. The two orbit supports therefore have sizes $60n^2+6, 99n^2+6$. The inserted entry is the pairing of these two points and has size at most $200n^2$. A frequency agreeing with λ as a character on that box gives exactly the same insertion. Apply Cauchy–Schwarz to the first orbit and (F) to the second, using $-\lambda$ when moving the complex conjugation. The squared double integral is bounded by $(c/\mathrm{md}(\lambda))^d$ times the product of the two squared $L^2(d\theta_j)$ norms and $\Lambda^{10^6 n^3}$.

For unequal last pivots put $j = a+1\le\varrho$ and assume that $\xi = (p_a^{-1} - q_a^{-1})/g$ is a unit of A. On the Heisenberg group of floor a write an element as $n(v,\beta)$, with $v\in V_{a+1}$ and β skew. Its measure is the image of $dv\,d\beta$. For $\sigma\in H_{a+1}$, the only possible nonzero term of the simple-root sum of $\sigma n(v,\beta)\sigma^{-1}$ is $-h(v,\sigma^{-1}f_{a+2})$. The Weyl ratio $w_q^{-1}w_p$ multiplies that simple-root entry by p_a/q_a. The translated kernel law at level g therefore inserts the frequency $\xi' = (p_a/q_a - 1)/g = -p_a\xi$ on this root sum.

Choose an active translate as origin. After translation and conjugation by the floor isometry from the support estimate, the active difference, modulo the preceding unipotent group, has its dual-frame column divided by $u = p_a^{-1}$ of size h_1. If the centered input translate has size h_0, the orbit point support has size

h_0+6. In these coordinates $v/u \in B_{h_1}$. At fixed v, two admissible skew coordinates have difference divided by u of size h_1+1, by subtracting the two central entries of that column. If S is the central slice and δ is the skew unit, the measure-preserving split $(\beta,\beta') \mapsto \tau\delta^{-1}\beta+\beta'$ shows that

$$\nu(S)^2 \le \mathrm{md}(u)\,\Lambda^{2(h_1+2)}.$$

Indeed $\tau\delta^{-1}$ has gauge size one. The differences of two points in the image of $S\times S$ belong to u times a scalar box of size h_1+3, whose volume before scaling is at most Λ^{h_1+4}. Changing $v = uv'$ contributes $\mathrm{md}(u)^{d_a}$, where $d_a = \dim_A V_{a+1} = k-1+2(\varrho-a)$. Thus the total scaling factor from this box is $\mathrm{fl}_a(u) = \mathrm{md}(u)^{d_a+1/2}$.

Expand the translated K'_j as the θ_j-integral of K'_{j+1}. Moving n through σ supplies the character just computed. Translation of the origin gives a fixed shift y_0 in (F) and the horizontal change supplies frequency $-\xi'\bar{u}$. Apply (F) with the input floor quadric and output ordinary volume of V_{a+1}. Since $\mathrm{md}(\xi'\bar{u}) = \mathrm{md}(\xi)$, its contribution to the squared integral, relative to the next $L^2(d\theta_j)$ norm, is at most

$$\left(\frac{c}{\mathrm{md}(\xi)}\right)^{d} \mathrm{fl}_a(p_a^{-1})\,\Lambda^{100n(h_0+h_1+31)}.$$

Both $u'y$ and $w_{p,j}u'y$ are near at $j=a+1$ on the active support. Apply the common-center calculation to the second point with transform $w_{p,j}^{-1}$. It gives $t_{u'} \in H_j$ with $t_{u'}u'y$ of size $99n^2$. For a fixed active u_0, the e_a^*-column of $(t_{u_0}u_0y)(t_{u'}u'y)^{-1}$ is that of $t_{u_0}u_0(u')^{-1}t_{u_0}^{-1}$ because the centering isometries fix e_a^*. Inversion and multiplication give size $198n^2+5$. Subtracting the identity costs one, and division by p_a^{-1} costs $10n$. Thus $h_0 = 99n^2$ and $h_1 = 198n^2+10n+6$ suffice. Translations of the origin and conjugation preserve the Heisenberg Haar. The sets of active points are the measurable near sets of the integral. These choices put the last gauge power below $\Lambda^{10^6n^3}$.

We record how the two estimates use the same weighted norms. For a fixed position write P_l for the product of the first l floor-value volumes to the local volume exponent. For either kernel put $A_l = P_l|K_{l+1}|^2$, $B_l(g) = \int (A_l + B_{l+1})(\sigma g)\,d\theta_l(\sigma)$ with $B_n = 0$, and Z_l for the integral of B_l against $\theta_{l-1} * \cdots * \theta_0$, with the identity mass at $l=0$. Since $P_l(\sigma g) = P_l(g)$ on H_l, the next squared orbit norm is at most $B_l(g)/P_l(g)$. Put $C_\Gamma = \Lambda^{10^6n^3}$ and $J_l = \prod_{b<l}\mathrm{fl}_b(p_b^{-1})$. The inverse-weight calculation bounds the integral of an invariant amplitude divided by $P_l(w_{p,l}g)$ against these preceding Heisenberg measures. The bound is $C_\Gamma J_l$ times the integral against the remaining measure, on the stated near support.

For the inserted entry, apply Cauchy–Schwarz to the two factors $(B_j/P_j)^{1/2}$ in the preceding translate. Applying the inverse-weight estimate in the two Weyl coordinates and using $Z_j \le Z_0$ gives $C_\Gamma^3 J_j^2 Z_0 Z'_j$ times the Fourier factor. For unequal pivots, write $h = w_{p,j}^{-1}g$. Before Cauchy–Schwarz the squared integrand on its support factors as

$$\frac{A_a(g)}{P_j(h)}\,\frac{P_j(h)}{P_a(w_{p,a}h)}\,|K'_j(h)|^2.$$

Here $P_a(w_{p,a}h) = P_a(g)$ by floor invariance. Retain the indicators that g,h are near at j. The first integral costs $C_\Gamma J_j$ and uses the A_a summand of Z_0. In the second, the function of h, with its near-point indicator, is left H_j-invariant. Left transport by $w_{p,j}^{-1}$ therefore preserves its integral against the ordered stack measure $(\eta_{j-1} * \theta_{j-1}) * \cdots * (\eta_0 * \theta_0)$, where η_l is the Heisenberg measure at l. In this integral isolate the last Heisenberg measure. The quotient of weights is constant along this integration. The horizontal estimate above costs $(c/\mathrm{md}(\xi))^d C_\Gamma\,\mathrm{fl}_a(p_a^{-1})$. The ensuing B'_j/P_j cancels the numerator P_j. The inverse-weight estimate at length a costs $C_\Gamma J_a$ and leaves Z'_j. Thus both bounds for the squared integral of the floor modulus have the required normalization

$$(c/\mathrm{md}(\lambda))^d\,\mathcal{D}_j(x), \qquad \mathcal{D}_j(x) = C_\Gamma^3 J_j^2 Z_0 Z'_j,$$

with $\lambda = \xi$ for unequal pivots. At the real places the level in this difference is $g=1$. For finite places all measures and weights refer to the same subgroup and local cut.

Integration against the preceding inverse weights. Fix a position x. Let $I_l(h)$ be $1/P_l(h)$ when $Q_l h$ meets the box K, and zero otherwise. Here P_l is the product of the first l floor-value volumes to the power s. Write $\omega_{<l}$ for the convolution of the first l Heisenberg measures. For a head p with all inverse initial products of size one, we prove, for $l \le \varrho$,

$$\int \Phi(h) I_l(w_{p,l} h)\, d(\omega_{<l} * \mu)(h) \le \Lambda^{10^6 n^3} \prod_{a<l} \mathrm{fl}_a(p_a^{-1}) \int \Phi \, d\mu.$$

Here $\Phi \ge 0$ is measurable, invariant on the left by Q_b and supported where $Q_b h$ meets K, for some $b \ge l$. The measure μ may be any countable sum of finite measures. The preceding unipotents lie in Q_b, so Φ passes through their integration. For each remaining h, some element of $Q_b \subset Q_l$ takes h to $g_0 \in K$. Factor that element in $H_l \mathcal{N}_{<l}$. The measure $\omega_{<l}$ is invariant under translations by $\mathcal{N}_{<l}$ and under conjugation by H_l. Also H_l commutes with $w_{p,l}$ and preserves I_l on the left. Thus it suffices to bound $\int I_l(w_{p,l} \nu g_0)\, d\omega_{<l}(\nu)$ uniformly on the box. At length zero the bound is one.

The Heisenberg convolution has the factorization $\omega_{<a+1} = \eta_a * \omega_{<a}$ by invariance of the earlier measure under the new floor group. Fix $\nu \in \mathcal{N}_{<a}$ and put $n_v = n(v, \beta)$ for the new Heisenberg element at $a < l$. The points $w_{p,a+1} n_v \nu g_0$ and $w_{p,a} \nu g_0$ differ by left multiplication by an element of H_a. Their first a weights agree, and nearness at $a+1$ for the first implies nearness at a for the second. On the active set $(n_v \nu)^{-1}$ returns $n_v \nu g_0$ to K. Apply the head-column bound with the head $\bar{p}$, whose Weyl element is $w_{p,a+1}^{-1}$. The scalar $\bar{p}_a$ and the column $(\nu^{-1} n_v^{-1}) e_a^*$ both have size $10n$. For $u = (-\bar{p}_a)^{-1}$, the column of $(\nu^{-1} n_v^{-1} - 1)/u$ therefore has size $j_0 = 20n + 1$. Thus v/u lies in the additive floor box of size j_0. At fixed v, differences of admissible skew coordinates divided by u have size at most $j_0 + 1$. The measure-preserving combination $(b, b') \mapsto \tau \delta^{-1} b + b'$ sends two copies of this skew slice into a translate of u times a box of size $j_0 + 3$. Its squared mass is at most $\mathrm{md}(u) \Lambda^{j_0+4}$. The change $v = uv'$ in $d_a = k - 1 + 2(\varrho - a)$ horizontal coordinates therefore costs, with the skew slice, at most $\mathrm{fl}_a(u) \Lambda^{j_0+2}$, where $\mathrm{fl}_a(u) = \mathrm{md}(u)^{d_a+1/2}$.

Let E_0 embed the d orthonormal coordinates. Define D by the ϱ row constraints $(\nu E_0 w)_{e_i} = 0$ and the condition that $(g_0^{-1} E_0 w, (g_0 x)^{-1} E_0 w)$ is a close pair. The latter requires the first vector and the difference of the two divided by the level g to be in the level ball. In fact $\nu E_0 w = E_0 w$ on D. A preceding unipotent fixes all the other coordinates of a vector with zero earlier dual coordinates. The Heisenberg and Weyl formulas now give

$$w_{p,a+1} n_v \nu E_0 w = E_0 w - p_a h(v, E_0 w) e_a^*.$$

This vector has zero earlier dual coordinates and its pullbacks at the two translated points give the prescribed close pair. Its e_a^*-coordinate is $h(v/u, E_0 w)$. The new floor-value set thus contains $z^* D$, where $z = E_0^* J(v/u)$.

We need a measurable $\phi \ge \mathrm{vol}(z^* D)^{-1}$ with $\int_{B_t^d} \phi(z)^s \, dz \le \Lambda^{99n(t+1)}$ at every size t. Over $\mathbf{C}$, $d = s + 1 + 3n$, $s = k - 3n - 2$. Add the row constraints forcing the negative projections of $g_0^{-1} E_0 w$ and $(g_0 x)^{-1} E_0 w$ to vanish. The negative space has dimension n, so there are at most $3n$ rows altogether. Choose $s + 1$ orthonormal columns O in their common kernel. On their image both pullbacks preserve Hilbert norm, since they preserve the Hermitian form and take values in the positive eigenspace. Thus $O\{v : \|v\| \le 1/9\} \subset D$, since the difference has norm at most $2/9$, the level is $g = 1$, and the ball has radius $1/4$. The image $z^* D$ contains the disk of radius $\|O^* z\|/9$. Take $\phi(z) = 36/\|O^* z\|^2$. After unitary completion of O, enclose the size-t box in the ball of radius $S = \sqrt{d}\, L^t$. In $s + 1$ complex dimensions the integral of $\|v\|^{-2s}$ over that ball is $\pi^{s+1} S^2 / s!$. The other $3n$ coordinates contribute at most $(\pi S^2)^{3n}$. The resulting moment is bounded by $\pi e^{36\pi} (\pi d)^{3n} d\, L^{2t(3n+1)}$, which is at most $\Lambda^{99n(t+1)}$ for $\Lambda = L^2$, $4 \le L$, $d \le L$.

At the finite places $d = k - 6\varrho - 13$, $s = k - 12\varrho - 23$. The finite inverse moment is on the open position set $\{y^{-1} v_0 y k_0 : v_0 \in U,\ k_0 \in K(g)\}$ at fixed y. Openness follows from $K(g)$ being open. The common-coset test puts its position in this set. Write $x = y^{-1} v_0 y k_0$ with $v_0 \in U$, $k_0 \in K(g)$, and $C = y g_0^{-1} E_0$. In addition to the first ϱ rows require the dual-frame coordinates of Cw, and the first ϱ coordinates of $(v_0^{-1} - 1) Cw$, to vanish. Then $v_0^{-1} Cw = Cw$, since a Bessel unipotent can change only the first frame

coordinates when the dual coordinates vanish. Clear the finitely many row denominators in each factor and use column operations over its discrete valuation ring. The resulting matrix in $\mathrm{GL}_d(R)$ has its last $h_f = d - 3\varrho = s + 3\varrho + 10$ columns Y in this common kernel. Hence $YR^{h_f} \subset D$. The first pulled-back vector is integral, and the second is its image under k_0^{-1}.

Take $\phi(z) = \mathrm{vol}\{c : cY^*z \in R^{h_f}\}$. For almost all z the entries of Y^*z generate a fractional ideal tR. Thus $\phi(z) = \mathrm{md}(t)^{-1}$ because $z^*YR^{h_f} = \bar{t}R$. Its inclusion in z^*D gives the required majorization. On R^d, integral changes of coordinates preserve volume. Summing over the possible content ideals bounds the moment by $\sum_{I \subset R} N(I)^{s-h_f} = \zeta_R(h_f - s) \le 2$. On $\alpha^{-t_0}R^d$ the cost of scaling is $\mathrm{md}(\alpha)^{-t_0(d-s)}$. Since $d - s \le 10n$, this fits the desired power of the finite Λ.

In either case use the volume-preserving floor basis for v/u. The box of ambient size j_0 has all floor coordinates of size j_0+4. The paired coordinates have total volume at most $\Lambda^{99n(j_0+6)}$, since each scalar box costs at most Λ^{j_0+5} and there are at most $2(1 + 9n + \varrho)$ such coordinates. On the orthonormal coordinates use the moment at size $j_0 + 4$. Together with the skew and horizontal factors this costs, in one step, at most

$$\mathrm{fl}_a(p_a^{-1})\Lambda^{j_0+2+99n(2j_0+11)} \le \mathrm{fl}_a(p_a^{-1})\Lambda^{10^4 n^2}.$$

Integrate the earlier weight and iterate for the $l \le n$ floors. This proves the asserted inverse-weight inequality on both local levels.

The scalar volume in the position shells. The same disks and integral columns also bound the positive floor weight. In the construction of D, allow any system of at most ϱ homogeneous row equations in w. Their number is the only property used to construct O and Y. For fixed z, let C_z consist of the scalars c for which each $w \in D$ admits two close pairs $(a, a - ge), (b', b' - gf')$ with

$$(cz)^* w = g^{-1}\big(h(a, b') - h(a - ge, b' - gf')\big) = h(e, b') + h(a, f') - gh(e, f').$$

Here $g = \bar{g}$, and a, b', e, f' lie in the level ball. At the finite places the last expression lies in R, since the Gram matrix and g are integral. Testing each column of Y and conjugating gives $cY^*z \in R^{h_f}$. The latter scalar set has volume $\phi(z)$, so $\mathrm{vol}(C_z) \le \phi(z)$.

Over $\mathbf{C}$, $g = 1$ and the Gram matrix has operator norm one. The three terms have total absolute value at most $3/16$. If $r_z = \|O^*z\| > 0$, the disk in z^*D contains $r_z/9$. Testing this value gives $|c| \le 27/(16r_z) \le 3/r_z$, whence $\mathrm{vol}(C_z) \le 9\pi/r_z^2 \le 36/r_z^2 = \phi(z)$. When $r_z = 0$ the majorant is infinite. Thus the same s-th moment controls the inverse floor volume and this scalar volume. The positive-weight integral uses the case of zero row constraints before integration over the positions.

Integration on a restricted set of positions. We prove the first of the three position estimates. The two gauges use the same level, scalar $g = \bar{g}$, box K, close pairs, floor measures and Haar measure M. Write $\mathcal{P} = \{x \in \text{positions} : (x - I)/g \text{ has size 1 in the second gauge}\}$. All sizes in the proof up to the final head calculation refer to this gauge, of constant Λ'. Put $d = d_{\mathrm{orth}}$, $d_a = k - 1 + 2(\varrho - a)$, $n = \varrho + 1$, and

$$C_a = \mathrm{fl}_a(g)c^d(\Lambda')^{4600n} \quad (a < \varrho), \qquad C_\varrho = 1.$$

Here $\mathrm{fl}_a(g) = \mathrm{md}(g)^{d_a+1/2}$. For two points h, y let $P_l(h, y)$ be the product of the first l floor-value volumes to the power s. The a-th value is $\zeta_{e_a^*}$ for a vector ζ with earlier dual coordinates zero for which $(h^{-1}\zeta, y^{-1}\zeta)$ is a close pair. These sets have jointly open membership in the points and the scalar, as one can vary the selected scalar coordinate in a witness vector and keep the others fixed. Their volumes are measurable, positive and finite, since each contains a neighborhood of zero and lies in the image of the bounded set $h\overline{\mathcal{B}}$ under the coordinate map.

Write η_a for the Haar measure of the Heisenberg group $\mathcal{N}_a$, with $\eta_\varrho = \delta_1$, and $\omega_l = \eta_0 * \cdots * \eta_{l-1}$, the measure of $\mathcal{N}_{<l}$. The coordinates $n(v, \beta)$ give $d\eta_a = dv\, d\beta$. Translation shears the scalar coordinate and translates the vector coordinate. Conjugation by H_{a+1} acts on v by an isometry of additive modulus one

and fixes β. Using the factorization $\mathcal{N}_{<a+1} = \mathcal{N}_{<a}\mathcal{N}_a$ and $\mathcal{N}_a \subset H_a$, induction shows that ω_l is invariant on both sides by $\mathcal{N}_{<l}$ and under conjugation by H_l.

The positive weight. For $Q_l h \cap K \neq \emptyset$ and $l \leq \varrho$ we claim

$$\int \mathbf{1}_{\mathcal{P}}(h^{-1} v\, y) P_l(h, v\, y)\, d\omega_l(v) \leq \prod_{a<l} C_a. \tag{1}$$

Factor a transporter taking h into K as γu with $\gamma \in H_l$, $u \in \mathcal{N}_{<l}$. Replace h, y by $\gamma u h, \gamma y$, and change v to $\gamma u v \gamma^{-1}$. The position is unchanged. Both points of the transported pair are multiplied by γu. This preserves the first l value sets, since Q_l preserves each successive dual coordinate on the subspace where the earlier ones vanish. It suffices to prove (1) for $h \in K$, by induction on l. For $l = 0$ its left side is an indicator.

In the step $a \longmapsto a+1$, a nonzero integral allows us to choose an element $u_0 \in \mathcal{N}_{<a+1}$ with $h^{-1} u_0 y \in \mathcal{P}$. Right invariance permits replacement of y by $u_0 y$, so write $x_0 = h^{-1} y \in \mathcal{P}$. Factor $\omega_{a+1} = \omega_a * \eta_a$. Whenever $x_1 = h^{-1} v\, n(v, \beta) y \in \mathcal{P}$, the matrix

$$\frac{(hx_0)(hx_1)^{-1} - I}{g} = h \left(\frac{x_0 - I}{g} - \frac{x_1 - I}{g} \right) (hx_1)^{-1}$$

has size 14. Indeed $x_1 = I + g((x_1 - I)/g)$ has size 3, products add the sizes plus one, and $X^{-1} = J^{-1} X^* J$ costs four. Thus the two factors after the first multiplication have sizes 4 and 9. The element $(hx_0)(hx_1)^{-1} = n^{-1} v^{-1}$ is an earlier unipotent times n^{-1}. Its e_a^*-column gives the horizontal coordinates of n^{-1} and its e_a-entry unchanged by that earlier factor. Consequently v/g has size 14. At fixed v, differences of admissible β's divided by g have size 15, by subtraction of the two central entries. The skew-slice calculation in the one-floor proof then gives, for a nonnegative measurable function b_1 supported on such slices and bounded there by $B(v/g)$ with B measurable,

$$\int b_1(n')\, d\eta_a(n') \leq \mathrm{fl}_a(g)(\Lambda')^{16} \int_{B_{14}} B(v')\, dv'. \tag{2}$$

In fact the translate of the image of two copies of the skew slice under $(b, b') \mapsto \tau \delta^{-1} b + b'$ lies in g times a scalar box of size 17. Thus its squared mass is at most $\mathrm{md}(g)(\Lambda')^{18}$. The horizontal change of variables costs $\mathrm{md}(g)^{d_a}$. These give (2), also for a smaller measurable support.

Apply the scalar-volume bound at h, x_0 with no row constraints. Thus $w \in D$ when $(h^{-1} E_0 w, y^{-1} E_0 w)$ is a close pair. Set $z = E_0^* J(v/g)$. For every value $b = \zeta_{e_a^*}$ at $(h, v\, n y)$ and every $w \in D$, unitarity gives

$$\begin{aligned} g^{-1}\big(h_J(h^{-1}\zeta, h^{-1} E_0 w) - h_J((v\, n y)^{-1}\zeta, y^{-1} E_0 w)\big) &= g^{-1} h_J(\zeta, E_0 w - v\, n E_0 w) \\ &= (b z)^* w, \end{aligned}$$

where $h_J(v, v') = v^* J v'$. Indeed $n E_0 w = E_0 w - h_J(v, E_0 w) e_a$, and the earlier unipotent adds only earlier frame vectors, orthogonal to ζ. The two pairs on the left are close. Hence the volume of the new value set is at most the same $\phi(z)$ for every such v. Integrating the earlier P_a with its position indicator costs at most $\prod_{b<a} C_b$ by induction, with ny as second point. Apply (2) to the resulting majorant ϕ^s on the remaining active slices. The additive floor box has paired coordinate fibre mass at most $(\Lambda')^{99n(14+6)}$. Its orthonormal coordinates have size $14 + 4$, where the moment of ϕ costs at most $(\Lambda')^{99n(14+5)}$. Since $16 + 99n(20 + 19) \leq 4600n$ and $c \geq 1$, this proves (1).

One integral of a squared pairing. Let $\Theta = \theta_n * \cdots * \theta_0$, $\lambda_l = \theta_l * \cdots * \theta_0$, and k be either raw pairing $k(h, y) = \langle W(h), W'(D_j y) \rangle$, $D_j = w_{q,j}^{-1} w_{p,j}$, or the self pairing with $D_0 = 1$ and $W' = W$. Use the corresponding starting floor j or 0. Write $K_l(h, y)$ for its diagonal average in $\theta_n * \cdots * \theta_l$, so $K_{n+1} = k$. The cut on the first vector implies $K_l(h, y) = 0$ unless $Q_l h$ meets K, since each higher transporter belongs to $H_l \subset Q_l$ and $U \subset Q_l$. These are measurable, absolutely convergent averages. To see this, at each floor $b \leq \varrho$ first translate a point near b into the box. The orbit point of an active translate then has size 7,

and its quadric mass there is at most $M_7\ \mathrm{vol}(B_{11}^{\mathrm{scalar}})^d$ by the fibre estimate. The change of representative preserves this integral by invariance of the quadric and its orbit parametrization. The terminal measure gives mass at most one to the translates meeting the cut. Iteration and boundedness of both vectors give uniform finite absolute integrals on every upper stack, so Fubini gives $K_l = \int\ K_{l+1}(\sigma\,\cdot, \sigma\,\cdot)\,d\theta_l(\sigma)$.

On $\mathcal{N}_{<l}$ the two unipotent transformations of $K_{l+1}(h, y)$ have the unitary phases coming from the two Bessel characters. Indeed one conjugates them by the averaging element of H_{l+1}, which fixes the first $l+1$ frame vectors, then uses the Bessel law in the raw pairing. In particular $|K_{l+1}(h, v\,y)| = |K_{l+1}(h, y)|$ for $v \in \mathcal{N}_{<l}$. For $b \geq j$ the twist commutes with H_b. Thus the amplitude $K_{b+1}(\sigma h, \sigma y)$, $\sigma \in H_b$, descends to the orbit $\sigma^{-1} f_{b+1}$. Here simultaneous translations in the stabilizer $H_b \cap Q_{b+1}$ preserve the amplitude. The unipotent phases cancel, and right change of representative in successive floor integrals preserves the invariant quadric measure, downwards from right invariance of the terminal measure. The orbit factor is measurable. In fact H_b, its ambient vector space and $\mathbf{C}$ are standard Borel spaces and the orbit map is Borel. For any rational sublevel set of the real or imaginary part of the amplitude, its image and the image of its complement are disjoint analytic sets because the amplitude is constant on fibres. Borel separation gives a Borel set agreeing with the sublevel set on the image. The countable family of these sets recovers each component there by rational cuts and extends it measurably to the ambient vector space. This supplies the measurable amplitude on the quadric in (F).

For $0 \leq l \leq n$ use the measurable functions

$$m_l(h, y) = \int \mathbf{1}_{\mathcal{P}}(h^{-1} v\,y)\,d\omega_l(v), \qquad A_l(h, y) = \mathbf{1}_{\{m_l(h,y)>0\}}\,|K_{l+1}(h, y)|^2.$$

The cut by positive mass avoids choosing a representative in a projected support. We prove for $j \leq a \leq \varrho$ that

$$\int A_a(h, n'y)\,d\eta_a(n') \leq C_a \int A_{a+1}(\sigma h, \sigma y)\,d\theta_{a+1}(\sigma). \tag{3}$$

If $m_{a+1}(h, y) = 0$, the left side vanishes by Tonelli and $\omega_{a+1} = \omega_a * \eta_a$. Otherwise $m_{a+1}(\sigma h, \sigma y) = m_{a+1}(h, y) > 0$ on H_{a+1} by conjugation invariance. Thus the right side uses the full square. At $a = \varrho$ the left integral is at the identity and (3) follows by Cauchy–Schwarz against θ_n on the part of its support meeting the cut, of mass at most one.

Suppose $a < \varrho$ and the left integral is positive. Choose one active $n_0 \in \mathcal{N}_a$ with $h^{-1} v_0 n_0 y \in \mathcal{P}$, $v_0 \in \mathcal{N}_{<a}$. Nonvanishing implies that h is near $a+1$. Factor a transporter taking h into K as $\xi = R u_*$, where $u_* \in \mathcal{N}_{<a}$ and $R \in H_a \cap Q_{a+1}$, and factor $R = tr$ with $t \in H_{a+1}$, $r \in \mathcal{N}_a$. For any active n' choose v with $h^{-1} v\,n'y \in \mathcal{P}$. Compare these two positions at the center ξh by the size-14 calculation above. In the variable $m = R n' n_0^{-1} R^{-1}$ their quotient is

$$(\xi v_0 n_0 y)(\xi v\,n'y)^{-1} = v' m^{-1}, \qquad v' \in \mathcal{N}_{<a},$$

since H_a normalizes the earlier unipotents. Therefore (2) holds for the active $m = n(v, \beta)$. This change of variables preserves $d\eta_a$. Its horizontal coordinate in n' is $t^{-1}v + v_0$, where v_0 is that of n_0.

Expand $K_{a+1}(h, n'y)$ at the next floor. Moving n' past $\sigma \in H_{a+1}$ inserts $\psi_A(-g^{-1} h_J(t^{-1}v + v_0, \sigma^{-1} f_{a+2}))$. Indeed the conjugated simple-root entry is the negative of this pairing before division by g. The twist acts trivially on $\mathcal{N}_a$ since $a \geq j$, and only the second argument was translated. Thus with $v = gv'$ the frequency in (F) is -1, its shift is tv_0/g, and its input orbit variable is $t\sigma^{-1} f_{a+2}$. This orbit carries the floor quadric measure. Where $K_{a+2}(\sigma h, \sigma y) \neq 0$, that variable has size 7. In detail $t^{-1}\xi \in \mathcal{N}_{<a+1}$, and conjugating this element by σ preserves that subgroup. Consequently $(\sigma t^{-1})(\xi h)$ is near $a+2$. Its translate into the box and the size-one matrix ξh give the column of $(\sigma t^{-1})^{-1}$ on f_{a+2} of size $1 + 5 + 1 = 7$ after projecting off the earlier frame vectors fixed by H_{a+1}. Applying (2) and then (F) with quadric input at size 7 and additive output at size 14 bounds the left side of (3) by

$$\mathrm{fl}_a(g) c^d (\Lambda')^{16+99n(7+6)+99n(14+6)} \int |K_{a+2}(\sigma h, \sigma y)|^2\,d\theta_{a+1}(\sigma).$$

The exponent is at most $4600n$. This proves (3), including the case of a null active slice.

Set $\mu_n = \Theta$ and $\mu_l = \eta_l * \mu_{l+1}$ downwards. In the ordered stack move the earlier Heisenberg measures past the higher floor factors using their conjugation invariance. Since $\eta_\varrho = \eta_n = \delta_1$, this gives $M = \omega_l * \mu_l$ for each l. When $s_1 \in H_l$, the measures μ_l and its left translate by s_1 therefore give the same integral to every nonnegative left $\mathcal{N}_{<l}$-invariant function. Indeed their convolutions on the left with ω_l are both M, by left invariance of M and conjugation invariance of ω_l. On invariant functions such a factor is unique by the quotient integral formula. Explicitly, choose a strictly positive integrable b on (G, M) and put $b_l(y) = \int b(v\,y)\, d\omega_l(v)$. This is positive, finite almost everywhere for either factor, and left $\mathcal{N}_{<l}$-invariant. Set $b/b_l = 0$ on the exceptional null set. For an invariant nonnegative F_1, integrating $F_1 b/b_l$ against M gives $\int F_1\, d\mu_l$ by the convolution identity, and likewise for the translated factor.

Write $T_l = \iint A_l(h, y)\, d\mu_l(y)\, d\lambda_l(h)$. Integrate (3) against $d\lambda_a(h)\, d\mu_{a+1}(y)$. For fixed σh, $A_{a+1}(\sigma h, \cdot)$ is invariant under $\mathcal{N}_{<a+1}$, by right invariance of ω_{a+1} and the character law for K_{a+2}. The preceding identity allows the simultaneous translate $\sigma\, y$ to be replaced by y in its integral. It follows that $T_a \le C_a T_{a+1}$, and $T_n \le \iint |k(h, y)|^2\, d\Theta(y)\, d\Theta(h)$.

Expanding the finite sum defining $Z_j(k; x)$ gives one term for each $j \le l \le \varrho$, namely the λ_l-integral of $P_l(h, hx)|K_{l+1}(h, hx)|^2$. In its integral over $x \in \mathcal{P}$, put $hx = v\,y$ using $M = \omega_l * \mu_l$ and Haar invariance. The square is constant as v varies. If $m_l = 0$, the indicator vanishes almost everywhere in v. A zero square also gives zero. Otherwise h is near $l + 1$, hence near l, and (1) bounds the term by $(\prod_{a<l} C_a) T_l$. Iterating (3) proves

$$\int_{\mathcal{P}} Z_j(k; x)\, dM(x) \le n \left(\prod_{a<\varrho} C_a \right) \iint |k(h, y)|^2\, d\Theta(y)\, d\Theta(h). \tag{4}$$

There are at most n summands. The choices of centers in (1) and (3) were made only for fixed outer variables in pointwise inequalities. The integrals use the measurable m_l, A_l.

The head factors. Put $\mathcal{N}(V, V') = \iint |\langle V(h), V'(y)\rangle|^2\, d\Theta(y)\, d\Theta(h)$. The raw self pairing in (4) has squared integral $\mathcal{N}(W, W)$. For the second pairing it is

$$J\, \mathcal{N}(W, W'), \qquad J = \prod_{a<j} \mathrm{md}(p_a/q_a)^{d_a+1}.$$

Indeed the a-th Weyl parameter in Θ has density $\mathrm{md}(b)^{d_a}\, db$, and its Weyl factor commutes with all higher floors. Write $D_a = w_{q,a}^{-1} w_{p,a}$ and $r_a = p_a/q_a$. In that factor $D_{a+1} w_a (r_a b)^{-1} = D_a w_a (b)^{-1}$ by the two-by-two hyperbolic block. Substituting $r_a b$ contributes $\mathrm{md}(r_a)^{d_a+1}$. Iterating proves the formula.

Let $N_0^2 \ge \mathcal{N}(W, W)\mathcal{N}(W, W')$. With $S \subset \mathcal{P}$, f nonnegative measurable and $f^2 \le E^2 \boldsymbol{D}_j$ there, apply Cauchy–Schwarz in x to the two square roots in $\boldsymbol{D}_j$, extending their squared integrals to $\mathcal{P}$. Formula (4) bounds $\kappa_j \int_S f$ by

$$E N_0\, n \left(\prod_{a<\varrho} C_a \right) \Lambda^{1500000 n^3} \kappa_j \left(\prod_{a<j} \mathrm{fl}_a(p_a^{-1}) \right) J^{1/2}.$$

Since $\kappa_j = \prod_{a<j} \mathrm{md}(p_a q_a)^{d_a/2}$, the last three factors starting with κ_j multiply to $\mathrm{md}((\prod_{a<j} q_a)^{-1})^{1/2}$. The inverse prefix has head-gauge size one, hence size two and modulus at most Λ^2. The argument requires only q to be a head. As $n \le 2^n \le (\Lambda')^n$ and $4600 n\varrho + n \le 50000 n^2$, the resulting estimate is

$$\kappa_j \int_S f\, dM \le E N_0\, c^{d\varrho} \left(\prod_{a<\varrho} \mathrm{fl}_a(g) \right) \Lambda^{1500000 n^3 + 1} (\Lambda')^{50000 n^2}.$$

This includes $j = 0$ with empty head product. For the finite pairing norms in the applications, a zero norm makes the corresponding square integral vanish by (4), and $E = 0$ already gives $f = 0$ almost everywhere on S. The proof also bounds the upper integral of $\mathbf{1}_S f$ for an arbitrary subset $S \subset \mathcal{P}$ and arbitrary nonnegative f satisfying the pointwise comparison, since $\mathbf{1}_{\mathcal{P}} E \sqrt{\boldsymbol{D}_j}$ is its measurable majorant in the same Cauchy–Schwarz calculation.

8.3.2 Which finite frequency survives the summation

For the same A, R, Z and j, write $D_l(p) = p_0 \cdots p_{l-1}$. On a nonzero finite orbital majorant all $D_l(p), D_l(q)$, $1 \le l \le j$, are integral: they are the leading minors of the two frames in integral representatives of the active cell. Let $u = \prod_{1\le l\le j} D_l(p)$ and $v = \prod_{1\le l\le j} D_l(q)$. Their quotients by any of the respective partial products or by any single pivot are integral. For example $u/p_i = (u/D_{i+1}(p))D_i(p)$. Thus $\alpha = g\overline{uv}\,uv$ is an integral self-conjugate scale clearing both heads, $1/g$, and all $1/p_i - 1/q_i$. For an exponent K_1 define $\Phi_{K_1}(p) = \mathrm{md}(u)^{-K_1/2}$ on this support. Set $\Phi_{K_1}(p) = 0$ when any required partial product is nonintegral. For $4E_1 \le K_1$, multiplicativity and $\mathrm{md}(g), \mathrm{md}(uv) \le 1$ give

$$\Lambda^{E_1} \le (2/\mathrm{md}(g))^{K_1}\Phi_{K_1}(p)\Phi_{K_1}(q).$$

Let $\xi \in A^\times$. The two local tests needed for the sieve are the additive character of ξ times the next diagonal Gram entry, and, at $j = i + 1$, the constant weight with $\xi = (1/p_i - 1/q_i)/g$. Write $\beta = \alpha^{200n^2}$. For the character test choose η componentwise to be either ξ or β, taking one of smaller valuation at each component. Then $\xi/\eta, \beta/\eta \in R$ and $\eta - \xi \in \beta R$. On the arguments z with $\beta z \in R$ the characters of ξz and ηz agree. Put $w = \beta/\eta$. For the other test take $\eta = \xi,\ w = \beta$. Then both α/g and $\alpha(1/p_i - 1/q_i)$ are integral. In both cases $w, \xi w, \xi/\eta \in R$, and, with $a_0 = \min(1, \mathrm{md}(\eta)^{-1/2})$,

$$a_0^2 \le \mathrm{md}(w)/\mathrm{md}(\beta).$$

The one-floor Plancherel bounds for these tests show that their absolute pivot integrand squared is at most $\mathrm{md}(\eta)^{-|\iota|}\mathcal{D}(x)$, where $|\iota|$ is the number of orthonormal coordinates. For the first test the Fourier transform is taken in the next frame coordinate, on which frequencies agreeing on $\beta^{-1}R$ give the same character. For the second it is taken in the Heisenberg coordinate, whose two translated characters differ by $(1/p_i - 1/q_i)/g$. Since $s_{\mathfrak{f}} \le |\iota|$, these are exactly the gains required by the shell bound whenever $\mathrm{md}(\eta)^{-1/2} < 1$.

Put $d(\xi) = [R : R \cap \xi^{-1}R]$. Since $wR \subset R \cap \xi^{-1}R$ and $\xi/\eta \in R$, we have $\mathrm{md}(w) \le d(\xi)^{-1}$ and $a_0 \le \mathrm{md}(\xi)^{-1/2}$. Use the shell estimate at $t + 12$, keeping t powers for the frequency and six pairs for its denominator. It gives for the floor integral I_j with either of these weights

$$\frac{|I_j|}{M_Z} \le \mathrm{md}(\xi)^{-t/2}d(\xi)^{-6}(2/\mathrm{md}(g))^{K_1}\Phi_{K_1}(p)\Phi_{K_1}(q).$$

Indeed the residual power of Λ is at most $E_1 = 10^9n^4 + 1200n^2$, including the six factors of $\mathrm{md}(\beta)^{-1}$. The assumptions $t + 10^{10}n^4 \le k,\ K_1 \ge 10^{10}n^4$ give $4E_1 \le K_1$ and the required shell range. In the sieve $t = k - n^7,\ K_1 = n^6$ satisfy these inequalities for sufficiently large n. The finite set contains all places over the rational level $g = g_0g_1$, so $\mathrm{md}(g) = g^{-2f}$. From $g \le b^{22}$ we obtain $(2/\mathrm{md}(g))^{n^6} \le b^{45fn^6} \le b^{fn^9}$. This gives the specified finite frequency bound in both uses of the floor recursion.

The contribution of the principal level. Put $\delta_R = \mathrm{vol}(R \setminus R^\times)$ and write the floor-zero form as

$$\mathcal{B}(Z) = \int \left| \int_Z \chi_g(yxy^{-1})\, dM(y) \right|^2 dM(x).$$

On $K(g)$ each conjugated character equals one, since $K(g)$ is normal in K. The contribution of this coset is exactly $M(Z)^2M(K(g))$.

For the finite bottom use $j = 0$ and the depth-one gauge in which every size set is R. The Weyl elements are the identity and Q_0 is the whole isometry group, so both points are near at zero. The head scaling is empty. A nonzero coset-vector pairing supplies the integral box representatives and integral divided difference as above. If every $L_a(z, x)$ contained a unit of R, those inverses would also be integral, and the reconstruction would give integral $(x - I)/g$. Since $g \in R$, this makes x integral and its unitary inverse is integral as well. This would put x in $K(g)$.

For a position outside $K(g)$, therefore, every active pair has some $L_a(z,x)$ with no unit of R. By the finite ideal calculation, it is an integral ideal d_aR disjoint from $R^\times$, with $[R : d_aR] \ge 2$. Thus the cap on its additive volume is $\min(\mathrm{vol}(R \setminus R^\times), 1/2)$. This is the cap for the exterior term. The contribution of $K(g)$ itself comes from the principal phase identity.

In the integral gauge, take $j = 0$, $p = q = 1$ and $e = \min(\delta_R, 1/2)$. The principal reconstruction gives the cap e at every nonzero pair outside $K(g)$. Here the two absolute pivots coincide. The coincident-pivot bound (8.8) gives $f_0^2 \le e^{2s_{\mathfrak{f}}} \mathcal{D}$ for the floor-zero majorant. The modulus of the phase integral is v_U times the modulus of the floor-zero upper pairing, so its square is at most $v_U^2 f_0$. Apply the shell bound to $1_{K(g)^c} f_0$ with $t' = t + 1 + n + 10^9 n^4$. If $t + 10^{10} n^4 \le k$, then $t' + 10^8 n^4 \le k$ and $2^{10^9 n^4} e^{t'} \le 2^{-n} \delta_R^{t+1}$. Restoring $v_U^2 M_Z = M(Z)M(K(g))$ yields

$$\mathcal{B}(Z) \le M(Z)M(K(g))\big(M(Z) + 2^{-n}\delta_R^{t+1}\big).$$

The zero cap also gives this estimate when $\delta_R = 0$.

8.4 Real correlations and their Haar measures

8.4.1 The real correlation forms for the open term

Work at one real place and write $F_\flat = W^\flat$. For $p \in (\mathbf{C}^\times)^{\varrho+1}$, let w_p be the Weyl element which sends e_a to $p_a e_a^*$ and e_a^* to $\bar{p}_a^{-1} e_a$ for $a < \varrho$. Let σ_c fix the hyperbolic pairs and act on the plane of a specified positive unit vector and z by

$$\begin{pmatrix} \bar{u} & \sqrt{|u|^2 - 1} \\ -\sqrt{|u|^2 - 1} & -u \end{pmatrix}, \qquad u = c/\min(1, |c|).$$

It is the identity on the other positive coordinates. For $\gamma = w_p \sigma$ with $\sigma \in U(X_0 \oplus \mathbf{C}z)$, write

$$J_F(\gamma) = \int_{G_1} F_\flat(y) F_\flat(\gamma^{-1} y)^* \, dy.$$

This Haar is the filling Haar on the group. To check absolute convergence, put $K_R = \{y : \|y\| \le R\}$. Both factors can be nonzero only on $UK_R \cap \gamma U K_R$. Conjugation by γ sends U to the opposite unipotent. In a factorization of a point of this intersection the product of the inverse lower and the upper unipotent factors lies in $\gamma K_R K_R^{-1}$. Elimination on lower-times-upper unitriangular matrices recovers both factors with pivots one. Thus the supports have compact intersection, locally uniformly in these parameters. This also gives continuity of the correlation by dominated convergence.

Define

$$c_G = \frac{\mathrm{vol}\, U(k)}{\mathrm{vol}\, U(r)} \frac{2\pi^k}{(k-1)!}, \qquad \Theta = \pi^{(2\varrho+1)k} \frac{\mathrm{vol}\, U(k)}{\mathrm{vol}\, U(r)},$$

$$d_a = k - 1 + 2(\varrho - a) \quad (a < \varrho), \quad \Xi(p) = \frac{\sqrt{c_G}}{d} (|p_\varrho|^2 - 1)_+^{(k-1)/2} \prod_{a<\varrho} |p_a|^{d_a} J_F(w_p \sigma_{p_\varrho}).$$

The top form is $J_n(p,q) = \langle \Xi(p), \Xi(q) \rangle_{\mathrm{HS}}$, with the Hilbert–Schmidt inner product linear in the second variable. It is continuous and positive semidefinite when the lists have a common terminal entry. For any one σ with $h(z, \sigma z) = p_\varrho = q_\varrho$, one may replace σ_{p_ϱ} in both arguments by σ. When the terminal modulus is one the scalar is zero. When it is larger than one, aligning the positive part of σz by H gives a factorization $\sigma = h_1 \sigma_{p_\varrho} h_2$ with $h_1, h_2 \in H$. Indeed the remaining isometry fixes z. Then J_F in both arguments is multiplied on the left by $\rho_H(h_1)$ and on the right by $\rho_H(h_2)$. These changes preserve the inner product.

For a rational packet with lower pivots p_a and terminal pairing $t = h(z, \sigma z)$ in the rational frame of norm $-N$, the input to this real form is the list

$$p'_{va} = s_{va}^2 \Phi_v(p_a) \quad (a < \varrho), \qquad p'_{v\varrho} = \Phi_v(t)/N.$$

Indeed the basis change conjugates each Weyl block by $\mathrm{diag}(s_{va}, s_{va}^{-1})$, with the inverse on the left, and scales the terminal pairing by $s_{v\varrho}^2$. Equivalently the leading minor of order j is multiplied by $\prod_{a<j} s_{va}^2$ under Φ_v. The condition $|\Phi_v(t)| > N$ makes the real terminal modulus larger than one. These substitutions apply to both members of each pair at the same t, with the same $F_\flat$.

Here are the lower forms and their estimates in these measures. For $i \le \varrho$, let G_i denote the group fixing the earlier hyperbolic pairs. At a lower floor $i < \varrho$, its measure θ_i is the law of $w_i(c)^{-1}u$, with $|c|^{2d_i}\,dc$ and Lebesgue Haar $d\eta_i(u)$ on the corresponding Heisenberg unipotent. At the terminal floor use the law of $\sigma_c^{-1}h$ with $2\pi^k(|c|^2-1)_+^{k-1}dc/(k-1)!$ and probability Haar on H. This is polar integration on the positive vector in $\sigma_c^{-1}z$, whose squared length is $|c|^2-1$. At floor n use probability Haar on H. Write M_i for the ordered convolution of $\eta_a * \theta_a$ for $a < i$, highest floor on the left, and $\mathcal{M} = M_{n+1}$. There are no Heisenberg factors at or above ϱ. Also write $B_i^D(z,z')$ for the average of $\langle F_\flat(vz), F_\flat(Dvz')\rangle$ against the convolution of $\theta_n,\dots,\theta_i$, in that order. For truncated Weyl products $w_p^{(i)}$, $w_q^{(i)}$ of length i, put $D = (w_q^{(i)})^{-1}w_p^{(i)}$. Define

$$\mathcal{I}_i(p,q;\omega) = \prod_{a<i}|p_aq_a|^{d_a}\int_{x,g}\int_{\sigma,\tau}\omega(\sigma\tau^{-1})B_{i+1}^1(\sigma g,\sigma gx) \cdot \overline{B_{i+1}^D\big(\tau(w_p^{(i)})^{-1}g,\tau(w_p^{(i)})^{-1}gx\big)}\,d\theta_i(\sigma)d\theta_i(\tau)\,dM_i(g)d\mathcal{M}(x).$$

The superscript 1 means the identity operator.

The real Gram cell and its Haar measure. Write $m=\varrho$, $n=m+1$, $W=X_0\simeq\mathbf{C}^k$ and $r=2m+k+1$. Use the Euclidean-orthonormal split basis in the order $(E,W,\mathbf{C}z,D)$, with $E=(e_a)_{a=1}^m$, $D=(e_a^*)_{a=1}^m$, and take x_0 to be the designated basis vector of W. The Hermitian involution J interchanges E,D, acts by 1 on W and by -1 on z. Let $F=(E,z)$, $F'=(D,x_0)$, $p=\mathrm{diag}(0_m,1)$, and $p_0=1-p$. The mixed Gram map is $Z(y)=F^*JyF$. Write $D_m(Z)$ for the leading principal minor of order m. The cofactor formula gives, for invertible Z,

$$\mathfrak{g}(Z)=|\det Z|^2-|D_m(Z)|^2=|\det Z|^2\delta(Z),\qquad \delta(Z)=1-|(Z^{-1})_{nn}|^2.$$

Thus the open Gram cell $\Omega_{\mathrm{Gr}}=\{\mathfrak{g}>0\}$ is the domain on which Z is invertible and $\delta>0$. Write $\Omega_G=\{y\in G_1 : Z(y)\in\Omega_{\mathrm{Gr}}\}$.

The section and its stabilizer. For $Z\in\Omega_{\mathrm{Gr}}$, put

$$K_Z=(Z^{-1})^*pZ^{-1},\qquad T_Z=-p-\tfrac12 p_0K_Z(1+p),\qquad \Lambda=p_0+\sqrt{\delta(Z)}\,p.$$

The section γ_Z used below sends F,F' to

$$X_Z=(FT_Z+F'\Lambda)Z,\qquad X_Z'=\big(F(Z^{-1})^*-X_ZT_{Z^*}\big)\Lambda^{-1},$$

and fixes $W\ominus\mathbf{C}x_0$ pointwise. These are the section matrices in the filling coordinates. Indeed $F^*JF=-p$, $(F')^*JF'=p$, $F^*JF'=p_0$. Consequently $F^*JX_Z=Z$. The Gram of X_ZZ^{-1} has blocks $-(K_Z)_{11},-(K_Z)_{1n},-(K_Z)_{n1}$ and $\delta(Z)-1=-(K_Z)_{nn}$, so $X_Z^*JX_Z=-p$. Using $pT_{Z^*}=-p$ gives $X_Z^*JX_Z'=p_0$. For the last block, $T_{Z^*}+T_{Z^*}^*=-2p-K_{Z^*}+pK_{Z^*}p$ gives

$$(X_Z')^*JX_Z'=\Lambda^{-1}(p-pK_{Z^*}p)\Lambda^{-1}=p.$$

The form on the span of F,F' is nondegenerate, so the Gram identities give $\gamma_Z\in G_1$. The displayed formulas are smooth on Ω_{Gr}.

The pointwise stabilizer of F is $S=U^0H$, with unique factorization $u(B,A)h$, where $H=U(W)$ acts trivially on the other blocks and

$$u(B,A)=\begin{pmatrix}1&-B^*&0&A-\frac12B^*B\\0&1&0&B\\0&0&1&0\\0&0&0&1\end{pmatrix},\qquad B\in M_{k,m}(\mathbf{C}),\quad A^*=-A.$$

In fact the fixed frame forces the D and z rows to agree with the identity. Pairings of the remaining columns force the W-block h to be unitary, the E, W-block to be $-B^*h$, and the Hermitian part of the E, D-block to be $-B^*B/2$. The free entries of $u(B, A)$ in the E, W-block are $-B^*$, and those in the E, D-block are the strictly upper entries and the imaginary diagonal. At fixed B the upper entries differ from those of A by a translation, and the imaginary diagonal is the same as in A. Thus $du = dB\, dA$ is the prescribed free-entry measure. The U^0 product adds the two B's and adds the two A's with a skew Hermitian shear. Left and right multiplication therefore preserve du, as does conjugation by h, which sends B to hB and fixes A. This also gives the bi-invariant product Haar on S.

The normalized frame density. Use independent-entry Lebesgue measure on skew Hermitian matrices, with complex off-diagonal entries and real imaginary diagonal entries. Write $d\omega_G$ and $d\omega_H$ for the raw Haar densities whose values at the identity use these coordinates on JY and Y, respectively. The filling measure is $dy = d\omega_G/\operatorname{vol} U(r)$. In fact the Cayley map of a skew Hermitian matrix V is $g = 2C^{-1} - 1$ with $C = 1 - J^{-1}V/2$. Its differential in left-translated coordinates is $Jg^{-1}dg = C^{-*}(dV)C^{-1}$, where $C^{-*} = (C^{-1})^*$. On Hermitian or skew Hermitian matrices of size j, congruence by a matrix C has absolute real determinant $|\det C|^{2j}$. Unitary factors preserve Euclidean volume, whose density is a constant times independent-entry Lebesgue measure. Positive diagonal singular values (s_i) contribute s_i^2 in the real diagonal and $(s_i s_j)^2$ in each complex off-diagonal position. This proves the determinant formula and gives the filling density on the Cayley chart with the prescribed divisor $\operatorname{vol} U(r)$. Its complement, where $g + 1$ is singular, has Haar measure zero. For each g there are finitely many phases t in one period with $1 + e^{it}g$ singular, so central invariance and Tonelli give this null-set assertion.

For $H = U(k)$, the last-column map has stabilizer $U(k-1)$. At the base column x_0, the off-diagonal entries and imaginary diagonal in that column give the Euclidean sphere tangent density with absolute determinant one. The other entries give the same raw Haar on $U(k-1)$. The unitary action preserves sphere area, so these densities disintegrate $d\omega_H$. The sphere area follows by polar integration of the Gaussian on $\mathbf{C}^k$. The compact volumes satisfy the resulting recurrence

$$s_k = \frac{2\pi^k}{(k-1)!}, \qquad \operatorname{vol} U(k) = s_k \operatorname{vol} U(k-1), \qquad dh = \frac{d\omega_H}{\operatorname{vol} U(k)}.$$

Start with $\operatorname{vol} U(0) = 1$. Here s_k is the area of the unit sphere of W and dh is probability Haar. Its image under $h \mapsto hx_0$ is unit sphere probability. These are the compact volumes in the filling normalization.

On full-rank $r \times n$ complex matrices X, the map $q(X) = X^*JX$ is a submersion onto its image in the space of Hermitian matrices. In its differential $X^*J\, dX$ can be arbitrary. On that Hermitian space use the density $d\lambda(q)$ in the independent complex off-diagonal entries and the real diagonal entries $q_{ii}/2$. Disintegrating ambient Lebesgue dX against this density gives a smooth density $d\mu$ on the fibre $q = -p$. The action of G_1 preserves q and dX, since $|\det y| = 1$, and therefore preserves $d\mu$.

At the base frame F, put $P = F^*J\, dX$. Then $dq = P + P^*$. For $i < j$, passage from $(P_{ij}, \overline{P_{ji}})$ to the ij entries of $(P - P^*)/2$ and dq has complex determinant 1. On the diagonal the real part of P_{ii} gives $dq_{ii}/2$ and the imaginary part gives the independent skew entry. Hence $d\mu$ at F is Lebesgue on these skew entries together with all entries of the $F^\perp$-block of $J\, dX$, where $F^\perp$ denotes the Euclidean orthogonal complement. The differential of $y \mapsto yF$ reads just these entries in the F-columns of JY. Its kernel is the remaining skew Hermitian block on $F^\perp = (D, W)$. On the tangent of $u(B, A)h$ this remaining block is $\left(\begin{smallmatrix} A & -B^* \\ B & Y_H \end{smallmatrix}\right)$. Its independent entries have density $du\, d\omega_H$.

It follows that $d\omega_G$ disintegrates over the frame orbit with quotient density $d\mu$ and right stabilizer density $du\, d\omega_H$. In detail, the differential just computed makes the orbit open in the full-rank fibre and supplies smooth local lifts y_X with $y_X F = X$. Integrating a test on $y_X S$ against the raw stabilizer Haar and then $d\mu(X)$ defines a left-invariant smooth density. The two tangent densities at the identity give $d\omega_G$, fixing its scale. Inversion on S preserves its Haar, so the right factor can be written $(vh)^{-1}$ with density $\operatorname{vol} U(k)\, dv\, dh$.

The density on the positive-discriminant cell. For $X = yF$, use $Z = F^*JX$ as part of the ambient coordinates. It records the D-rows and the negative of the z-row. At fixed Z, right-multiply X by Z^{-1}. Write the E- and W-blocks of XZ^{-1} as (L, a) and (B, x), with $a \in \mathbf{C}^m$ and $x \in W$ the last columns. The other two blocks are fixed, and the constraint is

$$q' = Z^{-*}q(X)Z^{-1} = \begin{pmatrix} L + L^* + B^*B & a + B^*x \\ a^* + x^*B & \|x\|^2 - 1 \end{pmatrix} = -K_Z.$$

Returning to X has ambient Jacobian $|\det Z|^{2(m+k)}$ on the $m + k$ free complex rows. Returning from q' to q has Jacobian $|\det Z|^{2n}$ on Hermitian matrices. Both are changes with Z fixed, so their ratio is also the Jacobian in the joint coordinates with dZ. In the upper-left block the change from L to its skew part $A = (L - L^*)/2$ and $L + L^*$ has the same determinant as at F. The off-block a is translated by B^*x. For the last real coordinate use $q'_{nn}/2 = (\|x\|^2 - 1)/2$. Polar integration at $\|x\|^2 = \delta(Z) > 0$ leaves $\delta(Z)^{k-1}d\sigma(\xi)$, where $x = \sqrt{\delta(Z)}\,\xi$ and $d\sigma$ is unit sphere area. Therefore the frame density in these coordinates is

$$d\mu(X) = |\det Z|^{2(k-1)}\delta(Z)^{k-1}\,dZ\,dB\,dA\,d\sigma(\xi) = \mathfrak{g}(Z)^{k-1}\,dZ\,du\,d\sigma(\xi).$$

All coordinates of X are recovered since

$$L = -\tfrac{1}{2}\big((K_Z)_{11} + B^*B\big) + A, \qquad a = -(K_Z)_{1n} - B^*x.$$

For a lift $k_\xi \in H$ of the column $\xi = k_\xi x_0$, the frame $u(B, A)k_\xi\gamma_Z F$ has exactly these blocks. Thus the coordinates cover the whole positive-discriminant part of the fibre. Two lifts differ by right multiplication by $H_0 = U(W \ominus \mathbf{C}x_0)$, which commutes with γ_Z. For $y = uk_\xi\gamma_Z s^{-1}$, the simultaneous change $(k_\xi, s) \mapsto (k_\xi\ell, s\ell)$, $\ell \in H_0$, is precisely the redundancy: the frame determines Z, u, ξ, and for a fixed lift the right factor is unique.

The two disintegrations, with sphere probability supplied by probability Haar dk_0 on the left H, now give

$$\begin{aligned} \int_{\Omega_G} f(y)\,dy &= \int_H \mathcal{I}(k_0)\,dk_0, \\ \mathcal{I}(k_0) &= \int_{\Omega_{\mathrm{Gr}}} c_G\,\mathfrak{g}(Z)^{k-1} \int_H \int_{U^0} \int_{U^0} f\big(uk_0\gamma_Z(vh)^{-1}\big)\,du\,dv\,dh\,dZ, \\ c_G &= \frac{\operatorname{vol} U(k)}{\operatorname{vol} U(r)}\frac{2\pi^k}{(k-1)!}. \end{aligned}$$

Here the right-factor integral is independent of the chosen lift by right invariance on S. For every non-negative Borel f, Tonelli permits these extended integrals. In particular $\inf_{k_0 \in H} \mathcal{I}(k_0) \le \int_{G_1} f\,dy$. The infimum is over a single compact variable outside the Gram and stabilizer integrals. This is the lower bound used in the real cell contraction.

The Cayley-coordinate Jacobian. The column lift in the compact chart takes $(t, b) \in \mathbf{R} \times \mathbf{C}^{k-1}$ to the Cayley transform of $\left(\begin{smallmatrix} 0 & b \\ -b^* & it \end{smallmatrix}\right)$. Its last column is

$$\xi = (b/\Delta,\ 2/\Delta - 1), \qquad \Delta = 1 + \|b\|^2/4 - it/2.$$

The inverse on the sphere off $-x_0$ reads $\Delta = 2/(1 + \xi_k)$, $b = \Delta\,\xi_{<k}$, $t = -2\operatorname{Im}\Delta$. The sphere relation gives $\operatorname{Re}\Delta = 1 + \|b\|^2/4$, so this is a smooth section on that chart. To compute its density, first use (b, Δ) as complex ambient coordinates in the displayed map. Their Jacobian is $4|\Delta|^{-2(k-1)-4}$. On the sphere the absolute derivative of $(\|\xi\|^2 - 1)/2$ in $\operatorname{Re}\Delta$ is $2/|\Delta|^2$. After dividing by this derivative and using $d(\operatorname{Im}\Delta) = dt/2$ in absolute density, the sphere measure is $|\Delta|^{-2k}\,db\,dt$. This is the left Cayley weight w_L, since Δ is the determinant of its Cayley denominator.

On the right use the full compact Cayley chart, with weight $w_R = |\det(1 - Y_R/2)|^{-2k}$ for skew Hermitian Y_R. Its inverse exists on the set where $h + 1$ is nonsingular. With these two choices the group chart is

injective. The frame recovers Z, u, ξ, the column chart recovers its left lift, and $y^{-1}uk_\xi\gamma_Z = vh$ recovers the right parameters by the stabilizer factorization. These inverse formulas are smooth on the chart image. The omitted sphere point and the omitted right compact set have fibre measure zero. The group chart thus pulls back $d\omega_G$ with density $\mathfrak{g}(Z)^{k-1}w_L w_R$ in the indicated Lebesgue coordinates, and dy with this density divided by $\operatorname{vol} U(r)$. Splitting a skew Hermitian matrix into independent entries, or isolating its last column, preserves these coordinate densities. Integrating the left weight gives s_k and integrating the right gives $\operatorname{vol} U(k)$, as in the normalized formula above.

The ordered Haar measure and its quotient. Write V_i for the space after removing the preceding Witt pairs and $G_i^{\mathrm{fl}} = U(V_i)$, so $G_0^{\mathrm{fl}} = G_1$ in the filling notation. Write $d\omega_i$ for the Haar density whose value at the identity is independent-entry Lebesgue measure on the skew Hermitian matrix J_iX, $X \in \mathfrak{u}(V_i)$. The split frame has $J_i^2 = 1$. The Cayley chart used in the filling measure identifies $dy = d\omega_0/\operatorname{vol} U(r)$. At the final compact group $H = U(X_0)$, the prescribed probability Haar is $d\omega_n/\operatorname{vol} U(k)$. These densities are both left and right invariant. For example, the conjugation action on the Lie algebra has absolute real determinant one since its complexification is conjugation on the full matrix algebra.

At a Witt floor write the fixed frame vector as e, its opposite as f, with $h(e,f) = 1$, and identify $V_i = \mathbf{C}e \oplus V_{i+1} \oplus \mathbf{C}f$. The stabilizer S_i of e in G_i^{fl} is $G_{i+1}^{\mathrm{fl}}N_i$, where N_i is the Heisenberg group of that floor. Indeed $g \in S_i$ has f-coefficient one in gf, and $h(gf, gf) = 0$ gives

$$gf = f + a + (it - h(a,a)/2)e \qquad (a \in V_{i+1},\ t \in \mathbf{R}).$$

There is exactly one $n(a,t) \in N_i$ with this value at f. Then $n(a,t)^{-1}g$ fixes the pair. The compact or indefinite unitary action of G_{i+1}^{fl} on a preserves Lebesgue measure and fixes t. The group law in N_i is translation in a and translation plus a linear shear in t. Thus $d\omega_{i+1} * d\eta_i$ is bi-invariant Haar on S_i.

The orbit map $g \mapsto g^{-1}e$ identifies $S_i\backslash G_i^{\mathrm{fl}}$ with the nonzero isotropic quadric. Use the invariant measure on this quadric obtained by disintegrating complex Lebesgue measure against $d(h(v,v)/2)$. This is the real place's skew coordinate $\beta = (i/2)h(v,v)$, with Lebesgue measure on $\operatorname{Im}\beta$. In the chart $v = c(f + a + (u + it)e)$, $c \neq 0$, one has $h(v,v)/2 = |c|^2(u + h(a,a)/2)$. The ambient Jacobian is $|c|^{2d_i+2}$, since $d_i = \dim_{\mathbf{C}} V_{i+1}$. Dividing by $|c|^2$ gives $|c|^{2d_i}\,dc\,da\,dt$ on the quadric. The omitted complex hyperplane has real dimension smaller than the smooth quadric and hence has zero quadric measure. For $g = w_i(c)^{-1}n(a,t)$ its orbit value is $cn(a,t)^{-1}f$. Inversion in these Heisenberg coordinates sends (a,t) to $(-a,-t)$, which preserves the density. Thus this invariant quadric measure is exactly the image of the specified θ_i. At the point e, use the chart with e, f exchanged. Its quadric density at $c = 1, a = t = 0$ is one.

Integrating over the stabilizer on the left and this quotient on the right gives a Haar measure. For right invariance, a right translation takes the quotient to itself by its unitary action and changes a chosen lift by left multiplication in S_i. The inner integral is unchanged by right translation within S_i. The measure has a smooth positive local density, by a smooth local section. To determine that density at the identity, put $Y = J_iX$. The differential of the orbit map there is $-J_iYe$. The complex entries in the e-column of Y and the imaginary diagonal entry give the quotient coordinates (c,a,t) at e with absolute real determinant one. Indeed J_i interchanges e, f and is unitary on V_{i+1}, and the e-diagonal is sent to the imaginary f-coordinate. The kernel has this column zero. Its remaining independent entries are exactly those of the middle skew Hermitian block, the column coupling the middle block to f, and the imaginary f-diagonal. These give the stabilizer density $d\omega_{i+1} * d\eta_i$, with absolute determinant one. The infinitesimal action of $n(a,t)$ on f is $a + ite$, so its last column in Y reads $J_{i+1}a + itf$. Hence

$$d\omega_{i+1} * d\eta_i * d\theta_i = d\omega_i \qquad (i < \varrho).$$

At the terminal floor use $V_\varrho = X_0 \oplus \mathbf{C}z$, $h(z,z) = -1$ and the stabilizer of z is H. On $h(a+cz, a+cz) = -1$, disintegration against $d(h/2)$ and ordinary polar integration in $a \in \mathbf{C}^k$ give the density $2\pi^k(|c|^2 -$

$1)_+^{k-1}\,dc/(k-1)!$ and the uniform sphere probability in a of squared length $|c|^2-1$. That sphere probability is the image of the right compact factor of θ_ϱ. The terminal boost sends z to a vector with z-coordinate $-c$ and first component of squared length $|c|^2-1$. The change $c\mapsto -c$ preserves the density, so this is precisely the orbit law. Near z, write $c=e^{i\phi}\sqrt{1+\|a\|^2}$. The ambient polar factor $|c|\,d|c|\,d\phi$ in the last coordinate becomes $d(h/2)\,d\phi$ in absolute density, giving $da\,d\phi$ on the hyperboloid. The differential $-J_\varrho Yz$ identifies these coordinates at z with the off-diagonal z-column and the imaginary z-diagonal with absolute determinant one. The remaining middle entries are the stabilizer density $d\omega_n$. This gives $d\omega_n * d\theta_\varrho = d\omega_\varrho$. In particular no extra scalar comes from either probability compact factor. Iterating the two floor formulas yields

$$\mathcal{M}=\frac{d\omega_0}{\operatorname{vol}U(k)}=\lambda_0\,dy,\qquad \lambda_0=\frac{\operatorname{vol}U(r)}{\operatorname{vol}U(k)}>0,\qquad \lambda_0\frac{\operatorname{vol}U(k)}{\operatorname{vol}U(r)}\le 1.$$

For the quotient formula, each later floor group fixes the preceding pair and conjugates $n(a,t)$ by a unitary change of a with t fixed. Therefore η_i commutes under convolution with the measures supported in the later floor groups. Moving these factors to the left in increasing order gives

$$\mathcal{M}=(\eta_0*\cdots*\eta_{\varrho-1})*(\theta_n*\cdots*\theta_0).$$

The first factor is Haar on U. Successively fixing the opposite Witt vectors writes each $u\in U$ uniquely in these ordered factors, with each later group acting on the preceding Heisenberg group by the same measure preserving conjugation. To fix the scale, the prescribed free entries are in the row of e_i, in the columns of the middle block and the later Witt pairs, together with the imaginary part in the column of f_i. At the identity the i-th Heisenberg factor gives the row $-a^*J_{i+1}$ in those complex entries and t in the real entry, with all other free rows zero. The real determinant has absolute value one, so the first factor is exactly du. Put $\Theta_0=\theta_n*\cdots*\theta_0$. Comparing the identity $\mathcal{M}=du*\Theta_0$ with $dy=du\,d\dot y$ gives the image measure $\lambda_0\,d\dot y$ of Θ_0 on $U\backslash G_1$. Explicitly, use a measurable section of this quotient and multiply a nonnegative quotient test by a function on U of Haar integral one in the section coordinates. Integration over U gives the test, so both iterated integrals give the asserted quotient identity.

The lower bound for the full Gram chart is used separately in the real cell contraction.

Set

$$a_H=\operatorname{vol}U(k)/\operatorname{vol}U(r),\qquad J_i(p,q)=\frac{a_H}{(\lambda_0 d)^2}\mathcal{I}_i(p,q;1).$$

Common-entry descent on the real level. Use the real cut vector $F_\flat$ and the ordered measures above. For $0\le i\le\varrho$, use $G_i^{\mathrm{fl}}=U(V_i)$ for the group fixing the first i hyperbolic pairs. Write f_i for the next frame vector, which is e_i for $i<\varrho$ and z for $i=\varrho$. Its stabilizer in G_i^{fl} is $S_i=G_i^{\mathrm{fl}}\cap Q_{i+1}$. Let $D=w_q^{-1}w_p$, with both Weyl products truncated at i. These products and D commute with G_i^{fl}. The two upper pairings B_{i+1}^1 and B_{i+1}^D are invariant under simultaneous left multiplication by S_i in their two arguments, by the stabilizer invariance of the upper averages.

The floor measure θ_i pushes forward to the invariant quadric measure under $\tau\mapsto\tau^{-1}f_i$. It therefore gives the same integral to a left S_i-invariant observable of τ when τ is replaced by τs, for any $s\in G_i^{\mathrm{fl}}$. Indeed the observable descends to the orbit and the substitution acts there by s^{-1}. The entry $Z_i(v)=h(f_i,vf_i)$ is invariant under both left and right multiplication by S_i, since S_i fixes f_i. Write the ordered disintegration

$$d\theta_i=\big(c\mapsto v_i(c)^{-1}\big)_*\big(\rho_i(c)\,dc\big)*d\zeta_i.$$

For $i<\varrho$, $v_i(c)$ is the Weyl element on the next pair with entry c, $\zeta_i=\eta_i$ is its Heisenberg measure, and $\rho_i(c)=|c|^{2d_i}$ with $d_i=k-1+2(\varrho-i)$. At $i=\varrho$, $v_i(c)$ is the terminal boost and $\zeta_i=\theta_n$ is probability Haar

on $H = U(X_0)$. The density from the skew-coordinate polar calculation in the ordered Haar comparison is

$$\rho_{\varrho}(c) = \frac{2\pi^{k}}{(k-1)!}(|c|^{2}-1)_{+}^{k-1}.$$

The skew base measure is $d(\beta/2)$ for the Hermitian level $\beta = h(v,v)$. On the support of ρ_i we have $Z_i(v_i(c)) = c$ almost everywhere. In both cases ζ_i is carried by S_i.

Put $\omega(v) = \mathrm{wave}(\xi Z_i(v))$. In the integral $\mathcal{I}_i(p,q;\omega)$, hold x, y, σ fixed and first integrate over τ. Its integrand has the three factors

$$\omega(\sigma\tau^{-1})\, B_{i+1}^{1}(\sigma y, \sigma y x)\, \overline{B_{i+1}^{D}(\tau w_p^{-1} y, \tau w_p^{-1} y x)}.$$

The observable of τ is left S_i-invariant, by the right invariance of Z_i under S_i and the simultaneous left invariance of the mixed pairing. Use the orbit change above to replace τ by $\tau_0\sigma$ inside the integral and then put $\tau_0 = v_i(c)^{-1}u$, $u \in S_i$, using the displayed disintegration. The phase argument becomes $\sigma\tau^{-1} = u^{-1}v_i(c)$, hence its entry is c. Put $y' = u\sigma y$. The first pairing is unchanged when $\sigma y, \sigma y x$ are replaced by $y', y'x$. Since both u and σ lie in G_i^{fl},

$$\tau w_p^{-1} y = v_i(c)^{-1} u \sigma w_p^{-1} y = (w_p v_i(c))^{-1} y'.$$

The same change gives $(w_p v_i(c))^{-1} y' x$ in the other argument. The phase is $\mathrm{wave}(\xi c)$. The new measure of y' is precisely $d(\zeta_i * \theta_i * M_i)$.

Apply the change first to the product of the absolute values of the two pairings, using nonnegative integrals. Its initial integral against $\mathcal{M}$, M_i and the two copies of θ_i is the integral of the absolute majorant $\mathcal{F}_i^{+}$, finite by the real shell estimate at $t_1 = 1$. Thus the changed integral is absolutely convergent, and Fubini also applies with its complex phase. Below the terminal floor the new measure is $M_{i+1} = \eta_i * \theta_i * M_i$. Appending the same c to the two heads changes the Weyl products to $w_p v_i(c)$, $w_q v_i(c)$, keeps the ratio D, and multiplies the head density by $|c|^{2d_i} = \rho_i(c)$. Write $p \oplus c$ for replacing the entry at i by c, with the suffix ignored at level $i+1$. The successive-average identity for the two upper pairings identifies their product with the form at the next floor and weight one. Consequently the normalized forms satisfy

$$\int_{\mathbb{C}} J_{i+1}(p \oplus c, q \oplus c)\, \mathrm{wave}(\xi c)\, dc = \frac{a_H}{(\lambda_0 d)^2}\, \mathcal{I}_i(p,q;\mathrm{wave}(\xi Z_i)). \tag{8.11}$$

For $i < \varrho$ the appended entry at $c = 0$ can be chosen as the unit 1, with the displayed phase still evaluated at $c = 0$. This point has measure zero. The terminal case of (8.11) uses the Gram computation below.

The terminal Gram in the same measures. At $i = \varrho$ both η_{ϱ} and η_n are point masses at the identity, and $\theta_n = \zeta_{\varrho}$. Thus $\zeta_{\varrho} * \theta_{\varrho} * M_{\varrho} = \mathcal{M} = \lambda_0\, dy$. The upper averages at floor $n = \varrho + 1$ are over probability Haar on H. The cut vector satisfies $F_{\flat}(hy) = \pi_H(h) F_{\flat}(y)$ with $\pi_H(h)$ unitary. As D commutes with H, both displayed upper averages therefore equal their raw pairings. For $|c| > 1$ put $\gamma_p = w_p v_{\varrho}(c)$ and $\gamma_q = w_q v_{\varrho}(c)$. The commutation with the terminal group gives $D\gamma_p^{-1} = \gamma_q^{-1}$. In the remaining integral in x, y' put $y'' = y'x$, using left invariance of Haar for the x-integral. The product of measures is $\lambda_0^2\, dy'\, dy''$. The two raw factors become

$$\langle F_{\flat}(y'), F_{\flat}(y'')\rangle \overline{\langle F_{\flat}(\gamma_p^{-1} y'), F_{\flat}(\gamma_q^{-1} y'')\rangle}.$$

For each terminal $\gamma = w_p v$, $v \in G_{\varrho}^{\mathrm{fl}}$, the products of coordinates used in $J_{F_{\flat}}(\gamma)$ have compact support in the integration variable. In fact, if both $F_{\flat}(y)$ and $F_{\flat}(\gamma^{-1}y)$ are nonzero, take $u \in U$ with uy in the cut box. Since $G_{\varrho}^{\mathrm{fl}} \subset Q_{\varrho}$, the second nonzero value makes $w_p^{-1} y$ near the box at ϱ. The head minor calculation bounds $u e_a^{*}$ for every $a < \varrho$ in any gauge in which p is a head. Use the test vectors e_a^{*} in (8.9), with $v = u$ and $g = 1$. The vectors and $(I-u)e_a^{*}$ are bounded, so the entries of u^{-1} are bounded in that gauge. Hence so are the entries of $y = u^{-1}(uy)$. The isometry group is closed in matrix space, so this bounds

the support in a compact set. The same gauge works locally in the nonzero head variables. This gives absolute integrability of the tensors and continuity by dominated convergence in the boost and heads.

Recall

$$J_{F_\flat}(\gamma)_{a,b} = \int F_{\flat,a}(y)\overline{F_{\flat,b}(\gamma^{-1}y)}\,dy.$$

Expanding the two raw factors, summing the finite coordinates, and applying Fubini on the compact support gives $\lambda_0^2\langle J_{F_\flat}(\gamma_p), J_{F_\flat}(\gamma_q)\rangle$, with the inner product linear in its second argument. The normalization of the form cancels λ_0^2. Write

$$c_G = a_H\,\frac{2\pi^k}{(k-1)!}, \qquad \Xi(p;c) = \frac{\sqrt{c_G}}{d}(|c|^2-1)_+^{(k-1)/2}\prod_{a<\varrho}|p_a|^{d_a}J_{F_\flat}(w_p v_\varrho(c)).$$

Define $J_n(p,q)$ on all unit heads by the inner product of $\Xi(p;p_\varrho)$ and $\Xi(q;q_\varrho)$. These top values form a positive semidefinite Gram kernel. The density $\rho_\varrho(c)$ multiplied by the earlier head density supplies exactly the two scalar factors in $\langle\Xi(p;c),\Xi(q;c)\rangle$. Thus the new inner integral at $i=\varrho$, with its density and normalization, is $J_n(p\oplus c, q\oplus c) = \langle\Xi(p;c),\Xi(q;c)\rangle$. Both sides vanish for $|c|\le 1$, since $k>1$. The padding unit 1 at $c=0$ also gives a zero terminal vector. The density and correlation formula prove (8.11) also at the terminal floor.

On the actual terminal group let $\gamma\in G_\varrho^{\mathrm{fl}}$ with $h(z,\gamma z)=c$. Its positive component of γz has squared length $|c|^2-1$, since $h(\gamma z,\gamma z)=-1$. For $|c|>1$, H acts transitively on the sphere of that length and fixes z. Take $h_1\in H$ sending $v_\varrho(c)z$ to γz. Then $h_2=(h_1v_\varrho(c))^{-1}\gamma$ fixes z and lies in H. Hence $\gamma = h_1v_\varrho(c)h_2$. Changing variables in the tensor integral and using the unitary law for $F_\flat$ gives for $a,b\in H$

$$J_{F_\flat}(a\gamma b) = \big(\pi_H(a)\otimes\overline{\pi_H(b^{-1})}\big)J_{F_\flat}(\gamma).$$

The w_p, w_q commute with h_1, so replacing the boost by γ applies the same unitary operator to both tensors. At $|c|=1$ the polar prefactors give zero for either choice. This is exactly the top Gram on the prescribed terminal element and head data. In particular the real inserted-frequency bound on the right of (8.11) applies to the same sequence of forms whose top belongs to the open packet. The unequal-head bound applies directly at its own floor, and the bottom estimate concerns J_0 in this same normalization.

Put $\rho=R=b^{6n+2}$ and $L=L_1\ge\rho$. The real level has $g=1$, with $k-1$ orthonormal coordinates and exponent $s_0=k-3n-2$. Use the ordered measure comparison above, $dM=\lambda_0\,dy$ and $dM_0^+=\lambda_0\,d\dot y$ on left U-invariant integrands, with $a_H\lambda_0\le 1$. For the first vector use $F_\flat$. For the second take a Bessel vector W' with the same character, supported near floor j, with squared quotient norm at most $\epsilon_1^2 d$. The pointwise inner-product inequality gives

$$\mathcal{N}(F_\flat,F_\flat)\le(\lambda_0 d)^2, \qquad \mathcal{N}(F_\flat,W')\le\epsilon_1^2(\lambda_0 d)^2,$$

since the squared quotient norm of $F_\flat$ is at most d. Thus the factor $N_0=\epsilon_1(\lambda_0 d)^2$ becomes $a_H\epsilon_1$ after multiplication by the form constant.

In Lebesgue coordinates $z=x+iy$, the oscillator character is $\mathrm{wave}(\xi z)=e^{2\pi i((\operatorname{Im}\xi)x+(\operatorname{Re}\xi)y)/\pi}$. The change to Euclidean frequency gives Plancherel constant π^2 in one complex coordinate. Multiplication by ξ has modulus $|\xi|^2$. Hence the shared one-floor estimates have squared factor $(\pi^2/|\xi|^2)^{k-1}$, both for the inserted diagonal frequency and, when $j=i+1\le\varrho$, for the Heisenberg frequency $\xi=p_i^{-1}-q_i^{-1}$. The floor-modulus product is one at $g=1$.

Fix $t_1\ge 1$ and $0\le f\le\mathcal{F}_j^+$. When $t_1>1$ assume in addition $f^2\le((\pi\rho^{24})^k/t_1^{s_0})^2\mathcal{D}$. Put $T_i=t_1L^{11ni}$ and let B_i be the positions where every entry of $x-I$ has absolute value at most T_{i+1}^{99n}. The sets B_i increase to the whole position set. In the annulus $A_i=B_i\setminus B_{i-1}$, with $B_{-1}=\emptyset$, the restricted integral uses a second gauge of base T_{i+1}^{99n} and constant T_{i+1}^{198n}. At the real place $A=\mathbf{C}$, $g=1$, the cut box is the operator-norm ball of radius R_{cut} and a gauge of base $L\ge R_{\mathrm{cut}}$ has scalar size i given by $|c|\le L^i$. With $R_{\mathrm{cut}}=R$ and $t=t_1$, let the nested position sets have entrywise radii T_{i+1}^{99n} for $x-I$. On annulus $i\ge 1$, a position is outside

the preceding set of radius T_i^{99n}. The head gauge has base L and the estimating gauge has base T_i. When $j > 0$, the head scaling holds because $L^{10n} \le T_i$. At an active pair both values of the cut vector are nonzero, so the two U-translates into the cut box exist. Their difference has size two. The required Weyl translates are near by the active-pair condition.

If each $L_a(z,x)$, $j \le a < \varrho$, had area greater than π/T_i^2, then in each there would be c_a outside the open disk of radius $1/T_i$. This is a nonzero complex number, with $|c_a^{-1}| \le T_i$. The reconstruction would give $|(x-I)_{bc}| \le T_i^{99n}$ at every entry, contradicting the annulus condition. The same implication rules out an active pair on that annulus when $j = \varrho$. On the initial set $i = 0$ the application instead uses the supplied pointwise frequency bound when $t > 1$, and the unrestricted cap when $t = 1$.

Thus the small cap on A_i, $i \ge 1$, is π/T_i^2. On A_0, take the supplied pointwise bound when $t_1 > 1$ and the bound with cap one when $t_1 = 1$.

The cut gauge Γ_ρ supplies the current floor mass. Its size-seven fibre and the $k-1$ size-eleven coordinates give

$$V \le \rho^{2574n}(\pi\rho^{22})^{k-1} \le \rho^{2574n+24(k-1)}.$$

The constant of the cut gauge is ρ^2, and $\rho \ge 4$ gives the last inequality. With cap π/T_i^2 the pointwise factor from (8.5) is bounded by $\pi^{s_0/2} T_i^{-s_0} \rho^{2574n+24(k-1)}$. All three cases therefore allow $E_i = (\pi\rho^{24})^k \rho^{9999n} T_i^{-s_0}$. In the restricted integral the head constant is L^2, and $T_{i+1}^{198n\cdot 50000n^2} = T_i^{9900000n^3} L^{108900000n^4}$. The normalized bound on A_i is thus

$$\Theta\epsilon_1 \rho^{24k+9999n} L^{108900000n^4+3000000n^3+2} T_i^{9900000n^3-s_0}, \qquad \Theta = a_H \pi^{(2\varrho+1)k}.$$

The power of π uses $k + 2(k-1)\varrho \le (2\varrho+1)k$. As $s_0 > 10^9 n^3$, summing the last power of T_i costs at most $2t_1^{9900000n^3-s_0}$. Use $\rho \le L$, $L \ge 2$, and $108900000n^4 + 3000000n^3 + 9999n + 3 \le 10^9 n^4$. This gives the bound, including convergence at $t_1 = 1$,

$$\frac{a_H}{(\lambda_0 d)^2} \kappa_j \int f \, dM \le \Theta L_1^{10^9 n^4} R^{24k} t_1^{9900000n^3-s_0} \epsilon_1. \tag{8.12}$$

Let $\ell_K(u) = u^K \min(1,u)^6$, and put

$$H_j(p) = \prod_{a<j} \ell_{n^6}(|p_0 \cdots p_a|^{-1}), \qquad \mathfrak{b}(\xi) = \Theta b^{2000n^9} \ell_{k-n^7}(|\xi|^{-1}).$$

Here $\xi \ne 0$. In the chosen real data $n \ge 10^8$, $k \ge n^7$, and $s_0 > 10^9 n^3$. Take $t_1 = \max(1, \rho^{24}|\xi|)$ and $\epsilon_1 = 1$. For $t_1 > 1$, $(\pi/|\xi|)^{k-1} \le (\pi\rho^{24})^k / t_1^{s_0}$, since $k - 1 \ge s_0$. If there is no active pair anywhere, the absolute majorant vanishes, so the desired bound holds. Otherwise an active pair passes both near-frame tests, and each direct partial product is a leading minor in the head calculation and has modulus at most $R_* = \rho^{9n}$. The common head radius is the maximum of ρ and the inverse moduli of the partial products. For $0 < a \le R_*$, $R_*^{n^6+6} \ell_{n^6}(1/a)$ dominates both 1 and a^{-n^6}. Applying this to each partial product gives

$$L^{10^9 n^4} \le L^{n^6} \le R_*^{(n^6+6)(2j+1)} H_j(p) H_j(q).$$

Now $s_0 - 9900000n^3 \ge k - n^7 + 6$ and $\rho^{24}/t_1 \le |\xi|^{-1}$, $t_1^{-1} \le \min(1, |\xi|^{-1})$. Therefore $\rho^{24k} t_1^{9900000n^3-s_0}$ is at most $\rho^{24n^7} \ell_{k-n^7}(|\xi|^{-1})$. With $\rho = b^{6n+2}$, the resulting power is bounded by b^{2000n^9}. This proves $\mathfrak{b}(\xi) H_j(p) H_j(q)$ for the inserted integral at floor j and for the unequal-head test at $j = i+1 \le \varrho$.

At the empty head write $Q_{A,B}(x) = \int_{U \setminus G_1} \langle A(y), B(yx) \rangle \, d\dot{y}$, and $E_\flat = F_\flat - W$ for the full vector W. Both upper averages give λ_0, and the position measure supplies a third factor. Hence

$$J_0 = \frac{a_H \lambda_0}{d^2} \int |Q_{F_\flat, F_\flat}(x)|^2 \, dx.$$

For a compactly supported scalar test h, the Schur estimate of W yields

$$\left| \int h(x) Q_{F_\flat, W}(x) \, dx \right| \le \|F_\flat\|_{L^2(d\dot{y})} (d/\delta)^{1/2} \|h\|_{L^2(dx)} \le d\delta^{-1/2} \|h\|_{L^2(dx)},$$

where $\delta = \deg_v$. Fubini on this support follows by Cauchy–Schwarz in the quotient and right invariance. Thus L^2 duality, together with $a_H \lambda_0 \le 1$, bounds the normalized integral of $|Q_{F_\flat,W}|^2$ by $1/\delta$. The inequality

$$|Q_{F_\flat,F_\flat}|^2 \le |Q_{F_\flat,W}|^2 + 2|Q_{F_\flat,F_\flat} Q_{F_\flat,E_\flat}|$$

leaves the mixed absolute majorant. The difference $E_\flat$ is bounded and has the same Bessel character. At floor zero every point is near, by translating to the identity in the box. In (8.12) take $L_1 = R$, $t_1 = 1$ and $\epsilon_1 = b^{-5r^2}$ from the squared cut norm. The mixed contribution is at most $2\Theta R^{10^9 n^4 + 24k} b^{-5r^2}$. The Weyl degree and the compact sphere factors give $\delta\Theta \le b^{4r^2}$: the factorials of indices k through $r-1$ cancel and π has nonpositive power. In the actual parameters $10^9 n^4 < nk$, $(6n+2)(n+24) < 7n^2$, and $7n^2 k + 11 \le r^2$. As $b \ge 2$, the mixed term is at most b^{-10}/δ. This proves $J_0 \le (1 + b^{-10})/\deg_v$ for the same forms whose highest member is the normalized top Gram.

Thus the actual top packet Gram has the positive top form, the diagonal descent and the numerical bottom and frequency bounds for the real part of the sieve.

8.5 Finite packets and the open orbital sum

8.5.1 Rational open orbits and their finite packets

The next three local arguments apply to the finite sets T and fixing subgroups Z used in the rational packet comparison below. We use a single coset vector for the terminal fibre and the two floor identities.

The finite orbital tensors and their terminal fibres. Fix the conjugation-stable set T of the finite packet calculation. Write $A_T = \prod_{w \in T} E_w$, $R_T = \prod_{w \in T} \mathcal{O}_w$, and z for the marked vector in the filling lattice basis. Use $h(v, w) = v^* J w$, so $h(z, z) = -N$. The additive measure gives R_T mass one, and R_T is self-dual for $\psi_T(x) = \psi_f(x/\varpi)$ restricted to T. Take $K_0 = K(g)_T$, $K = K(1)_T$, and $K_0 \subseteq Z \subseteq K$. For $Q = K_0 \backslash Z$ use counting measure. Its vector is

$$W_a(y) = \chi_g(y a^{-1}) \quad (a \in Q), \qquad \chi_g(uk) = \psi_T\big(g^{-1} \operatorname{sr}(u)\big) \quad (u \in U,\ k \in K_0),$$

with zero value off UK_0. Here sr is the sum of simple-root entries. It is additive on U and belongs to gR_T on $U \cap K_0$, so the value is well defined. Right K_0-invariance of χ_g makes the formula independent of the chosen representative of a. Normality of K_0 in K also gives $W(yk) = W(y)$ for $k \in K_0$. In particular W is locally constant, bounded by $\sqrt{|Q|}$, and $W(y) \ne 0$ implies $uy \in K$ for some $u \in U$. These are the same coordinates used in the finite shell bound.

Let G be the isometry group over A_T, let $G_i \subset G$ fix the first i hyperbolic pairs for $0 \le i \le \varrho$, and let $H \subset G_\varrho$ fix z. Put $t(l) = h(z, l^{-1} z)$, $e(l) = l^{-1} z$ on G_ϱ. The terminal floor measure θ_ϱ pushes forward by e to the norm-$-N$ quadric measure δ. Let dh be right Haar on H with $H \cap K$ of mass one. The terminal group measure and the full group measure are

$$\nu_T = dh * \theta_\varrho, \qquad \mu = \nu_T * M_\varrho, \qquad M_i = (\eta_{i-1} * \theta_{i-1}) * \cdots * (\eta_0 * \theta_0), \quad M_0 = \delta_1.$$

Here η_i is the Heisenberg measure and θ_i the next-vector floor measure. Thus the terminal stabilizer in ν_T is integrated with its actual Haar measure. Indeed l, l' over e differ by a left factor of H. For a nonnegative test F, the fibre integral $\int_H F(hl)\, dh$ is independent of l by right invariance. Right multiplication of l by $s \in G_\varrho$ acts by s^{-1} on e, preserving δ. This proves the right invariance of ν_T. For compact $C \subset G_\varrho$, any nonempty fibre in C has Haar mass at most $\operatorname{vol}_H(H \cap CC^{-1})$, and its image lies in the compact set $e(C)$. The quadric measure is locally finite there by the regular norm chart. Hence $\nu_T(C) < \infty$.

The same orbit argument at each lower floor yields the ordered Haar μ. In detail, the stabilizer of the next isotropic vector in G_i is generated by the Heisenberg group at i and G_{i+1}. The latter preserves horizontal volume by its isometry action and fixes the skew centre. Its right Haar followed by η_i is therefore

right invariant on this stabilizer. Integrating that Haar along θ_i gives right Haar on G_i by the invariant quadric pushforward. Compact sets project to compact sets of regular orbit points. The same fibre bound as above gives local finiteness, and an open set meeting the group has positive measure by the local quadric chart. The full group is unimodular. Indeed its adjoint action preserves the trace pairing $\operatorname{tr}(XY)$ on its skew-adjoint Lie algebra over the local fixed field. This pairing is nondegenerate, as extension to a splitting field gives the trace pairing on all matrices. Thus the absolute adjoint determinant is one, also in a split pair. The left translate of a right invariant volume has this same density, so μ is left Haar as well.

For unit heads p, let w_p use the first ϱ entries, $d_i = k - 1 + 2(\varrho - i)$ for $i < \varrho$, and put

$$\kappa_p = \prod_{i<\varrho} \operatorname{md}_T(p_i)^{d_i/2}, \qquad C_p(l)_{ab} = \kappa_p \int_G W_a(y)\overline{W_b(l w_p^{-1} y)}\, d\mu(y).$$

This is the tensor used in the rational packet calculation below. We record the support needed in these integrals. For fixed heads, the triangular reconstruction bounds y uniformly in $l \in G_\varrho$ whenever $W(y)$ and $W(l w_p^{-1} y)$ are both nonzero. Indeed take $u \in U$ with $uy \in K$. The second condition puts $w_p^{-1} y$ in $Q_\varrho K$, where Q_ϱ is the group unitriangular on the earlier frame. The head-minor calculation bounds all columns $u e_i^*$. Applying the triangular inverse formula to e_i^* bounds u^{-1}, hence y. One gauge works in a neighbourhood of a finite list of heads and their conjugates, since there are only finitely many inverse partial products to put in its first size ball. Its bounded matrix set in the closed isometry group is compact. Now write $y' = l w_p^{-1} y$. The Weyl blocks give $w_{\tilde p} = w_p^{-1}$, and G_ϱ commutes with them. Thus $y = l^{-1} w_{\tilde p}^{-1} y'$. The same bound with the conjugate head bounds y'. Finally $l = y' y^{-1} w_p$ lies in a compact set as well. At a split pair the involution interchanges the two components when forming $\tilde p$. The conjugation-stable gauge controls both.

In particular the four-vector product used to compute $\langle C_p(l), C_q(l)\rangle$ is bounded and has common compact support in l, y_1, y_2 for nearby p, q. There is also uniform local constancy. For p' near p the head modulus is unchanged. For $s \in G_\varrho$ sufficiently near 1, $y^{-1} w_p s w_{p'}^{-1} y \in K_0$ throughout the compact set just obtained. This follows from continuity and compactness, since K_0 is open. The right K_0-invariance of W now gives $C_{p'}(ls) = C_p(l)$ for every $l \in G_\varrho$. Outside that compact y-set both integrands vanish whenever their first vector is nonzero. For $a \in H \cap K_0$ we also have $W(ay) = W(y)$, since a normalizes U and fixes the simple-root character. Consequently

$$C_p(al) = C_p(l) = C_p(la).$$

For the second equality change the integration variable by left multiplication by a^{-1} and commute a with w_p.

For a finite family p with a common last entry t, let R be any positive semidefinite matrix on its index set. For $d = (a, b) \in Q^2$, write $v_d(l)_p = C_p(l)_{ab}$. With the inner product linear in its second argument,

$$\sum_{p,q} R_{pq}\langle C_p(l), C_q(l)\rangle = \sum_{d\in Q^2} v_d(l)^* R\, v_d(l) \;\geq 0.$$

Integrate this expression on $t(l) - t \in L$ for a compact open additive subgroup $L \subset A_T$. For arbitrary heads set

$$\mathcal{K}_L(p, q) = \int_{t(l)-p_\varrho \in L} \langle C_p(l), C_q(l)\rangle\, d\nu_T(l).$$

Thus the integral of the positive form is $\Re \sum_{p,q} R_{pq} \mathcal{K}_L(p, q)$. At a common last entry this is the Gram matrix of the functions $1_{t+L}(t(l)) C_p(l)$ in $L^2(d\nu_T; \mathbb{C}^{Q^2})$. The integral is jointly locally constant in the two unit heads, since the tensors agree uniformly for nearby heads and changing p_ϱ inside the same coset of L preserves the aperture.

On a terminal fibre put

$$E_R(e) = \inf_{l^{-1}z=e} \int_H \sum_d v_d(hl)^* R\, v_d(hl)\, dh.$$

The infimum ranges over G_ϱ, with value $+\infty$ for an empty fibre. On the orbit its value equals every displayed integral by right invariance of dh. The function on that orbit is measurable by the existence of continuous local sections. At a fixed point e_0, for v on its quadric near e_0, put $a = e_0 + v$, $D = h(a, e_0)$, and

$$r_v x = x - a\, h(a, x)/D, \qquad s(v) = r_v r_{e_0}^{-1} \quad (x \in V_\varrho).$$

The element D is a unit of A_T near $2h(e_0, e_0) = -2N$. Since $D + \bar{D} = h(a, a)$, direct expansion shows that r_v preserves h. It sends e_0 to $-v$, so $s(v)e_0 = v$ and $s(e_0) = 1$. Extend by the identity on the hyperbolic pairs. The same construction works about any point of the norm quadric, so each orbit is open in the quadric. Integration along a section and then against δ now gives

$$\int_{h(z,e)-t\in L} E_R(e)\, d\delta(e) = \mathfrak{R} \sum_{p,q} R_{pq} \mathcal{K}_L(p, q).$$

The pushforward description of δ gives full measure to the orbit in this identity.

Here is the consequence on individual labels. Suppose a finite list (l_j, d_j) lies over e_0, and $l_j l_{j'}^{-1} \in K_0$, $d_j = d_{j'}$ imply $j = j'$. For v near e_0 on the quadric, $v = a_0 y$ with $a_0 \in H \cap K_0$, take $s = s(v)^{-1}$. The elements $l_j s a_0$ lie over y and have exactly the original $C_p(l_j)$ for every head on the list. This uses the uniform constancy for right multiplication by s, followed by the right $H \cap K_0$-invariance. For any l over y put $h_j = l_j s a_0 l^{-1} \in H$. In the integral over H, retain coordinate d_j on $(H \cap K_0)h_j$. It is constant there by left invariance of the tensor under $H \cap K_0$. For a given coordinate these cosets are disjoint, since $h_j h_{j'}^{-1} = l_j l_{j'}^{-1}$. Positivity on every coordinate gives throughout this neighbourhood

$$E_R(y) \;\geq\; \mathrm{vol}_H(H \cap K_0) \sum_j v_{d_j}(l_j)^* R\, v_{d_j}(l_j).$$

For y off the orbit the same inequality holds with infinite left side. This provides the neighbourhood and the individual coset mass for the one-vector count. No averaging factor occurs in the coordinates of W.

The Fourier multiplier of the finite terminal aperture. Set $\mu_0 = \mu(K_0)$, $v_U = \mathrm{vol}_U(U \cap K_0)$, and $m_Z = \mu(Z)\mu_0 / v_U^2$. With the filling modulus, $v_U = \prod_{i<\varrho} \mathrm{md}_T(g)^{d_i+1/2}$ and $\mathrm{md}_T(g) = g^{-2f}$. Thus $v_U^{-1} = g^{f\varrho(2k+2\varrho+1)} = U_0$. Write $\mathcal{I}_i(p, q; \omega)$ for the finite floor form, including its head density $\kappa_i(p, q) = \prod_{a<i} \mathrm{md}_T(p_a q_a)^{d_a/2}$. Write $Z_i(x) = h(f_i, x f_i)$, with $f_i = e_i$ for $i < \varrho$ and $f_\varrho = z$. We claim

$$\int_{G_\varrho} \psi_T(\xi t(l))\, \langle C_p(l), C_q(l) \rangle\, d\nu_T(l) = \mathcal{I}_\varrho(p, q; \psi_T(\xi Z_\varrho)). \tag{8.13}$$

To prove the identity expand the two tensors. With $w_p = w_{p,\varrho}$, $w_q = w_{q,\varrho}$, the four factors give

$$\kappa_p \kappa_q\, \langle W(y_1), W(y_2) \rangle \overline{\langle W(l w_p^{-1} y_1), W(l w_q^{-1} y_2) \rangle}.$$

The integrations in y_1, y_2 use $d\mu\, d\mu$. The compact support proved above allows all three integrations in l, y_1, y_2 in either order, even with absolute values. Put $y_2 = y_1 x$, using left Haar in x. In the y_1-integral put $y_1 = h_1 \sigma y$ by $d\mu = dh_1 * d\theta_\varrho * dM_\varrho$. For fixed h_1, σ replace l by $h_2 \tau (h_1 \sigma)^{-1}$, using right invariance of ν_T and then $d\nu_T = dh_2 * d\theta_\varrho$. Both h_1, h_2 fix z, giving

$$t(l) = h(z, \sigma\tau^{-1} z), \qquad l w_p^{-1} y_1 = h_2 \tau w_p^{-1} y, \qquad l w_q^{-1} y_2 = h_2 \tau w_q^{-1} y x.$$

With $D = w_q^{-1} w_p$, the last expression is also $D h_2 \tau w_p^{-1} y x$. Integrating the first raw pairing over h_1 gives the upper pairing of $(u, v) \mapsto \langle W(u), W(v) \rangle$ at $(\sigma y, \sigma y x)$. Integrating the second over h_2 gives the conjugate of the upper pairing of $(u, v) \mapsto \langle W(u), W(Dv) \rangle$ at $(\tau w_p^{-1} y, \tau w_p^{-1} y x)$. These upper integrals are over dh at this terminal stage. The two remaining θ_ϱ-integrals, with the multiplier at $\sigma\tau^{-1}$, are precisely the defining

two translates of $\mathcal{I}_\varrho$. Its density here is $\kappa_\varrho(p,q)=\kappa_p\kappa_q$. This proves (8.13) with the specified conjugation and finite stabilizer measures.

Use $p\oplus_i c$ for replacing the entry at i by c. In a unit head use 1 when c has a zero component. Such c's form a null set for dc, being a finite union of coordinate hyperplanes in A_T. For the terminal form put $J_{\varrho+1}=\mathcal{K}_L/(|L|m_Z)$, where $|L|=\mathrm{vol}(L)$. Fubini is applicable when integrating it over the common last entry c. Indeed integration over c of the absolute majorant $1_L(t(l)-c)|\langle C_p(l),C_q(l)\rangle|$ costs $|L|$ times an integrable function of l. The substitution $u=t(l)-c$ yields

$$\int_{A_T}\psi_T(\xi c)\,J_{\varrho+1}(p\oplus_\varrho c,q\oplus_\varrho c)\,dc \\ =\frac{\mathcal{I}_\varrho(p,q;\psi_T(\xi Z_\varrho))}{m_Z}\,\frac{1}{|L|}\int_L\psi_T(-\xi u)\,du. \tag{8.14}$$

The last factor is exactly $1_{L^\perp}(\xi)$, where $L^\perp=\{\xi:\psi_T(\xi L)=1\}$. If the character is nontrivial on L, translation by a point where it differs from one forces the integral to vanish. In particular the factor equals one at zero frequency and has modulus at most one at every frequency.

Descent of the finite common entry. For $i\le\varrho$ put $J_i(p,q)=\mathcal{I}_i(p,q;1)/m_Z$. For $i<\varrho$, write $P=w_{p,i}$, $P'=w_{q,i}$ for the truncated Weyl elements and $D=(P')^{-1}P$. Each commutes with G_i. Denote the two upper pairings just used, now at floor $i+1$, by B^1,B^D. The quadric pushforward and the stabilizer invariance of upper averages apply to this finite level. Both pairings are invariant under simultaneous left translation by $S_i=\mathrm{Stab}_{G_i}(f_i)$. For the finite floor, θ_i is the law of $w_i(c)^{-1}u$ for $\mathrm{md}_T(c)^{d_i}\,dc\,d\eta_i(u)$. Here u fixes f_i, and $Z_i(w_i(c))=c$.

In $\mathcal{I}_i(p,q;\psi_T(\xi Z_i))$ the upper pairings occur as

$$\psi_T(\xi Z_i(\sigma\tau^{-1}))\,B^1(\sigma y,\sigma yx)\,\overline{B^D(\tau P^{-1}y,\tau P^{-1}yx)}.$$

Keep σ fixed. The expression in τ descends to $\tau^{-1}f_i$, since B^D is invariant under simultaneous left translation of both arguments by S_i, and Z_i is invariant under right multiplication by S_i. Thus τ may be replaced in the integral by $\tau_0\sigma$ using invariance of the quadric measure. Put $\tau_0=w_i(c)^{-1}u$ and $y'=u\sigma y$. Then $Z_i(\sigma\tau^{-1})=Z_i(u^{-1}w_i(c))=c$. The first pairing becomes $B^1(y',y'x)$ and the second uses $(Pw_i(c))^{-1}y',(Pw_i(c))^{-1}y'x$. Appending c to both heads leaves D unchanged because $w_i(c)$ commutes with the preceding blocks. The resulting measure in y' is $\eta_i*\theta_i*M_i=M_{i+1}$. The new head density is $\kappa_{i+1}(p\oplus_i c,q\oplus_i c)=\kappa_i(p,q)\,\mathrm{md}_T(c)^{d_i}$. With weight one, each upper pairing is the successive average of the next upper pairing, by the ordered convolution of its floor measures. Their product is therefore the two-translate integral defining $\mathcal{I}_{i+1}$ at these heads.

For the absolute values the initial integral is bounded by the finite denominator-shell estimate with gain 1 and exponent 0. At fixed inputs each successive upper average converges absolutely by the measured floor-mass bound on its support. The same change of variables for the nonnegative integrands gives integrability in c, then Fubini applies to the complex integral. We have proved

$$\int_{A_T}\psi_T(\xi c)\,J_{i+1}(p\oplus_i c,q\oplus_i c)\,dc=\frac{\mathcal{I}_i(p,q;\psi_T(\xi Z_i))}{m_Z}\qquad(i<\varrho). \tag{8.15}$$

At the bottom the Weyl elements are 1, the preceding stack is δ_1, and the two upper pairings coincide. The unfolded phase pairing gives

$$J_0=\frac{1}{v_U^2m_Z}\int_G\left|\int_Z\chi_g(yxy^{-1})\,d\mu(y)\right|^2d\mu(x)=\frac{\mathcal{B}(Z)}{\mu(Z)\mu_0}.$$

Thus the bound from each prescribed standard or dual vector stabilizer acts on this very bottom form. In (8.15) and (8.14) the right side at $\xi=0$ is J_i. At a unit frequency ξ the inserted-character bound applies, with the two head weights at i. For $i<\varrho$ the unequal-pivot bound applies directly to $J_{i+1}(p,q)$ when $g^{-1}(p_i^{-1}-q_i^{-1})\in A_T^\times$, with both head weights at $i+1$. Together with common-last-entry positivity and continuity, these are all the finite inputs of the adelic sieve for this Z,L.

The residue conditions and their finite levels. Use $n = \varrho + 1$ and $r = k + 2\varrho + 1$ on the given filling datum. Let h be the number of column classes and d_B the number of product Bessel coordinates, so $C = hd_B$. For proper A, a residue in Ω_A fixes the prescribed vectors and has all leading frame minors of orders $1, \dots, n$ nonzero at one chosen place above g_0. The class representatives fix the frame, and taking a minor of the integral term position $h_c^{-1}\gamma_f h_{c'}$ commutes with reduction in the filling lattice basis. Hence every term over this residue has all these minors nonzero over E. The last minor includes the marked anisotropic column. The full fixing set and the two complementary residue sets use the empty choice of retained links in the crude comparison.

For a finite calculation take a conjugation-stable set T containing the places of g and N and the supports of the denominators of the filling basis, the finitely many term positions, and their inverses. Enlarge T so that the selected pivot values, terminal entries and their inverses are integral off T. Write K for the integral group at T and $K_0 = K(g)_T$ for its principal level. Membership in the global principal level means membership in K_0 at T and integrality with integral inverse off T. At a split pair $\{w, \bar{w}\}$ the unitary isometry is determined by its first matrix. In dual coordinates its other component is the inverse adjoint. The reduction of the first matrix is onto $\mathrm{GL}_r(\mathbf{F}_q)$, since entrywise lifts of an invertible residue matrix have unit determinant. At a prescribed pair $q = g_0$, let Z_w be the inverse image of the specified vector stabilizer. The frame is injective on residue vectors and both actions are transitive on nonzero vectors. Thus $[K_w : Z_w] = q^r - 1$. Take $Z = \prod_{w\in A} Z_w \times K_{\mathrm{rem}}$ over the distinct pairs and the remaining conjugation-stable factor. We have $K_0 \subseteq Z \subseteq K$, and the components at T of all lifts of residues in Ω_A belong to Z.

Write H_T for the terminal stabilizer H in the local packet measure. Use the normalized finite bottom

$$J_0^T = \frac{\mathcal{B}(Z)}{\mu(Z)\mu_0}, \qquad \mu_0 = \mu(K_0), \qquad \mathcal{B}(Z) = \int \left| \int_Z \chi_g(yxy^{-1})\, d\mu(y) \right|^2 d\mu(x).$$

For $x \in K_0$ the phase in the inner integral is one by normality. Its contribution to $\mathcal{B}(Z)$ is $\mu(Z)^2\mu_0$. On a prescribed split pair, the unit-pivot charts have integral horizontal and skew coordinates of volume one. At the last floor use the marked pair above and the unit coordinate $a = h(u, e)$ on the norm-$-N$ fibre. Scale v by a and u by $\bar{a}^{-1}$, then apply the Heisenberg translation fixing u. Division by the unit a keeps the horizontal coordinates integral, and the norm equation leaves an integral skew coordinate. The same integral chart therefore has mass at least $(1-q^{-1})^2$ on each of the n floors. The integral terminal stabilizer has mass one. It follows that $\mu(K_w) \geq (1-q^{-1})^{2n} \geq 2^{-n}$ for the chosen q. The nonunits of the split ring have additive measure $\delta_w \leq 2/q \leq b^{-9}$. Apply the finite bottom estimate with $t_0 = k - 10^{10}n^4$ to Z_w and the principal lower bound to K_w. This gives

$$\frac{\mathcal{B}(Z_w)/\mu(Z_w)}{\mathcal{B}(K_w)/\mu(K_w)} \leq \frac{1}{q^r - 1} + \frac{2^{-n}}{\mu(K_w)}\delta_w^{t_0+1} \leq 2b^{-10r} + b^{-8k} \leq b^{-6k}.$$

Here $t_0 + 1 \geq 8k/9$ on the sufficiently large selected squares. The split pairs and complementary factor give a direct product of groups, and the phase is a product. Rescaling the product Haar by a positive scalar rescales $\mathcal{B}$ by its cube and each subgroup mass by that scalar. Thus the quotient of the two ratios is unchanged, and Fubini multiplies the individual bounds. On the whole semilocal group the finite bottom estimate at exterior exponent zero uses $\delta \leq 1$ and $\mu(K) \leq 1 + 2^{-n}$. For this upper bound unload the preceding integral unipotents, of volume one. At each orbit choose box representatives at that floor and the next. The indicated column of their quotient modulo the earlier frames is integral, so the translated orbit vector remains in the unit box. The finite cone calculation with e split planes bounds this fibre by $\zeta_R(e-1)$ uniformly in its norm level. Indeed the content divisor bound uses trace one in the skew coordinate after the orthonormal coordinates are fixed. There are at least $m_0 + 1$ planes at each floor, including the auxiliary terminal planes, and the orthonormal boxes have mass one. The flag disintegration therefore gives $\mu(K) \leq \zeta_R(m_0)^n$ with $\mu_H(H \cap K) = 1$. This is at most $1 + 2^{-n}$. Since $2^{1-n} \leq b^{-6}$, this yields

$$0 \leq J_0^T \leq e_T = (1 + b^{-6})b^{-6k|A|}.$$

The estimate holds throughout the finite exhaustion.

Unfolding on the labelled cosets. For the rational open-cell parametrization write $F = E^{+}$, $\gamma_p = w_p\sigma$, $t = h(z, \sigma z)$ and $y = (\sigma z)_{X_0}$. The nonzero leading minors provide the successive pivots. At each floor the middle and skew coordinates of the image of the pivot vector, divided by its nonzero opposite coefficient, specify the first Eichler translation. The opposite column and the unit pairing specify the second. The link in the Weyl position is the graph of σ^{-1} on $X^{\perp}/X$. For a fixed rational open term and its terminal representative σ, this elimination determines the two unipotents and all the pivots uniquely. The block fixes σ and can have many heads. The block labels (c, s, c', Λ) inject into (c, c', s, σ). The identities $\sigma z = y - (t/N)z$, $Ny^{*}y = t\bar{t} - N^2$ and $\mathrm{disc}(\gamma_p) = |\prod_{a<\varrho} p_a|^2(t\bar{t} - N^2)$ give $|t_v| > N$ at every real place on the open links. We keep the sum over heads inside each block and group the block labels by t.

For a block choose a lift $\tilde{s} \in K_r$, put $D_\gamma = h_c^{-1}\gamma_f h_{c'}$, and let $l = (D_\sigma|_T)^{-1}$. The class representatives commute with the Weyl operators, giving $D_{\gamma_p}|_T = w_p l^{-1}$. In the filling basis $l^{-1}z$ is the lattice-coordinate vector of $h_c^{-1}y - (t/N)z$. Use the finite correlation $C_p(l)$ of the actual vector W_Z, with factor $\varkappa(p) = \prod_{a<\varrho} \mathrm{md}_T(p_a)^{e_a/2}$, $e_a = k - 1 + 2(\varrho - a)$. For each $b \in Q = K_0\backslash Z$ choose $a = K_0\tilde{s}_T b$, the unique coset with $\tilde{s}_T^{-1}ab^{-1} \in K_0$. Right multiplication of the integration variable by a in the (a, b) entry gives $\chi_g(x_T)\overline{\chi_g((D_{\gamma_p}^{-1}x\tilde{s})_T)}$. Here x can be taken to equal one off T. Write $\zeta = D_{\gamma_p}^{-1}x\tilde{s}$, which is integral off T with integral inverse.

For the two nonzero phases the rational unipotents solving

$$(h_c^{-1}u_f h_c)x \in K(g), \qquad (h_{c'}^{-1}u'_f h_{c'})\zeta \in K(g)$$

form one left coset of U_c and $U_{c'}$, respectively. Indeed the local support gives finite adelic unipotent solutions, taking the identity off T. The rational points of U are dense in its finite adeles by the affine unipotent coordinates and approximation with denominators supported at finitely many places. The conditions are open and the ratio of two solutions is in the respective class unipotent. A vanishing phase has no solutions. The two finite arguments lie in the inverse unipotent cosets, so their respective χ_g values are $\chi_f(u)^{-1}$ and $\chi_f(u')^{-1}$. This uses triviality on the principal level, invariance of the simple roots under the middle stabilizer and triviality on the indicated integral positions off T. The product formula gives

$$\chi_g(x_T)\overline{\chi_g(\zeta_T)} = \chi_\infty(u)\overline{\chi_\infty(u')}.$$

In the cut real correlation $B_\infty(\gamma_p)$ insert the partitions $\sum_u \beta_c(uY) = 1$ and $\sum_{u'} \beta_{c'}(u'\gamma_p^{-1}Y) = 1$ on these two cosets. Multiplying the real integrand by the displayed phase replaces its two Bessel factors by the values at uY and $u'\gamma_p^{-1}Y$. With $Y_1 = uY$ this is the amplitude of $\delta = u\gamma_p u'^{-1}$, since $\delta^{-1}Y_1 = u'\gamma_p^{-1}Y$. For fixed u, u' the two finite conditions say that $D_\delta \in K_r$ has residue s, $h_c^{-1}u_f h_c$ is integral off T, and $x_T \in (h_c^{-1}u_f h_c)_T^{-1}K_0$. In fact if k_1, k_2 are the products in the two conditions then $D_\delta = k_1\tilde{s}k_2^{-1}$. Conversely the condition on x_T defines $k_1 \in K(g)$, and $k_2 = D_\delta^{-1}k_1\tilde{s} \in K(g)$ by its reduction. The finite coset contributes μ_0. Open-cell uniqueness and the absolute integral majorant for amplitudes justify the interchange of the sum and the two integrals, since the absolute integrals sum to at most μ_0 times that majorant. Thus

$$C_p(l)_{ab}B_\infty(\gamma_p) = \varkappa(p)\mu_0\mathcal{A}_{p,T}(c, s, c', \sigma),$$

where $\mathcal{A}_{p,T}$ is the indicated restricted sum of rational cut amplitudes. This holds for every $b \in Q$ with its specified a.

The labelled multiplicity at the terminal quadric. For fixed (c, y, t) the terminal vector e is common to the corresponding labels. The pairs $(K_0 l, (a, b))$ are distinct as the label and $b \in Q$ vary. Indeed equality of the cosets for labels i, j gives $h_{c_i'}^{-1}\sigma_i^{-1}\sigma_j h_{c_j'}$ in K_0 at T. The two term positions are integral off T, so this element is in the global principal level. Equality of the terminal vector and of the fixed frames places $\sigma_i^{-1}\sigma_j$ in $H_0(F)$. The column level is the pullback of the principal row level to H_0. The class double cosets

agree, and freeness at that level gives $\sigma_i = \sigma_j$. Finally equality of (a, b) forces equality of the two residues, both reductions of ab^{-1}. This gives the injection into the labelled cosets with all $|Q|$ choices per block.

Use $\ell_{va} = |\kappa_v|^{2(\varrho-a)}/N$, d_B the number of product Bessel coordinates, and $G_t = \prod_v (|t_v|^2 - N^2)^{k-1}$. The real terminal Gram formula on the scaled head and entry t_v/N gives the same Gram matrix R for all labels with that t. Its vectors are $\eta_p B_\infty(\gamma_p)$, where

$$\lambda_0 = \frac{c_G^{f/2}}{d_B} \prod_v \prod_{a=0}^{\varrho} \ell_{va}^{e_a}, \qquad \eta_p = \lambda_0 \sqrt{G_t} \prod_v \prod_{a<\varrho} |\Phi_v(p_a)|^{e_a}.$$

The terminal scale is $\ell_{v\varrho} = 1/N$ and contributes $N^{-(k-1)}$ at each v. The product formula for each T-unit gives $\varkappa(p)\eta_p = \lambda_0 \sqrt{G_t}$. Thus at the labelled coordinate $v_p = C_p(l)_{ab}$ the form $v^* R v$ equals $\mu_0^2 G_t \lambda_0^2 \| \sum_p \mathcal{A}_{p,T} \|^2$.

The one-vector count on the marked terminal quadric. Write $V = \mathrm{vol}(\widehat{\mathcal{O}}_E)$ for the global finite additive measure. The constant of the marked quadric in these coordinates is

$$c_{\mathrm{quad}} = |D_E|^{1/4} V^k \left(\frac{\pi^k}{(k-1)!\, N^k} \right)^f.$$

For a finite list of distinct (c, y) with $N y^* y = t\bar{t} - N^2$ and $h_c^{-1} y$ integral off T, let $a_{c,y} \geq 0$ be lower bounds for a nonnegative function F_0 on $H_T \cap K_0$-orbits of a neighbourhood of the corresponding terminal points. For all sufficiently small L,

$$|L| \sum_{(c,y)} a_{c,y} \;\leq\; (1 + b^{-6}) h c_{\mathrm{quad}} G_t \int_{h(z,e) - t_T \in L} F_0(e)\, d\delta_{-N}(e). \tag{*}$$

To see the constant, choose a compact locally constant test attaining the specified lower bounds and minorizing F_0 on the norm quadric. Take the maximum of the finitely many $a_{c,y}$ on sufficiently small compact open neighbourhoods saturated under $H_T \cap K_0$. Such saturations are again compact and open. Use Weil's formula on a regular one-vector fibre. For the unitary group H_0 of a k-dimensional Hermitian space with $k > 2$, suppose at one place the full-rank null orbit is nonempty and the group is transitive on each full-rank fibre, and suppose the integral of the absolute Schwartz theta sum converges. Write μ_α for disintegration of the covolume-one additive Haar of $X_0(\mathbf{A}_F)$ against the covolume-one additive Haar of the norm coordinate on $\mathbf{A}_F$. For $\alpha \in F^\times$ and an adelic Schwartz test f_0 the formula with probability quotient Haar is

$$\int_{H_0(F)\backslash H_0(\mathbf{A}_F)} \sum_{\substack{y \in X_0(F) \\ q_X(y) = \alpha}} f_0(h^{-1} y)\, dh = \int_{q_X(y)=\alpha} f_0(y)\, d\mu_\alpha(y).$$

This is the type (I$_2$) specialization with $\epsilon = \frac{1}{2}$ and one input vector, so the inequality on the dimension in that theorem is $k > 2$. [27, §27, pp. 40–41, §44, Theorem 2, pp. 62–63, and §52, Theorem 5 and (41), pp. 76–77]

Here $q_X(y) = y^* y$ on the definite space X_0. At a split place over g_0 the two halves of y are a vector and a covector. For each fixed contraction, nonzero vectors and nonzero covectors form one GL_k orbit. For zero contraction such pairs exist for $k \geq 2$. Complete the vector to a basis of the kernel of the covector and then add a vector on which the covector takes value one. This proves the relevant transitivity for zero contraction. For nonzero contraction use the vector and a basis of the kernel. The quotient is compact by definiteness and the Schwartz theta sum is bounded there by the finite class decomposition and the lattice theta sums. The compact-level test is the saturated finite test at T times the integral boxes outside T. At infinity take a Schwartz test equal to one on the compact spheres of norm $\alpha_v = (|t_v|^2 - N^2)/N > 0$. With probability Haar on the quotient each of the h classes has mass $1/h$ because the rational column stabilizers are trivial.

In Weil's (41), use the fibre gauge given by disintegration of the ambient additive gauge against the fixed-field additive gauge. With real Lebesgue in the chosen coordinates, the finite integer boxes have covolume-one additive measures $V = 2^f|D_E|^{-1/2}$ over E and $|D_F|^{-1/2}$ over F. These give the normalized rational gauges, since multiplying a gauge by a rational scalar has adelic product one. The quotient of the normalization factors is $V^k\sqrt{|D_F|} = V^k|D_E|^{1/4}$, using finite unramifiedness of E/F. At each real place polar integration against $d\alpha$ gives $\pi^k\alpha_v^{k-1}/(k-1)!$. The class sum thus has scalar $hV^k|D_E|^{1/4}$ times the real fibre masses and the finite fibre integral whose integer boxes have mass one. Off T, Fourier inversion can first be performed on a congruence slab around α. At a place of F of residue size Q, square the absolute integral for a character of conductor Q^a and subtract the two variables. The perfect unramified or split trace pairing, including above 2, forces the difference into the a-th prime-power lattice. The squared integral is at most Q^{-2ak} with the integer box of mass one. Since there are at most Q^a frequency classes of this level, absolute summation bounds the density by $(1-Q^{1-k})^{-1}$, uniformly in α. Their product is at most $\zeta_F(k-1) \le \zeta_E(k/2-1)$, and $\zeta_E(s) \le \zeta(s)^{2f}$ makes its relative error $o(b^{-6})$ on these data. Let B_{lat} have the class lattice columns in the rational middle space. The chart $(y,\tau) \mapsto B_{\mathrm{lat}}^{-1}(y-\tau z/N)$ has determinant $\bar\pi/N = 1/\pi$ up to sign, of absolute modulus N^f. In volume-one finite coordinates the ambient chart measure is therefore $N^f\,dy\,d\tau$. Dividing by the unchanged fixed-field norm measure for $q_X(y)-\tau\bar\tau/N$ gives $d\tau\,dy|_{q_X=\tau\bar\tau/N-N} = N^{-f}\,d\delta_{-N}$. The skew coordinate is changed by an integral skew unit and keeps the integer-lattice measure. Together with the $N^{-(k-1)f}$ from α this gives precisely c_{quad}. For a compact locally constant minorant the slab integrals at the regular norm fibre tend to its quadric integral. Indeed a finite partition of the compact support into local norm-coordinate charts shows that the minorant's fibre integral is locally constant as τ varies near t_T. The outside-density bound is uniform in this perturbation. Its slack within $1+b^{-6}$ gives $(*)$ on all sufficiently small slabs for the finite list. [27, §53, p. 81, for the quotient measures]

The normalized packet comparison on the common fibre. The neighbouring-fibre lower bound for these pairs supplies the factor $\mu_H(H_T \cap K_0)$ on each occupied coset. Apply the one-vector quadric count with these lower bounds, grouping by distinct (c,y) for fixed t. In rational coordinates $D_\sigma z = h_c^{-1}y-(t/N)z$. Since D_σ, the basis change and its inverse, and t/N are integral off T, the required $h_c^{-1}y$ is integral there. The norm and strict positivity conditions were verified in the rational open cell. For the test F_0 in that count use the nonnegative fibre form E_R from the local argument. Its integral over the aperture is the integral in the terminal comparison. For sufficiently small compact open additive L, the integral of the positive fibre form over the aperture is at most $\mathfrak{R}\sum_{p,q} R_{pq}\mathcal{K}_L(p,q)$ by the terminal disintegration. A compact open additive subgroup inside the intersection of the finitely many neighbourhoods works for all selected t.

Normalize $\mu_H(H_T \cap K) = 1$ as in the local terminal measure. Identity extension outside T identifies $\bar H$ with $(H_T \cap K)/(H_T \cap K_0)$. Thus $\mu_H(H_T \cap K_0) = |\bar H|^{-1}$. At the level primes the integral Bessel subgroup has integral unipotent and middle factors in the flag basis. Their reductions intersect trivially by their action on the diagonal frame blocks. The local factors extend integrally by the identity off those primes, so $|J| = |\bar H|U_0$, with the unipotent index $U_0 = v_U^{-1} = g^{f\varrho(2k+2\varrho+1)}$ in the local normalization. Also $\mu(Z) = \mu_0|Q|$. These identities give

$$m_Z = \mu(Z)\mu_0 U_0^2 = |J|U_0\,\mu_H(H_T \cap K_0)\mu_0^2|Q|.$$

In the one-vector inequality divide by $|L|$, cancel $G_t > 0$ and apply this identity. The result for the chosen labels is

$$\lambda_0^2 \sum_{\text{labels}} \Big\| \sum_p \mathcal{A}_{p,T} \Big\|^2 \le (1+b^{-6})\Pi\,\mathfrak{R}\sum_{p,q} R_{pq}\frac{\mathcal{K}_L(p,q)}{|L|m_Z}, \qquad \Pi = |J|U_0 h c_{\mathrm{quad}}.$$

The terminal kernel here uses the same unnormalized coset vector and Haar as the finite floor identities.

The rational heads and the frequency budget. We give the summation over rational pivots, since their weights vary with the actual scale and with the places of T. Put $c_i = V/\prod_v \ell_{vi}^2$. Let $O_T = \prod_{w\in T} \mathcal{O}_w$, $[y]_T = \prod_{w\in T} |y|_w^{-1}$, and $D_T(y) = [O_T : O_T \cap y^{-1}O_T]$. Put $L_h(s) = s^h \min(1,s)^6$. For frequencies $y \in \mathcal{O}_{E,T}$, integral away from T, we shall use

$$\sum_{0\neq y\in\mathcal{O}_{E,T}} \frac{[y]_T^{h/2}}{D_T(y)^6} \prod_v L_h(|c_v y_v|^{-1}) \;\le\; C_1^f B^{6f} \prod_v X_v^h \qquad (|c_v|^{-1} \le X_v \le B,\ B \ge 1).$$

Here C_1 is absolute. The product formula cancels the first h powers, since $[y]_T^{1/2} \le \prod_v |y_v|$. The remaining real factors are at most $B^{6f}\prod_v \max(1,|y_v|)^{-6}$. Write $\mathfrak{d}(y)$ for the denominator ideal. Its norm is at most $D_T(y)$. For a fixed denominator $\mathfrak{d}$, its elements lie in at most $N\mathfrak{d}$ cosets of $\mathcal{O}_E$. In each coset a product of complex squares of side $1/2$ contains at most one element, since every conjugate of a nonzero integral difference would otherwise have modulus less than one. Summing the real weights over the grid costs at most an absolute constant per place. Summing the remaining denominator weights costs at most $\zeta_E(2) \le \zeta(2)^{2f}$. This also proves the bound for all finite subsums uniformly in T.

Remove the two head weights from the real frequency bound. At a nonzero frequency the remaining factor is

$$\Theta_\infty b^{2000n^9} L_{k-n^7}(|\xi_v|^{-1}), \qquad \Theta_\infty = \pi^{(2\varrho+1)k}\frac{\operatorname{vol} U(k)}{\operatorname{vol} U(r)}.$$

At the finite places the two local Fourier bounds use the next diagonal Gram entry and the difference of two last pivots, respectively. With their denominator summation, the finite frequency factor is bounded by

$$b^{n^9 f}[y]_T^{(k-n^7)/2}/D_T(y)^6.$$

Here the frequency exponent is $k-n^7$, the head exponent is n^6, and $(2/\operatorname{md}_T(g))^{n^6} = (2g^{2f})^{n^6} \le b^{n^9 f}$ since $g \le b^{23}$. It gives the same bound in the off-diagonal step at $y = (p^{-1}-q^{-1})/g$ for distinct last rational pivots p, q. In this case the real frequency is gy_v/ℓ_{vi}. In both cases the inverse real frequency scale is at most

$$X_v = \frac{|\kappa_v|^{2\varrho+1}}{gN}, \qquad \max(1,X_v) \le b^{O(n)}.$$

Indeed $|\kappa_v| = g|\Phi_v(\varpi)|/(2\pi) \ge 1$. The packing bound therefore bounds the sum of the frequency factors, or one such off-diagonal factor, by $\epsilon_* = \Theta_\infty^f \prod_v X_v^{k-n^7} b^{O(fn^9)}$. All implied constants here are absolute.

For a head of length i, its successive cumulative products $d_j = p_0\cdots p_j$, $j < i$, determine the head. The finite head weight is zero unless every d_j is integral at T, and equals $\prod_j [d_j]_T^{n^6/2}$ otherwise. Away from T these d_j are units. For each j apply the packing bound with $h = n^6$, $c_v = \prod_{a\le j}\ell_{va}$ and $X_v = B = N^i$. These choices are valid since $\ell_{va} \ge N^{-1}$. The real head weights are exactly the corresponding products of L_{n^6}. As $N \le b^{60n}$, summing the heads independently via the d_j's costs $b^{O(fn^9)}$, also after squaring. Moreover $\ell_{vi} \le b^{46n}$ and $V = (2/\operatorname{rd}_E)^f$. Consequently the squared sum of head weights divided by $\prod_{a<i} c_a$ is bounded by a common $\Gamma = b^{O(fn^9)}$.

Let $\bar u$ be the mean of the logarithms of the nominal weights and put $G_* = \log(g\operatorname{rd}_E/(2\pi))$. The different generator has $\prod_v |\Phi_v(\varpi)| = \operatorname{rd}_E^f$. Since $m = 2\varrho+2$, the specified window for N yields

$$A_* = \frac{1}{f}\sum_v \log X_v = (m-1)G_* - \log(gN) \le -m\bar u - \tfrac{25}{2}\log b, \qquad |A_*| = O(m\log b).$$

Here $g \ge b^{20}$. In Weyl's dimension product the factors for the unchanged initial $k-1$ parameters cancel. For each last index $l = k-1,\dots,r-1$ there are l remaining differences, each at most $w(v)e^{1/x}$, with denominator $l!$. Thus

$$f^{-1}\log\deg_{\mathrm{Bes}} \le mr(\bar u + 1/x) - \sum_{j=k-1}^{r-1}\log(j!).$$

The sphere product for vol $U(r)/\operatorname{vol} U(k)$ gives $\log \Theta_\infty \le \sum_{j=k}^{r-1} \log(j!)$. It follows that

$$\frac{1}{f}\log\left(\deg_{\mathrm{Bes}} \Theta_\infty^f \prod_v X_v^{k-n^7}\right) \le -\tfrac{25}{2} k \log b + O\big(m^2 \log b + mr/x + n^7 m \log b\big) \le -10k \log b$$

for all sufficiently large selected squares. Indeed $m \asymp \sqrt{x}$, $k \asymp x^5$ and $0 \le \bar{u} \le 2\log b$. Also $n^9 = o(k)$. Hence $(2\varrho+1)\Gamma\epsilon_* \le b^{-6kf-10}/\deg_{\mathrm{Bes}}$.

Ordered Poisson summation on the actual heads. Multiply the normalized finite forms by the real forms on the scaled heads. Denote the product forms by J_i, $0 \le i \le n$. At the top with common terminal value the finite form is $\mathcal{K}_L/(|L|m_Z)$ and the real form is R. Their entrywise product is positive semidefinite. Descent integrates one common coordinate at frequency zero by (8.14) and (8.15) and the real descent. The multiplier $1_{L^\perp}$ is one at zero and its modulus is at most one at all frequencies. The product system thus uses the same two local frequency bounds.

Let $\mathcal{P} \subset E^\times$ be a finite set of T-units including the terminal values. Use the factors c_i defined in the frequency budget. Let $\Lambda_i(p)$ be the product of the finite and real head weights for $p \in \mathcal{P}^i$ with the real scales applied. Write $H_i = \sum_{p\in\mathcal{P}^i} \Lambda_i(p)$. Use the bounds from the denominator packing and frequency budget, $H_i^2 \le \Gamma \prod_{a<i} c_a$ and a common ϵ_* for the sum of nonzero frequency gains and for each unequal-last-entry gain. These are uniform in T and L.

For each i let $F_i(z)$ be the sum of $J_{i+1}(p \oplus z, q \oplus z)$ over the two length-i heads, with the real scalings on those heads. This function is real nonnegative by the top Gram and successive zero-frequency integrations. It is integrable by the common-entry identities. At a rational sample the heads and varied entry are units in the product of fields. Continuity at the top and Fatou's lemma at each descent give lower semicontinuity at each such sample. The zero-coordinate sets in the integrating variables have measure zero.

Pull F_i back to the adeles using the real scalings ℓ_{vi} and projection at T, extending by zero outside those adeles integral off T. The complex scaling has real determinant $\prod_v \ell_{vi}^2$. The finite box off T in the product adelic measure has volume V, since the local measures used at T give the local integers mass one. The pullback has integral $c_i \int F_i$. The adelic quotient by E has volume one because the real covolume of $\mathcal{O}_E$ is $2^{-f}\sqrt{|D_E|}$ and the finite integral box has measure $V = 2^f/\sqrt{|D_E|}$. Index the Fourier coefficients by the inverse characters $c \mapsto \psi(-yc/\varpi)$, for $y \in E$. Their integrals use $\psi(yc/\varpi)$. The surviving frequencies have $y \in \mathcal{O}_{E,T}$ and occur at

$$\xi_T = y_T, \qquad \xi_v = \frac{2\pi \mathrm{i}\, \Phi_v(y)}{\Phi_v(\varpi)\ell_{vi}}.$$

Indeed the integrals use $\psi_T(\xi_T c_T)$ at T and $\mathrm{wave}(\xi_v c_v)$ at the scaled complex coordinate c_v. Both factors are those of the same global character. If y fails integrality off T, an integral translation at those places leaves the pullback invariant and acts by a nontrivial character, so its coefficient vanishes. At nonzero surviving y the product frequency is a unit in the product of local fields. The absolute coefficients sum to at most $c_i\epsilon_* H_i^2$ by the product bounds.

Periodize the nonnegative pullback on $\mathbb{A}_E/E$. Tonelli's theorem gives an L^1 function with these Fourier coefficients. Their absolute convergence gives a continuous function agreeing with the periodization almost everywhere by uniqueness of Fourier coefficients. Near zero each finite sum over the translates at the sampled rationals is bounded by the periodization. Lower semicontinuity of this sum at zero and density of the full-measure equality set show that its value at zero is at most the continuous value at zero. For $V_i \subseteq \mathcal{P}$ this gives

$$\sum_{t\in V_i} F_i(t) \le c_i\left(\Re \sum_{p,q\in\mathcal{P}^i} J_i(p,q) + \epsilon_* H_i^2\right),$$

where $F_i(t)$ is evaluated at the scaled rational point. At $i = 0$ the zeroth coefficient uses the bottom $E_0 = e_T(1+b^{-10})^f/\deg_{\mathrm{Bes}}$. At $i > 0$ the summands of J_i with equal last pivots give the preceding sum

with common entry set $\mathcal{P}$. Unequal last pivots u, u' use $y = (u^{-1} - u'^{-1})/g$, nonzero and integral off T. The real frequencies are $g\Phi_v(y)/\ell_{v,i-1}$ and the finite frequency is y_T. Their bound costs a further $\epsilon_* H_i^2$. Induction now gives

$$\sum_{t\in V_i} F_i(t) \le \left(\prod_{a=0}^{i} c_a\right)\left(E_0 + (2i+1)\Gamma\epsilon_*\right).$$

The numerical budget on this datum gives $(2\varrho+1)\Gamma\epsilon_* \le b^{-6kf-10}/\deg_{\mathrm{Bes}}$. At $i = \varrho$ the left side is $\Re\sum_{t,p,q} R_{pq}\mathcal{K}_L(p,q)/(|L|m_Z)$.

8.5.2 The packet constant and the row mass

Let mass_j be the sum of the inverse finite stabilizer orders for the standard self-dual positive definite genus in rank j over E/F. At every finite place E/F is split or unramified. In the split case a unimodular pairing identifies a basis on the first factor with a dual basis on the second. In the unramified field case a self-dual integral Hermitian lattice has an integral vector of unit norm after rescaling. Indeed reduction yields a vector with nonzero norm by the nondegenerate residue pairing and the surjective residue trace. The norm on unramified units is surjective. Split off that vector and repeat on its unimodular orthogonal complement. These self-dual lattices are maximal for integral norms. Polarization against every $\mathcal{O}_E$-multiple of a lattice vector puts an enlargement of integral norms in the Hermitian dual, by the perfect local trace pairing, including in residue characteristic two.

For such a maximal lattice in a totally definite rank-j Hermitian space over a CM extension of a totally real degree-f field, the mass theorem of Gan, Hanke and Yu gives

$$\mathrm{mass}_j = 2\,2^{-jf}\prod_{s=1}^{j} |L_F(1-s,\eta_{E/F}^{s})|. \tag{8.16}$$

In the odd-rank case the local correction set consists of ramified and unramified nonnorm-determinant places. In the even-rank case it consists of nonsplit nonnorm-determinant places. Both sets are empty here, since the determinant is a unit norm at each unramified place. The motive contributes the product and the Tamagawa number is 2. [28, (4.2)–(4.3) and Propositions 4.4–4.5, pp. 109–110]

The quotient volume on a Hermitian inner form with these finite lattices and the compact-dual real measure is also mass_j. The invariant densities correspond under inner conjugation, which has determinant one on the Lie algebra. The convergence factors depend on the common character motive. Thus the two forms have the same scale relative to Tamagawa measure and both have Tamagawa number 2. On the definite group with finite stabilizers of volume one and real Haar of total mass one, each class contributes its inverse stabilizer order. [28, Proposition 2.13 and (2.2)–(2.6)]

For a finite-order Hecke character χ of a totally real field, with conductor $\mathfrak{f}_\chi$ and local parities a_v, the completed function

$$\Lambda(z,\chi) = \left(|D_F|\,\mathrm{N}\mathfrak{f}_\chi\right)^{z/2}\prod_{v\mid\infty}\Gamma_{\mathbf{R}}(z+a_v)L_F(z,\chi)$$

satisfies $\Lambda(z,\chi) = w_\chi\Lambda(1-z,\overline{\chi})$, with $|w_\chi| = 1$. [29, §§4.4–4.5, pp. 338–346] For $\chi = \eta_{E/F}^{s}$ the conductor is the unit ideal and $a_v = a = s \bmod 2$ at every real place. The character equals its conjugate. Taking absolute values at $z = s$ and using the gamma identities $\Gamma_{\mathbf{R}}(z) = \pi^{-z/2}\Gamma(z/2)$ and

$$\frac{\Gamma_{\mathbf{R}}(s+a)}{|\Gamma_{\mathbf{R}}(1-s+a)|} = \frac{2(s-1)!}{(2\pi)^s}$$

give

$$|L_F(1-s,\eta_{E/F}^{s})| = |D_F|^{s-1/2}\left(\frac{2(s-1)!}{(2\pi)^s}\right)^{f} L_F(s,\eta_{E/F}^{s}). \tag{8.17}$$

The gamma ratio follows from reflection and duplication for Γ at the integers and half-integers. At $s = 1$ the character is nontrivial, so both values are finite.

Normalization of the packet sum. Set $s_U = \varrho(k+\varrho)$ and $F_{\mathrm{sc}} = \prod_{v,a} \ell_{va}^{k+2(\varrho-a)}$. For the actual column classes and the class-unipotent covolume,

$$h = |\bar{H}|(N^k-1)^f \,\mathrm{mass}_k, \qquad \nu F_{\mathrm{sc}} = U_0 N^{-kf} V^{-s_U} |D_E|^{\varrho/4}.$$

The first identity follows by refining the integral genus of rank k. At the places of N the lattice basis prescribes a nonzero dual vector in the residue space of dimension k. Each of the f split pairs has index N^k-1. At g the principal refinement has index $|\bar{H}|$ and its stabilizers are trivial, so the refined mass counts classes. The second identity is the class-unipotent covolume in the real frame coordinates used for P.

Every factor in the packet allowance can now be evaluated. With $\mathscr{M} = C^2/(\deg_{\mathrm{Bes}} \mathrm{mass}_r)$ one has

$$\frac{\Pi V^n}{\mathscr{M}\,\deg_{\mathrm{Bes}}\, c_G^f (\nu F_{\mathrm{sc}}/d_B)^2} = (1-N^{-k})^{-f} \frac{\mathrm{mass}_r}{\mathrm{mass}_k} \prod_{j=k}^{r-1} \left(\frac{(2\pi)^{j+1}}{\mathrm{rd}_E^{j+1/2}\, j!} \right)^f .$$

Indeed $C = h d_B$, $\Pi = |\bar{H}| U_0^2 h c_{\mathrm{quad}}$, and the powers of N reduce to $(1-N^{-k})^{-f}$. The remaining factor before the unitary volumes is $|D_E|^{1/4-\varrho/2} V^{2s_U+n+k}$. Substitute $V = (2/\mathrm{rd}_E)^f$ and $|D_E| = \mathrm{rd}_E^{2f}$. The power of the root discriminant is $-f(r^2-k^2)/2$. The powers of 2 before the sphere product use $2s_U + n + k - 1 = \sum_{j=k}^{r-1} j$. Finally $\mathrm{vol}\, U(r)/\mathrm{vol}\, U(k) = \prod_{j=k}^{r-1} 2\pi^{j+1}/j!$, giving the displayed identity.

The masses in this ratio are those of the standard positive definite self-dual genera in ranks j. The relative discriminant of E/F is the unit ideal by the covolume calculation. Apply the self-dual mass formula (8.16) and its consecutive critical-value factor (8.17). Since $|D_F| = \mathrm{rd}_E^f$, their quotient gives

$$\frac{\mathrm{mass}_{j+1}}{\mathrm{mass}_j} = \left(\frac{\mathrm{rd}_E^{j+1/2}\, j!}{(2\pi)^{j+1}} \right)^f L_F(j+1, \eta_{E/F}^{j+1}).$$

The absolute logarithm of the last factor is at most $f \log \zeta(j+1)$, by the absolutely convergent Euler product. This is at most $f\, 2^{-j-1}(1+2/j)$ by bounding the zeta tail beyond its first term by the integral from 2. For $j \ge 4u+12$ and $8f \le 2^u$ we obtain

$$\left| \log \frac{\mathrm{mass}_{j+1}}{\mathrm{mass}_j} - f \log \frac{\mathrm{rd}_E^{j+1/2}\, j!}{(2\pi)^{j+1}} \right| \le 2^{1-u}. \tag{8.18}$$

Take $u = \varrho$ throughout $k \le j < r$. The logarithm of the packet ratio above is bounded above by

$$2f/N + 2(r-k)\, 2^{-\varrho} = o(b^{-9}).$$

Here $\log N \ge m \log b$, f grows polynomially in x, $\varrho = m/2 - 1 \asymp \sqrt{x}$, and $\log b = x^{\varepsilon}$. Thus the ratio is at most $1 + b^{-9}$ on all sufficiently large selected squares.

Combining the packet and sieve inequalities requires

$$\frac{1}{\lambda_0^2} \prod_{a=0}^{\varrho} c_a = \frac{V^n d_B^2}{c_G^f F_{\mathrm{sc}}^2}.$$

This uses the exponent $k-1+2(\varrho-a)$ in λ_0 at every a, including the terminal index. The mass comparison now bounds the finite packet sum, divided by $\nu^2 \mathscr{M} b^{-6k|A|}$, by

$$(1+b^{-6})(1+b^{-9})\big((1+b^{-10})^f (1+b^{-6}) + b^{-6k(f-|A|)-10}\big) \le 1 + b^{-4}.$$

Indeed $f \le b/1000$, $b \ge 1000$, and $|A| \le f$.

The comparison on the original column space. The packet bound is independent of the finite choices. Exhaust T with the finitely many labels and heads fixed. The off-T indicator in $\mathcal{A}_{p,T}$ tends to one, with the absolutely summable rational amplitudes as dominator. Next exhaust the pivot values inside each block using that absolute sum. Finally take the supremum over finite lists of blocks, after squaring within each link. We obtain

$$\sum_{s\in\Omega_A,\, c,c',\,\Lambda} \left\|A^{\mathrm{cut}}_{\{\mathrm{disc}\neq 0,\ \mathrm{link}=\Lambda\}}(c,s,c')\right\|_{\mathrm{HS}}^2 \le (1+b^{-4})\nu^2 b^{-6k|A|}\mathscr{M}.$$

The crude energy comparison below uses this to bound its retained-link term before accounting for the cutoff error.

For the complementary energies the finite comparison matrix on the C original column coordinates is

$$(T_S(s))_{(c,i),(c',j)} = \nu^{-1}\overline{A_{S(F)}(c,s,c')_{ij}},$$
$$(K(s)-T_S(s))_{(c,i),(c',j)} = \nu^{-1}\overline{A_{S(F)^c}(c,s,c')_{ij}}.$$

These are the full-vector blocks and the conjugation is that of the original Petersson unfolding. The finite tensor $C_p(l)$ has coordinates on $Q\times Q$ and occurs with the cut real correlation in the open-link inequality. In the original unfolding the finite row condition contributes the principal coset volume $\mu_f(K(g))$, which cancels the same scalar in $P=(\mu_f(K(g))\nu)^{-1/2}\mathcal{Q}$ for the Poincaré map $\mathcal{Q}$. The same Petersson unfolding gives the identity block on $S(F)$ and the unitary transport matrices $T_S(t)$ for $t\in J$. The complementary energy comparison uses these same coordinates to bound $\Delta_i=\mathscr{E}_{S(F)^c}(JG_i)/(|J|C)$.

The two energy ratios on the selected datum. Here $r=k+m-1$ and $\bar{w}=f^{-1}\sum_v \log w(v)$. Use $U_0=g^{f\varrho(2k+2\varrho+1)}$ for the order of the reduced unipotent group and d_B for the number of Bessel coordinates per column class. The class count and the decomposition of J give the exact ratio

$$\frac{\mathscr{M}}{\mathscr{D}}=\frac{(N^k-1)^f\,\mathrm{mass}_k\;d_B}{U_0\;\deg_{\mathrm{Bes}}\,\mathrm{mass}_r}.$$

Indeed the number of column classes is $|\bar{H}|(N^k-1)^f\,\mathrm{mass}_k$, and $|J|=|\bar{H}|U_0$. These are the class and unipotent counts used in the open packet argument.

For the real Bessel coordinates at v, the $U(X_0)$-type in the construction is $\mathrm{Sym}^{\ell_v}\mathbf{C}^k$ with $0\le\ell_v\le a$. Thus its dimension multiplied by $(k-1)!$ is $\prod_{j=1}^{k-1}(\ell_v+j)\le(a+k)^{k-1}$. Here $a+k\le 2xb\le b^{5/4}$ for all sufficiently large selected squares, using $a<xb+N_1$ and the polynomial sizes of N_1,k. Consequently

$$0\le f^{-1}\log d_B+\log((k-1)!)\le\tfrac{5}{4}k\log b.$$

In the Weyl product for $\deg_v$, the factors involving only the first $k-1$ parameters cancel against their denominators. Here $k=r+1-m$ is even. For $i\le k-2$ in zero-based indices, $\lambda_i=k/2-1-i\ge 1-k/2$. For $l\ge k-1$ put $j=l+2-k$, so $\lambda_l=h_2-w_j(v)$. The parameter bounds $w_j(v)-h_2\ge w(v)+2jh_2$ and $h_2=r^8\ge k$ yield

$$\lambda_i-\lambda_l\ge w(v)+2h_2+1-k/2\ge w(v).$$

Every remaining difference is at least one and at most $w(v)e^{1/x}$, including the differences from the initial parameters. The denominator contributed by the parameter of index j, starting from zero, is $j!$. It follows that

$$f^{-1}\log\deg_{\mathrm{Bes}}=mk\bar{w}-\sum_{j=k-1}^{r-1}\log(j!)+O(m^2\log b+mr/x).$$

For the lower bound before taking this error, use $m(k-1)\bar{w}$. For the upper bound use $mr(\bar{w}+x^{-1})$.

Sum the consecutive mass estimate (8.18) over $k \le j < r$. The exponents of rd_E and 2π in its logarithm are both $(m-1)k + O(m^2)$, with opposite signs. Hence

$$f^{-1}\log\frac{\mathrm{mass}_r}{\mathrm{mass}_k} = (m-1)k\log\frac{\mathrm{rd}_E}{2\pi} + \sum_{j=k}^{r-1}\log(j!) + O(m^2\log b + m2^{-\varrho}/f).$$

Also $f^{-1}\log U_0 = (m-2)k\log g + O(m^2\log b)$. Insert these expressions in the exact ratio and use

$$\log N = m\bar{w} + (m-1)\log\frac{g\,\mathrm{rd}_E}{2\pi} - \tfrac{15}{2}\log b + O(\sqrt{\log b})$$

for the selected prime N. The weight terms cancel and the two factorial sums leave $\log((k-1)!)$. Thus

$$f^{-1}\log(\mathscr{M}/\mathscr{D}) = k\big(\log g - \tfrac{15}{2}\log b\big) + f^{-1}\log d_B + \log((k-1)!) + \mathcal{E},$$

where

$$|\mathcal{E}| \ll k\sqrt{\log b} + m^2\log b + mr/x + m2^{-\varrho}/f + N^{-k}.$$

Since $m \asymp \sqrt{x}$, $k \asymp x^5$, and $\log b = x^{\varepsilon}$, this error is at most $\frac{1}{8}k\log b$ on all sufficiently large selected squares. The lower bound $g \ge b^{20}$ and the Bessel dimension estimate therefore give

$$b^{11kf} \le \frac{\mathscr{M}}{\mathscr{D}} \le (g/b^{49/8})^{kf} \le \tfrac{1}{2}(g/b^{6})^{kf}.$$

The last comparison uses $b^{kf/8} \ge 2$. This proves $\tau^{44}\mathscr{D} \le \mathscr{M}$ and $2(b^6/g)^{kf}\mathscr{M} \le \mathscr{D}$, with $\tau = b^{kf/4}$ on the same datum. The auxiliary factors used in the crude bound also satisfy $2\tau^2(b/g)^{kf} \le 1$, $N^f \le \tau$, and $\tau \ge b^{n+31}$, where $n = \varrho + 1$. For these use $g \ge b^{20}$, $N \le b^{60n}$ and $k \ge 1000(2n)^9$.

8.6 Compatible orbits, cutoff and real frame bounds

8.6.1 Compatible frame orbits

For a rational term γ over (c, s, c'), put $B = U \cap \gamma U \gamma^{-1}$ and $\Lambda = U_c \cap \gamma U_{c'} \gamma^{-1}$. After joining the cutoff integrals in its class-unipotent double orbit, the quotient of integration is $\Lambda\backslash G_\infty$. The intersection B is a unipotent group over F, and Λ is its rational intersection with an open compact subgroup at the finite places. Unipotent strong approximation gives a compact quotient $\Lambda\backslash B_\infty$. The integrand transforms on B_∞ by the ratio of the two Bessel characters. This character is trivial on Λ. If the two finite characters disagree on any rational element of B, the product formula gives a nontrivial ratio on B_∞. Disintegrating the absolutely convergent integral over $\Lambda\backslash B_\infty$ then gives zero by translation on this compact quotient. We can therefore restrict the rational orbits to the pairs with agreeing characters.

For such a pair let $f_0, \dots, f_\varrho$ be $e_0, \dots, e_{\varrho-1}, z$, and write $f_i' = \gamma f_i$. Let X_i, X_i' be their spans through index $i-1$. Suppose that a common vector has leading coefficients $c, c' \neq 0$ at indices $i, i' < \varrho$. Denote the vector by u. For y perpendicular to $X_{i+1} + X'_{i'+1}$, the isometry

$$v \longmapsto v + h(v, y)u - h(v, u)y + \ell\, h(v, u)u, \qquad \ell + \bar{\ell} = -h(y, y),$$

belongs to both flag unipotents. Here h is linear in its first argument. The sum of its simple entries for the first flag is $c\,h(f_{i+1}, y)$. Earlier simple entries vanish by the orthogonality imposed on y, and later entries vanish by the level of u. The second flag gives $c'\,h(f'_{i'+1}, y)$. The common finite character is χ_f. The character equality holds on replacing y by every E-multiple of y. The field E is dense in its finite adeles, so a nonzero difference of the two simple-entry sums would contradict the nontriviality of this additive character. Thus

$$c f_{i+1} - c' f'_{i'+1} \in X_{i+1} + X'_{i'+1}.$$

Thus a common vector at the next pair of levels has the same leading coefficients. A common vector with the terminal index on only one side would have negative norm on that side and zero norm on the other. Iteration shows that the levels are equal. At the terminal index both distinguished vectors have norm $-N$, so $c\bar{c} = c'\bar{c}'$. A common vector already at the terminal level gives that equality directly.

Write $X = X_\varrho$, $X' = X'_\varrho$ and $I = X \cap X'$. At a real place $(X' + Ez')^\perp$ is positive semidefinite with radical X', where $z' = \gamma z$. If a common vector occurs, let j be the first level with one. The lifted common vectors from j through $\varrho - 1$ are independent, and the intersection I meets X_j and X'_j trivially. Hence $X = X_j \oplus I$ and $X' = X'_j \oplus I$. The terminal common vector gives $z' - \zeta z \in X + X'$ with $\zeta\bar{\zeta} = 1$. A vector of X_j perpendicular to X'_j is also perpendicular to I and hence to X'. The terminal relation makes it perpendicular to z'. Its isotropy then puts it in X', by the positive semidefinite description. Thus it is zero and the leading frame minor of size j is nonzero. The convention for size zero is 1.

If $X + Ez$ and $X' + Ez'$ have zero intersection, their cross pairing Z is nonsingular. A vector in its left kernel has nonpositive norm as an element of $X + Ez$. Since it is in $(X' + Ez')^\perp$, it must have zero norm and belong to X'. It also belongs to X, and hence is zero. For a pair with zero discriminant the equation $|\det Z|^2 = N^2 |\det Z_\varrho|^2$ now gives a nonsingular pairing between X and X'. Their joint span is hyperbolic. Project z, z' orthogonally onto its complement. We subtract only an element of X from z and an element of X' from z', so the projections have norm $-N$. The vanishing discriminant makes their joint Gram matrix singular by (8.4). The complement has negative index one. Subtracting a multiple of the first projected vector from the second then gives an isotropic vector in a positive definite space, so the two projections are proportional. This gives a nonzero intersection of the augmented spans. Consequently every compatible pair with zero discriminant has a first common level $j \le \varrho$, with common leading coefficients $(\zeta, 1)$ at every level from j on and with the leading minor of size j nonzero.

At depth zero the relations make γ upper triangular on the entire augmented frame, with constant frame diagonal ζ. The class representatives fix the frame, and the lattice change preserves its lines with the same scale on both sides of a term position. Thus the multiplier on its first vector is still ζ. Integrality of the position gives $\zeta \in \mathcal{O}_E$. For a position in JG_i, its component away from G_i fixes that first vector, so ζ reduces to one at a chosen place above g_{1-i}. At this split place the completion is $\mathbf{Q}_{g_{1-i}}$ and its principal units are torsion-free since g_{1-i} is odd. The norm-one algebraic integer ζ is a root of unity, so it equals one. The upper unitriangular frame condition then puts γ in the Bessel subgroup $S(F)$. This is the required depth-zero exclusion for the complementary energy.

The class cutoffs in the frame box. Write $a = \varrho$, $n = a + 1$, $B = b^{2n^2}$, and $\theta = b^{6n+2}$. Let $\Gamma_c = U_c$ be the rational class-unipotent group at level g, and $\Gamma_c^0 = U_c^0$ its subgroup fixing the frame $\mathbf{f} = (e_0, \dots, e_{a-1}, z)$ pointwise. At the real places use the normalized frame with scales $s_{vh} = \kappa_v^{a-h}/\sqrt{N}$ for $0 \le h \le a$. The set Box consists of the points y whose inverse frame vectors have length at most B at every place and for which the local size is at most θ. Both conditions are left invariant under $U_\infty^0 H_{0,\infty}$.

The rectangular coordinates in the class-unipotent fibre give an element of Γ_c with rational frame action $1 + gNxE_{ih}$ for every $i < h \le a$ and $x \in \mathcal{O}_E$. Use the single free entry gNx at column f_h in row i. The hyperbolic coordinates are integral in the class basis and the marked coordinate has denominator dividing N, so the row divided by g is integral in that basis. The skew entry can be chosen as in the fibre computation. In the real frame the same entry is multiplied by s_{vh}/s_{vi}, of absolute value at most one. In the ordinary Euclidean norm on $\mathbf{C}^f$, the trace dual of $\mathcal{O}_E$ is $2\mathfrak{D}_E^{-1}$. Its shortest length is at least $2/b$, since $\mathfrak{D}_E = (\varpi)$, $|\Phi_v(\varpi)| \le b$, and a nonzero algebraic integer has some conjugate of absolute value at least one. The elementary covering-radius estimate $\mu(L) \le (2f)^2/\lambda_1(L^*)$ therefore gives a covering radius at most b^3 for $\mathcal{O}_E$, using $f \le b/1000$.

At an arbitrary y, choose $u_v \in U_v$ with $\|u_v y_v\|_{\mathrm{op}} \le \mathrm{size}(y_v)$, by the definition of the size with the additive one. Starting with the frame matrices of u_v, multiply on the right by the preceding rational shears, processing row indices from largest to smallest. A shear at (i, h) changes that entry by $gN(s_{vh}/s_{vi})\Phi_v(x)$, and changes no entries in lower rows. Approximation in $\mathcal{O}_E$ at once over all v leaves absolute value at

most $gNb^3 \le b^{60n+26}$ in each processed entry. For the resulting $q \in \Gamma_c$, the inverse frames of $q^{-1}y$ have lengths at most b^{60n+27} size(y_v) at v. This follows by applying $(u_v y_v)^{-1}$, whose operator norm equals that of $u_v y_v$, to the columns of the frame matrix of $u_v q_v$.

Let $\mathcal{F}_c(y)$ be the left Γ_c^0-cosets of $t \in \Gamma_c$ for which $(ty)_v^{-1}\mathbf{f}$ obeys these last length bounds. This is a finite nonempty set. Finiteness follows by applying y_v to the bounded columns. The resulting matrices have bounded frame coefficients at every embedding, and the rational coefficients along the isotropic rows are integral in the class chart, which fixes those rows and the frame. Equal frame matrices give the same coset. Put

$$\omega_c(y) = \frac{\mathbf{1}_{[1]\in\mathcal{F}_c(y)}}{|\mathcal{F}_c(y)|}.$$

This function is measurable and belongs to $[0,1]$. Every element of $U_\infty^0 H_{0,\infty}$ fixes the inverse frame and the size in this test, since $t_v^{-1}\mathbf{f}$ is in the frame span. For $t \in \Gamma_c$ the cosets counted at ty are the cosets counted at y right translated by t^{-1}. Thus $\sum_{t\in\Gamma_c^0\backslash\Gamma_c} \omega_c(ty) = 1$. When $\omega_c(y) > 0$ and the cut holds at v, the inverse frame lengths at that place are at most $b^{66n+29} \le B$. Choose a measurable partition cutoff p_c for the discrete left action of Γ_c^0 on G_∞. For example, the indicator of a measurable fundamental section is such a cutoff. Take $\beta_c = p_c\omega_c$. Summing first within the left cosets shows $\sum_{t\in\Gamma_c} \beta_c(ty) = 1$. This is the partition of unity in the Poincaré unfolding, so it gives the same periodization and the same fibre integral ν.

Separation of the rational frames in the cutoff. Let $n = \varrho + 1$, $B = b^{2n^2}$, and write $f_0, \dots, f_\varrho$ for the augmented frame, with $f_\varrho = z$. Consider two rational terms in the same residue count after their class-unipotent reductions into the cutoff box. Write γ for their relative rational isometry, and y_v for the real coordinates of the first reduced term. Use the supremum, over v, a, of Euclidean distance between corresponding vectors $y_v^{-1} f_a$ and $y_v^{-1}\gamma_v f_a$. If this distance is less than $(NB)^{-1}$, the two rational frames coincide.

Indeed the class representatives preserve the frame. The relative integral term position consequently makes every frame pairing $h(f_b, \gamma f_a)$ an algebraic integer: in the class lattice chart these are entries of the product of its integral Gram matrix and an integral isometry. The entries of γf_a along the first ϱ coordinate rows are integral as well, since the class change acts on the middle block and fixes those coordinate rows and the marked frame. In the normalized real frame, the scale at f_a is $s_{va} = \kappa_v^{\varrho-a}/\sqrt{N}$, with $1 \le |\kappa_v| \le b^{23}$. The real Gram matrix has operator norm one. Pairing the difference of the two vectors with $y_v^{-1} f_b$, of length at most B, bounds each conjugate of $h(f_b, \gamma f_a) - h(f_b, f_a)$ by NB times their distance. The factor N covers $1/|s_{va}s_{vb}|$. An algebraic integer whose absolute values at all embeddings are less than one vanishes, by its norm. All these pairings agree.

The pairings with the isotropic frame and with z fix the dual frame coordinates and the z-coordinate. Unitarity now makes the component of $\gamma f_a - f_a$ in the definite middle space have norm zero. Thus this difference lies in the isotropic span. For fixed a, recover its coordinates there starting with the last one. In the normalized frame the coefficient at e_i is $(s_{va}/s_{vi})\Phi_v((\gamma - I)_{i,a})$, where the index a on the right denotes the frame column. As y_v is in the cut, choose $u_v \in U_v$ with $\|u_v y_v\|_{\mathrm{op}} \le b^{6n+2}$. The i-th row of the upper unitriangular action of u_v reads precisely that coefficient once the later coefficients vanish. Hence

$$|\Phi_v((\gamma - I)_{i,a})| \le b^{23\varrho+6n+2}\,\|y_v^{-1}\gamma_v f_a - y_v^{-1} f_a\| < 1.$$

Here $b^{23\varrho+6n+2} \le NB$ for the selected large squares. Integrality and the same norm argument finish the descending induction. This proves the asserted separation.

Translate one frame to the origin in the real subspace containing the differences, of dimension $D \le 2rnf$ in the unrestricted count. The resulting centres lie in a norm ball of radius $2B$. Their open balls of radius $1/(2NB)$ are disjoint, so comparison of Lebesgue volumes bounds their number by $(1 + 4NB^2)^D \le b^{6n^2 D}$, since $N \le b^{60n}$ and n tends to infinity. For two reduced terms with the same source class and residue, the relative finite position is the quotient of their two integral positions and belongs to $K(g)$.

When their frames coincide, class freeness gives the same target class and the relative term is a class unipotent fixing the frame. The boxing corrections are themselves class unipotents.

For the version with a right translate $x \in G_\infty$, fix the source class, residue and t. Group the eligible rational terms by right multiplication by $U_{c'}$ for each target class c'. This leaves the finite position modulo $K(g)$ fixed and leaves the test $\gamma^{-1}tx \in \mathrm{InCut}$ invariant. Within each such orbit the weights $\beta_{c'}(\gamma^{-1}t)$, evaluated at the unshifted arguments, sum to one, or at most one for a restricted term set. Choose a representative of the full orbit with $\omega_{c'}(\gamma^{-1}tx) > 0$, using the partition at the shifted argument. Its inverse frame $(tx)^{-1}\gamma\mathbf{f}$ lies in the box. Distinct orbits have distinct boxed frames by the class freeness and the relative-position argument above. The separation and volume comparison on these frames bounds the total weight by b^{12rn^3f} for every x.

The full and cut relative traces. Use the Bessel vectors and the class cutoffs of the normalized unfolding. Write $a = \varrho$, $n = a + 1$, $r = k + 2a + 1$, and $d_W = \prod_v |I_v|$, so $C = |\mathrm{Cl}|d_W$. For a family B of matrix blocks put

$$\|B\|_\Omega = \nu^{-1}\Big(\sum_{s\in\Omega,c,c'} \|B(c,s,c')\|_{\mathrm{HS}}^2\Big)^{1/2}.$$

All the residue sums use counting measure. The norm of the tensor formed from two product Bessel vectors is the product of their norms. For the cut pair $W_v, W_v^\flat$ use the real majorant and decay (5.14)–(5.15). Write $M_v = A_v$ and $m_v = \mu_v$ for the U_v-invariant majorant and its finite positive measure at v. The norm in these bounds is the one on the full vectors and their cut, and the integral uses the same quotient Haar. Here $\mathrm{InCut}_v = \{\mathrm{size} \le R\}$ as in the absolute periodization.

Write $M(t) = \prod_v M_v(t_v)$ and let $K_\gamma(t)$ and $K_\gamma^\flat(t)$ be the full and cut integrands before their two class cutoffs. If the two tensors differ, some local vector differs at t_v or at $(\gamma^{-1}t)_v$. Hence

$$\|K_\gamma(t) - K_\gamma^\flat(t)\|_{\mathrm{HS}} \le 2b^{-3r^3f}M(t)M(\gamma^{-1}t).$$

At fixed c, s, t and right translate $x \in G_\infty$, the sum of $\beta_{c'}(\gamma^{-1}t)$ over (c', γ) with $\gamma^{-1}tx \in \mathrm{InCut}$ is at most b^{12rn^3f} by the weighted frame count. This holds for each right translate x in the positive measure. Integrate this count against $\prod_v m_v$ using the majorant of $M(\gamma^{-1}t)$. The integral of $\beta_c(t)M(t)$ is $\nu\prod_v \int_{U_v\backslash G_v} M_v$ by the class cutoff partition and the unipotent covolume. Thus for every c, s and any term set T,

$$\sum_{c'} \|A_T(c,s,c') - A_T^\flat(c,s,c')\|_{\mathrm{HS}} \le 2\nu\, d_W\, b^{(12rn^3-r^3)f} \le \tfrac{1}{2}\nu\, b^{-k^3f}.$$

For the last inequality use $d_W \le b^{2kf}$, $r^3 - k^3 \ge r^2$, and $12rn^3 + 2k + 2 \le r^2$. The latter follows from $r \ge 24n^3$, $k + 1 \le r$ and $r \ge 4$. The square of the sum in c' bounds the sum of the squares. Since $|G| \le g^{r^2f} \le b^{23r^2f}$ and $|\mathrm{Cl}| \le \mathscr{D} = |J|C$, we get

$$\|A_T - A_T^\flat\|_\Omega^2 \le b^{-40}\mathscr{D}. \tag{8.19}$$

Here $2k^3 - 23r^2 \ge 2k^2(k - 46) \ge 40$ by $r \le 2k$ and $k \ge 47$. The bound also proves convergence of the difference with these weights. Absolute convergence of the two amplitude sums permits their use on every term set in the unfolding.

The Haar measures on the two real unipotent groups. Work at a fixed infinite place in the normalized frame. Write $m = \varrho$, $n = m + 1$, $k = \dim X_0$, and $r = k + 2m + 1$, and order the blocks as (e, X_0, z, e^*). Use the Euclidean operator norm, so $g^{-1} = Jg^*J$ gives $\|g^{-1}\| = \|g\|$. Write $F = (e_0, \dots, e_{m-1}, z)$ for the frame matrix. In the Gram section coordinates the frame stabilizer is U^0H, with $H = U(X_0)$ and U^0 parametrized there by (B, A_{Gr}). For the tube write $A = -B^*$ and $S = A_{\mathrm{Gr}}$. This linear change has absolute real determinant one, and the e, e^* block is $S - AA^*/2$.

Multiplication of $u(A,S)$ and $u(B,V)$ adds A,B and replaces the second coordinate by $S+V+(BA^*-AB^*)/2$. Inversion negates A,S. Thus $du_0 = dA\,dS$ is left and right Haar, with dS Lebesgue in the entries $S_{ab} \in \mathbf{C}$ for $a<b$ and $\operatorname{Im} S_{aa} \in \mathbf{R}$. To compare the normalization, the real coordinates of the skew matrix are sent to these coordinates by

$$(c_{ab}, c_{ba}) \longmapsto \big((c_{ab}-c_{ba})/\sqrt{2},\ (c_{ab}+c_{ba})/\sqrt{2}\big) \quad (a<b), \qquad c_{aa} \longmapsto \operatorname{Im} S_{aa}.$$

All these maps have absolute determinant one. At fixed A the free entries of $S-AA^*/2$ differ from those of S by a translation. Consequently the prescribed Haar measure on U^0 is $du_0 = dA\,dS$.

The frame action $uF = FQ(u)$ of U has kernel U^0 and image the upper unitriangular n-by-n group. For $Q = \left(\begin{smallmatrix} a & v \\ 0 & 1 \end{smallmatrix}\right)$ use the lift fixing X_0, taking e to ea, z to $z+ev$, and e^* to $(e^*+zv^*+evv^*/2)a^{-*}$. Here $C^{-*} = (C^*)^{-1}$ for any invertible matrix C. These columns have the original J-pairings, as follows from $a^*a^{-*} = 1$ and the cancellation $vv^*/2 + vv^*/2 - vv^* = 0$ in the dual Gram matrix. The lift belongs to U by its triangular form. Its conjugation action on U^0 sends (A,S) to (aA, aSa^*). Both maps preserve the displayed Lebesgue measures since a is unipotent and the corresponding linear maps are unipotent. Multiplication on either side in the upper unitriangular group is triangular in the entries ordered by superdiagonal distance and has diagonal one. The product $dQ\,du_0$ under $(Q,u_0) \mapsto u_0\widetilde{Q}$ is therefore Haar on U and agrees with the prescribed product measure. It also gives Lebesgue in the full free entries. In fact the e, e^* block of $u(A,S)\widetilde{Q}$ is $(S-AA^*/2+vv^*/2)a^{-*}$. For fixed A, Q, right multiplication by a^{-*} is a shear, triangular in descending column order in each free upper row, with coefficient one on each complex entry and on the imaginary diagonal. The remaining free entries are A and the strict upper entries of Q. Here dQ is $d\mathfrak{R}\,d\mathfrak{I}$ in each free complex entry.

Write dy, dh, c_G for the filling Haar, the probability Haar on H, and the constant in the Gram density formula. Write $d\dot{y}$ for the quotient disintegration of dy by the just computed du on U. The local data are read after the diagonal basis change. At place v the frame scales are $s_{va} = \kappa_v^{m-a}/\sqrt{N}$, with the reciprocal scales on e_a^* and the unscaled X_0. The integral coordinates and rational term are transported before taking operator norms, free entries, and Gram pairings. In particular the term γ of the real calculation is the rescaled image of the rational term. Each leading principal Gram determinant of order j gains $\prod_{a<j}|s_{va}|^2$ from row and column scaling. The last scale is $1/\sqrt{N}$, so the rational discriminant $|D_n|^2 - N^2|D_m|^2$ has common factor $(\prod_{a<n}|s_{va}|^2)^2$ in the rescaled real discriminant. This is also the frame and Haar used by the full vector in the normalized Poincaré map.

A compact tube with prescribed fiber bound. For $u = u(A,S)$ set $L(u) = \max(\|A\|, \|S\|^{1/2})$. The group law gives $L(uu') \le L(u)+L(u')$, since the central correction has norm at most $\|A\|\|B\|$. Also $L(u^{-1}) = L(u)$. The dilations $(A,S) \mapsto (\lambda A, \lambda^2 S)$ multiply du_0 by λ^{D_0}, $D_0 = 2m(k+m) \le r^2$, and multiply L by λ. This includes $m=0$, when $U^0 = \{1\}$ with mass one.

For $T \ge 8r$, put $R = r^2T^2$, $\Omega = \{L \le R\}$ and $A_T = \Omega B(T)$, where $B(T) = \{y \in G_1 : \|y\| \le T\}$. The two factors are compact. The first is a closed bounded box in the displayed coordinates and the second is a closed matrix norm ball in $U(J)$. The box has finite positive du_0-measure. On this box, $\|u\| \le 1 + 2R + 3R^2/2 \le 5R^2$, so

$$A_T \subset B(5r^4T^5) \subset B(T^9) \subset B(T^{20n}).$$

If the fiber over y is nonempty, fix $u_0y = \omega_0 z_0$ with $\omega_0 \in \Omega$ and $z_0 \in B(T)$, and set $q = \omega_0^{-1}u_0$. For another $uy = \omega z$ in the fiber, $uq^{-1} = \omega z z_0^{-1}$. Here $zz_0^{-1} \in U^0$ and its norm is at most T^2. For any $u(A,S)$ of norm at most T^2, its A-block and the skew part S of its e, e^*-block have norms at most T^2, so $L(u) \le T^2$. Right translation by q^{-1}, followed by homogeneous scaling of the box, gives

$$\int_{U^0} \mathbf{1}_{A_T}(uy)\,du_0 \le (1+1/r^2)^{D_0}\,\mathrm{vol}_{U^0}(\Omega) \le e\ \mathrm{vol}_{U^0}(\Omega) \le T\ \mathrm{vol}_{U^0}(\Omega).$$

The inequality holds also for an empty fiber.

Geometry of the actual real frame box. Fix γ with $Z = F^*J\gamma F$ and $\mathfrak{g} = |\det Z|^2 - |D_m(Z)|^2 > 0$. Use the notation $\mathrm{Box}(\rho) = \{y : \|y^{-1}f_a\| \le \rho \text{ for all } a\}$ and $\mathrm{size}(y) = \inf_{s\in UH}(\|sy\| + 1)$. Suppose $\mathrm{Box}(\rho) \cap \gamma\,\mathrm{Box}(\rho) \neq \emptyset$, with $\rho \ge b \ge 8r$. This gives $|Z_{ab}| \le \rho^2$ by taking the pairings after a common isometry. The determinant bound $\mathfrak{g} \le (n\rho^2)^{2n}$ implies $M_2 \ge \rho$ for

$$M_2 = \rho^{9n}/\mathfrak{g}, \quad P_2 = M_2^{9n}, \quad T_2 = P_2^5, \quad P = P_3 = T_2^n, \quad c_2 = c_G(\mathfrak{g}/2)^{k-1}.$$

Also $M_2^{-1} \le \mathfrak{g}$ and $P \ge b^5$. The sign condition holds at each real place of a rational term of nonzero discriminant. Indeed the Gram matrix on $(F, \gamma F)$ has determinant $(-1)^n\mathfrak{g}$ by swapping the two block rows and using the rank-one Schur complement. The identity extends to singular Z by continuity. When the pair Gram is nonsingular its negative index is n: the ambient negative index is n, and inside the pair span the negative vector z and the m-dimensional totally isotropic subspace in $z^\perp$ force at least n.

More generally, let $M \ge 8n$, $|Z_{ab}| \le M^2$, $\mathfrak{g}(Z) \ge M^{-1}$ and $L_* \ge M^{9n}$. For $\|Z' - Z\| < L_*^{-6}$, the two matrix norms are at most M^3. By column multilinearity the determinant of a principal block of order $j \ge 1$ changes by at most $jM^{3(j-1)}\|Z' - Z\|$. The total change in $\mathfrak{g}$ is at most $4nM^{6n-3}L_*^{-6} \le M^{-1}/2$ (the order-zero minor is constant). Hence $\mathfrak{g}(Z') \ge \mathfrak{g}(Z)/2$. The singular values and the last cofactor now show that

$$\|Z'\|, \quad \|(Z')^{-1}\|, \quad \left(1 - |((Z')^{-1})_{nn}|^2\right)^{-1/2} \le L_*.$$

Indeed $|\det Z'| \ge (2M)^{-1/2}$, which bounds the inverse norm by $\sqrt{2M}\,M^{3n-3}$. The last quantity is $|\det Z'|/\sqrt{\mathfrak{g}(Z')}$, at most $\sqrt{2M}\,M^{3n}$. The same inverse estimate applies without a neighborhood whenever the entries and the discriminant of the tested matrix satisfy the original bounds with M.

Use the explicit Gram section, with $E = p$, $D = p_0$, $Q = Z^{-1}$, $\Delta = \Lambda$, $A = T_Z$, $A' = T_{Z^*}$, and $S_Z = X_Z = \gamma_Z F$. Its column identities give

$$\gamma_Z F' = (FQ^* - S_Z A')\Delta^{-1}, \qquad \gamma_Z^{-1}F(Z^*)^{-1} = FA' + F'\Delta.$$

We use these identities to bound the section and its derivative.

On a region where the three displayed bounds hold with $L_* = P \ge 8$, we have $\|A\|, \|A'\| \le 2P^2$, $\|S_Z\| \le 3P^3$, and $\|\gamma_Z\| \le 11P^6 \le P^{12}$. For a real matrix coordinate direction, $dQ = -Q(dZ)Q$ has norm at most P^2. Differentiation gives $\|dA\|, \|dA'\| \le 2P^3$, $|ds| \le P^3$, $|d(s^{-1})| \le P^5$, and $\|dS_Z\| \le 6P^4$. In the formula for $\gamma_Z F'$, the derivative before differentiating Δ^{-1} has norm at most $(P^2 + 18P^6)P$, and differentiation of Δ^{-1} costs at most $(P + 6P^5)P^5$. Thus the derivative of γ_Z itself is at most $32P^{10} \le 4P^{18}$ in each real coordinate.

The pair $(F, \gamma F)$ has the same Gram as $(F, \gamma_Z F)$. Their Gram is nonsingular and has signature (n, n). The isometry matching the two pairs extends by a unitary map between the positive orthogonal complements. It thus gives $s_1 \in U^0 H$ with $s_1\gamma_Z F = \gamma F$, and $s_2 = \gamma^{-1}s_1\gamma_Z$ also fixes the frame. Hence $\gamma = s_1\gamma_Z s_2^{-1}$. Choose h, v independently with laws dh and $du_0|_\Omega \,/\, \mathrm{vol}_{U^0}(\Omega)$, where Ω is the tube box for $T = T_2$. Set $t = s_2(vh)^{-1}$ and $A_2 = A_{T_2}$.

For these choices, suppose $y \in \mathrm{Box}(\rho)$ and $\mathrm{size}(\gamma t y) \le \rho$. Put $y_0 = (vh)^{-1}y$. Then $y_0^{-1}F = y^{-1}F$, and $\mathrm{size}(\gamma_Z y_0) = \mathrm{size}(\gamma t y)$ by the left UH-invariance of size. By its definition, some $u \in U$ satisfies $\|u\gamma_Z y_0\| \le \rho$: approximate the infimum of $\|s\gamma_Z y_0\| + 1$ to get $\|s\gamma_Z y_0\| < \rho$, then remove the Euclidean-unitary left H-factor. If $u^{-1}F = FC'$, then $Z' = F^*Ju\gamma_Z F = (C')^*Z$. Since C' is upper unitriangular, left multiplication by $(C')^*$ preserves every leading principal determinant. By the two frame bounds, $|Z'_{ab}| \le M_2^2$, and therefore $\|(Z')^{-1}\| \le P_2$. We have

$$y_0^{-1}\gamma_Z^{-1}F(Z^*)^{-1} = (u\gamma_Z y_0)^{-1}F((Z')^*)^{-1}.$$

The right side has norm at most M_2P_2 because F is an orthonormal column inclusion. By the inverse section formula and $\|A'\| \le 2P_2^2$, each $\|y_0^{-1}e_a^*\|$ is at most $M_2P_2 + 2\sqrt{n}M_2P_2^2 \le P_2^3$. In the negative eigenspace of J the orthonormal basis $(e_a - e_a^*)/\sqrt{2}, z$ gives $\|y_0^{-1}|_{J=-1}\| \le \sqrt{2n}P_2^3$. For any J-unitary matrix $X = \left(\begin{smallmatrix} a_1 & b_1 \\ c_1 & d_1 \end{smallmatrix}\right)$ in the positive-negative splitting, $XJX^* = J$ gives $a_1a_1^* = 1 + b_1b_1^*$ and $c_1c_1^* = d_1d_1^* - 1$. Thus if the negative restriction has norm at most $L \ge 1$, its full norm is at most $\sqrt{1 + L^2} + 3L \le 5L$. It follows that $\|v^{-1}y\| = \|hy_0\| = \|y_0^{-1}\| \le 5\sqrt{2n}P_2^3 \le T_2$. In particular $y \in A_2$.

Evaluation at the section and the full-vector contraction. Use the full normalized W, $d = |I_v|$, and $\delta = \deg_v$. The input from the Bessel estimates is $\int_{U\backslash G_1} \|W\|^2 d\dot{y} = d$, with the unitary left UH-law for W and its right derivative words W_w. For $P \geq b^5$ their Schur estimate, in the same Haar, is

$$\int_{U\backslash G_1} \Big\| \int_{G_1} \psi(y) W_w(gy) dy \Big\|^2 d\dot{g} \leq P^{2|w|}(d/\delta) \int_{G_1} |\psi|^2 \, dy, \qquad |w| \leq 2n^2.$$

For the geometric calculation use the Gram inequality with a single left compact variable outside all integrations. The density on the open cell is $c_G \mathfrak{g}(Z')^{k-1} dZ'\, dh\, dv_0\, du_0$, where $dv_0 = du_0(v)$ and dZ' is coordinate Lebesgue measure.

First note an unfolding bound. Let $\Phi \geq 0$ be Borel, with $\int_{U^0} \Phi(ug) du_0 \leq K$ for each g. Suppose that whenever $\Phi(g)$ and $\Phi(sg)$ are nonzero, with $s \in U$, every free entry of $Q(s)$ has modulus at most $R_* \geq 1$. For any nonnegative left-U-invariant w one has

$$\int_{G_1} \Phi(g) w(g)\, dg \leq K(2R_*)^{2n^2} \int_{U\backslash G_1} w(g)\, d\dot{g}.$$

On each orbit meeting the support, use a point in the support as origin by right translation in its U-integral. In the decomposition $s = u_0 \widetilde{Q}$ the U^0-fiber has integral at most K and the upper entries range over discs of radius R_*. Their product volume is $(\pi R_*^2)^{n(n-1)/2} \leq (2R_*)^{2n^2}$. Apply this on the quotient using the stated disintegration.

Let ψ be integrable and supported in $B(P^{20})$. For $v \in V_P = U^0 \cap B(P^{20})$, $h \in H$, evaluate the convolution at $g = \gamma_{Z'}(vh)^{-1}$. The cube $Z + [0, P^{-8}]^{2n^2}$ in the real and imaginary entries lies in $\mathcal{Q} = \{Z' : \|Z' - Z\| < P^{-6}\}$, since its operator radius is at most $2nP^{-8}$. The discriminant throughout this region is at least $\mathfrak{g}/2$. In a cube direction $X = g^{-1}\partial g$ has norm at most $4P^{70}$ from the section bounds and $\|vh\| \leq P^{20}$. The derivative of $W_w(gy)$ uses $y^{-1}Xy$. For each term in $(y^{-1}Xy)_{ab}$, absorb $(y^{-1})_{ac} y_{db}/P^{40}$ in ψ and retain the coefficient $P^{40} X_{cd}$, of modulus at most P^{112}. The absorbed factor has modulus at most one on the support.

Include all such modified tests together with their corresponding derivative words in the family of convolutions. A first derivative of any family member is a sum of at most r^4 members of one higher word length, with the stated coefficient bound. Differentiation under the integral follows from the continuous derivatives on the compact sets in this calculation and the integrability of the test. For a one-dimensional interval of length $\varepsilon = P^{-8}$, the fundamental theorem of calculus and Cauchy–Schwarz give

$$\|f(0)\|^2 \leq \tfrac{2}{\varepsilon} \int_0^{\varepsilon} \|f(x)\|^2 dx + 2\varepsilon \int_0^{\varepsilon} \|f'(x)\|^2 dx.$$

Integrate in h, v and the other current coordinates, and at each step bound the resulting integrals by their maximum over the permitted members. This costs at most $2P^8(1 + r^8 P^{224}) \leq P^{241}$ per coordinate, and uses only first derivatives of family members and words of length at most $2n^2$.

To bound each resulting cube integral, insert a nonnegative Borel factor $\beta(g)$ with $\int_{U^0} \beta(ug) du_0 = 1$. For instance take $e^{-\|g\|^2}$ divided by its U^0-orbital integral. This denominator is positive and finite, since $\|ug\| \geq \|u\|/\|g^{-1}\|$ and $L(u) \leq \|u\|$, with the polynomial volume growth proved above. It is constant on the left U^0-orbit. Multiply by $\mathbf{1}_{\{g : F^* J g F \in \mathcal{Q}\}}$ to define Φ. The Gram and the squared norm of each convolution are left $U^0 H$-invariant. Thus the u-integral in the Gram lower bound removes β. Every chart integral for fixed h_0 is at least c_2 times the corresponding cube integral over $H \times V_P$, so the same lower bound holds for its infimum and therefore for the Haar integral.

For the reverse estimate on this Haar integral, the unfolding bound applies with $K = 1$ and $R_* = P^2$. Indeed, for $s \in U$ and two Gram matrices from $\mathcal{Q}$,

$$F^* J g F = Q(s)^* F^* J s g F.$$

The Gram norm and inverse norm bounds give $\|Q(s)\| \le P^2$. Hence unfolding and Schur bound the Haar integral by $(2P^2)^{2n^2} P^{4n^2}(d/\delta) \int |\psi|^2$. With the coordinate evaluation cost this is bounded by the power P^{2000n^2}. We obtain

$$c_2 \int_H \int_{V_P} \left\| \int \psi(y) W(\gamma_Z(vh)^{-1} y) dy \right\|^2 dv_0\, dh \le P^{2000n^2}(d/\delta) \int |\psi|^2\, dy.$$

Now $A_2 \subset B(P^{20})$ and $\Omega \subset V_P$. For a bounded measurable ϕ, apply the section estimate to $\psi_j(y) = \mathbf{1}_{A_2}(y)\phi(y)\overline{W_j(y)}$ and sum over the d columns. The left Bessel actions of s_1 and t are unitary and independent of y, so

$$\mathbf{E}_{h,v} \left\| \int_{A_2} \phi(y) W(\gamma t y) W(ty)^* dy \right\|_{\mathrm{HS}}^2 \le \int |\phi|^2\, d\nu_2,$$
$$d\nu_2 = \frac{P^{2000n^2} d}{\delta c_2 \operatorname{vol}_{U^0}(\Omega)} \mathbf{1}_{A_2}(y) \|W(y)\|^2 dy.$$

Use the same unfolding bound with $\Phi = \mathbf{1}_{A_2}$. Now $K = T_2 \operatorname{vol}_{U^0}(\Omega)$ by the tube estimate, and $R_* = P^{40}$ since $\|s\| \le \|sy\|\|y^{-1}\|$ for two points in A_2. Thus $\int_{A_2} \|W(y)\|^2 dy \le T_2 \operatorname{vol}_{U^0}(\Omega)(2P^{40})^{2n^2} d$. The tube volume cancels. Using $P = M_2^{45n^2}$ and $M_2 \ge 2$ gives

$$\nu_2(G_1) \le M_2^{93600n^4+2n^2+45n} \frac{d^2}{\delta c_2} \le \left(\frac{\rho^{9n}}{\mathfrak{g}} \right)^{10^5 n^4} \frac{|I_v|^2}{\deg_v c_G(\mathfrak{g}/2)^{k-1}}.$$

In the classwise application below γ is read in its normalized frame at every infinite place. The radius is $\rho = b^{2n^2}$, which is at least the analytic cutoff radius b^{6n+2} since $\varrho \ge 10^5$. Use the local tubes and sampling laws above in the class unfolding with a large discriminant.

8.7 Bounds on the cut orbital sums

8.7.1 The sizes of cut blocks

At fixed (c, s, c') choose a representative of each $U_c^0 \gamma U_{c'}^0$ class, where U_c^0 fixes the whole frame. Only classes with both real arguments in the frame box contribute. Write j for the level, so $0 \le j \le n$, and set $d_j = \min(j, a)$. For the position $\kappa = h_c^{-1} \gamma h_{c'}$ in the adelic lattice chart let $I_j(\kappa)$ be the index of the $\widehat{\mathcal{O}}_E$-span of the adjoints of the first d_j rows of the frame pairing in the span of its first a rows. Write

$$D_j(\gamma) = |N_{E/\mathbb{Q}} \det Z_j(\gamma)| / N^{2f(j-a)_+},$$
$$x(\gamma) = \prod_{v|\infty} |\delta(\gamma)|_v,$$
$$B_j(\gamma) = \frac{\max(x(\gamma), N^f D_j(\gamma))}{I_j(\kappa) I_j(\kappa^{-1})},$$
$$M_0 = |J|\,|\mathrm{Cl}|\, N^{-kf},$$
$$\mathcal{D}_j = \begin{cases} \tau^{13-32j} \mathscr{D} & j \le a, \\ \mathscr{M} & j = n. \end{cases}$$

Here δ is the frame discriminant and a leading minor of size zero is one. The indices are finite and positive. For a stable set of level j, the cut amplitude norm of its block is bounded by the sum, over its retained near representatives in the block, of $v\, u_j(\gamma)$, where

$$u_j(\gamma) = \tau^2 \left(\frac{\mathcal{D}_j}{M_0 g^{2f d_j k} B_j(\gamma)^{k+1}} \right)^{1/2}. \tag{8.20}$$

Only compatible residues are included in this sum. This means that the reduced unipotent phase takes equal values at $\bar{u}, \bar{u}'$ whenever $\bar{u}s = s\bar{u}'$. The two-sided covariance multiplies the block by the ratio of these values and gives zero when they differ. This is a condition on the residue in addition to the character test on a rational stabilizer.

Integral frame translations within a class. Use $Z = Z(\gamma) = (h(f_i, \gamma f_h))_{i,h}$ in the rational frame, and let Y have the first j rows of Z followed by the corresponding rows of the identity. It is invertible since its determinant is the leading minor of order j. For $h \geq j$, the shared flag vectors in the level condition can be combined so that the coefficients on the moved frame from index j onward are exactly those of e_h in the coordinate space of the frame. Indeed these trailing coefficient vectors are unitriangular. The resulting vector belongs to the original frame span, so its first j pairings vanish, since these first j frame vectors are isotropic. Its moved-frame coefficients are therefore $Y^{-1}e_h$. In particular $\gamma\mathbf{f}Y^{-1}e_h$ has zero positive middle and dual-isotropic components. If $j = n$ and $h = a$, the dual component still vanishes and the z component is $-1/N$, by the equation $ZY^{-1}e_a = e_a$. Its positive middle component has squared length at most $1/N$ at every embedding. In fact, preservation of the Hermitian square gives that length as $1/N - N|\Phi_v((Y^{-1}e_a)_a)|^2$.

Let $C_{c,0}$ be the finite coordinate map from X_0 into the middle block of the class lattice of c. The intersection lattice $L_c = \{x \in X_0(E) : C_{c,0}x \in \widehat{\mathcal{O}}_E^{k+1}\}$ has integral Hermitian squares because the full class lattice is self-dual. A vector of L_c whose squared length is strictly less than one at every infinite place must be zero by the norm of its Hermitian square.

Here is the translation calculation in a form also used for the frame-value count. Let p, p_0 be terms of level j with the same Z and source class c, and let D_{mid} be the positive middle block of $(p - p_0)\mathbf{f}$. Choose isotropic row indices J_0 containing all indices less than $d_j = \min(a, j)$, and write $\xi = Z_{J_0,:}$. Suppose that in the finite chart $C_{c,0}D_{\mathrm{mid}} = g(A_0\xi + B_0Y)$ with A_0, B_0 integral. In each finite valuation ring the intersection of the image of $C_{c,0}$ with the integral middle module is saturated, hence admits an integral projection onto it. Apply this projection to the two coefficient matrices. By density of E in its finite adeles, choose a rational matrix L_0 close enough to the preimage of the projected A_0 that L_0 and all columns of $Q = g^{-1}(D_{\mathrm{mid}} - gL_0\xi)Y^{-1}$ belong to L_c. The two integrality conditions define an open neighbourhood of that preimage. For $h \geq j$ the column of Q vanishes, since both shared vectors have zero positive middle block and $\xi Y^{-1}e_h = 0$ by their dual-isotropic blocks. The omitted index below j can only be $h = a$ when $j = n$, and then $\xi Y^{-1}e_a = 0$ as well. Its column is integral with squared length at most $4/(g^2N) < 1$ at every place, by the two bounds just proved, and is zero as well. Since $Y_{h,:} = \xi_{h,:}$ for the selected rows $h < j$, we obtain $D_{\mathrm{mid}} = gL\xi$ with columns of L in L_c. All the other rows below the isotropic block in $(p - p_0)\mathbf{f}$ vanish since the frame pairings give its dual-isotropic and marked rows.

For the thin class $C = \Gamma_c^0\gamma\Gamma_{c'}^0$, both unipotent factors fix the frame, so every term has the same level, Z and residue. The right factor fixes the moved columns before applying the left factor. In the finite class chart, the change in each middle row of a moved frame is a linear combination of the dual-isotropic rows of the original moved frame, with coefficients in $g\widehat{\mathcal{O}}_E$. These are the rows of the isotropic frame pairing. Let R_a and R_j be the finite spans of the adjoints of the first a and d_j rows of Z, respectively. They have rank d_j. For $j \leq a$ the shared vectors give zero columns from j onward after multiplying the isotropic rows by Y^{-1}. At $j = n$ this rank follows from the invertible full pairing. In either case the first d_j rows are independent. The quotient gR_a/gR_j has order $I_j(\kappa)$. Colour the terms by the differences of their $k+1$ adjoint middle rows in this quotient, giving at most $I_j(\kappa)^{k+1}$ colours. In one colour choose a base term p_0. The translation calculation with $J_0 = \{0, \dots, d_j - 1\}$ and $B_0 = 0$ gives $D_{\mathrm{mid}} = gL(p)Z_{J_0,:}$. For equal $L(p)$ the moved frames therefore differ only in the isotropic rows.

The real volume integral. At a fixed real place write J_1 for the Hermitian involution defining the normalized form, and use $\mathbf{f}$ for the normalized frame matrix. The columns of $F_- = (1 - J_1)\mathbf{f}$ span the maximal negative space. Put $\Delta(y) = |\det(F_-^* y^{-1}\mathbf{f})|$. It is positive. A vector in $y^{-1}\operatorname{span}(\mathbf{f})$ has nonpositive

Hermitian square, whereas a nonzero vector whose projection to the negative space is zero has positive square. Moreover Δ is left invariant under U_v^0, which fixes $\mathbf{f}$.

Fix y of size at most θ, and let ζ be an isometry with the given normalized pairing matrix Z_v and level j normal form. In the following calculation ι includes the first j frame coordinates, $F_j = \mathbf{f}\iota$, and $z_j = \iota^* Z_v \iota$. Define

$$B_0 = (F_-^* y^{-1}\mathbf{f})^{-1} F_-^* y^{-1}, \qquad t = \zeta F_j z_j^{-1}, \qquad p = Z_v \iota z_j^{-1}.$$

Then $B_0\mathbf{f} = 1$, $p = \mathbf{f}^* J_1 t$ and $\iota^* p = 1$. Write $M = B_0 t z_j - \iota$. For the first j standard columns, $B_0\zeta\mathbf{f}$ and $1 + Mp^*$ have equal values. For every remaining index, the shared flag vector takes the form $\mathbf{f}x = \zeta\mathbf{f}x'$ with both coefficient vectors supported through that index and last coefficients of modulus one. Here $p^* x = t^* J_1 \zeta \mathbf{f} x' = z_j^{-*} F_j^* J_1 \mathbf{f} x' = 0$, because the first j vectors are isotropic when there are remaining indices. Thus $B_0\zeta\mathbf{f}x' = (1+Mp^*)x = x$. The two triangular change-of-columns matrices have determinant of modulus one. The identity $\det(1 + Mp^*) = \det(1 + p^* M)$ now gives

$$\frac{\Delta(\zeta^{-1} y)}{\Delta(y)} = |\det B_0 \zeta \mathbf{f}| = |\det z_j|\,|\det(p^* B_0 t)|.$$

At $j = 0$ this is one, with empty determinants.

Choose $u \in U_v$ with $\|uy\|_{\mathrm{op}} \le \theta$. Let T be θ times the first j square block of its frame matrix, and set

$$C = Tp^* B_0, \qquad D = Cy, \qquad Y_1 = J_1 t T^* - C^*, \qquad W_1 = \mathrm{pr}_{X_0} Y_1, \qquad \Theta = -CtT^*.$$

Since $\mathbf{f}^* Y_1 = 0$, we have $Y_1^* J_1 Y_1 = W_1^* W_1$. The rows of D lie in the negative subspace, so $CJ_1C^* = DJ_1D^* = -DD^*$. Also $Q = \theta^{-1}(uy)^{-1}F_j$ is a contraction and $QT = y^{-1}F_j$, whence $DQ = 1$. Thus $DD^* \succeq 1$. Expansion gives

$$\Theta + \Theta^* = W_1^* W_1 + DD^* - Tt^* J_1 t T^* \succeq 1 + W_1^* W_1,$$

as $t^* J_1 t = z_j^{-*} F_j^* J_1 F_j z_j^{-1} \preceq 0$. If $H = (\Theta + \Theta^*)/2$, conjugating Θ by $H^{-1/2}$ gives the identity plus i times a Hermitian matrix. Its determinant has modulus at least one, so $|\det\Theta| \ge \det H$. For an inclusion $e : \mathbf{C}^d \to \mathbf{C}^j$ of coordinate columns, $W_1 W_1^* \succeq (W_1 e)(W_1 e)^*$. The preceding determinant identity and $|\det T| = \theta^j$ give

$$\frac{\Delta(y)}{\Delta(\zeta^{-1}y)} \le \frac{(2\theta^2)^j}{|\det z_j|\det(1 + (W_1 e)^*(W_1 e))}.$$

Suppose the specified positive middle block of ζF_j has the form $V_0 + ARe^* z_j$, where V_0 is fixed and R is an invertible d-by-d matrix. Then $W_1 e = AR(e^* T^* e) + C_0$ for a fixed C_0, since J_1 is the identity on X_0. The affine map on each row has real Jacobian $|\det R|^2\theta^{2d}$. To integrate the determinant to the power $-(k+n)$, integrate one row w at a time. For the positive matrix H_0 from the remaining rows, $\det(H_0 + w^* w) = \det H_0(1 + wH_0^{-1}w^*)$. Changing variables by $H_0^{1/2}$ lowers the exponent of $\det H_0$ by one, with factor at most $\pi^d/d!$, since for an integer $q \ge d+1$

$$\int_{\mathbf{C}^d} (1 + \|w\|^2)^{-q}\, dw = \frac{\pi^d}{(q-1)\cdots(q-d)} \le \frac{\pi^d}{d!}.$$

Here the equality is the radial beta integral, with the empty product at $d = 0$. The successive exponents are $k+n, \dots, n+1$. We obtain a measurable majorant $K_v(A)$ for $(\Delta(y)/\Delta(\zeta^{-1}y))^{k+n}$ with

$$\int K_v(A)\, dA \le \left(\frac{(2\theta^2)^j}{|\det z_j|}\right)^{k+n} \left(\frac{\pi^d}{d!\theta^{2d}|\det R|^2}\right)^k. \tag{8.21}$$

The measure in each complex coordinate is ordinary Lebesgue measure, with unit disk area π. The case $j = 0$ uses the unit mass of the zero-dimensional parameter space.

The one-sided Schur estimate. For cut vectors put $\beta_c^0 = \mathbf{1}_{\mathrm{Box}}\beta_c$ and $V(y,z) = \prod_v(\Delta(y_v)/\Delta(z_v))^{k+n}$. For the class $\mathcal{C} = \Gamma_c^0\gamma\Gamma_{c'}^0$ define

$$F(y) = \sum_{p\in\mathcal{C}} \beta_{c'}^0(p^{-1}y)V(y,p^{-1}y), \qquad \sum_{p\in\mathcal{C}} S(c,c',p) = \int \beta_c^0(y)\|W^\flat(y)\|^2F(y)\,dy.$$

Here $S(c,c',p)$ denotes the integral with the single summand of F. All rational arguments in real integrals mean their normalized real images. The function F is invariant under left Γ_c^0, by reindexing p and the invariance of Δ. In the last integral use $\beta_c^0 = \mathbf{1}_{\mathrm{Box}}p_c\omega_c$ and disintegrate along U_∞^0. At a section point y, all factors except p_c and F are constant on the fibre. The cutoff p_c unfolds that fibre to $\Gamma_c^0\backslash U_\infty^0$. Right translation by any $s \in U_\infty^0$ preserves its quotient Haar, since a unipotent group is unimodular. Consequently, if for every boxed y and every $u \in U_\infty^0$ one has $\int_{\mathcal{X}} F(us(x)y)\,dx \le M_*$, where $s(x) \in U_\infty^0$ and $\mathrm{vol}\,\mathcal{X} \ge 1$, Tonelli bounds the last displayed integral by $M_*\int \beta_c(y)\|W^\flat(y)\|^2\,dy \le M_*\nu\,d_W$. This uses the full class partition on U_∞ and the cut squared norms $\int_{U_v\backslash G_v}\|W_v^\flat\|^2 \le |I_v|$.

Take $d = d_j$, $\mathcal{X} = \prod_v\{x_v \in \mathbf{C}^{k\times d} : \max|x_{v,ih}| < 1\}$, of volume π^{kdf}. In the v-factor put

$$R_v = \mathrm{diag}_{h<d}\,\frac{\pi g}{b\,s_{vh}}.$$

Let $s_v(x_v)$ fix z and have blocks on $X \oplus X_0 \oplus X^*$

$$\begin{pmatrix} 1 & Q_v & -Q_vQ_v^*/2 \\ 0 & 1 & -Q_v^* \\ 0 & 0 & 1 \end{pmatrix}, \qquad Q_v = \iota_{a,d}(x_vR_v)^*,$$

where $\iota_{a,d}$ includes the first d indices of the isotropic block. The matrix preserves the split form and fixes the frame. Its inverse adds Q_v^* times the dual-isotropic block to the positive middle block. For a colour with base term p_0, the isometry $s_v(x_v)^{-1}u_v^{-1}p_v$ has pairing Z_v and the same shared flag relations. On its first j frame columns the positive middle block is $V_0 + (x_v + \lambda_v(p))R_ve^*z_j$, with

$$V_0 = \mathrm{pr}_{X_0}(u_v^{-1}p_{0,v}F_j), \qquad \lambda_v(p) = (b/\pi)\Phi_v(L(p)), \qquad e = \iota_{j,d}.$$

Indeed the rational translation $gL(p)Z_{J_0,:}$ is rescaled on its frame columns by s_{vh}, and the corresponding row of Z_v has the additional factor s_{vi} at row i. Hence its coefficient after rescaling is $g\Phi_v(L(p))\,\mathrm{diag}(s_{vi}^{-1})$. This is $\lambda_v(p)R_v$. Also u_v^{-1} does not change the positive middle block of the difference, whose dual-isotropic coordinates are zero. The volume ratio at $us(x)y$ is therefore bounded by $\prod_v K_v(x_v+\lambda_v(p))$ with the same kernels throughout the colour.

Translate the integration variable for the summand indexed by p by $\lambda(p)$. At a fixed translated argument, at most one matrix $L(p)$ can occur with a nonzero boxed weight. In fact an intersection of two translated unit boxes would imply that every coordinate of $\Phi_v(L(p) - L(q))$ has modulus less than $2\pi/b$. In each column the squared length would be less than $64k/b^2 \le 64/b < 1$ at every place. The lattice norm argument above then gives equality of the matrices. For that fixed matrix all nonzero summands have the same original x, and their boxed inverse frame tuples differ in $\prod_v(y_v^{-1}X_v)^n$, of real dimension at most $2anf$. This follows from the isotropic row difference and the fact that $us(x)$ fixes X pointwise. The weighted frame count bounds the sum of their $\beta_{c'}^0$ weights by $b^{12an^3f} \le \tau$. Equal frames give relative elements in $\Gamma_{c'}^0$ by class freeness at the same residue, so this count uses the partition weights on one class lattice. Now integrate the common product kernel and sum over the colours. The required estimate for M_* is at most $I_j(\kappa)^{k+1}\tau\prod_v\int K_v$.

Write $\epsilon = j - d$, $D_v = |\Phi_v(\det Z_j(\gamma))|$, $P_0 = \prod_v D_v$, and $w_v = \prod_{h<d}|\kappa_v|^{a-h}$. We have $\epsilon \in \{0,1\}$, $w_v \ge 1$ and $P_0 \ge 1$ because the nonzero rational leading minor is an algebraic integer. The normalized local block satisfies

$$|\det z_j| = w_v^2N^{-j}D_v, \qquad |\det R_v| = (\pi g/b)^dN^{d/2}/w_v, \qquad \sqrt{D_j(\gamma)} = P_0/N^{\epsilon f}.$$

Thus (8.21) is a product of the factor $((2\theta^2)^j/|\det z_j|)^n$ and the k-th power of

$$\frac{2^j\theta^{2\epsilon}}{\pi^d d!}\frac{b^{2d}N^{\epsilon}}{g^{2d}D_v}.$$

The first fraction is at most $(g/b^6)^d$. If $\epsilon = 0$, use $2 \le \pi$ and $g \ge b^{19}$. If $\epsilon = 1$, then $d = a$, and $2\theta^2 \le b^{12a+17} \le b^{13a} \le (g/b^6)^a$. The n-th power factors have product at most

$$\frac{N^{\epsilon f}}{P_0}(2\theta^2)^{jnf}N^{(jn-\epsilon)f} \le \frac{\tau}{\sqrt{D_j(\gamma)}}.$$

For the last inequality the leftover exponent of b is at most $(72n+5)n^2 f$ by $N \le b^{60n}$ and $j \le n$, below $kf/4$ since $k \ge 1000(2n)^9$. Together with packing and disintegration this proves

$$\sum_{p\in C} S(c,c',p) \le \nu\, d_W \tau^2 (b^4 g)^{-d_j kf}\left(\frac{I_j(\kappa)}{\sqrt{D_j(\gamma)}}\right)^{k+1}. \tag{8.22}$$

All the estimates also hold for the zero level with the empty minor and the zero-dimensional shift.

The cut amplitude in the small discriminant range. Both arguments of a nonzero cut integrand with nonzero cutoff belong to Box. For a single term p, insert the reciprocal weights $V(y, p^{-1}y)$ and $V(p^{-1}y, y)$ in Cauchy–Schwarz and change variables by p^{-1} in the second integral. This bounds the square of its absolute cut amplitude by $S(c,c',p)S(c',c,p^{-1})$. Apply Cauchy–Schwarz to the sum over p, and then use (8.22) for the two orientations. Inversion sends Z to Z^*, so it preserves j, D_j, and the finite position becomes κ^{-1}. If $x(\gamma) \le N^f D_j(\gamma)$, divide the resulting squared bound by $\nu^2 u_j(\gamma)^2$ from (8.20). With $M_0 = |J|\,|\mathrm{Cl}|\,N^{-kf}$ and $\mathscr{D} = |J|\,|\mathrm{Cl}|d_W$, the ratio is at most

$$\frac{d_W \mathscr{D} N^f b^{-8d_j kf}}{\mathcal{D}_j} \le 1.$$

Use $d_W \le \tau^8$, $N^f \le \tau$, and $b^{-8d_j kf} = \tau^{-32d_j}$. For $j \le a$ take $\mathcal{D}_j = \tau^{13-32j}\mathscr{D}$, and at $j = n$ take $\mathcal{D}_n = \mathscr{M} \ge \tau^{44}\mathscr{D}$. If the cut absolute amplitude in this range is nonzero at a term, there is a point witnessing that the term is near. Nearness is preserved by both frame-fixing class-unipotent factors, so only near representatives enter the bound.

The real cell mass on a top-level class. Let γ belong to the large range $x(\gamma) > N^f D_j(\gamma)$ of the fixed block. Its rational frame discriminant δ is nonzero. Put $a = \varrho$, $n = a+1$, $B = b^{2n^2}$, $s_0 = 10^5 n^4$. Write $x = \prod_v |\Phi_v(\delta)|$ and $L_v = \prod_{i=0}^{a} \ell_{vi}$, $L = \prod_v L_v$, where $\ell_{vi} = |\kappa_v|^{2(a-i)}/N$. The two frame-row indices at level n are one. Put $u = u_n(\gamma)$ from (8.20) in the large range $B_n(\gamma) = x$.

To see the sign of the discriminant, multiply the rational Gram matrix Z by $y = \mathrm{adj}(Z)e_n$, where e_n is the last coordinate vector. This gives $(\det Z)e_n$, and the last entry of y is the leading minor D_a. In the isometric image of the frame combination Fy, all dual isotropic coordinates therefore vanish, and its terminal coordinate is $-\det Z/N$. Equality of its norms in the two frames gives $\Phi_v(\delta) = N\|(\gamma F y)_{X_0,v}\|^2$. The frame products on the retained finite coset are integral. Thus x is the absolute norm of a nonzero algebraic integer from the totally real field, and $x \ge 1$. Scaling each row and column of Z to the real frame gives $\mathfrak{g}_v = L_v^2|\Phi_v(\delta)|$. In this calculation the last scale satisfies $\ell_{va} = 1/N$, which accounts for the N^2 on the rational minor D_a.

For the compact-cell estimates take γ to be a retained near representative. Let M_v denote the upper bound for the mass in the local contraction at γ_v, formal degree $\deg_v$, and $d_v = |I_v|$. It is

$$M_v = \left(\frac{B^{9n}}{\mathfrak{g}_v}\right)^{s_0}\frac{d_v^2}{\deg_v c_G(\mathfrak{g}_v/2)^{k-1}}.$$

The product of the d_v is d_B. Use the packet normalization inequality for $\Pi, V, F_{\mathrm{sc}}, \nu, \mathscr{M}$. Substitution of $\prod_v \mathfrak{g}_v = L^2 x$ gives

$$\frac{4\prod_v M_v}{\nu^2 u^2} \le \frac{4(1+b^{-9})}{\tau^4}\, x^{2-s_0} B^{9nfs_0}\, \frac{F_{\mathrm{sc}}^2}{L^{2(k-1+s_0)}}\, \frac{2^{(k-1)f} M_0 g^{2fak}}{\Pi V^n}.$$

Here $M_0 = |J| h N^{-kf}$ and $\Pi = |J| g^{fa(2k+2a+1)} h\, c_{\mathrm{quad}}$. The last fraction is

$$g^{-fa(2a+1)} \left(\frac{(k-1)!\, \mathrm{rd}_E^{k+n-1/2}}{2^{n+1}\pi^k} \right)^f.$$

It is at most $b^{(k+n)f/1000}$, by $(200 f k\, \mathrm{rd}_E)^{1000} \le b$. The quotient containing F_{sc} is the product of $\ell_{vi}^{2(2(a-i)+1-s_0)}$. All these exponents are negative and $\ell_{vi} \ge N^{-1}$, so the quotient is at most N^{2nfs_0}. By $N \le b^{60n}$ and $x \ge 1$, the exponent of b^f contributed by these factors and B^{9nfs_0} is at most $(120n^2 + 18n^3)s_0$. This is at most $k/2$ for $n \ge 10$ and $k \ge 1000(2n)^9$, as on the chosen scales. Since $\tau^4 = b^{kf}$, the displayed ratio is at most one.

For the cutoff on the same representative, the two boxed frames give real pairings of modulus at most B^2 by Cauchy–Schwarz. Passing to the rational frame multiplies an entry by at most N, since $|\kappa_v| \ge 1$. Thus each entry is at most $NB^2 \le b^{5n^2}$ on the chosen scales. The determinant expansion, $n \le b$, and $0 \le \Phi_v(\delta) \le |\Phi_v(\det Z)|^2$ give $x \le b^{11n^3 f}$. Using $\mathscr{M} \ge \tau^{44} |J| h d_B$ in the same class size now yields

$$u^2 \ge \frac{\tau^{48} N^{kf} d_B}{g^{2fak} x^{k+1}} \ge b^{-2k^3 f}.$$

Indeed $g \le b^{22}$ and $44ak + 11n^3(k+1) \le 2k^3$. Consequently $\sqrt{\prod_v M_v} \le \nu u/2$ and $\nu b^{-k^3 f}/2 \le \nu u/2$ for the two estimates on the same class.

The cut amplitude in the large discriminant range. For $x(\gamma) > N^f D_j(\gamma)$ one has $j = n$, since the discriminant vanishes at other levels. Both row indices at n are one, and $B_n(\gamma) = x(\gamma)$. The discriminant computation in the cell mass comparison shows that every normalized local discriminant $\mathfrak{g}_v$ is positive and $x(\gamma) \ge 1$. The sign and integrality arguments there apply to all terms in the range, before using nearness. In particular the full normalized pairing and hence Z are invertible.

Fix a stable set of terms in this range and a block (c, s, c'). Stability holds for this restriction since the two leading determinants and the discriminant are unchanged by the upper unitriangular frame actions on the left and on the right. Choose one representative ρ of each double class for Γ_c on the left and $\Gamma_{c'}^0$ on the right. Choose it among the fixed representatives for the smaller left group Γ_c^0, preferring a near one whenever it exists. The larger left group acts freely on these right cosets. To see this, if $u \in U$ fixes $\gamma \mathbf{f}$, it fixes the first a columns of $\gamma \mathbf{f} Z^{-1}$, whose dual-isotropic block is the identity. In $X \oplus (X_0 \oplus Ez) \oplus X^*$ an element of U is upper block triangular, acts trivially on $X^\perp/X$, and has contragredient diagonal blocks on X, X^*. Fixing these columns forces both those diagonal blocks to be one and the block from X^* to the middle to vanish. Preservation of the form then makes the middle-to-X block vanish. Fixing the columns now also makes the X^*-to-X block vanish. This proves $\Gamma_c \cap \gamma \Gamma_{c'}^0 \gamma^{-1} = \{1\}$.

On the pair (γ, y) let $(u, v) \in \Gamma_c \times \Gamma_{c'}$ act by $(u\gamma v^{-1}, uy)$. Counting measure in the terms and real Haar in y are preserved. The weight

$$\phi(\gamma, y) = \mathbf{1}_{\gamma \text{ chosen}}\, \omega_{c'}(\gamma^{-1} y)$$

sums to one along the action. For each left $\Gamma_{c'}^0$-coset of v there is exactly one pair mapping γ to the chosen representative, by the preceding freeness. Its weight is $\omega_{c'}(v\gamma^{-1} y)$, constant on that coset, and their sum is one. The original cutoff product also sums to one under the action. Both kernels are invariant by their class-level Bessel characters. Inserting the sum of either weight into the other weighted integral, using Tonelli for absolute values, and changing variables by the action proves equality of the two integrals for each invariant absolutely integrable integrand. This applies to the full and cut kernels and to the norm of

their difference. In particular the sum of absolute errors with the new weights is at most $\nu\, b^{-k^3 f}/2$. Here the absolute error bound is the integral of the majorant in the full-to-cut estimate, before the triangle inequality on the individual blocks.

If a chosen representative has a nonzero cut integrand, its right argument is boxed by the property of $\omega_{c'}$, and the left argument is in the cut. The left partition gives a $u \in \Gamma_c$ moving that left argument into Box. Thus the larger double class contains a near term $u\rho$, and its smaller double class has a near representative. By the choice rule the selected representative is near.

For a near selected γ apply the real cell contraction at every place with frame bound B and $\mathfrak{g}_v > 0$. It gives compact sets A_v, probability laws on the permitted translates $t_v \in U_v^0 H_{0,v}$, and positive measures of masses at most M_v from the cell mass comparison. For a product translate change variables to the left argument $\gamma t y$. The weight is now $\omega_{c'}(ty) = \omega_{c'}(y)$. If the cut integrand is nonzero, y_v has inverse frame of size at most B and $\operatorname{size}(\gamma_v t_v y_v) \le \theta \le B$, hence $y_v \in A_v$. The cut integral is therefore the integral on $A = \prod_v A_v$. Tensoring the local contractions and applying them to $\omega_{c'}|_A$ shows that the average squared norm of the full-vector integral on A is at most $\prod_v M_v$, since $0 \le \omega_{c'} \le 1$. To tensor, first use finite sums of simple tensors. Approximate the bounded weight in L^2 for the product of the local sums of cell measure and restricted Haar. This gives approximation in both product measures separately. Convergence in restricted Haar identifies the displayed integral for each fixed translate by continuity of the full vectors on the compact domain. The product contraction gives convergence in the averaged L^2 norm and hence the claimed bound. Choose a translate satisfying the support assertion and this squared norm bound. The change from the full vector to the cut on A costs at most the absolute error for this representative under ϕ, since the latter integrates the norm on the whole group. By the cell mass comparison, the full integral on A has norm at most $\nu\, u_n(\gamma)/2$ and the error at most $\nu\, b^{-k^3 f}/2 \le \nu\, u_n(\gamma)/2$. This proves the large-range bound by summing over the selected near representatives, which are a subset of the representatives for the smaller double classes in (8.20). Together with the small-range bound and the residue test, it gives the cut block estimate on every stable term set of a fixed level.

A finite-module count for inverse indices. Write $\widehat{O} = \widehat{\mathcal{O}}_E$, and for an integral finite idele d write $[d] = |\widehat{O}/d\widehat{O}|$. For a finite adelic lattice M and a submodule K put $I_K(M) = [M : M \cap K]$ when the index is finite. Let Y_α be integral matrices of size $\sigma \times (r_0 + s)$, and let m be an integral matrix whose r_0 columns span a primitive direct summand of $\widehat{O}^{r_0+s}$. Suppose K contains $d\widehat{O}^\sigma$. For each α let u_α be a nonnegative integer divisible by $I_K(Y_\alpha m \widehat{O}^{r_0})$. If at most x indices α have any one matrix of column classes modulo K, then for $p_0 \ge r_0$

$$\sum_\alpha u_\alpha^{-p_0} \le [d]^{\sigma(\sigma+s)} x. \tag{8.23}$$

The weight at $u_\alpha = 0$ is zero.

Complete m to a basis of $\widehat{O}^{r_0+s}$ and express the columns in this basis. Put $Q = \widehat{O}^\sigma/(\widehat{O}^\sigma \cap K)$, of order at most $[d]^\sigma$. Group the matrices by the submodule $A \subset Q$ generated by their first r_0 column classes and by the ordered last s column classes. Each submodule of Q can be generated by σ elements. Indeed, over each discrete valuation ring in the finite support of Q, lift the submodule to the free module of rank σ, use elementary divisors, and project a basis back to the submodule. Combining the local generators gives the desired σ generators. There are therefore at most $|Q|^{\sigma+s}$ groups. In a group the index $I_K(Y_\alpha m \widehat{O}^{r_0})$ is $|A|$. There are at most $|A|^{r_0}$ ordered first-column classes. Each occurs at most x times, and its summands have weight at most $|A|^{-p_0}$. This proves (8.23).

The sum of inverse discriminants on the marked class lattices. Put $a = \varrho$, $n = a + 1$, $k = \dim_E X_0$, and $f = [E : \mathbf{Q}]/2$. Use one complex embedding at each of the f infinite places and the usual complex Euclidean norm in the orthonormal coordinates of $X_0 = E^k$. For a finite adelic lattice Λ in $X_0 \otimes_E \mathbb{A}_{E,f}$ and an E-subspace V of X_0, let $d_\Lambda(V)$ be $|\det(m_V^* m_V)|_f^{-1/2}$, where m_V is a basis of $\Lambda \cap V_{\mathbb{A}_f}$, $*$ is conjugate transpose, and the finite absolute value is the product of the normalized local absolute values. A change

of basis over $\widehat{O}$ leaves this quantity unchanged. With $r_0 = \dim_E V$, the covolume of $V \cap \Lambda$ in its real span in the f complex embeddings is

$$(2^{-f}\sqrt{|D_E|})^{r_0} d_\Lambda(V).$$

To see this, write $m_V = xA$ with x an E-basis of V and A invertible over the finite adeles. Finite elementary divisors multiply the covolume of $\mathcal{O}_E^{r_0}$ in the coefficient space by $|\det A|_f^{-1}$. At a complex embedding the real volume multiplier of x is $\det(x_v^* x_v)$. The product formula gives $\prod_v \det(x_v^* x_v) = |\det(x^* x)|_f^{-1/2}$. The normalized discriminant is their product with $|\det A|_f^{-1}$. In particular passing to a sublattice of finite index I in the same subspace multiplies the discriminant by I.

First let $\Lambda = u\widehat{O}^k$ with u a finite adelic isometry of X_0. This lattice is self-dual. For the primitive lattice $M = \Lambda \cap V_{\mathbb{A}_f}$, projection from Λ to $V_{\mathbb{A}_f}$ has image $M^\vee$, where the dual is taken for the Hermitian form. Indeed a functional on M extends to Λ by a finite adelic integral complement, and self-duality realizes that functional by a member of Λ. The projection has kernel $\Lambda \cap V_{\mathbb{A}_f}^\perp$. It follows that the index of $(\Lambda \cap V_{\mathbb{A}_f}) \oplus (\Lambda \cap V_{\mathbb{A}_f}^\perp)$ in Λ is $[M^\vee : M] = d_\Lambda(V)^2$. Interchanging $V, V^\perp$ gives $d_\Lambda(V) = d_\Lambda(V^\perp)$.

Let $s = \operatorname{codim}_E V$. In the real dimension $2fs$ of $V^\perp$, denote the Euclidean successive minima of $\Lambda \cap V^\perp$ by λ_h. If $s > 0$, choose independent E-vectors $x_1, \dots, x_s$ there with $\|x_i\| \le \lambda_{2f(i-1)+1}$, where $\|\cdot\|$ is the Euclidean norm summed over the complex embeddings. At step i the preceding E-span has real dimension $2f(i-1)$, so one of the $2f(i-1)+1$ independent lattice vectors at that minimum lies outside it. The Euclidean bound for the product of the minima follows by stretching the unit ball along an orthonormal flag with scales λ_h. Choose the flag so that its $(h-1)$-space contains all lattice vectors of norm strictly below λ_h. A lattice point in the open stretched ball with last nonzero coordinate h would have norm $< \lambda_h$. Half of the stretched ball thus injects, up to sets of measure zero, into the quotient by the lattice. Its volume is at most the lattice covolume. Writing $\mathrm{rd}_E = |D_E|^{1/(2f)}$ and using the cube of side $2/\sqrt{2fs}$ in the unit ball gives

$$\left(\prod_{i=1}^{s} \|x_i\|\right)^{2f} \le \prod_{h=1}^{2fs} \lambda_h \le (8fs\,\mathrm{rd}_E)^{fs} d_\Lambda(V).$$

For $s = 0$ the empty product has value one.

Every nonzero $x \in X_0 \cap \Lambda$ has $\|x\| \ge 1$, since its squared Hermitian length is a nonzero totally positive algebraic integer. The same bound for differences gives unit spacing. Packing half-unit balls yields at most $(1+2R)^{2fk}$ points of norm at most R. Integration of this count for $R \ge 1$ gives

$$\sum_{0 \ne x \in X_0 \cap \Lambda} \|x\|^{-2f(k+1)} \le 3^{2fk}(k+1) \le 5^{2f(k+1)}.$$

The ordered tuple chosen for V determines $V^\perp$ and hence V. Raising the minima bound to $k+1$ and summing the products of inverse vector lengths over all tuples now gives

$$\sum_{\operatorname{codim}_E V = s} d_{u\widehat{O}^k}(V)^{-k-1} \le (200fk\,\mathrm{rd}_E)^{f(k+1)s}. \tag{8.24}$$

Here $s \le k$. For $s > k$ the sum is empty.

Now use the actual class lattices. Let $p = \pi$ be the integral finite idele of the marked class basis. Then $p\bar{p} = N$ and $[p] = N^f$. Let o be the marked coordinate of X_0, K^0 its definite unitary standard lattice stabilizer, and K_H the column level. For the class represented by h_c the definite part is

$$\Lambda_c^+ = h_c\{v \in \widehat{O}^k : v_o \in p\widehat{O}\}.$$

Indeed in the marked basis the new vector is $e_o/\bar{p} + z/N$. On X_0 its coefficient is $\bar{p}\,v_o$, and the coefficient of z is $-v_o/p$. Thus $d_c(V) = d_{\Lambda_c^+}(V)$. The preimage of the full ambient lattice stabilizer in K^0 is $K_p = \{h \in K^0 : e_o^* h \equiv e_o^* \bmod p\}$. In the marked coordinates the differences $h - I$ must be integral after dividing row o by p and column o by $\bar{p}$. The row congruence for h^{-1} follows from the row congruence for h. Conjugate

transpose therefore gives the required column congruence for h. At each finite place at least one of $p, \bar{p}$ is a unit, so both divisions preserve the required integrality. Conversely the marked chart recovers the row congruence by multiplying back the column scale $\bar{p}$. The column level satisfies $K_H \subset K_p$. The admissible datum gives

$$[K^0 : K_p] \geq (N^k/b)^f.$$

Fix a representative u at level K^0, and write y_ϕ for representatives of K^0/K_p. On $u\widehat{O}^k$ the row defining the marked lattice of uy_ϕ is $v \mapsto e_o^* y_\phi^{-1} u^{-1} v \bmod p$. These rows are distinct modulo p in the u-basis. Equality of the rows for ϕ, ϕ', followed by right multiplication by $y_{\phi'}$, would give $y_\phi^{-1} y_{\phi'} \in K_p$. For a fixed V of codimension s, its primitive lattice in $u\widehat{O}^k$ has $k - s$ basis vectors. The index $I_\phi(V)$ of the marked lattice in this primitive lattice is the order of the image of the row on it modulo p. Apply (8.23) with one output coordinate, modulus p, $x = 1$, and $p_0 = k + 1$. It gives $\sum_\phi I_\phi(V)^{-k-1} \leq N^{f(s+1)}$. Since $d^+_{uy_\phi}(V) = d_{u\widehat{O}^k}(V) I_\phi(V)$, (8.24) and the index bound give

$$\frac{1}{[K^0 : K_p]} \sum_\phi \sum_{\operatorname{codim}_E V = s} d^+_{uy_\phi}(V)^{-k-1} \leq N^{-kf} b^f N^{f(s+1)} (200 f k \operatorname{rd}_E)^{f(k+1)s}.$$

This average is also the uniform average over the column-level classes above the class of u. In fact the finite rational stabilizer $u^{-1}U(X_0)(E^+)u \cap K^0$ acts on K^0/K_H with trivial stabilizers by freeness of the column level. The sum over subspaces is invariant under rational transport. Thus all orbits have the same size, and passing from K^0/K_H to K^0/K_p has constant multiplicity $[K_p : K_H]$. This proves the assertion about the average, including its class multiplicity.

Use $N \leq b^{60n}$, $(200 f k \operatorname{rd}_E)^{1000} \leq b$, and $k \geq 1000(2n)^9$. The extra exponent of b divided by f is at most $1 + 60n(s+1) + (k+1)s/1000 \leq k(s+1)/4$. With $\tau = b^{kf/4}$, summing over the column-level classes gives

$$\sum_c \sum_{\operatorname{codim}_E V = s} d_c(V)^{-k-1} \leq |\mathrm{Cl}|\, N^{-kf} \tau^{s+1}. \tag{8.25}$$

The inverse frame-value sum on a fixed label. Fix a source class c, a level j with $0 \leq j \leq a + 1$, and the J-coordinate t in the fixed set bijection $G \simeq G/J \times J$. Take any set $\mathcal{A}$ of pairs (c', γ) of this label represented by near terms of level j chosen in their U_c^0-$U_{c'}^0$ double cosets. Write $\mathsf{F} = (e_0, \dots, e_{a-1}, z)$ as a column matrix and $Z = \mathsf{F}^* h \gamma \mathsf{F}$. For a fixed $V \subset X_0$ of codimension at most s, let

$$M_\gamma(V) = \{((\gamma\mathsf{F})^* h) v : v \in \Lambda_c^+ \cap V_{\mathbb{A}_f}\}, \qquad R_\gamma = \langle Z^*_{0,:}, \dots, Z^*_{a-1,:} \rangle_{\widehat{O}}.$$

For any nonnegative integer u_γ on each pair in $\mathcal{A}$ divisible by $I_{gR_\gamma}(M_\gamma(V))$, with the infinite index recorded as 0, we prove

$$\sum_{(c',\gamma) \in \mathcal{A}} u_\gamma^{-k} \leq b^{2000 n^6 (s+1) f}. \tag{8.26}$$

The zero convention for the weight is as in (8.23).

In the class chart the paired frames are integral and the class basis preserves the frame vectors. Consequently Z has algebraic integer entries. At a near witness y_v, both y_v and $\gamma_v^{-1} y_v$ are in the real frame box, with $\|y_v^{-1} f_h\|$ and $\|(\gamma_v^{-1} y_v)^{-1} f_h\| \leq B = b^{2n^2}$. The signed frame scales are $s_{vh} = \kappa_v^{a-h}/\sqrt{N}$ with $1 \leq |\kappa_v| \leq b^{23}$. After undoing the two scales the Gram pairing at v is bounded by $NB^2 \leq b^{5n^2}$. Two distinct integral Z's differ by an algebraic integer with absolute value at least one at some complex embedding and coordinate. Packing in the product of the entry disks thus bounds the number of Z's by $(1 + 2b^{5n^2})^{2fn^2} \leq b^{12fn^4}$.

Fix one such Z. Let P be its first a rows and Y its first j rows padded with the remaining rows of the identity. Put $\delta = \det Y \neq 0$, $M = g_0^{n^4}$, and

$$d = Mg\delta\bar{\delta}, \qquad K = gR_\gamma + gY^* \widehat{O}^n.$$

These objects are constant in the fixed Z family. Multiplication by the adjugate of Y^* shows $d\widehat{O}^n \subset gY^*\widehat{O}^n$. Also $I_K(M_\gamma(V))$ divides u_γ, because $gR_\gamma \subset K$ and the index modulo K is finite. The weight is zero in the infinite-index case of the original gR_γ index. The Leibniz bound at the real places gives $|\delta|_v^2 \le b^{11n^3}$. Using $g \le b^{23}$, $g_0 \le b^{11}$, $n \ge 10^5$, and the absolute norm of a principal integral idele gives

$$[d] = \prod_v |\Phi_v(d)|^2 \le b^{24fn^4}. \tag{8.27}$$

For the application of (8.23), use the adjoint of the middle part of the moved frame in the class chart as the matrix Y_γ, of size $n \times (k+1)$. Use the chart's integral unimodular middle Gram matrix times an integral basis of $\Lambda_c^+ \cap V_{\mathbb{A}_f}$, in that chart, as the matrix m. Its columns generate a primitive direct summand. Their ambient codimension is at most $s+1$, and $Y_\gamma m$ generates precisely $M_\gamma(V)$.

Consider a nonempty set of the matrices Y_γ whose columns agree modulo K. Take a fixed comparison term in the set. The integral translation calculation with $\xi = P$ writes the change needed in the positive middle rows as gLP, with the columns of L in the definite class lattice. Its congruence hypothesis is agreement modulo $gR_\gamma + gY^*\widehat{O}^n$. Extend L by zero on the z-row. If h_M is the middle Gram matrix and $\omega = -\zeta_3$ has trace one, the rational unipotent matrix

$$\begin{pmatrix} I & -gL^*h_M & -\omega g^2 L^* h_M L \\ 0 & I & gL \\ 0 & 0 & I \end{pmatrix}$$

is unitary and fixes the augmented frame. The first property follows from the trace of ω, and the second from the orthogonality of L to z. Self-duality of the class lattice makes the three off-diagonal blocks integral multiples of g in its chart. Thus the matrix gives a left translate in U_c^0. It adds gLP in the middle and leaves the dual rows fixed. The marked row is fixed by the pairings, so after this translate the moved frame agrees below the e-rows with the comparison frame. Both Z and the vectors $\gamma \mathsf{F}\,\mathrm{adj}(Y)e_h$ for $h \ge j$ are unchanged by any such translate, since the latter vectors belong to the standard augmented frame span.

We bound the remaining a-by-n integral e-rows by grouping their residues modulo d, after the middle shears. To prove injectivity of these residues, we first bound the shared vectors. For $h \ge j$ write $\gamma \mathsf{F}\,\mathrm{adj}(Y)e_h = \mathsf{F}w_h$, where $w_h \in E^n$. At the witness y_v we have

$$y_v^{-1}\mathsf{F}(\Phi_v(w_h)/s_v) = (\gamma_v^{-1}y_v)^{-1}\mathsf{F}(\Phi_v(\mathrm{adj}(Y)e_h)/s_v),$$

where division by s_v is entry by entry. The right side has length at most $nB\sqrt{N}\,b^{5n^3}$ by the adjugate bound and the box. Choose a left unipotent v_0 with $\|v_0 y_v\|_{\mathrm{op}} \le b^{7n}$, using the cutoff radius b^{6n+2}. Its action on the augmented frame is upper unitriangular with entries of size at most $b^{7n}B$, since $\|v_0 f_h\| \le \|v_0 y_v\|_{\mathrm{op}}\|y_v^{-1}f_h\|$. The inverse on this span has entries bounded by $(n-1)!(b^{7n}B)^{n-1} \le b^{3n^3}$. Multiplying by $v_0 y_v$ and this inverse, then undoing the scales $|s_{vh}| \le b^{23n}$, gives $|\Phi_v((w_h)_e)| \le b^{10n^3}$ in every e-coordinate. In particular twice this bound is strictly less than M, since $g_0 \ge b^{10}$.

Suppose two aligned frames of the same e-row residue differ in their e-rows by dX_1 with X_1 integral. For $h \ge j$ the difference of the two shared vectors in the e-rows is $dX_1\,\mathrm{adj}(Y)e_h$. Its quotient by M is integral and has absolute value strictly less than one at every infinite place. Hence $X_1\,\mathrm{adj}(Y)e_h = 0$. The frames have equal internal Gram matrix and agree below the e-rows, including the e^*-block P. As $d = \bar{d}$, the equality of the internal Gram matrices gives $X_1^*P = -P^*X_1$. Let α consist of the first a columns of $\mathrm{adj}(Y)$ and put $A_1 = X_1\alpha$. For $j \le a$, the vanishing of the other columns and $\mathrm{adj}(Y)Y = \delta I$ give $A_1P = \delta X_1$. For $j = a+1$, let β be the last column of the adjugate. Then $P\beta = 0$, so $P^*X_1\beta = 0$ by the Gram relation. Since P has full row rank, $X_1\beta = 0$. This gives $A_1P = \delta X_1$ also at the top level. The integral matrix

$$S_1 = \bar{\delta}A_1 - \delta A_1^* - (P\alpha)^*A_1$$

is skew Hermitian by the Gram relation. Using $A_1P = \delta X_1$ and $A_1^*P = -(P\alpha)^*X_1$ gives $S_1P = \delta\bar{\delta}X_1$. The unipotent shear adding MgS_1 times the e^*-rows to the e-rows therefore matches the two frames. It fixes the standard frame, and is integral and congruent to one modulo g on the class lattice, since the class chart is the identity on these hyperbolic pairs. This puts it in U_c^0. For the resulting equal frames, the relative position $\kappa^{-1}\kappa'$ of the two integral positions fixes the standard frame. Its residue is in J, so the G/J-coordinates of the two residues agree. Their fixed t gives equal exact residues, hence $\kappa^{-1}\kappa' \in K(g)$. The relative rational isometry η fixing the frame has upper triangular blocks in $X \oplus (X_0 \oplus Ez) \oplus X^\vee$. On its middle quotient it acts on X_0 as an isometry h_η and fixes z. In the finite class chart of $\kappa^{-1}\kappa'$, the diagonal middle block is the marked lift of $h_{c'}^{-1}h_\eta h_{c''}$, where c', c'' are the two target classes. Since the position is integral and congruent to one modulo g, that lift belongs to the principal level, so $h_{c'}^{-1}h_\eta h_{c''} \in K_H$. The class condition gives $c' = c''$, and freeness of the column level gives $h_\eta = 1$. The remaining rational relative term belongs to $U_{c'}^0$. The original pairs therefore have the same chosen double-coset representative.

There are at most $[d]^{an}$ e-row residue matrices. Thus in the application of (8.23) we can take $x = [d]^{an}$, $\sigma = n$, $p_0 = k$, and complementary rank at most $s+1$. Summing over Z, the exponent of b is at most

$$f\{12n^4 + 24n^4[n(n+s+1) + an]\} \le 2000n^6(s+1)f.$$

This proves (8.26).

8.8 Summing over pairs of frame terms

8.8.1 Pair weights for the congruent frame terms

For two near terms γ, γ' of level j in the same block let L be the intersection of their links. Each link is a maximal isotropic for $h_M \oplus (-h_M)$ on M^2 with $M = X_0 \oplus Ez$. Put $i = k + 1 - \dim_E L$. Let V be the first projection of $L \cap X_0^2$ and put $l = k - \dim_E V$. The first projection is injective on $L \cap X_0^2$ by positivity of X_0. The terminal pairings on either link have rank at most one in the degenerate levels and at most two in the top level. Thus $l \le i + (j-a)_+$, and $i = 0$ means that the two links coincide.

Let $d_c(V)$ be the Hermitian discriminant of the intersection of V with the definite lattice of c. If m is an adelic basis of this primitive lattice in the coordinates of X_0, set $d_c(V) = |\det(m^*m)|_{\mathbb{A}_f}^{-1/2}$. Define $F_\gamma(V)$ as the index $[M_\gamma : M_\gamma \cap R_\gamma]$, where M_γ consists of the frame values of the integral vectors of V and R_γ is the span of the adjoints of the first a rows of the position's frame pairing. Define d_γ as the index of the lattice generated by the two integral isotropic frames in its saturation. Use the corresponding definitions for γ'. These indices are taken in the lattices of c, c' in the given term position. The two ordered weights are $\omega(\gamma, \gamma')$ and its transpose, with

$$\omega(\gamma, \gamma') = \frac{d_c(V)^{-k-1}(g^{2fd_j}F_\gamma(V))^{-k}d_{\gamma'}^{-k}}{M_0(2\tau^2)^l}.$$

Use zero weight off the congruent near pairs of level j.

Summing the ordered pair weights. For a congruent pair in the same block (c, s, c') let V be the definite projection of the common link, as in the definition of the pair weight. Write $l = \operatorname{codim}_E V$. For either link and $y \in V$, its lift in $X^\perp \cap (\gamma X)^\perp$ can be written $y + x$, with $x \in X = \langle e_0, \dots, e_{a-1}\rangle$, and its pullback lies in $y' + X$ for $y' \in X_0$. Both projections have zero z-coordinate. Thus the lift is orthogonal to $\gamma\mathsf{F}$, including its terminal vector. It follows that $M_\gamma(V)$ takes values in the adelic span of R_γ, so $F_\gamma(V) = I_{R_\gamma}(M_\gamma(V))$ is finite and positive. The rank of R_γ is $d_j = \min(j, a)$. The nonzero leading minor gives the lower bound, and for $j \le a$ the shared columns give $PY^{-1}e_h = 0$ for $h \ge j$. Therefore $[R_\gamma : gR_\gamma] = g^{2fd_j}$, and

$$I_{gR_\gamma}(M_\gamma(V)) \mid g^{2fd_j}F_\gamma(V).$$

For the second sum, keep γ fixed and restrict to $V_0 = V \cap (\gamma \operatorname{span} \mathsf{F})^{\perp}$, of codimension at most $l + n$. If $y \in V_0$, then for the lift $y + x$ just used, $\gamma^{-1}x$ is isotropic and orthogonal to F. Its image in the positive definite quotient $X^{\perp} \cap z^{\perp}/X$ must therefore be zero. It follows that $\gamma^{-1}y \in y' + X$. The second link has a lift of the same pair (y, y'), so $(\gamma'\gamma^{-1} - 1)y \in X + \gamma' X$. For integral y in the source chart, equality of the exact residues makes $g^{-1}(\gamma'\gamma^{-1} - 1)y$ integral. Pairing it with the moved augmented frame gives $-g^{-1}(\gamma'\mathsf{F})^{*}hy$ by the defining orthogonality of V_0. The unsaturated two-frame lattice has image precisely $R_{\gamma'}$ under this pairing. Its saturated quotient has order $d_{\gamma'}$. Thus

$$I_{gR_{\gamma'}}(M_{\gamma'}(V_0)) \mid d_{\gamma'}.$$

The assertions hold on the finite adelic lattices by extending the rational subspaces and using the same congruence there.

For fixed c, V, t sum first over γ' with γ held fixed, then over γ, using the two displayed divisibilities in (8.26). The second sum uses V_0 on that same γ, and the first-term sum uses V. The product of bounds is

$$b^{2000n^6(2l+n+2)f} \le \tau^{l+1}.$$

The inequality follows from $k \ge 1000(2n)^9$ and $2l + n + 2 \le 3n(l+1)$. All pairs in the same block carry one label t, and there are $|J|$ labels. Each ordered pair is counted with its actual V and no further multiplicity. With $M_0 = |J|\,|\mathrm{Cl}|\,N^{-kf}$, insertion of (8.25) bounds the sum of the ordered weights at codimension l by

$$\frac{|J|}{M_0(2\tau^2)^l}\tau^{l+1}|\mathrm{Cl}|N^{-kf}\tau^{l+1} = 2^{-l}\tau^2.$$

Summing $l \ge 0$ gives, at each level j,

$$\sum_{c,c',\gamma,\gamma'} \omega(\gamma, \gamma') \le 2\tau^2. \tag{8.28}$$

This applies to the entire family of congruent near pairs, and hence to each restricted residue set in the fixing-set and complementary estimates. Apply the bordered lattice calculation below to these same pairs for the gain from their common-link codimension and the terminal divisor on equal top links.

The bordered Gram determinant on the class lattice. Let all residues in Ω fix nonzero frame vectors at places with total norm at least $L_* \ge 1$, counting chosen and conjugate places separately. Keep the notation of the pair count. Let X be the isotropic frame span, $M = X_0 \oplus Ez$ the middle space, h the ambient Hermitian Gram matrix, and Λ_c the self-dual adelic class lattice. Gram entries are written $v^{*}hw$. Intersections of a rational subspace with an adelic lattice use its adelic span. For a term γ of level j, put $d = \min(j, a)$, $t = a - d$, $A = X + \gamma X$, and $I = X \cap \gamma X$. The columns of

$$F = (e_0, \dots, e_{a-1}, \gamma e_0, \dots, \gamma e_{d-1})$$

are a basis of A. For $j \le a$, the leading minor of size j implies that the first j isotropic vectors are independent modulo I. The shared vectors at positions $j, \dots, a-1$ give an independent family in I with unit triangular last t coordinates. The same holds in the moved frame after rescaling the common leading coefficient. For $j = a + 1$, invertibility of the full frame pairing gives $I = 0$, since a vector in I would pair to zero with the entire moved frame. Thus $\dim I = t$ in both cases. The map

$$\ell_\gamma : A^{\perp} \longrightarrow M^2, \qquad v \longmapsto ([v], [\gamma^{-1}v])$$

has kernel I and image $L_\gamma = \mathrm{link}(\gamma)$. Its image is isotropic for $h_M \oplus (-h_M)$ and has dimension $k + 1$, so it equals its orthogonal complement for that form.

Write $S_0 = (e_d^{*}, \dots, e_{a-1}^{*})$, $p = (z, \gamma z)$ and $P_0 = (S_0, p)$. The bordered matrix is

$$B_\gamma = (F, P_0)^{*}h(F, P_0) + \mathrm{diag}(0, N1_{t+2}).$$

The notation $[u]$ for a nonzero finite idele means its inverse finite modulus, so $[u] = |\mathrm{Nm}_{E/\mathbf{Q}}(u)|$ on $E^\times$. Group the columns as $(\gamma e_{<d}, S_0, \gamma z)$ and $(e_{<a}, z)$. The second diagonal block is zero, since the added N cancels $h(z,z) = -N$. The cross block R has the first d rows of $Z(\gamma^{-1})$, the last t isotropic identity rows and the last row of $Z(\gamma^{-1})$. Hence $[\det B_\gamma] = [\det R]^2$. For $j = a+1$, $R = Z(\gamma^{-1})$. Otherwise replace the last row of R by the last identity row to obtain Y. Then $\det Y = \det Z_j(\gamma^{-1})$. A shared terminal vector can be given terminal coordinate one in the first frame and zero in its last t isotropic coordinates by subtracting shared isotropic vectors. Its leading coefficient in the other frame has absolute value one at every complex place. Its first-frame coordinate vector is $Y^{-1}e_{\mathrm{last}}$, because it is orthogonal to $\gamma e_{<d}$. The row replacement therefore multiplies $\det Y$ by $-N$ times a scalar of norm one. In the notation of the cut block sizes, this gives

$$[\det B_\gamma] = (N^{2f} D_j(\gamma))^2. \tag{8.29}$$

Let $Q = \frac{1}{2}\operatorname{diag}(1_k, N, 1_k, N)$ on M^2, and let $\mathcal{L}_\gamma = \ell_\gamma(\Lambda_c \cap A^\perp)$. For an adelic basis v of a lattice in this positive form write $\mathrm{disc}_Q(v) = [\det(v^* Q v)]^{1/2}$. We compute the discriminant of $\mathcal{L}_\gamma$. Choose a rational basis (u_I, u_L) of $A^\perp$ with u_I a basis of I, $S_0^* h u_I = 1_t$ and $S_0^* h u_L = 0$. The projection on the last t isotropic coordinates of I is invertible by the shared vectors, so this choice is possible. The columns $\ell_\gamma(u_L) = v$ are a basis of the link. In the flag $I \subset A^\perp$, choose an integral adelic basis of the saturated intersection with Λ_c, of the form $\mathcal{U} = (u_I, u_L)\left(\begin{smallmatrix} A_0 & B_0 \\ 0 & E_0 \end{smallmatrix}\right)$. Choose the basis over each valuation ring by extending a basis of the primitive first lattice. The associated changes are unimodular outside a finite set. The columns vE_0 are a basis of $\mathcal{L}_\gamma$.

Let $\iota = [\Lambda_c \cap A \,:\, \langle F \rangle]$, with the spans taken over $\widehat{\mathcal{O}}_E$. The complementary-minor identity on this lattice gives

$$[\det B_\gamma] = [N]^{t+2} \iota^2 [\det(\mathcal{U}^*(h + N^{-1} h P_0 P_0^* h)\mathcal{U})].$$

To see the identity, first replace F by a primitive basis F_{sat}, extracting ι^2. Primitivity and self-duality make $F_{\mathrm{sat}}^* h$ a surjection on the integral ambient lattice. In the unimodular basis $(T, \mathcal{U})$ provided by a section $F_{\mathrm{sat}}^* h T = 1$, let J_0 be the ambient Gram matrix and B the coordinates of P_0. The bordered Gram matrix for F_{sat} is the principal minor deleting the $\mathcal{U}$-indices of

$$\begin{pmatrix} J_0^{-1} & B \\ B^* & B^* J_0 B + N 1 \end{pmatrix}.$$

The Schur complement gives the full determinant and upper inverse block, respectively,

$$\frac{N^{t+2}}{\det J_0}, \qquad J_0 + N^{-1} J_0 B B^* J_0.$$

Jacobi's identity gives the formula since $[\det J_0] = 1$.

In the rational basis (u_I, u_L) the last Gram matrix before the adelic basis change is $\operatorname{diag}(N^{-1}1_t, v^* Q v)$. Indeed the S_0-pairings give the first block. The two middle components of ℓ_γ have the same h_M-pairing on the link, and the z, γz terms change their negative coordinates to positive ones. The factor $[N]^{t+2}[N^{-1}]^t = [N]^2$ cancels N^{4f} in (8.29), giving

$$\mathrm{disc}_Q(\mathcal{L}_\gamma)[\det A_0]\iota = D_j(\gamma).$$

The class chart fixes e_i, z, e_i^*. Its other columns use the middle lattice basis with its marked coordinate. The map $x \mapsto (\gamma f)^* h x$ on the integral unmoved isotropic lattice has image R_γ and kernel $I \cap \Lambda_c$. For the kernel, $\gamma^{-1}x$ is isotropic and orthogonal to the unmoved augmented frame. Its class in the positive definite quotient of $X^\perp \cap z^\perp$ by X is zero. This proves the kernel identity over E and hence in the adelic chart. Modulo the image of the first d unmoved isotropic vectors, the quotient is the last t coordinates modulo $A_0 \widehat{\mathcal{O}}_E^t$. Its index is $I_j(\kappa) = [\det A_0]$. In the other direction, insert the whole two-frame lattice between $\langle F \rangle$ and $\Lambda_c \cap A$, extracting its saturation index d_γ. Pairing with the unmoved augmented frame kills X, and

maps the remaining index onto the quotient of the columns for all γe_i by those for $i < d$. Those columns are the adjoints of the isotropic rows at κ^{-1}. Thus $I_j(\kappa^{-1})d_\gamma$ divides ι, and

$$\operatorname{disc}_Q(\mathcal{L}_\gamma) I_j(\kappa) I_j(\kappa^{-1}) d_\gamma \le D_j(\gamma). \tag{8.30}$$

Separating the common-link and terminal factors. Let $L = L_\gamma \cap L_{\gamma'}$, $L_0 = L \cap X_0^2$, and let V be its first projection. Use Q-orthogonal complements for the rational flag $L_0 \subset L \subset L_\gamma$. An adapted adelic basis of $\mathcal{L}_\gamma$ is upper triangular in this flag. Its discriminant is the product of the discriminants of the three diagonal lattices. The first projection identifies the form on L_0 with the definite form on V. An integral lift of a vector of L_0 has the form $x + x_0$, with $x \in \Lambda_c \cap V$, x_0 in the integral isotropic lattice, and is orthogonal to the entire moved frame. Hence $(\gamma f)^* hx = -(\gamma f)^* h x_0 \in R_\gamma$. The first diagonal lattice therefore has discriminant at least $d_c(V) F_\gamma(V)$, by the finite index of this frame-value condition on $\Lambda_c \cap V$.

Let z_0 take pairings with z on the two components of the link. On L_γ it is represented for Q by two vectors $\epsilon_1 = (u - z, yz)$, $\epsilon_2 = (\bar{y}z, u' - z)$ in L_γ. At top level obtain ϵ_1 by applying γf to $NZ(\gamma)^{-1} e_{\text{last}}$ and taking its link image. Then $y = N \det Z_a / \det Z$ and $\delta(\gamma) = \det Z \, \overline{\det Z}(1 - y\bar{y})$. At lower levels use a shared terminal vector, giving $y\bar{y} = 1$. The inverse construction gives the other vector with terminal multiplier $\bar{y}$, by isotropy for the difference form. This isotropy also gives $\epsilon^* Q = z_0$ on the link. Their Gram is $N \left(\begin{smallmatrix} 1 & -\bar{y} \\ -y & 1 \end{smallmatrix} \right)$. By positivity its rank is the rank of z_0 on the link, one below the top level and two at the top.

For the middle diagonal lattice, of rank $r_0 = \dim(L/L_0)$, this shows that $l = k - \dim V = i - 1 + r_0 \le i + (j - a)_+$. Its terminal pairings agree with those of the full integral lifts. Thus they are integral and have a nonzero minor of size r_0. Let $L' = L_*$ when $j = a + 1$ and $L = L_\gamma$, and $L' = 1$ otherwise. In the first case $r_0 = 2$, and the terminal determinant has index at least L_* by the fixed-vector determinant test below. In the other cases the nonzero integral minor has index at least one. Gram Cauchy–Schwarz on the corresponding r_0 columns of ϵ bounds that index by their discriminant times the discriminant of the middle diagonal lattice. The former factor is 1 for $r_0 = 0$, N^f for $r_0 = 1$, and $x(\gamma)/D_j(\gamma)$ for $r_0 = 2$. The last equality uses $r_0 = 2$ only at top level and the factor N^{2f} in $D_j(\gamma)$.

For the determinant test, work at one of the specified split places and use the reduction of the lattice chart. If the terminal determinant were a unit, the two forms defined by the integral lifts, which vanish on X and κX, would force both terminal coefficients of every common vector of the two augmented frames to be zero. Suppose a nonzero common vector has last nonzero isotropic coefficients x_b, $y_{b'}$ on the two sides. For any linear form λ vanishing on the first $b + 1$ and moved first $b' + 1$ isotropic vectors, the rank-one translation $1 + u\,v\lambda$, u in the residue field, belongs to both flag unipotents. At the conjugate place the corresponding inverse adjoint fixes the augmented frame, so its simple-root sum is zero. Lift the two translations separately, using the corresponding flag vectors and linear forms in their respective coordinates. The lifts can be supported on this one split pair. The reductions intertwine at both places by inverse adjoint, including when the fixed vector is on the conjugate side. The locally self-dual phase has conductor $\mathcal{O}_{E,w}$ at the working place w, where g has valuation one. Division by g therefore gives a nontrivial additive character of its residue field. Comparing the phases for all u gives $x_b \lambda(f_{b+1}) = y_{b'} \lambda(\kappa f_{b'+1})$. The annihilator relation gives a common vector with these same leading coefficients at the next indices. Iterating reaches a nonzero terminal coefficient. This contradicts the unit determinant since a fixed nonzero frame vector is already a common vector. This applies at each chosen or conjugate place separately, so the product of their norms divides the ideal index of the determinant.

Write $\mathcal{Y}_\gamma$ for the third diagonal lattice in the complement of L in L_γ. Combining the first and the middle factors with (8.30) gives exactly

$$L' d_c(V) F_\gamma(V) d_\gamma \operatorname{disc}_Q(\mathcal{Y}_\gamma) \le B_j(\gamma).$$

The pairing for $h_M \oplus (-h_M)$ between the complements of L on the two sides is nondegenerate. It equals the pairing of the integral lifts by $h(1 - \gamma\gamma'^{-1})$, since components in the common link pair to zero. Congruence

of the two term positions makes this pairing divisible by g. The difference form is twice a Q-unitary change, so the nonzero pairing determinant in codimension i gives $\operatorname{disc}_Q(\mathcal{Y}_\gamma)\operatorname{disc}_Q(\mathcal{Y}_{\gamma'}) \geq (g/2)^{2if}$ by Gram Cauchy–Schwarz. Multiplication and taking square roots give the pair charge on the actual term positions

$$\sqrt{B_j(\gamma)B_j(\gamma')} \geq (g/2)^{if} L' d_c(V) \sqrt{F_\gamma(V)F_{\gamma'}(V) d_\gamma d_{\gamma'}}. \tag{8.31}$$

The charge for a pair of cut terms. To use this gain, retain a set $\mathcal{R}$ of links in the top level, and retain none at a lower level. An omitted pair either lies at a lower level, has distinct top links, or has both terms in a top link outside $\mathcal{R}$. For such a pair inserting (8.31) into (8.20) and keeping the extra factor $d_c(V)^{-1}$ gives

$$u_j(\gamma)u_j(\gamma') \leq \frac{\tau^4 \boldsymbol{D}_j (2\tau^2)^l (2/g)^{ikf}}{(L')^k} \sqrt{\omega(\gamma,\gamma')\omega(\gamma',\gamma)}.$$

The extra powers of the positive integer indices and of $(g/2)^{if}L'$ can be dropped because these quantities are at least one. Put $\eta = 2\tau^2(b/g)^{kf}$. The energy ratios on this datum give $\eta \leq 1$ and $\tau^{18}\eta\mathscr{M} = 2(b^6/g)^{kf}\mathscr{M} \leq \mathscr{D}$. Fix $C_* \geq \mathscr{D}$, with $C_* \geq \tau^{20}\mathscr{D}$ if depth zero occurs, and with $C_* \geq \tau^{12}\mathscr{M}/L_*^k$ if top links are omitted. Then the coefficient of the geometric mean is at most C_*/τ^3 in every omitted pair.

For $j = 0$ the bound follows from $\tau^{17}\mathscr{D} \leq C_*/\tau^3$. For $1 \leq j \leq a$, the factor $\eta^i \leq 1$ leaves at most $\tau^{17-32j}\mathscr{D}$. For $j = n, i \geq 1$, use $l \leq i+1$ and $2\tau^6\eta\mathscr{M} \leq 2\tau^{-12}\mathscr{D}$. For $j = n, i = 0$ use $2\tau^6\mathscr{M}/L_*^k \leq 2\tau^{-6}C_*$. The inequalities use $b \geq 1000$ and $\tau \geq 2$.

At fixed (c, s, c') sum the amplitudes within each link before expanding the squared norm. Retained diagonal link squares stay in that sum. Bound the remaining inner products by the products of the norms, then use (8.20) and the bound for omitted pairs. A pair with fixed c, c' determines its residue s, so it enters the sum over Ω once. By AM–GM and (8.28), the total omitted energy at level j is at most

$$\frac{C_*}{2\tau^3} \sum_{c,c',\gamma,\gamma'} \big(\omega(\gamma,\gamma') + \omega(\gamma',\gamma)\big) \leq \frac{2C_*}{\tau} \leq \frac{b^{-30}C_*}{4^{n-j}}.$$

Here $\tau \geq b^{n+31}$ on the datum. This calculation uses the same near representatives and class charts throughout.

For a stable set T, partition the degenerate terms with agreeing rational characters by their least level $j < n$. The least-level sets are stable because each level and the character-agreement condition are invariant under the two rational unipotent actions. Terms with disagreeing rational characters have zero amplitude. The norm of the entire cut degenerate part is at most $b^{-15}\sqrt{C_*}\sum_{j=0}^{n-1} 2^{-(n-j)} \leq b^{-15}\sqrt{C_*}$. If the retained cut link squares at the top level sum to at most $v^2 w_*$, the top norm is at most $\sqrt{w_* + b^{-30}C_*}$. The triangle inequality and (8.19) therefore give

$$\|A_T\|_\Omega \leq \sqrt{w_*} + 3b^{-15}\sqrt{C_*}. \tag{8.32}$$

Young's inequality with parameter b^{-4} gives the stated full energy bound $(1+b^{-4})^2 w_* + b^{-12}C_*$. Indeed the two coefficients supplied by Young are $1 + b^{-4}$ and $9(1+b^4)b^{-30}$, which are no larger than the stated bounds.

8.8.2 The full energy on the required residue sets

On Ω_A take $C_* = \mathscr{M}_A$, using $\tau^{44}\mathscr{D} \leq \mathscr{M}$ from the datum ratios and $|A| \leq f$. For proper A retain all top links and set $w_* = (1+b^{-4})\mathscr{M}_A$ by the cut open-link estimate. When A is the full set, retain no links, take $w_* = 0$ and use $L_* \geq g_0^f$, which gives $\tau^{12}\mathscr{M}/L_*^k \leq \mathscr{M}_A$. The crude estimate with these constants gives

$$\mathscr{E}_{\mathrm{all}}(\Omega_A) \leq (1 + 4b^{-4})\mathscr{M}_A. \tag{8.33}$$

The constant follows from $(1+z)^3 + z^3 \leq 1 + 4z$ at $z = b^{-4}$.

The complement on the two required residue sets. Set $\Omega_i^{\perp} = JG_i$ and $T = S(F)^c$, and retain no links. The coordinate away from G_i reduces to a Bessel-subgroup element. It fixes the first isotropic frame vector at every chosen place and its conjugate over g_{1-i}, so $L_* = g_{1-i}^{2f}$. The depth-zero calculation for compatible frame orbits puts each such term in $S(F)$.

The ratio of the two level primes is at most b. Thus

$$\frac{\tau^{12}\mathscr{M}}{L_*^k} \le (b^4/g)^{kf}\mathscr{M} \le \tfrac{1}{2}b^{-2kf}\mathscr{D} \le \mathscr{D},$$

The middle inequality uses $2(b^6/g)^{kf}\mathscr{M} \le \mathscr{D}$ from the datum ratios. Take $C_* = \mathscr{D}$ and $w_* = 0$ in (8.32). In the normalized unfolding, the amplitude is that of $P^*\rho(s)P - T_S(s)$ on the given map P. Thus

$$\mathscr{E}_{S(F)^c}(\Omega_i^{\perp}) \le 9b^{-30}\mathscr{D} \le b^{-12}|J|C. \tag{8.34}$$

8.9 The two projected trace moments

8.9.1 The finite Bessel phase

Write $\chi_U^{\natural}(s) = |G|^{-1}\sum_{u\in G}\chi_U(usu^{-1})$. We bound the finite phase in counting measure. At a chosen place above g_i, put $q = g_i$, $a = \varrho \ge 1$, $n = k + 2a + 1$, and $V = \mathbf{F}_q^n$. At the split place the residual group is $\mathrm{GL}(V)$. Order the frame and its dual as

$$e_1, \dots, e_a, \quad [z, X_0'], \quad e_a^*, \dots, e_1^*.$$

Here X_0' is the orthogonal complement of z in the middle block over the residue field. The lattice change preserves z, of norm $-N$, and N is a unit at this place. The image of the integral unipotents is $1 + \mathfrak{u}$, with $\mathfrak{u}$ the strictly block upper triangular matrices for block sizes $(1^a, k+1, 1^a)$. To lift a residue, lift the strictly upper blocks in the chosen split component. At its conjugate component use the inverse transpose transported by the lattice Gram matrix. These integral lifts are block unitriangular and can be set to the identity away from the chosen pairs of places.

The unipotent character has the simple entries joining $e_1, \dots, e_a, z$ at the chosen place. At the conjugate place the inverse transpose moves those entries to the successive links joining $z, e_a^*, \dots, e_1^*$. In fact the adjacent block of the inverse is the negative of the original adjacent block. The inverse Gram matrix interchanges the paired isotropic vectors, and the middle boundary uses the covector $h(z, \cdot)$, which takes the nonzero value $-N$ on z. Both residue additive characters are nontrivial because $\psi_f(x/\varpi)$ is self-dual on the local integers, and the level prime has valuation one. Relative to one fixed nontrivial additive character ψ of $\mathbf{F}_q$, each link has a nonzero coefficient. Conjugating by a diagonal change of basis along the chain makes them all one. Such a change preserves $\mathfrak{u}$ and the squared norm of the conjugacy average. With J_0 the resulting lower shift along this chain, the phase is $\psi(\mathrm{tr}(J_0X))$ at $1 + X$, $X \in \mathfrak{u}$, and zero elsewhere.

Let $\mathfrak{p}$ be the block upper triangular algebra including its diagonal blocks, and let $\chi^{\natural}$ denote the local conjugacy average. For $F(X) = 1_{\mathfrak{u}}(X)\psi(\mathrm{tr}(J_0X))$ on $\mathrm{End}(V)$, use the counting transform $\widehat{F}(A) = \sum_X F(X)\psi(-\mathrm{tr}(AX))$. An entry in a strict upper block of X pairs with the transposed entry of A, so $\mathfrak{u}^{\perp} = \mathfrak{p}$ for the trace pairing. Thus $\widehat{F} = |\mathfrak{u}|1_{J_0+\mathfrak{p}}$. Fourier transform commutes with conjugacy averaging. The average of F is supported on nilpotents, and $X \mapsto 1 + X$ identifies it there with $\chi^{\natural}$, with $\chi^{\natural}(1) = F(0) = 1$. Plancherel, with A uniform in $\mathrm{End}(V)$, gives

$$\|\chi^{\natural}\|_{\ell^2(\mathrm{GL}(V))}^2 = |\mathfrak{u}|^2\,\mathbb{E}_A p(A)^2, \qquad p(A) = \mathbb{P}_{g\in\mathrm{GL}(V)}(g^{-1}Ag \in J_0 + \mathfrak{p}).$$

In an ordered basis $(b_1, \dots, b_a, z, x_1, \dots, x_k, t_1, \dots, t_a)$ write $b_{a+1} = z$, $c_i = t_{a+1-i}^*$ for $1 \le i \le a$, and $c_{a+1} = z^*$, using the dual basis. The lower blocks of $J_0 + \mathfrak{p}$ prescribe the first a columns and last a rows below the diagonal blocks. Equivalently,

$$Ab_i - b_{i+1} \in \langle b_1, \dots, b_i\rangle, \qquad c_iA - c_{i+1} \in \langle c_1, \dots, c_i\rangle \quad (1 \le i \le a).$$

Starting at $v = b_1$ and $\lambda = c_1$, the two lists are unit triangular changes of $v, Av, \dots, A^a v$ and $\lambda, \lambda A, \dots, \lambda A^a$. Such changes preserve the evaluation matrix with only its last diagonal entry equal to one. Hence the dual-basis conditions $c_i(b_j) = 0$ except for $c_{a+1}(b_{a+1}) = 1$ are equivalent to

$$\lambda A^h v = 0 \quad (0 \le h < 2a), \qquad \lambda A^{2a} v = 1.$$

These moments force independence of both sequences through power $2a$. A dependence with highest index h would give a contradiction after multiplying by A^{2a-h} and evaluating against the other initial vector or covector. For each v with independent sequence there are q^k choices of λ, since the moments prescribe $2a+1$ independent evaluations. Each pair gives $q^{a(a+1)}$ choices of the two unit triangular changes. The x_i lift an ordered basis of $\big(\bigcap_{j=1}^{a+1} \ker c_j\big)/\langle b_1, \dots, b_a\rangle$, which has dimension k, with q^a choices of lift per vector. Each t_i can be any vector in its prescribed affine c_j-evaluation fibre of size q^{n-a-1}. These choices complete the basis and their counts are constant. It follows that $p(A)$ is proportional to the fraction $d(A)$ of vectors with independent sequence through power $2a$. The affine codimension gives $\mathbb{E}_A p(A) = |\mathrm{u}|^{-1}$. For a uniform A and uniform v, expose $v, Av, \dots, A^{2a}v$ in order. After independence up to any given step, the next vector is the image of a new independent input to the uniform linear map, and hence is uniform in V. Thus

$$\mathbb{E}_A d(A) = \prod_{j=0}^{2a} (1 - q^{j-n}), \qquad \|\chi^{\natural}\|^2_{\ell^2(\mathrm{GL}(V))} = \frac{\mathbb{E}_A d(A)^2}{(\mathbb{E}_A d(A))^2} \le \prod_{j=0}^{2a} (1 - q^{j-n})^{-1}.$$

All $2f$ split factors are independent in the residual group G. Both the phase and its conjugacy average factor over them. The product of the denominators in this bound is at least $1 - 2f/(q_{\min}^k(q_{\min} - 1)) > 1/2$, using $q_{\min} = \min(g_0, g_1) \ge b^{10}$ and $f \le b$ for large x. We have proved

$$\sum_{s \in G} |\chi_U^{\natural}(s)|^2 \le 2.$$

The projected covariance. To use (8.33), let $\Theta(s) = \|K(s)\|^2_{\mathrm{HS}}$. It is nonnegative and of positive type since it is the Hilbert–Schmidt pairing of PP^* with $\rho(s)PP^*\rho(s)^*$. Choose A and fixed vectors to maximize $b^{6k|A|} \sum_{s \in F_A} \Theta(s)$. If A is full, (8.33) suffices. Otherwise take one missing place, with residue field $\mathbf{F}_q$. Let $M_A = \sum_{s \in F_A} \Theta(s)$. Maximality bounds the Θ-mass of the stabilizer in F_A of any prescribed nonzero vector of the frame or dual frame there by $Y = b^{-6k} M_A$.

For a leading l-by-l frame block of s at this place, write it as $A_1^{\mathsf{t}} \pi(s) A_0$, where π acts on $\mathbf{F}_q^r$ and A_0, A_1 are the frame and paired-frame inclusions of $\mathbf{F}_q^l$. For a nontrivial additive character ψ of $\mathbf{F}_q$, the left kernel size equals $\sum_y \mathbb{E}_z \psi(y^{\mathsf{t}} A_1^{\mathsf{t}} \pi(s) A_0 z)$ and is at least q for a singular block. The two axes of this character sum cost at most $2M_A$ after multiplying by Θ and summing. For $y, z \ne 0$, use the delta basis for $\pi(s)A_0 z$ and the orthonormal Fourier basis for the contragredient orbit of $A_1 y$. Tensor each with vectors u_s having $\langle u_a, u_b\rangle = \Theta(a^{-1}b)$. The squared norm of either sum of tensor vectors over $s \in F_A$ equals $|F_A|$ times the Θ-mass of the corresponding stabilizer, by expanding the squared norm and grouping by $a^{-1}b$. The mixed pairing of the two sums equals $|F_A| q^{-r/2}$ times the corresponding phase sum. Cauchy–Schwarz therefore bounds that sum by $q^{r/2} Y$. Summing in y and averaging in z gives

$$q \sum_{\substack{s \in F_A \\ \det(A_1^{\mathsf{t}} \pi(s) A_0) = 0}} \Theta(s) \le 2M_A + q^{l + r/2} Y \le 3M_A.$$

The last bound uses $l \le \varrho + 1$, $r = k + 2\varrho + 1$, $q \le b^{11}$ and $k \ge 1000(2\varrho + 2)$, which gives $11(l + r/2) \le 6k$ through $l = \varrho + 1$. The union bound over all these minors costs at most $3(\varrho + 1)/q \le b^{-4}$ of M_A, since $\varrho + 1 \le k \le b$ and $q \ge b^{10}$. Every residue in $F_A \setminus \Omega_A$ is singular at one of these minors at the missing place. Thus $(1 - b^{-4})M_A \le (1 + 4b^{-4})\mathscr{M}_A$. The maximal choice bounds the unrestricted sum by $b^{6k|A|} M_A$. In particular $\sum_s \Theta(s) \le (1 + 9b^{-4})\mathscr{M}$.

For the rank argument use the projected covariance

$$H = |G|^{-1} \sum_s P_W \rho(s) P P^* \rho(s)^* P_W.$$

For a finite family of maps L_s into the same Hilbert space, cyclicity of trace gives

$$\operatorname{tr}\left(\left(\sum_s L_s L_s^* \right)^2 \right) = \sum_{s,t} \operatorname{tr}(L_t^* L_s L_s^* L_t) = \sum_{s,t} \|L_s^* L_t\|_{\mathrm{HS}}^2.$$

Here $L_s = \rho(s)P$, so $L_s^* L_t = K(s^{-1}t)$. Each relative translate occurs $|G|$ times. The ambient covariance, which averages $L_s L_s^*$ with weight $1/|G|$, has squared norm $|G|^{-1} \sum_s \Theta(s)$. Multiplication on both sides by the orthogonal projection P_W decreases this norm. Applying (8.3) to the same W gives

$$\frac{R \operatorname{tr}(H^2)}{C^2} \le 1 + 9b^{-4}. \tag{8.35}$$

The two complementary estimates needed for rank. The small share of a bad finite type reduces the precision needed in one of the complementary orbital estimates. Put

$$\Delta_i = \frac{1}{|J|C} \sum_{s \in J G_i} \|K(s) - T_S(s)\|_{\mathrm{HS}}^2.$$

Fix i and put $D_i = J \cap G_i$, $\tau_i = T_S|_{D_i}$ and $N_i = [G_i : D_i]$. On D_i the comparison term is a unitary representation on the column space C. Its extension by zero to G_i is $T_S|_{G_i}$. The covariance $K(ts) = \tau_i(t)K(s)$ implies $P^*\rho(t) = \tau_i(t)P^*$ on the span of the $\rho(s)PC$, and both sides vanish on its orthogonal complement. Taking adjoints and replacing t by t^{-1} gives $\rho(t)P = P\tau_i(t)$.

Let π be the induced representation on the orthogonal sum of the column fibres over G_i/D_i. We use $[ht, v] = [h, \tau_i(t)v]$ and left translation. The map $S_i[h, v] = \rho(h)Pv$ intertwines this representation with $\rho|_{G_i}$. Put $L_i = S_i^* S_i$. The $(hD_i, h'D_i)$ block of $L_i - I$ in coset representatives is $K(h^{-1}h') - T_S(h^{-1}h')$. The Hilbert–Schmidt norm of this block is unchanged by left and right multiplication of its argument by D_i. Each summation over G_i/D_i is consequently a sum over G_i divided by $|D_i|$. Left J-covariance gives

$$\frac{\|L_i - I\|_{\mathrm{HS}}^2}{N_i C} = \frac{1}{|D_i|C} \sum_{h \in G_i} \|K(h) - T_S(h)\|_{\mathrm{HS}}^2 = \Delta_i.$$

For any self-adjoint central idempotent $a \in \mathbf{C}[G_i]$, the identity-coset block of $\pi(a)(L_i - I)$ is $\sum_{h \in G_i} a(h)(K(h) - T_S(h))$. The diagonal coset blocks are unitarily conjugate by equivariance. Their squared norms sum to at most $\|L_i - I\|_{\mathrm{HS}}^2$, since $\pi(a)$ is an orthogonal projector. Thus

$$\Big\| \sum_{h \in G_i} a(h)(K(h) - T_S(h)) \Big\|_{\mathrm{HS}}^2 \le C\Delta_i.$$

Take $a_i = 1 - e^{(i)}$, $A_i = \pi(a_i)$ and $p_i = \operatorname{rank}(A_i)/(N_i C)$. The compression of $\pi(h)$ to the identity coset is $T_S(h)$. As all diagonal coset blocks of A_i have the same trace,

$$p_i = C^{-1} \sum_{h \in G_i} a_i(h) \operatorname{tr} T_S(h) = \sum_{h \in G_i} a_i(h) \chi_i^\natural(h),$$

where the one-factor trace formula gives the phase and $\chi_i^\natural$ is its conjugacy average on G_i. Likewise the identity-coset trace of $A_i L_i = S_i^* \rho(a_i) S_i$ is $\operatorname{tr}(P^* \rho(a_i) P)$. Hilbert–Schmidt Cauchy–Schwarz gives

$$\left| \frac{\operatorname{tr}(P^* \rho(a_i) P)}{C} - p_i \right| = \frac{|\operatorname{tr}(A_i(L_i - I))|}{N_i C} \le \sqrt{p_i \Delta_i}.$$

This uses the counting normalization of the actual complementary orbital sums and the phase of the Bessel term at the same place.

The identity contribution to the finite share. Put $q_{\min} = \min(g_0, g_1)$ and $u = 2f/[q_{\min}^k(q_{\min} - 1)]$. On G_i, $\chi_i^{\natural}(1) = 1$. The product of the chain probabilities in the finite phase calculation is at least $1-u$. The product over just the factors of G_i has the same lower bound. In counting measure, with δ_1 the indicator of the identity, this gives

$$\|\chi_i^{\natural} - \delta_1\|_2^2 = \|\chi_i^{\natural}\|_2^2 - 1 \le \frac{u}{1-u}.$$

Since a_i is a self-adjoint idempotent, $\|a_i\|_2^2 = a_i(1) = \beta_i$. Separate the identity in the formula for p_i and apply Cauchy–Schwarz to the remaining terms. By (8.2),

$$p_i \le \beta_i + \sqrt{\beta_i u/(1-u)} \le b^{-8}$$

for sufficiently large selected squares, since $\beta_i \le b^{-8}/2$, $q_{\min} \ge b^{10}$ and $f \le b$.

The projected diagonal bound. Taking a to be the unit in the induced block estimate gives $\|P^*P - I_C\|_{\mathrm{HS}}^2 \le C \min_i \Delta_i$. The two commuting bad projections dominate $I - P_W$ in sum. The positive-part trace is the supremum of the trace against $0 \le Q \le I$, and the trace norm of $I_C - P^*P$ is at most $\sqrt{C}\|I_C - P^*P\|_{\mathrm{HS}}$. Hence

$$d_+ := \frac{\operatorname{tr}((I_C - P^*P_W P)_+)}{C} \le \min_i \sqrt{\Delta_i} + \sum_{i=0}^{1} \left(p_i + \sqrt{p_i \Delta_i}\right) \le 5b^{-4}, \tag{8.36}$$

where the last bound uses $p_i \le b^{-8}$ and the supplied $\Delta_i \le b^{-12}$ for both factors.

8.10 Selection of a common family of rows

8.10.1 Selection using the projected maps

Let $L_j : V_j \to \mathcal{H}$ be a finite family of maps between finite-dimensional Hilbert spaces, and let Q_j be orthogonal projections with $L_j(\operatorname{im} Q_j) \subseteq Z$. Put $A = \sum_j L_j L_j^*$, $A_Q = \sum_j L_j Q_j L_j^*$, and $n_j = \operatorname{tr}((I - L_j^* L_j)_+)$. If P_Z projects orthogonally onto Z, then

$$\sum_j (\operatorname{rank} Q_j - n_j) \le \operatorname{tr} A_Q \le \operatorname{tr}(P_Z A) \le \sqrt{\dim Z \operatorname{tr}(A^2)}. \tag{8.37}$$

For the first inequality, take the trace of $I - L_j^* L_j$ against Q_j and bound it by the trace of its positive part. For the second, $A_Q \le A$ and $\operatorname{tr} A_Q = \operatorname{tr}(P_Z A_Q)$ by the range condition. The last is Hilbert–Schmidt Cauchy–Schwarz.

Use d_+ from (8.36) and put $\epsilon = (R \operatorname{tr}(H^2)/C^2 - 1)_+$. If $d_+ < 1$, then $R > 0$, since $P_W P = 0$ would give $d_+ = 1$. Let $T_s = P_W \rho(s) P$, and for $F \subseteq W$ let Q_s project onto $S_s = T_s^{-1} F$. In the positive trace compression estimate (8.37), take $L_j = T_s$, $Z = F$ and $A = |G|H$. Centrality of P_W gives $T_s^* T_s = P^* P_W P$ for each s, so $n_j = C d_+$. Dividing by $|G|$ and using the arithmetic–geometric mean gives

$$|G|^{-1} \sum_s \dim S_s - C d_+ \le \sqrt{\dim F \operatorname{tr}(H^2)} \le C \sqrt{(1+\epsilon)\frac{\dim F}{R}} \le \frac{C}{2}\left(1 + \epsilon + \frac{\dim F}{R}\right).$$

If $k_{\mathrm{ram}} = C^{-1} \dim \ker \mathcal{R}$, then $\dim \mathcal{R}^{-1} S_s \le \dim S_s + C k_{\mathrm{ram}}$ by rank–nullity. Thus the needed hypothesis is the combined bound $d_+ + \epsilon/2 + k_{\mathrm{ram}} \le \eta_0$, where we take $\eta_0 = b^{-3/2}$. Equations (8.36) and (8.35) and the central orbit bound $k_{\mathrm{ram}} \le f/(N-1)$ imply it for sufficiently large x. Here $N - 1 \ge b^9$. On the ramified columns the nonzero scalars from normalization, the seesaw and the integral scaling of Section 7 preserve these preimages. Consequently for every $F \subset W$ there is a translate with

$$\dim(B_s \mathcal{R})^{-1} F \le C\left(\frac{1 + \dim F/R}{2} + \eta_0\right). \tag{8.38}$$

Start with $F = 0$. While the codimension $D_F = R - \dim F$ exceeds $4\eta_0 R$, add the image of the selected translate. The rank gain is $\Delta \ge C(D_F/(2R) - \eta_0) > 0$. The potential $5R \log(1 + D_F/(\eta_0 R))$ decreases by at least $5R\Delta/(D_F + \eta_0 R) \ge C$. Thus a set of translates $\mathcal{S}$ has $|\mathcal{S}|C \le 5R \log(1 + 1/\eta_0)$ and images spanning a space of dimension at least $(1 - 4\eta_0)R$. Since $\eta_0 = b^{-3/2}$, the budget is at most $5R \log(1 + b^2)$.

The common family of rows. Choose a basis $y = (y_1, \dots, y_n)$ for the selected image span from the vectors $B_s\phi$, with $s \in \mathcal{S}$ and ϕ in the fixed short spanning set. For each i retain one such presentation $y_i = B_{s_i}\phi_i$. For a coefficient prime p, an embedding j over p, and a twist θ of level norm prime to p, group this fixed list by s_i. The integral short-form bound gives the two recombinations on each group of preimages at j and $j \circ \mathrm{conj}$. Taking their block-diagonal sums in the row order gives U_+, U_-, with integral lifted rows for the same θ and with the paired cost

$$\begin{gathered} -\log|j(\det U_+)|_p - \log|(j\circ\mathrm{conj})(\det U_-)|_p \le F_p, \qquad F_p = 2|\mathcal{S}|Ce_p\log p, \\ \sum_p F_p \le 2\delta, \qquad \delta = |\mathcal{S}|C\mathcal{B}. \end{gathered} \tag{8.39}$$

At infinity (7.9) gives $\|\Theta_1 y_i\|_\sigma^2 \le e^{h_{\mathrm{sh}}}$ at every σ. Applying Hadamard to the Gram matrix of these lifted rows gives

$$n \ge (1-4\eta_0)R, \qquad \delta \le 5R\log(1+b^2)\mathcal{B}, \qquad \mathrm{covol}_\sigma^2(\Theta_1 y) \le e^V \quad \text{for every } \sigma, \qquad V = nh_{\mathrm{sh}}. \tag{8.40}$$

9 Theta maps on the two grids

The period maps will be compared for a family of norm-one Hecke characters of weight zero. A trial stage character ψ has tame conductor Q_ψ, a squarefree product of chosen halves of split pairs with norm prime to gN. Its residue characters are nontrivial at every place in this support. The character is unramified on the norm-one group off these split pairs. The empty support includes the trivial character. A trial chain character ϑ is a stage character for which ϑ^j has the same conductor for $1 \le j \le m$. All levels at the chosen primes remain in force. The base spaces, real types, and the lifts at the trivial character are fixed by the arithmetic datum. Write υ_E for the splitting character of the arithmetic datum and υ for its restriction $\upsilon_{U(1)}$ to the norm-one torus.

9.1 The triangles and their first arrows

On the good row space W, write Θ_ϑ for the weighted chain on the V-column, read in the terminal class-sum metric, and put $Z_\vartheta = \det(\Theta_\vartheta^*\Theta_\vartheta)$. The determinant uses the fixed Petersson metric on W.

The determinant along the triangle. For a fixed m-fold twist consider the source triangle. Write $a_{i,n} : F_{i,n} \to F_{i,n+1}$ for the lifts, $d_{i,n} : F_{i,n+1} \to F_{i+1,n}$ for the descents, and e_i for the isometric inclusions of principal-level forms in $F_{i,0}$. Denote the positive triangle weights by q_n and the successive weighted lifts in row i by $M_{i,n}$, with $M_{i,0} = I$ and $M_{i,n+1} = q_n a_{i,n} M_{i,n}$.

The arrows on the two triangles. Let $z = 0$ for the grid V and $z = 1$ for Z. At (a, b) use the stem of rank $r + z$, followed by a blocks of size $t + z$ and b blocks of size $t - z$. We use the same index order in both grids. An interval of length n in the triangle of grid z starts at offset i. Its space is the node $(0, n)$ of that grid with the blocks numbered $i+1, \dots, i+n$. Its rank is $r+z+n(t-z)$. At length zero the space is the i-fold opposite of the original space. The isometry e_i on base-level forms is pullback by this identification of the unitary groups, with their quotient measures. The base level is the chosen principal level for V, and the principal level of $\mathfrak{b} = (\bar\pi_G)^{b_1}$ for Z. Write $\mathfrak{M} = (\bar\pi_G)$. The finite ring of the triangle is $\mathcal{O}_E/(\mathfrak{M}Q_\vartheta)$ for V and $\mathcal{O}_E/\mathfrak{b}$ for Z. In the latter all the arrows in the triangle are untwisted. The order q_n of the general linear group of the node of length n is the weight of its outgoing lift. It turns the quotient measure on its source into the measure with principal level of volume one.

At each real place use the row entries

$$\lambda_s = \begin{cases} \kappa/2 - 1 - s, & 0 \le s < \kappa - 1, \\ h_2 - w_{s+2-\kappa}, & s \ge \kappa - 1. \end{cases}$$

The stem uses $0 \le s < r + z$. If o is the offset, the doubled parameter has stem entries $2\lambda_s + z$, blocks centered at $2w_1, \dots, 2w_a$ of length $t+z$, and blocks centered at $-2w_{o+1}, \dots, -2w_{o+b}$ of length $t-z$. A block of length d has displacements $d-1, d-3, \dots, 1-d$. Lower every entry by one in V if $a+b$ is odd. On the stem the sign at position s is $(-1)^{o+A+\min((A-o)_+, b)}$, where $A = (s+1-\kappa)_+$. All block signs are positive. In V a step increasing b across the l-th block uses $\mu_l^0 \vartheta^l v^{\lfloor N/2 \rfloor}$ at a space of rank N. A step increasing a uses $\mu_l^0 v^{\lceil N/2 \rceil}$. In Z the twist is on the step increasing a, and both roundings equal $N/2$.

For an edge increasing b, subtract the weight of the character at the opposite end in the two parameters. In doubled coordinates the half-rank terms cancel the lowering, and the common coordinates are translated by $2w_l$. The new block is consequently centered at zero. Among existing stem coordinates the signs flip precisely for $A \ge l$, namely the coordinates below $-2w_l$. All earlier blocks in this column are above that center. For an edge increasing a, use the negative of the parameters. The ceiling cancels the lowering and the common coordinates become $2w_l - \widetilde{\lambda}$, where $\widetilde{\lambda}$ is a doubled entry before lowering. They are all positive, and their signs agree at the two ends. For the triangle edge deleting its first block, compare the negative parameters at $(i+1, n)$ and $(i, n+1)$, with character $b_i(N) = \alpha_{i,0}(N)^{-1} v^N$, where $\alpha_{i,0}(N)$ is the character of the first upward arrow at rank N. The shifted common entries are $-\widetilde{\lambda} - 2w_{i+1}$. The signs at offsets $i+1$ and i are opposite on the stem for $A \le i$, where this entry is negative. For $A \ge i+1$ use $\min(A-i, n+1) = 1 + \min(A-i-1, n)$ to see that they agree. The surviving blocks lie below $-2w_{i+1}$, so their shifted entries are positive. This argument is symmetric in the ends of that edge. The gaps between opposite signs exceed $500(N+N')^3$, since a sign change separates stem positions with different A, or occurs across distinct block centers. The minimum separation in either case is bounded below using $h - t - 2$ with $h = 2r^8$, and $N, N' \le 2n_m$. Use the relative Fock weight calculation with the actual excess $d = |N - N'|$ on the edge. In that calculation the entries of the central string $d-1, d-3, \dots, 1-d$ have relative weight zero. The centered entries on each matched row and column are opposite in the Fock coordinates. The sign conditions above give opposite relative weights. The product of the minors in the matched rows and columns has these weights, is killed by the mixed-sign Laplacians, and is nonzero on the matching matrix. This gives the prescribed joint highest line on each edge. The harmonics depend only on the infinite types and weights.

The finite datum of every upward arrow $a_{i,n}$ is the indicator of the standard inclusion $\binom{I}{0}$ in the triangle ring. Choose the polynomial for a left edge once for each (i, n). At $n = 0$, conjugating the upward kernel and changing its source to the opposite space changes its characters to b_i. On the residue datum this sends $p(x)$ to $\overline{p(\epsilon^{-1}x)}$, up to a scalar, with ϵ the integral unit used in the opposite-space identification. The polynomial is conjugated. In a Petersson pairing the convolution projection onto the chosen left-edge type replaces that polynomial by its left-edge harmonic up to a scalar independent of the forms. Absorbing the scalar into this first descent datum $p_{i,0}$ gives $\langle x, a_{i,0} e_i y \rangle = \langle d_{i,0} x, e_{i+1} y \rangle$. This is an identity on the realized form spaces. It follows by substituting the conjugate kernel into the Petersson integral, then interchanging the two integrations. The convolution replacement uses the same joint type at the two ends.

Compact support of a split matrix coefficient. At a split place write the source group as $\mathrm{GL}_r(F_v)$ in the chosen polarization. Its action on the matrix argument is right substitution, up to the scalar in the Weil action. Suppose that a finite Schwartz vector φ is supported on integral n-by-r matrices X, with $LX \bmod \varpi^e \in H \subseteq \mathrm{GL}_r(\mathcal{O}_v/\varpi^e)$ for a fixed integral row projection L, an integer $e \ge 1$ and a subgroup H. If X and Xg both lie in the support, then LX and $(LX)g$ are invertible integral matrices. Consequently $g = (LX)^{-1}((LX)g)$ is integral invertible and reduces to an element of H. Thus the oscillator matrix coefficient of φ vanishes outside the inverse image J_H of H in $\mathrm{GL}_r(\mathcal{O}_v)$. The same argument applies when the substitution uses g^{-1}. If φ transforms on J_H by a unitary character, write its coefficient there as $\xi(g)\|\varphi\|^2$, with ξ accounting for the conjugation in the pairing convention. If another coefficient m satisfies $m(g) = \xi(g)^{-1} m(1)$ on J_H, integration of the product gives $\mathrm{vol}(J_H)\|\varphi\|^2 m(1)$.

The point datum of the first arrows. In either triangle let q be the base rank and D' the residue ring of the first arrow. Thus $q = r$ and $D' = \mathcal{O}_E/(\mathfrak{M}Q_\vartheta)$ in V, while $q = r+1$ and $D' = \mathcal{O}_E/\mathfrak{b}$ in Z. The finite point datum is the indicator of the inclusion $J = \binom{I_q}{0}$ in this ring. At each of its split places the projection onto the marked rows is a common left inverse of all matrices in the residue support. The split support calculation with the identity subgroup at the stated congruence depth gives the principal subgroup as the only possible support of the coefficient. That subgroup fixes the residue indicator and its tame phase is one, so the coefficient there equals the squared norm of the indicator. With probability Haar on the compact source its integral is the squared norm divided by $|\mathrm{GL}_q(D')|$. Off the ramified places the datum is the vacuum.

A compact element k of the base level can be matched on the target by lifting the block diagonal matrix with blocks k mod D' and I_{r+t-q}. This follows by inserting J in the residue action on both sides. The resulting intertwining of translations is up to a unitary Weil phase, which cancels in the Gram pairing. Testing the Gram image of a k-fixed vector against all translated vectors shows that the Gram operator preserves the base level.

9.2 The doubling pairing for the fixed section

Let A have rank n and let B have rank $M > n$ and global Witt index q. Use a, b for the two characters of the original pair, acting by $a(\det h)b(\det g)^{-1}$ at $(g,h) \in U(A) \times U(B)$ in the fixed oscillator model. The doubled source is $A^{\square} = A \oplus (-A)$. Its Siegel parabolic stabilizes the diagonal isotropic subspace of dimension n. In the neutral doubled model the two characters are $v^n, 1$. If a Siegel Levi element acts on the stabilized isotropic subspace by the matrix D, its Weil action at zero has factor $v_E(\det D)^M |\det D|_{\mathbb{A}_E}^{M/2}$. Indeed, the source action on the mixed-model term uses the inverse of the source splitting character v_E^{-M}. The normalized induced section thus uses $\chi_{\square} = v_E^M$ on this Levi coordinate, and $\chi_{\square}|_{\mathbb{A}_F^\times} = \epsilon_{E/F}^M$. The original b enters through g_1 in the unconjugated kernel and g_2 in the conjugated kernel. Restoring b in both original kernels multiplies their product by

$$b(\det g_1)^{-1}\overline{b(\det g_2)^{-1}} = b(\det(g_2^{-1}g_1))^{-1}.$$

The two original a-factors at h cancel. In forming the doubled tensor, opposite-source conjugation changes the first splitting character in the conjugate factor to $a^{-1}v^n$, producing v^n on the B-side of the doubled oscillator. For a Schwartz test Φ in the doubled Weil representation $\omega^{\square}$ that is a polynomial Fock vector at infinity, let Φ_s be the standard Siegel section with value $\Phi_{s_0}(g) = (\omega^{\square}(g)\Phi)(0)$ at $s_0 = (M-n)/2$, in the normalized induction by $\chi_{\square} \circ \det |\det|^s$. If $M - q > n$, Yamana's convergent Siegel–Weil formula gives

$$E_{\mathrm{Sieg}}(g, s_0; \Phi_{s_0}) = I(g, \Phi), \qquad I(g,\Phi) = \int_{[U(B)]} \theta^{\square}(g,h;\Phi)\, dh_1.$$

Here dh_1 gives the quotient volume one. The theta integral converges absolutely and the Eisenstein series is holomorphic at s_0. For $M > 2n$ the defining Eisenstein series itself converges absolutely there. [30, Theorem 2, pp. 235–237] For $M > 2n$, the volume-one convention is stated explicitly in [31, Theorem 2.1(1), p. 859].

For the inner product of two theta lifts, take the doubled test $\delta(\varphi_2 \otimes \overline{\varphi_1})$. Here δ is the polarization change. Its evaluation at zero is $\langle \varphi_1, \varphi_2 \rangle$. Write $\omega_0(g)$ for the neutral one-copy action $\omega^{1,1}(g,1)$. Under δ, the doubled action at the block embedding $\iota_{\square}(g_1, g_2)$ corresponds to

$$\omega_0(g_1) \otimes \big(\chi_{\square}(\det g_2)\overline{\omega_0(g_2)}\big).$$

In the opposite-source model, conjugating the neutral one-copy action gives the characters (v^n, v^M). For the doubled action, the second block carries $(v^n, 1)$, so it acquires the factor $v(\det g_2)^M = \chi_{\square}(\det g_2)$. This

is the polarized tensor identity of [32, §11.3] in our determinant coordinates. Writing $G = U(A)$, insert the inverse of this second-factor character and the two b-factors in the pairing. Its value is

$$\operatorname{vol}[U(B)] \int_{[G]^2} E_{\mathrm{Sieg}}(\iota_{\square}(g_1, g_2), s_0; \Phi_{s_0}) \chi_{\square}(\det g_2)^{-1} \cdot b(\det(g_2^{-1} g_1))^{-1} f_2(g_1) \overline{f_1(g_2)}\, dg_1\, dg_2.$$

The volume factor converts the volume-one integral in the Siegel–Weil formula to the target Petersson measure.

Here is the boundary-orbit calculation for cuspidal source vectors. For a maximal isotropic subspace L in $A^{\square}$ write $X = L \cap A$ and $X' = L \cap (-A)$. Projection to the first copy gives $\mathrm{pr}_A L \subseteq X^{\perp}$ and has dimension $n - \dim X'$. Projection to the second gives the reverse inequality for $\dim X, \dim X'$. Thus both intersections have dimension d and $\mathrm{pr}_A L = X^{\perp}$. When $d > 0$, the unipotent radical N_X of the parabolic fixing X acts trivially on $X^{\perp}/X$. Consequently $N_X \times 1$ preserves L. It is normal in the stabilizer of L in $U(A) \times U(A)$ because that stabilizer preserves X in the first factor. Its action on L has determinant one, so the Siegel section is invariant under this radical in the unfolded orbit. The integral on $N_X(F)\backslash N_X(\mathbb{A}_F)$ is the cusp constant term of the first source form. This proves the boundary vanishing in the orbit unfolding of [31, (2-2)–(2-3), pp. 861–862] for the ranks and signatures here. If A is definite, there is no positive-dimensional rational isotropic intersection.

We also use the unfolding with arbitrary rapidly decreasing source vectors and a primitive finite test in the unconjugated kernel. In the polarization attached to L, the zero-term functional is invariant under the Heisenberg action of $B \otimes_E L$. Indeed, those operators act by multiplication in that polarization, with multiplier one at the origin. For $0 \neq x \in X$, take the part of $B \otimes_E Ex$ supported at the chosen split place in the summand dual to the retained matrix variable. It acts on that test by $\psi_v(\ell(Yx_v))$ on the matrix Y, with ℓ ranging over all linear forms on the target vector space. This is the split Heisenberg pairing given by contraction against the chosen column x_v. A source translate replaces x_v by an invertible translate of it. Average these modulations over increasing dual lattices. The averages restrict the support to the corresponding annihilator lattices for Yx_v. That image of the compact support stays away from zero because each retained matrix is injective. For a large enough dual lattice the average is zero. Since the zero-term functional is invariant under every modulation in the average, the boundary term is zero even before integrating the source forms. The conjugated test may have a different finite restriction.

For $d = 0$, L is a graph of a rational isometry of A. Unfold the graph orbit by the diagonal copy of $U(A)$. The inverse $\chi_{\square}$ factor cancels the Levi character on the diagonal and is one at the representative $(g, 1)$. With $R(g)$ denoting right translation and $\omega^{1,1}$ the oscillator action with the two determinant twists removed, the resulting pairing is

$$\operatorname{vol}[U(B)] \int_{U(A)(\mathbb{A}_F)} b(\det g)^{-1} \langle \varphi_1, \omega^{1,1}(g, 1)\varphi_2 \rangle \langle f_1, R(g) f_2 \rangle\, dg.$$

The polarization change evaluates the graph coefficient at $(g, 1)$ as the oscillator matrix coefficient displayed here. In all these applications $M > 2n$, so $s_0 > n/2$. Absolute convergence of the Siegel series at this point and rapid decrease justify the unfolded pairings.

Injectivity of the sharp-point lift. Write $a = r + t = 6r$ for the target rank in this calculation and $n \in \{r, r+1\}$ for the source rank. Every target here has global Witt index $w \leq a/2$, so $a - w \geq a/2 > n$. In Yamana's notation the parabolic isotropic rank of the doubled source is n and the evaluation point is $(a - n)/2$. His convergent Siegel–Weil formula therefore applies. The source forms are cuspidal, so the boundary-orbit calculation leaves the graph pairing. The primitive point support computes its marked finite factors. With the same highest polynomial in both copies of the section it expresses the mixed Petersson pairing of two lifts as the source pairing with the product of the local vacuum convolutions off

the modulus, times the real integrals, positive point masses at the specified split places and the target quotient volume. At the split places in the modulus the principal subgroups act trivially on both source forms, so the coefficient gives exactly the mass computed above. At the other finite places use the volume-one compact and include the prescribed splitting characters in the coefficient. Its compact-coset contribution is then the identity on the retained source vectors.

At a real place let h be the actual nonzero joint harmonic and put $f(g) = \langle h, \omega(g)h\rangle$. Let m be the corresponding unit highest coefficient of the oscillator-side discrete series. This series is contragredient to the bilinear automorphic input, so the compression to its lowest-K highest line pairs f against $\overline{m}$. Its restriction m_K to the maximal compact has $c = \int_K |m_K|^2 > 0$ with probability Haar. Formula (5.1) and compact Schur orthogonality give, for positive Hermitian z in the Cartan decomposition, the $K \times K$ average of $f\overline{m}$ as

$$c \int_K M_{-\mu}(u^* z u)^{-1} f(u^{-1} z u)\, du.$$

Indeed substitute $p_1 u, p_2 u$ for the two compact variables in (5.1). The two compact integrals project onto the respective copies of the highest line of h, each with factor c. The polar formula carries the normalization $1/c$, so the total factor is $c^2/c = c$. The real integral is absolutely convergent by the gap on these edges and boundedness of f. For such z, $H = I_a \otimes z^{-\mathrm{t}}$ is positive Hermitian and $\det z = 1$. In the normally ordered Fock Gaussian, each pure quadratic contraction couples opposite signs in the same row. It therefore kills the chosen harmonic polynomial P on either side. The remaining coefficient is the positive scalar $\det(H_{--})^{-1}$ times the Fock pairing of $P(C^{-1}\cdot)$ with P, where $C = P_+ H P_+ + P_- \overline{H} P_-$ and $P_\pm$ project onto the two signs. This pairing is the squared norm obtained by substituting the positive square root of C^{-1} into P. It is positive because the substitution is invertible. The same argument applies to $u^{-1} z u$, and the real integral is positive. [11, Theorem 6.3 and (25)–(26)]

For the finite convolutions let Λ be the self-dual oscillator lattice at a place off the modulus, with unit vacuum v, and put $v_g = \omega(g)v$. The Weyl translates of v_g by representatives of $\Lambda/(\Lambda \cap g\Lambda)$ are orthonormal. For two distinct representatives the difference is outside $g\Lambda$, so a shift in $g\Lambda$ with nontrivial commutator makes their pairing zero. Their pairings with v all have the same modulus because v is an eigenvector for translations in Λ. Bessel's inequality yields

$$|\langle v, v_g\rangle| \le [\Lambda : \Lambda \cap g\Lambda]^{-1/2}.$$

At a split place of residue size q in one half, let L be the source lattice, $L' = gL$, and set $u = \mathrm{length}(L/(L \cap L'))$, $w = \mathrm{length}(L'/(L \cap L'))$. The opposite half has intersection length w by duality. The oscillator lattice has a copies in each half, so the index is $q^{a(u+w)}$. For each u, w there are at most $q^{n(u+w)}$ possibilities for L'. Choose the intersection by u successive index-q sublattices and its extension by w index-q overlattices. Each step costs at most $(q^n - 1)/(q - 1) \le q^n$. At a nonsplit place let Q be the residue size over E. The two lengths are equal by self-duality and equal covolumes. The oscillator index is Q^{au}, with at most Q^{2nu} choices for L'. The compact stabilizer is the source lattice stabilizer, so each such right coset has mass one.

Consequently the convolution is $I + T_v$ on the corresponding source fixed space. Unitarity bounds $\|T_v\|$ by the total absolute coefficient mass of the nonidentity cosets, at most

$$\left(1 - q^{n-a/2}\right)^{-2} - 1 \quad \text{or} \quad \frac{Q^{2n-a/2}}{1 - Q^{2n-a/2}},$$

at split and nonsplit places, respectively. Both bounds are less than one and summable over the places since $a/2 - 2n \ge r - 2 \ge 2$. The coefficients are bi-invariant under the compact subgroup because the off-modulus datum is its vacuum and the splitting characters are unramified. The convolutions preserve the source fixed space and commute across different places. Use the Petersson Hilbert closure of the fixed finite level and real type. If $\delta_v = \|T_v\|$, then $\|(I+T_v)^{-1} - I\| \le \delta_v/(1-\delta_v)$, and these bounds are summable. The operator products and their inverse products converge in norm and remain inverse. The mixed Petersson relation tested against all forms of the fixed level and type, which are dense in this closure, shows that a

vector whose lift is zero is killed by the invertible product. This proves injectivity on the full base level for each of the two arrows, including the deep core level in Z. With the realized first adjoint this gives $d_{i,0}a_{i,0}e_i = e_{i+1}B_i$ with B_i injective.

9.3 The actual auxiliary stage lift

The stage lift and its coefficient space. At stage j, W_j is the rank-r signed space whose first $\kappa + j - 1$ stem positions are positive and whose other $m - j$ positions are negative. It has the row Harish–Chandra entries λ_s. The target is the definite space I_{6r}. The two Weil characters are v^{m-j} and $\mu_j^0 \psi v^{3r}$. Number the marked source columns by decreasing λ_s on the positive block and increasing λ_s on the negative block. The centered target parameter is obtained from the centered source entries $\lambda_s + w_j$ by inserting $(t-1)/2, (t-3)/2, \dots, (1-t)/2$ between the two sign blocks. The common entries have the signs of W_j and absolute values at least r^8, so both blocks on the target carry positive labels. The highest harmonic uses the initial minors on the positive block and the initial minors in conjugate variables on the oppositely ordered negative block. These minors give the positive and negative column diagrams of the Weyl coefficient space. The step at height h is the adjacent weight difference minus one for the centered parameter, including the adjacent end of the inserted string at the end of each block.

Let C_ψ be the theta descent from the definite coefficient space V_ψ at the auxiliary level to the full principal-level space on W_j with that harmonic. The finite section fixes the inclusion at $\mathfrak{M}$. At each new split place the marked source submatrix is invertible and contributes its inverse determinant character. The unmarked rows are free. Outside these places the section is the vacuum. The target compact level fixes the inclusion columns at $\mathfrak{M}$ and preserves the marked coordinate subspace at the twist places with the corresponding determinant character. The adjoint $u_\psi = C_\psi^*$ uses Petersson on W_j and the class sum of the compact coefficient Gram pairing weighted by $1/|\operatorname{Aut}|$ on V_ψ. On the good source space U_j of dimension R put $L_\psi = u_\psi U_j$ and $\lambda_j(\psi) = \det(u_\psi^* u_\psi|_{U_j})$.

The first-arrow pairings on the two inner forms. Put $j = i + 1$, $\psi = \vartheta^j$, and $D_\psi = \mathcal{O}_E/(\mathfrak{M}Q_\psi)$. The support of Q_ψ is that of Q_ϑ for $1 \le j \le m$. Both first-arrow pairs in the comparison have ranks r and $r + t = 6r$. The target dimension minus its global Witt index is at least $3r > r$ on the grid, and the auxiliary target is definite. Thus both are in Yamana's convergent case. Both pairs have the joint harmonics of the corresponding real edge, with weights independent of ϑ, and the source vectors in these types are cuspidal. Apply the convergent Siegel–Weil formula and the graph-orbit unfolding. In Yamana's notation the doubled skew Hermitian space has dimension $2r = 2n_{\mathrm{Yam}}$ and the target has $m_{\mathrm{Yam}} = 6r$. The normalized point is $s_0 = (m_{\mathrm{Yam}} - n_{\mathrm{Yam}})/2 = 5r/2 > n_{\mathrm{Yam}}/2$. Use Theorem 2 with the convergence bound above on the mixed-sign target. The fixed polynomial Fock tensor is the Schwartz datum. The doubled section has character $\chi_\square = v_E^{6r}$. On the surviving graph the actual character $b = \mu_j^0 \psi v^{3r}$ contributes $b(\det \gamma')^{-1}$ to the source matrix coefficient. The same-vector real calculation gives a positive scalar on the joint harmonic. Include the target quotient volume in that scalar. The opposite-space identification of the source is isometric and intertwines finite translations.

For the first arrow in V use the principal support and covariance just proved with $D' = D_\psi$ and $q = r$. Its point indicator has squared norm $|D_\psi|^{-(r+t)r}$ in the residue measure. The corresponding integral over the compact source is $q_0^{-1}|D_\psi|^{-(r+t)r}$, where $q_0 = |\mathrm{GL}_r(D_\psi)|$.

For the auxiliary lift u_ψ, Petersson adjunction with the definite class-sum metric realizes its values by the highest-vector theta lift. The projector on the Q_ψ-part of the section is a sum over source translates by $\prod_{v \in Q_\psi} \mathrm{GL}_r(k_v)$. They act trivially on the chosen-level source forms when lifted integrally and congruent to the identity at $\mathfrak{M}$. For each invertible marked submatrix exactly one translate of the identity condition contributes, with the determinant factor from the Weil action. Move each translate into the source form by right invariance. The resulting section is the indicator with marked rows equal to the identity multiplied

by $\gamma = \prod_{v\in Q_\psi} |\mathrm{GL}_r(k_v)|$. The unmarked entries at each of these places give q_v^{rt} possibilities. The auxiliary mass is therefore

$$\frac{\gamma^2\ \mathrm{N}(Q_\psi)^{rt}}{q_0|D_\psi|^{(r+t)r}}.$$

The conversion from Petersson to the definite class-sum norm multiplies by the auxiliary level index and by the fixed dimension of the coefficient representation. The resulting factor is $A_j\Delta_\vartheta|D_\psi|^{-(r+t)r}$, where $A_j \neq 0$ is fixed and $\Delta_\vartheta = [K_{\mathrm{full}} : J_{\mathrm{aux}}]\gamma^2\,\mathrm{N}(Q_\psi)^{rt}/q_0$. At a new place the auxiliary level J_{aux} has index $G_{r+t}(q)/(G_r(q)G_t(q)q^{rt})$, by counting the fixes of the ordered two-block flag, where $G_k(q) = |\,\mathrm{GL}_k(\mathbf{F}_q)|$. The finite count consequently gives

$$\Delta_\vartheta|D_\psi|^{-(r+t)r} = C_{\mathfrak{M}}\,\mathrm{N}(Q_\psi)^{rt}\prod_{v\in Q_\psi}\prod_{h=t+1}^{t+r}(1-q_v^{-h}),$$

since the local Gram factor $G_{r+t}(q_v)/G_t(q_v)$ contains $q_v^{r^2+2rt}$, with $G_s(q) = |\mathrm{GL}_s(\mathbf{F}_q)|$. This is the factor when the common Hecke sum is written without its residue-matrix denominator. The character of the $r+t$ side in both lifts is $\mu_j^0\psi\,v^{(r+t)/2}$, since $r+t$ is even. At the places away from the modulus and its conjugate, the vacuum matrix coefficient for a source element γ' is consequently the same scalar

$$\rho_{j,\psi}(\gamma') = \left((\mu_j^0\psi v^{(r+t)/2})(\det\gamma')\,\Delta_f(\gamma')^{r+t}\right)^{-1}.$$

Here Δ_f is the finite Weil multiplier in the original vacuum convention. The change from the realized level-J_m chain to a definite class-sum metric occurs after this comparison.

Let K be the chosen level on either rank-r group. With x, y fixed by K, the matrix coefficient of $\rho_V(\gamma')$ can be replaced by that of

$$\frac{1}{[K : K\cap\gamma' K(\gamma')^{-1}]}\sum_{k\in K/(K\cap\gamma' K(\gamma')^{-1})}\rho_V(k\gamma' k^{-1}).$$

Indeed each summand has the same pairing with x, y, and changing the chosen representative of k changes the action on y by an element of K. This operator on K-fixed vectors is the normalized double-coset operator. The common finite adelic identification uses γ' supported off the chosen places and the twist places. The coefficients $\rho_{j,\psi}(\gamma')/|D_\psi|^{(r+t)r}$ of the normalized double cosets agree in the two expressions. The unfolded integrals converge absolutely on each matrix entry. Since both base spaces are finite dimensional, the series lies in the closed paired image of the off-level Hecke algebra. Its operator $P_{j,\psi}$ thus acts on the good spaces in both inner forms, commutes with the chosen finite group and preserves the good subspaces. The two Gram forms $q_0\langle a_{i,0}e_i x, a_{i,0}e_i y\rangle$ and $\langle u_\psi x, u_\psi y\rangle$ are the actions of $P_{j,\psi}$ with scalars C_i and $A_j\Delta_\vartheta$, respectively, for $C_i \neq 0$ independent of ϑ. The good trace transfer gives equal determinants of $P_{j,\psi}$ on the two good spaces. In particular, if S_i is the good first-arrow Gram operator with the factor q_0 included, then $S_i = q_0 B_i$ on that good base and

$$\det S_i = (C_i/(A_j\Delta_\vartheta))^R\lambda_j(\psi).$$

At the trivial twist the auxiliary lift $W_j \to I_{6r}$ has this same point datum at the specified modulus $\mathfrak{M}$, as follows by Petersson adjunction with the definite class-sum metric. Its finite average at the trivial character is the single indicator with modulus $\mathfrak{M}$. Its off-modulus datum is the vacuum, its source rank is r, and the real scalar uses the actual joint harmonic for this edge. The argument just proved gives injectivity on the full principal-level source, which supplies the descent of a good form in the holomorphic peak argument.

9.4 The auxiliary lifts from the norm-one torus

The rank-one norm at a twist place. For a rank-one pair p of point z_p, denote the lift of the constant source form on $U(1)$ by F_ψ. It has the compact jet and vacuum scales of this pair. The definite target has rank $n_p = 2z_p \geq 4$. The splitting characters are 1 and $\mu_p(\psi)v^{z_p}$, with the point character $\mu_p(\psi)$ in the signed catalogue. The source $U(1)$ is compact modulo rational points, the constant source form is cuspidal, and $n_p > 2$, so the convergent doubling norm formula applies to the prescribed finite section and compact Fock vector. The split support calculation applies to the first-coordinate projection. At a split place newly in the conductor, the section has first coordinate 1 in the residue field and q^{n_p-1} choices for the other residues. Its squared norm relative to the vacuum is q^{-1}. The matrix coefficient is supported on source units reducing to 1, as both first coordinates must equal 1. The measure of this support is $(q-1)^{-1}$ with the maximal compact of volume one. The finite projection on the constant source multiplies the lift by $q-1$, because its $q-1$ summands are equal after translating the constant function. The target level preserves a hyperplane of codimension one and has index $(q^{n_p}-1)/(q-1)$, by transitivity on hyperplanes. Consequently the factor in the normalized class norm is

$$\frac{q^{n_p}-1}{q-1}\,(q-1)^2 q^{-1}\,\frac{1}{q-1} = q^{n_p-1}(1-q^{-n_p}).$$

Off the chosen places and the twist, the split vacuum has overlap $q^{-n_p|a|/2}$ under a source translation of valuation $a \in \mathbb{Z}$. This follows by intersecting the standard lattice with its scalar translate in the split Schrödinger model and inserting the square-root determinant in the Weil action. Its character factor is u^a for the unitary parameter u in the local Rallis splitting. The two geometric series give

$$\sum_{a\in\mathbb{Z}} u^a q^{-n_p|a|/2} = \frac{1-q^{-n_p}}{(1-uq^{-z_p})(1-u^{-1}q^{-z_p})}.$$

The two denominator factors come from the inverse of the base change of the weight character $\mu_p(\psi)$ of this rank-one pair. The Weil splitting in the oscillator overlap has already been included in $q^{-n_p|a|/2}$ and in u. At an unmarked nonsplit place the norm-one group is compact and the unramified vacuum coefficient is one on it. Thus its normalized local factor is one. At a fixed chosen place the section has all n_p residues prescribed, the target index is the number $q^{n_p}-1$ of nonzero vectors, and the coefficient has support in source units reducing to one. Changing the source integral to its principal-level class sum multiplies the lift by $q-1$. The fixed place thus contributes $(q-1)(1-q^{-n_p})$.

The product of the spherical numerators, the compact real volume and the real harmonic norm gives a common nonzero factor. In the unscaled kernel normalization the real harmonic factor is the hook number times the coefficient dimension. The jet and vacuum scales multiply it by the same positive quantity for every twist. The relative norm consequently has the ratio of the incomplete Hecke values of $\mu_p(\psi)^{-1}$, with the factor q^{n_p-1} at each new conductor place. Absolute convergence at $z_p > 1$ gives nonzero Hecke products. In particular $F_\psi \neq 0$ and

$$\rho_p(\psi)^{-1} = \frac{\|F_1\|^2}{\|F_\psi\|^2}\,\mathrm{N}(Q_\psi)^{\,n_p-1}.$$

The same calculation applies to the flat pairs of ranks 4 and $2z_p - 20$ used for the rank-one test.

9.5 The spherical polynomial on the auxiliary lifts

We use the same split spherical polynomial on the stage lift and on the lift of the constant function. For stage j, let the auxiliary source be W_j of rank $r' = r$ and the target I_n with $n = r + t$. Put $\mu = \mu_j^0\vartheta$, where ϑ is a trial finite-order norm-one character of the allowed added conductor. On the sharp line take the

constant source on I_1, so $r' = 1$, the target I_{2z}, and μ equal to the point character times the trial character. In each case write $t' = n - r'$, let r_- be the number of negative source signs, and use the Weil characters

$$a = v_{U(1)}^{r_-}, \qquad b = \mu\, v_{U(1)}^{n/2}.$$

The stage calculation uses the adjoint of the stage descent. Indeed, if its kernel is $K(g,h)$ and the scalar product on the definite target is the class sum with weights $1/|\operatorname{Aut}(h)|$, then the adjoint sends u to $h \mapsto \int K(g,h)^* u(g)\, dg$, with the source Petersson measure. Its transposed conjugate coefficient and the stage highest-vector datum give the bilinear theta pairing. The sharp kernel has the constant source and the one-row jet coefficient. At a split pair $w, \bar{w}$ off the support of the finite test data and the characters, both kernels therefore have the integral-box vacuum and volume-one spherical compact subgroups used in the rectangular vacuum identity.

Put $q = q_w$ at such a place. The two splitting scalars and the argument change for the source polynomial are

$$A_w = q^{r'/2} v_E(\varpi_w)^{2r_- - r'}, \qquad B_w = q^{n/2} \mu^{\mathrm{bc}}(\varpi_w), \qquad C_w = q^{t'} A_w / B_w.$$

Here $(v_{U(1)})^{\mathrm{bc}} = v_E^2$, since base change evaluates at $x/\bar{x}$ and $v_E(\bar{x}) = v_E(x)^{-1}$. For the monic spherical source and target polynomials P_w, Q_w, write $\check{P}_w(T) = T^{r'} P_w(1/T)$ and $\check{Q}_w(T) = T^n Q_w(1/T)$. On a nonzero lift of a spherical source eigenvector the rectangular identity gives

$$\check{Q}_w(T) = \prod_{i=0}^{t'-1} (1 - A_w q^i T)\, \check{P}_w(C_w T).$$

For the constant source $\check{P}_w(T) = 1 - T$. A fixed scalar in the realization by the class sum multiplies both sides of the operator identity by the same number. In particular the operator $H_w(\xi_w)$ given by the target reciprocal Hecke polynomial at $\xi_w = (A_w q^{t'-1})^{-1}$ kills the entire stage image and the sharp line. This last statement follows from the operator identity on every spherical source vector before taking a Hecke eigenbasis.

10 Complementary periods and the two chains

10.1 An exchange identity for complementary square integrals

Complementary square integrals and the realized commutations. In a square let the two exterior spaces have orthogonal sum W_* of even rank N, and let the intermediate spaces have ranks a, b with $a + b = N$ and $e = a - N/2$. Their negative sign counts add to the count q_* of W_*. To define partners over the residue ring R_* in the order $b \oplus a$, take $k \in \mathrm{GL}_N(R_*)$ and write $Y_k = k\binom{I_b}{0}$, $X_k = (0\ I_a)k^{-1}$. The two finite functions are supported on these frames, with $\Phi(X_k) = \chi(\det k)\Psi(Y_k)$, where χ is the residue character of ν. The product of the edge characters at spaces whose ranks add to N must be $\nu^{-1} v^{N/2}$. The weight μ of ν satisfies $N \le 2\mu + 2q_*$ and $N \ge 2q_* + 2$. Let S be the nonempty split support of the square modulus, disjoint from its conjugate. The local determinant characters on S are nontrivial of conductor one. For $e \ge 0$ put

$$\Pi_{\nu,S}(e) = \Big(\prod_{v \in S} q_v\Big)^{e^2} \prod_{l=1}^{e} |L^S(l, \nu)|^2.$$

Complementary integration interchanges the two kernels on matched data with scalar c satisfying $|c|^2 = C\Pi_{\nu,S}(e)$, where $C > 0$ is fixed in the twist. The constant for the map between form spaces is κ_*/c, with κ_* independent of the twist. The product for $e = 0$ is one. For $e < 0$ take $\Pi_{\nu,S}(e) = \Pi_{\nu,S}(-e)^{-1}$.

Primitive support and integration in the split model. All measures in this calculation are the measures of the fixed Weil datum. In the finite named model take a nonempty set S of split places disjoint from $\bar{S}$. Fix the finite Bruhat–Schwartz function Ψ in that model and the vacuum at the other places. Write G for the external hermitian space, H for the integrated one, and $Y : G \to H$ for a matrix in the named coordinates. Suppose Y is row-surjective at one named place wherever $\Psi(Y) \neq 0$. Put $h_G(g), h_H(h) \geq 1$ for the heights from the moved lattices and write $M_j(v) = \int |v(u)|(1+\|u\|)^j \, du$ for the Fock moment, with the normalized coherent-state measure. Then the kernel is bounded by

$$|\Theta_{\Psi,v}(g,h)|\, h_H(h)^B \leq C M_j(v) h_G(g)^A$$

for every prescribed $B \geq 0$, for suitable A, j, C. The constants here may depend on the fixed finite function. For the initial bound use reduction with the inverse first minimum as height. A real move of operator norm at most two halves heights above the fixed cutoff by expanding the short isotropic subspace and contracting its dual. After $O(\log h_G + \log h_H)$ moves, Mahler compactness gives rational translates in fixed compact sets. Operators of the compact set of real moves have a common norm bound in the coherent model, so the norm of the reduced operator is bounded by a fixed power of the heights. The rational translates leave the theta functional unchanged. This gives the initial bound for Λ on the coherent datum norm.

Here is the short-shift argument for this estimate. For large $h_H(h)$, take a nonzero short rational vector d in the moved integral lattice of H, of real length $O(h_H(h)^{-1})$. Its hermitian norm is integral and has all real absolute values less than one, so d is isotropic. Take an E-basis C of G in the moved integral lattice with real lengths $O(h_G(g)^{A_0})$. Indeed the moved lattices have fixed covolume, and the product bound on the successive Euclidean minima bounds the last minimum by a fixed power of the inverse first minimum. A full-rank set over $\mathbb{Q}$ contains an E-basis. In the original rational coordinates the shifts are dzC^*J_G, with z an adelic row. All lie in one isotropic space of $\mathrm{Hom}(G,H)$. The pullback $\Lambda = \Theta \circ \omega(g,h)$ is invariant under rational parameters z, and the shift operators T_z satisfy $T_z T_{z'} = T_{z+z'}$, since the Heisenberg pairing on this space is zero. Define $\mathcal{P}$ on these invariant functionals by $\mathcal{P}\Lambda(f) = \int_{[z]} \Lambda(T_z f)\, dz$ with probability Haar on the additive quotient. It is idempotent and $\mathcal{P}\Lambda$ is invariant under every adelic parameter.

At a selected split pair, restrict z to the half conjugate to the matrix coordinates. In the fixed named model this multiplies the finite function by an additive character of the pairing with Y. In the transported coordinates d', C' its linear coefficient is $d'^* Y C'$, up to the fixed invertible Weil and frame factors. It is nonzero because $d' \neq 0$, C' is invertible and Y is surjective at the selected place. In particular $\mathcal{P}\Lambda(f) = 0$ for the datum in question. To see this directly, each point of the compact support has a modulation whose character is different from one on an open neighborhood. Refine a finite subcover to a locally constant partition. The invariant functional vanishes on each summand by that modulation. This argument also works with the columns C restricted to a frame on which YC' is surjective.

Choose representatives for $[z]$ in a fixed bounded real set with finite part in $D\widehat{\mathcal{O}}_E$, where D is a fixed integer large enough to fix the given finite vector under the shifts. Such representatives exist by additive approximation and the compact real quotient of the resulting fractional lattice. The moved finite columns of d, C are integral before this multiplication by D. Their real shifts have length at most $C_0 h_G(g)^{A_0}/h_H(h)$. Since $(1-\mathcal{P})^m = 1 - \mathcal{P}$ for $m \geq 1$, the value $\Lambda(f)$ equals the integral of $\Lambda((1-T_{z_1})\cdots(1-T_{z_m})f)$ over these representatives. For shifts of length at most one, expand each coherent vector $\rho_\infty(u)v_{\mathrm{vac}}$ using the Heisenberg law and the fundamental theorem of calculus in each shift. The norm of the mixed derivative of order m is bounded by $C_m(1+\|u\|)^m$. The bound follows by differentiating the Gaussian translate and its phase, with translation lengths at most m and all directions normalized. Integrating the coherent expansion gives the product of the shift lengths times $C_m M_m(v)$. Use the moderate-growth bound for the theta functional in this coherent-state norm, with a fixed height exponent A_1. Taking $m > A_1 + B$ absorbs the height of h. If the bound on a shift exceeds one, then $h_H(h) \ll h_G(g)^{A_0}$, and the initial moderate-growth bound gives the assertion with a larger A.

For a Fourier coefficient along a fixed isotropic frame K with retained frame I, one also needs the sum of absolute values of the rational slice terms. Use finite sums of the slice functionals with unit phases chosen to make their values nonnegative on the fixed datum. The initial operator bound for these sums is uniform in the finite subset and the phases. In fact an adapted coherent vacuum has slice norms bounded by a Gaussian in the real variables and a box indicator in the finite variables. Their translates have uniformly bounded sums on the rational lattice, by periodizing the Gaussian and the compact box. The coherent integral gives the same bound times the datum norm for an arbitrary vector. Rational reduction in the integrated variable permutes the rational slices, giving the initial height bound for the sum of their absolute values. For this functional take C to be a common integral multiple of $(K \mid I)$, in rational coordinates. The shifts kill K, so the rational invariance of each slice functional follows from rational invariance of the retained theta functional. The preceding argument bounds every finite sum uniformly in h if the restriction of Y to this transported frame is row-surjective. This gives the absolute slice sum bound on the finite-volume quotient. For the top row datum this restriction is invertible by the stacked determinant. For the top column datum its first square block is invertible at the conjugate places. At a parabolic point the frame changes by an invertible block triangular matrix. For the lower row test the full row rank follows from invertibility of the stack of P and Y. These are the tests on which termwise integration and the cone estimate below are used.

The two realized line models. In the ordered frames $W^+ = L \oplus W \oplus L^*$ and $V^+ = B \oplus V \oplus B^*$, dual isotropic generators pair to one. The center parameter of the outer line radical multiplies the term indexed by $x \in V^+(E)$ by $\psi(z\, x^* J_{V^+} x)$. The radical coordinates here are those with central parameter twice the matrix parameter of the general center formula. The central mean is therefore the indicator $x^* J_{V^+} x = 0$. For $x = 0$ the remaining mean is the retained theta functional at the zero slice. For $x = \xi b_0 \neq 0$, where b_0 generates B in a rational Witt frame, the other radical coordinate u translates the first slice by $\xi b_0 u^*$ in the retained hom space. The map $u \mapsto \xi u^*$ is an automorphism of the corresponding additive quotient. The normalized zero term of the seesaw calculation therefore identifies this mean with the second zero projection, along $\mathrm{Hom}(W, B)$. At fixed integrated point the slice norms are summable and the radical phases have modulus one, so the expansion is termwise integrable over the compact radical. The primitive-support bound is uniform in the integrated group on this compact family of external points.

For clarity the two zero projections can be read in the same hom space before using the finite tests. The trace hermitian form pairs a map from block L_1 to block B_1 with the map from its dual block to the dual of B_1. In the lines just chosen, the common isotropic summand is

$$\mathrm{Hom}(W \oplus L^*, B) \oplus \mathrm{Hom}(L^*, V).$$

Its dual is $\mathrm{Hom}(W \oplus L, B^*) \oplus \mathrm{Hom}(L, V)$, as one sees by pairing the matrix units. The orthogonal retained piece is $\mathrm{Hom}(W, V)$ together with the hyperbolic plane $\mathrm{Hom}(L, B) \oplus \mathrm{Hom}(L^*, B^*)$. After the common projection the last two models keep opposite coordinates of this plane. Thus they are related by the self-dual Fourier transform on $\mathbb{A}_E$, with the dual variable read through conjugation. This is the plane on which the Tate calculation below takes place. The split basis is the original adapted integral frame, so its finite scaling and the real Gaussian are the ones in the normalized zero-term formula.

For the standard row datum PI'_A, the zero lower matrix annihilates the added first line together with the fixed rows, so its full stacked determinant is zero. Nonzero isotropic vectors form the one orbit of rational isotropic lines by Witt extension. The stabilizer of B has radical and Levi $E^\times \times U(V)$ on rational points. The line terms are invariant under its adelic radical because the two zero terms are invariant under the corresponding unipotent operators. The line-orbit sum unfolds to the quotient by that stabilizer. In Iwasawa coordinates the Levi sends the generator of B to λ^{-1} times itself. Its action on the radical column has modulus $|\lambda|^{-a}$ and that on the central E^+-coordinate has modulus $|\lambda|^{-1}$. This gives the quotient factor $|\lambda|^{a+1}$, with one fixed positive change of Haar measure. The two line arguments are respectively λu and $\overline{\lambda u}^{-1}$ at an external line scalar u. The two determinant characters and the zero-term dilation on the

ambient side give the Weil multipliers of the next calculation. In particular the formula at arbitrary u is used here only when the lower two ranks are equal, so that the extra modulus $|u|^{a+1}$ combines with the Iwasawa factor and u can be removed by multiplicative translation.

On the column side the upper function is read on transposes at $\bar{S}$, with the W-projection of the adapted frame. In this reading the transpose of the dual projection at S is the inclusion at $\bar{S}$, by the hermitian frame relations. The finite column datum after a slice at x is consequently the lower $\chi^{-1}(\det(PY^{\mathrm{t}}))$ on integral Y, multiplied by the vacuum factor if the finite component of the scaled slice argument is integral, and zero otherwise. In the nonzero isotropic term the remaining theta functional is invariant under translations along that line of the integrated space, after transport by the same integrated point h. Its local action at a split pair is a modulation with coefficient $x'^{*}Y$, where $x' \neq 0$ is the moved line vector and Y has full row rank on the retained named datum. The partition argument above makes this functional zero there. Thus precisely the zero slice remains. The rank twist loses the splitting character υ for the removed plane, and the constant term at $(1, g)$ is the lower integral on the real zero slice times the fixed nonzero factor κ_0 of the adapted vacuum.

The signed coefficient at the top of the tower. Use $W_1 = \mathbb{H} \perp W_*$, of even rank $n_1 = 2j$, and $X = \mathbb{H}^b \perp W_1$, with adapted splitting D over W_1. For the small space at the top take $U = V_{\mathrm{small}}^{-}$ with its form negated. Thus the other space at the top is $U \perp W_1$. If V_{small} has q_b negative signs, U has $b - q_b$, and this top intermediate space has $b - q_b + q_* + 1$ negative signs. Removing $b + 1$ planes on the large side leaves $q_* - q_b$, the sign count of V_{large}. The rank difference at the top is n_1, so the signed excess is j.

In the integral frame (L, K, I) of D, put $P = L'$, the dual row block. Right multiplication of the stacked row matrix by the frame makes its determinant the determinant on the (K, I) columns, times an integral unit from the frame and the prescribed row order. Call its character phase ϵ. On the small side the same L' reads the K columns at the conjugate split places when the data are transposed. At a point $n(c, z)m(q, g)$, the central character of the rational slice at a frame x reads its hermitian Gram matrix $x^{*}Jx$. The coefficient with index ${}_JT$, for the chosen purely imaginary Weil element J, is the sum of the orbit integrals of the slices with this Gram matrix. Here $\psi(\frac{1}{2}\operatorname{tr}(zT))$ is the common radical phase in the unscaled matrix parameter. Its normalization is fixed by the normalized mixed-model zero term. The absolute slice estimate just proved gives the needed integrability.

Write J_U for the form of U, x_U for its adelic frame and α for the fixed scale idele. A self-dual test is

$$T = J_U, \qquad q = \alpha x_U B, \qquad B_f k \in U(1_b)(\mathbb{A}_f) \quad \text{for some } k \in \mathrm{GL}_b(\widehat{\mathcal{O}}_E).$$

Here $U(1_b)$ is the group of the finite form 1_b. Witt extension gives a single rational orbit on each side, represented by $\left(\begin{smallmatrix} 1_b \\ 0 \end{smallmatrix}\right)$ and 1_b. Their stabilizers are $U(W_1)$ and the trivial group, respectively. On the large side unfold first through $U(U)\times U(W_1)$ in $U(U \perp W_1)$. If a finite translate does not lie in the product of this block group with the standard integral subgroup, the term vanishes. Indeed a nonzero value requires that the moved first columns be integral. Right multiplication by k then gives an integral orthonormal frame. Integral Witt extension carries it to the first standard columns, and an isometry fixing those columns acts in their orthogonal complement. This is the stated block-group support.

After the allowed finite translate we can use a block point $y \oplus y_1$. Put $B_y = y^{-1}B$ and $x = (B_y)_f$. Both slices require integral x. Since xk is unitary when x is integral, the determinants of x and x^{-1} are both integral. In particular $\det x$ is a finite unit. On the large side the determinant of $\left(\begin{smallmatrix} x & * \\ 0 & Y_1 \end{smallmatrix}\right)$ factors as $\det x \det Y_1$. The retained finite vector is the vacuum for (W_1, U) times the signed determinant datum for (W_1, W_1), multiplied by $\epsilon\chi^{-1}(\det x)$. On the small side it is the vacuum times $\chi^{-1}(\overline{\det x})$. The finite unit identity relating the two readings is

$$\chi^{-1}(\det x)\widetilde{\nu}_f(\det x) = \chi^{-1}(\overline{\det x}), \qquad \widetilde{\nu}(\lambda) = \nu(\lambda/\bar{\lambda}).$$

The integrated characters on $U(U)$ differ by $\nu(\det y)^2$. In the product Weil action this is the factor from $\widetilde{\nu}(\det y)$, since $\det y$ has norm one. The common radical phase and the retained Heisenberg translate $x_U^{-1}c^*J_{W_1}$ agree on the first block.

Real slices. Let v be the prescribed real vector on (W_1, U) and w a real vector on (W_1, W_1). In the real chart put $J = \operatorname{diag}(J_U, J_{W_1})$ in sign coordinates and $Z_0 = (1_b, 0)^{\mathsf{t}}$. Near its nondegenerate b-plane use the smooth projection

$$P(Z) = Z(Z^*JZ)^{-1}Z^*JZ_0.$$

It depends only on the plane of Z. Near Z_0, choose a smooth unitary $k(P)$ taking the plane of Z_0 to that of P. This follows by the local section of the orbit submersion on nondegenerate planes. Take a nonnegative smooth compactly supported bump σ near Z_0, with $\sigma(Z_0) = 1$. For the small vector prescribe the real slice at B to be

$$\kappa\,\sigma(Z_0B)\left|\prod_\tau(\det B_\tau)^{j+\mu_\tau}\right|^2 v.$$

The scalar κ is chosen using the integral of the displayed bump and determinant weight, which is positive. With the Fourier measure of this mixed model the slice at zero in the opposite splitting is then v. For the large vector put $Z' = k(P(Z))^{-1}Z$ and prescribe the slices

$$\kappa\,\sigma(P(Z))\,\sigma(Z')\prod_\tau\overline{\det Z'_{1,\tau}}^{\,2\mu_\tau}\,\omega_\infty(1, k(P(Z)))(v\boxtimes w).$$

Here Z'_1 is the first square row block and the real target weight in ω_∞ is $j-\mu_\tau$. Define the slices by zero outside the chart. The two cutoffs make the support compact in Z, with rank bounded away from degeneration, so the extension is smooth. In the mixed model these are compactly supported smooth families of retained Schwartz vectors. Inverse partial Fourier transformation produces the required vectors with all Fock moments finite. Every nonzero large slice has real Gram matrix congruent to J_U.

To compare the periods, set $Z = h_\infty^{-1}Z_0$. If the bump at $P(Z)$ is nonzero, $h_\infty k(P(Z))$ preserves the plane of Z_0 and is block diagonal. A left translate by its inverse in $U(U)\times U(W_1)$ cancels it. This translate also aligns the slice at ZB_y with the slice at $Z_0(B_y)_\infty$, since $P(ZB_y) = P(Z)$, with the first block of the left translate included in the former slice argument. Hence the two real factors have ratio

$$\frac{\rho(h_\infty)}{\widetilde{\nu}_\infty(\det(B_y)_\infty)\,|\det(B_y)_\infty|_{E,\infty}^{j}},\qquad \rho(h_\infty) = \sigma(P(h_\infty^{-1}Z_0)).$$

The determinant relation follows place by place from

$$\overline{d}^{\,2\mu_\tau} = \frac{|d|^{2(j+\mu_\tau)}}{(d/|d|)^{2\mu_\tau}|d|^{2j}}.$$

In the remaining block the integral of w is the constant equal-rank period $I_\nu(w)$ on W_1. Its constancy uses the one-plane formula as explained below.

The quotient stages supply a fixed positive factor R. Indeed the integral over the block subgroup, after quotient by the rational stabilizer, is unchanged by its left block translates. The weight ρ depends only on the plane of $h_\infty^{-1}Z_0$, and the finite indicator is also left block invariant. On the real quotient the support of ρ is in the image of the compact set of representatives $k(P)^{-1}$ over the bump. The finite support is the image of the standard compact subgroup. Consequently the quotient weight has finite mass, positive because it is positive near the identity coset. Its mass is R for every character, including a zero value of the common small coefficient. Writing $d_0 = \det(\alpha x_U)$ and κ_D^+, κ_D^- for the two fixed vacuum factors, the multiplier in the original coefficient identity is

$$C_{\mathrm{top}} = R\,\frac{\kappa_D^+}{\kappa_D^-}\,|d_0|_E^j\,\epsilon\,\widetilde{\nu}(d_0)\,I_\nu(w).$$

In fact the Levi modulus ratio is $|\det q|_E^j$ and its character ratio contributes $\widetilde{\nu}(\det q)$; $\det q = d_0 \det B$. The real determinant quotient, the finite unit identity and the $U(U)$-character adjustment cancel the factors containing B_y and y. This gives the displayed scalar including the conjugate finite reading. Its squared modulus is a fixed positive multiple of $\Pi_{\nu,S}(j)\mathcal{R}_\nu(w)$ with $\mathcal{R}_\nu(w) = |I_\nu(w)|^2/\Pi_{\nu,S}(j)$.

For an invertible T at a real place outside the class of U, every small orbit is empty. With the supported large vector just constructed, every large slice at that index is zero. One also needs the all-negative auxiliary case for arbitrary real input during the equal-rank nonvanishing argument. At a place where the Gram form of a frame of rank b has a positive direction, move that direction to a positive coordinate by a real isometry. The projection of the frame to the negative rows has rank at most $b-1$. There are $b+q(W_1)$ negative rows. For a coherent-state label, require in addition orthogonality to its $q(W_1)$ columns indexed by negative signs of W_1. A unit negative row remains. Rotation of its line by $\bar{\eta}$ fixes the frame. Its vacuum exponent on the remaining oscillator is

$$n_1 - q(W_1) - (j - \mu_\tau) = j - q(W_1) + \mu_\tau > 0$$

since $q(W_1) \le j$ and $\mu_\tau > 0$. The coherent label has zero coordinates in the terms on which the rotation would contribute negative powers of η. Its circle mean therefore vanishes after that vacuum factor. The compact stabilizer average is zero on all coherent vectors and hence on every Schwartz vector by the coherent-state integral. Haar on the stabilizer is right invariant under its compact subgroup, so the unfolded orbit integral vanishes. This also covers the alternative sign hypothesis in the top coefficient calculation when the real bump restriction is absent.

Propagation from the self-dual tests. Let D be the difference of the large integral and C_{top} times the small integral. Write Λ and I_{small} for those integrals. At a place $w \in S$ let $\tau_w(x) = \nu((x, x^{-1})) = \widetilde{\nu}_w(x)$, in the chosen split coordinate. Its restriction to integral units is χ_w. Use the full character, with its value at the uniformizer, and put

$$\eta_w(h) = \tau_w(\det h)|\det h|_w^{N/2}.$$

Here N is the ambient rank for the two complementary integrals being compared. Multiplying over the selected pairs gives $\eta = \prod_{w\in S}\eta_w$. The factor $|\det h|_w^{N/2}$ is the square root of the additive Jacobian of $X \mapsto hX$ on $M_{a,N}(E_w)$. In the split vacuum convention the unitary factor of the underlying action on the position space T_w is $\upsilon_{E,w}(\det T_w)^{-1}$. Indeed, the determinant valuations of the intersection lattices for T^{-1} in the two halves differ by $-\upsilon_w(\det T_w)$, and the two local values of υ_E are inverse under the split identification. On the large integrated group, postcomposition by h' has phase $\upsilon_{E,w}(\det h')^{-N}$. Multiplication by $(\nu^{-1}\upsilon^{N/2})(\det h')$ leaves $\tau_w(\det h')^{-1}$, since $\upsilon((x,x^{-1})) = \upsilon_{E,w}(x)^2$. The action with pullback $\Phi(hX)$ uses $h' = h^{-1}$. On the small group the character adjustment by ν^2 in the opposite coordinates gives $\nu\,\upsilon^{N/2}$ on the integrated group. The element with coordinate l in the chosen half has coordinate $l^{-\mathrm{t}}$ in the other half. Transposing the pullback in that half gives $\Psi(Yl)$, with factor $\eta_w(l)$. Right invariance of the two quotient measures therefore gives

$$\eta(h)\Lambda(\Phi(h\,\cdot)) = \Lambda(\Phi), \qquad \eta(l)I_{\mathrm{small}}(\Psi(\,\cdot\, l)) = I_{\mathrm{small}}(\Psi).$$

These identities use the original quotient measures.

Write $E_S = \prod_{w\in S} E_w$ and $\mathcal{O}_S = \prod_{w\in S}\mathcal{O}_w$, with $K_j = \mathrm{GL}_j(\mathcal{O}_S)$. Use Haar measures of compact mass one in the following pushforwards of the finite sections:

$$f_X(k) = \int_{\mathrm{GL}_a(E_S)} \eta(h)\Phi(hX_k)\,dh, \qquad f_Y(k) = \int_{\mathrm{GL}_b(E_S)} \eta(l)\Psi(Y_k l)\,dl.$$

The integrals are over compact subsets for each k, since the supports are compact in the full-rank loci. On primitive frames the pushforwards give the primitive reductions. For example $X\mathcal{O}_S^N$ is the lattice generated by the columns of X. On the support it equals $h\mathcal{O}_S^a$ for finitely many right cosets hK_a. The

part on each such coset is moved to primitive X by $X \mapsto hX$ and weighted by $\eta(h)$, by the measured row covariance. The compact average with the same covariance yields f_X; the column calculation uses right multiplication on the basis of the column frame. Thus the integrals with Φ, Ψ equal the integrals with their primitive reductions, with no extra measure factors.

For the block upper triangular parabolic $P_{b,a}$ stabilizing the first b coordinates, let u, v be the two diagonal blocks of p. Then $X_{kp} = v^{-1}X_k$, $Y_{kp} = Y_k u$, so

$$f_X(kp) = \eta(v) f_X(k), \qquad f_Y(kp) = \eta(u)^{-1} f_Y(k).$$

Put $\tau = \prod_{w \in S} \tau_w$ and $|\det g|_S = \prod_{w \in S} |\det g_w|_w$. Equivalently, the sections

$$|\det g|_S^{a/2} f_X(g^{-1}), \qquad |\det g|_S^{-b/2} \tau(\det g)^{-1} f_Y(g^{-1})$$

have the same left transformation law $\delta_{P_{b,a}}(p)^{1/2} \tau(\det v)^{-1}$, where $\delta_{P_{b,a}}(p) = |\det u|_S^a |\det v|_S^{-b}$. Consequently the integral partner relation on K_N extends, by $\mathrm{GL}_N(E_S) = K_N P_{b,a}(E_S)$, to $f_X(k) = \eta(k) f_Y(k)$. For any full-rank column test the relation specifies its primitive row partner from f_Y on K_N, and defines $\mathcal{D}$ linearly in that test.

To see covariance under the outer group, let β_X, β_Y be the exterior characters before the opposite-space reading, with $\beta_X \beta_Y = \nu^{-1} v^{N/2}$. In the latter reading the column exterior character is $\beta_Y^{-1} v^b$. For g at a chosen half the pullbacks $X \mapsto Xg$ and $Y \mapsto g^{-1}Y$ have Weil multipliers of modulus $|\det g|_w^{a/2}$ and $|\det g|_w^{-b/2}$, respectively. The ratio of their unitary factors is $v_{E,w}^{a-b}(\beta_Y^{-1} v^b)/\beta_X$, evaluated in the split coordinate. This is $\tau_w(\det g)$. Their full ratio is therefore $\eta_w(g)$, which is exactly the factor obtained by replacing k by $g^{-1}k$ in the normalized partner relation. The kernel of $\mathcal{D}$ is stable under the outer split group. Its image retains the standard integral level off the selected pairs. The determinant-one shears below have Weil multiplier one. For the scalar translate by $-\varpi^{-1}$ and the dilations $D_d k$, divide each normalized column translate by its nonzero Weil multiplier before summing. The multiplier may depend on k. The linear kernel therefore contains the bare pullbacks used in the support and multiplicity calculations.

On the exterior blocks of the triangle the character product uses the opening character $\alpha_{i,0}^{-1} v^{N'}$ at rank N' and the upper character of step $i + n + 2$. The exponent from their rank terms is $\lfloor a/2 \rfloor + \lceil b/2 \rceil$, or the same with a, b interchanged according to the exterior block. The exponents equal $N/2$ since $a + b = N$ and a, b have the same parity. This leaves $\nu^{-1} = (v^{-h} \chi_0 \vartheta)^{n+1}$, with $\vartheta = 1$ on the core triangle. On a square of the full bordered grid the horizontal and vertical rank exponents give the same floor and ceiling in V. In Z both ranks are even, so their integral halves also sum to $N/2$. The remaining product is $\mu_{j_1}^0 \mu_{j_2}^0 \vartheta^\ell$, for horizontal and vertical step indices j_1, j_2, with ℓ the twisted step. This calculation holds on each exterior block. For the full ambient group take exterior characters of product $\nu^{-1} v^{N/2}$. The difference between these characters and the original block characters contributes the same multiplier on both paths after the opposite-space reading. It therefore cancels on return to the original periods. In the auxiliary comparisons, put $\delta_N = \nu^{-1} v^{N/2}$. For the row pair $W_N \to H_a$ write its two characters (A_R, B_R); for the column pair $H_b \to W_N$ before taking the adjoint write (A_C, B_C). Use $A_R = B_C = \delta_N$ on the integrated groups and $A_C B_R = \delta_N$ on the exterior. On raising the rank by a plane use $\delta_{N+2} = v \delta_N$, increase B_R by the factor v, and keep A_C fixed. The zero-term action on descent restores the lower characters. Taking the opposite small space gives integrated character $v^N B_C^{-1} = \nu v^{N/2}$ and exterior character $v^b A_C^{-1}$. Thus at the top, with $a - b = 2j$, the row exterior character divided by the opposite small one is $\delta_N v^{-b} = \nu^{-1} v^j$. For the two rank-zero tests use $(B_R, A_C) = (\delta_N, 1)$ when $b = 0$ and $(1, \delta_N)$ when $a = 0$, giving the same characters (δ_N, δ_N) on the common equal-rank kernel. At an old chosen pair the unramified factor of the trial character still takes its full value on a uniformizer. At a new pair the unramified factors of the original stage characters are likewise retained. On K_N all modulus factors are one and only the specified residue character remains.

Retain the extra plane of $W_1 = \mathbb{H} \perp W_*$ outside the frame of D. Write $P = L'$ and let $\Psi_0(Y) = 1_{\mathrm{int}}(Y) \chi^{-1}(\det(PY))$, with its zero extension off the unit determinant locus. Let F be its image under $\mathcal{D}$.

At a point that fixes the D-frame at the selected pairs consider the two partial means of a coefficient F_T, one for each line of the extra plane. Partial averaging along the opposite line gives invariance under $Y \mapsto Y + uzPY$ in the linear functional of the finite datum. Here u is the extra line column in the moved frame at S, and z is any row. Only determinant-one shears occur. This implies invariance of the value at Ψ_0 under $Y \mapsto Y + urY$ for $Pu = ru = 0$. In detail, average both tests over $Y \mapsto Y + urHPY$, where H varies over integral matrices. If the affine space $Y + uZ$ meets the support, take $Y_0 = Y + uZ_0$ in that support. Then $H_0 = Y_0(PY)^{-1}$ is integral, $rY_0 = rY$, and the two averages differ by $H \mapsto H + H_0$. Otherwise both are zero. The averaging is a finite sum after passage to a sufficiently small open subgroup. The actual local unitary shear has its operation $1 + ur$ at S and its adjoint inverse at $\bar{S}$. The column test reads PY at the indicated half and $Pu = 0$. This proves the partial-mean shear invariance on both lines of the extra plane.

Write $\Xi = \jmath T$ for the skew-hermitian Fourier index and $\delta = \jmath\alpha\bar{\alpha}$ for the inverse-different idele in these coordinates. Consider an invertible index for which $\delta^{-1}\Xi$ is unimodular at S. It is then unimodular at $\bar{S}$, since this normalized matrix and its inverse are hermitian. Use the standard compact level over the complementary finite places. Let $\phi_g^+(a)$ be the mean of $F_T(n_-(b)n_+(a)g)$ over $b \in E^b\backslash\mathbb{A}_E^b$, where $n_\pm$ are the two shifts of the added line by a row linking it to the D-frame. Define ϕ_g^- with the lines exchanged. The commutator of the shifts is in the D-center. Moving g by $n_+(a)$ translates ϕ_g^+ by a and multiplies $\phi_g^-(b)$ by $\psi(-\operatorname{tr}(\Xi a^* b))$ in the coefficient convention. Rational translates of a read all Fourier coefficients in b, because Ξ is invertible. Thus ϕ_g^+ determines $F_T(n_-(b)n_+(a)g)$ and then ϕ_g^-, and scalar identities pass between the two means.

To straighten a point with a nonzero mean, use Iwasawa for the extra line and the D-frame. Right multiplication by an integral unitary matrix away from the selected pairs makes the action on this flag block upper triangular, with the D-block denoted by γ. The outer radical factor may be removed from the partial mean. For the center of the extra line this follows from invariance under a split local shear and additive approximation with its rational center. The partial mean then equals the coefficient at $0 \oplus \Xi$ of the enlarged frame. The rational radical acts on that frame by $\left(\begin{smallmatrix} 1 & \lambda \\ 0 & 1 \end{smallmatrix}\right)$, which preserves $0 \oplus \Xi$. The radical over the selected split half acts through the shear invariance just proved. Apply additive approximation in the radical column and then in its center to remove the whole factor. In particular straightening preserves a nonzero mean and gives a representative which scales the two extra lines by mutually dual scalars $t, t' = \bar{t}^{-1}$ and acts on the original frame by γ away from the selected pairs. At those pairs it still fixes the original frame.

Among such representatives with a nonzero partial mean in either orientation, choose one minimizing the denominator index of $\det Q_\gamma^{-1}$ away from the selected pairs, where $Q_\gamma = \delta^{-1}\gamma^*\Xi\gamma$. The matrix Q_γ is integral there. Indeed a nonzero mean gives a nonzero coefficient at a shifted point with the same original frame. Invariance under integral central shifts of the moved frame forces integrality by the integral additive conductor. The support of ϕ^+ now satisfies $\gamma^{-1}a^*t$ integral away from these pairs. To check this, straighten $n_+(a)g$. Its triangular frame before straightening has lower-left block $-a^*t$ and lower-right block γ. If γ' is the new lower-right block, comparing the two frames by the inverse integral right factor gives

$$-a^*t = \gamma' H_{21}, \qquad \gamma = \gamma' H_{22},$$

with H_{21}, H_{22} integral at these places. Nonintegrality of $\gamma^{-1}a^*t$ would force H_{22} to be nonunimodular. Both normalized Gram matrices are integral and $Q_\gamma = H_{22}^* Q_{\gamma'} H_{22}$, so the denominator index would strictly decrease. By symmetry the support of ϕ^- satisfies the corresponding condition on $\gamma^{-1}b^*t'$. For each column x of Q_γ^{-1}, the shear with $a = (\gamma x\ \overline{t'})^*$ supported at the complementary finite places has character one on this support of ϕ^-. This follows from the preceding support condition and the integrality of $Q_\gamma x$. It is therefore a period of ϕ^+, by the transfer of scalar identities between the two means. Comparing one nonzero point of ϕ^+ with its translate by that period shows x is integral. This makes Q_γ unimodular.

There is then a point with a nonzero coefficient and this unimodular Gram block, since $\phi^{+} \neq 0$ is a mean of coefficients at unchanged original frames. Glue the unimodular matrices at all finite places using the fixed frame at the selected pairs. For the field in the grid, which contains a primitive 5040-th root of unity, the finite unimodular hermitian form admits an integral orthonormal frame. At split pairs this uses the perfect integral pairing, and at the unramified fixed places it uses the residue norm and integral hermitian diagonalization. This is applied after removing the inverse-different factor δ. If the real class is wrong, the off-class vanishing already applies. Landherr's local–global theorem says that two nondegenerate Hermitian forms over a quadratic extension of number fields are isometric over that extension if they are isometric over every completion. The real class supplies the signatures, and the integral orthonormal frames give the finite local isometries with J_U. Apply that rational isometry to the original frame. The new coefficient is a self-dual test and was proved zero. This contradiction proves vanishing at all points trivial at the selected pairs for indices primitive there.[1]

Now use the integral opposite frame at the selected places. On the translates $\Psi_z(Y) = \Psi_0(Y + KzL'Y)$, a central Fourier functional at a point trivial at those places is multiplied by the character of z with the normalized Fourier index. On integral z the standard datum is unchanged. Thus nonzero functionals can only have integral normalized index, by the conductor. Put $\varpi = (\varpi_v)_{v \in S}$ and combine the tests $\Psi_{\varpi^{-1}\widetilde{r}}$ with weights $\chi^{-1}(\det r)$, over square matrices r in the product of residue fields. Use zero extension of the determinant character and take the lift of 1 to be 1. The Fourier multiplier of this sum is a matrix Gauss sum. At a singular frequency it is zero, since a stabilizer on a complementary line realizes every determinant and the character is nontrivial. At the remaining frequencies the primitive coefficient test gives zero. Fourier uniqueness on the center and weak approximation away from the selected places put the combination in the kernel.

The scalar translate by $-\varpi^{-1}$ of this combination is $r(Y)$ times the standard datum for the opposite row block, with $r(Y)$ a nonnegative integer. Indeed, if X is the sheared variable of a nonzero summand, then $Y = K\widetilde{r}(L'X) - \varpi X$. The opposite row block of Y reduces to r times the row block $L'X$. Each nonzero summand is therefore the value of that opposite standard datum. The term $r = 1$ is nonzero at $K - \varpi L$, whose opposite standard row block is 1.

Such a test generates all the required locally constant column functions in the stable kernel. Let Y_0 be its witnessing column and $P_0 Y_0 = 1$ its standard row block. A column partner is covariant by $\chi^{-1}(\det h)$ for right $h \in \mathrm{GL}_b(\mathcal{O}_S)$. Choose a depth d at which the given column function Ψ is translation invariant on integral matrices. The map $D_d = Y_0 P_0 + \varpi^{-d}(1 - Y_0 P_0)$ fixes Y_0 and dilates the kernel of P_0. Average the test $p(D_d k Y)$ over integral k with the weight $\Psi(k^{-1} Y_0)$. On its support

$$Y \equiv k^{-1} Y_0 (P_0 D_d k Y) \pmod{\varpi^{d}},$$

and covariance cancels the determinant character. Thus the average is $\Psi(Y)$ times the integral of $|p(D_d k Y)|$. The latter is constant on primitive columns by transitivity, and is positive at Y_0. Compactness and local constancy make this average a finite linear combination of translates. This proves the extension to standard rows and columns. The last averaging argument, starting from an ordinary standard column test in the kernel after descent, proves the identity on all partners. The depth d concerns the matrix functions and can exceed the conductor of the determinant character.

The normalized Tate factor for one plane. Write W for the lower ambient space of even rank N and V for the lower intermediate space of rank a. The unchanged intermediate space has rank $b = N - a$. Let A, X be the adapted plane splittings of $W^{+} = \mathbb{H} \perp W$ and $V^{+} = \mathbb{H} \perp V$. Use the prescribed row orders $\sigma : b \oplus a \longrightarrow N$ and $\sigma^{+} : b \oplus (a+2) \longrightarrow N+2$, where the integers in these orders denote their coordinate sets. The fixed b rows at the lower stage are an integral matrix P at the selected split places,

[1] W. Landherr, *Äquivalenz Hermitescher Formen über einem beliebigen algebraischen Zahlkörper*, Abh. Math. Sem. Univ. Hamburg 11 (1935), pp. 245–248, especially p. 247.

and the upper fixed rows are PI'_A, where I'_A projects onto the W-coordinates of the adapted frame. The standard row datum at either stage is the inverse determinant character of the stacked, ordered rows on the integral invertible locus, extended by zero.

Let ν be the square character of positive infinite weight μ, put $\widetilde{\nu}(\lambda) = \nu(\lambda/\bar{\lambda})$, and let S be the one-sided, nonempty split support of the tame part, disjoint from $\bar{S}$. Put $Q_S = \prod_{w\in S} q_w$. For $w \in S$ let $\chi_w \neq 1$ be the residue-field character read on the w-component, and identify $E_{\bar{w}}$ with E_w via conjugation. The finite character on the two unit factors is (χ_w, χ_w^{-1}). Thus $\widetilde{\nu}$ has conductor exponent one on both factors and is unramified off $S \cup \bar{S}$. Set $e = a - N/2$, so the upper difference is $(a+2) - (N+2)/2 = e+1$. Put $x = e+1$ for $e \geq 0$ and $x = -e$ for $e < 0$. For the two line models use $\beta = +$ in the first case and $\beta = -$ in the second. With the cone vector formed using this (β, x), the normalized one-plane multiplier γ has

$$|\gamma|^2 = C_{\beta,x}\, Q_S^{2x-1} |L^S(x,\nu)|^2 = C_{\beta,x} \frac{\Pi_{\nu,S}(e+1)}{\Pi_{\nu,S}(e)}, \qquad C_{\beta,x} > 0.$$

The constant can depend on the fixed spaces, splittings, infinite weight and measures. It is independent of the character of weight μ, the ramified set, the integral P, the real vector and the adelic point. This scalar evaluation uses the mixed-model line and Levi formulas and the primitive-data lower theta kernel bound stated below.

The split line calculation. In an adapted frame, a line term fixes a scalar s on one pivot and partially Fourier transforms the other removed coordinates. At $w \in S$, let M be the N-square matrix of the rows of P and of the variable lower matrix Y, stacked in order σ. The ambient and intermediate isotropic line coordinates put the upper stacked matrix, up to fixed integral invertible frame changes on the left and right, in the form

$$T_+ = \begin{pmatrix} s & z_W & z_0 \\ 0 & M & (0, z_V)^{\mathsf{T}} \\ 0 & 0 & t \end{pmatrix}.$$

Here z_W is a row of length N, z_0 is a scalar, and the last column of the middle block has b zero entries and a entries from z_V, in order σ. The other model uses T_- obtained by interchanging the two diagonal scalars s, t. These are the two mixed coordinates on the same removed plane. Indeed the rows of PI'_A vanish on both plane coordinates, the lower variable block maps W to V, and the remaining integrated maps range from W to one removed line, from the other removed line to V, and between the removed lines. The two models exchange the fixed and integrated maps at the two diagonal pivots. The full frames and the permutations σ, σ^+ change the determinant by an integral unit independent of the matrix arguments. On either model $\det T_\beta = s\, t\, \det M$. Integrality is entrywise and unit determinant requires each of $s, t, \det M$ to be a unit.

Write $D_w(t) = 1_{\mathcal{O}_w^\times}(t)\chi_w(t)^{-1}$ and use the additive Haar measure of volume one on $\mathcal{O}_w$ for the additive character ψ_w of conductor $\mathcal{O}_w$. In the line term at $\alpha = \alpha_0\lambda$, where α_0 is the fixed Weil scale idele, $s = \lambda_w$ and the Fourier frequency of t is $\bar{\lambda}_w$. All other Fourier frequencies are zero. Those variables give $\mathcal{O}_w$-volumes one; the determinant phase from the fixed frames has modulus one. The remaining transform is

$$\int_{\mathcal{O}_w} D_w(t)\psi_w(bt)\, dt = g_w\, 1_{\mathcal{O}_w^\times}(\varpi_w b)\chi_w(\varpi_w b), \qquad g_w = q_w^{-1} \sum_{u \in k_w^\times} \chi_w(u)^{-1}\psi_w(\varpi_w^{-1}\tilde{u}).$$

In Tate's local calculation the ramified test is the inverse of the multiplicative character on the units and vanishes off them. Its Fourier transform is supported at the complementary conductor, and there it is a Gauss sum times the multiplicative character. Here the multiplicative conductor is exactly $\varpi_w\mathcal{O}_w$ and the additive conductor is $\mathcal{O}_w$, so the complementary support has valuation -1. [29, §§2.4–2.5] For $b \in \mathcal{O}_w$ the character sum of D_w vanishes. If $v_w(b) < -1$, orthogonality on each class modulo ϖ_w gives zero. If $v_w(b) = -1$, summing the residue classes and changing variables by $\varpi_w b$ gives the formula. For the unnormalized finite Fourier transform of χ_w^{-1} extended by zero at zero, the value at frequency zero is zero

and all nonzero frequencies have the same modulus. Plancherel gives $(q_w-1)q_w^2|g_w|^2 = q_w(q_w-1)$, hence $|g_w|^2 = q_w^{-1}$.

For $p = (p_w)_{w\in S}$, define $F_p(\lambda_f)$ on the finite adeles by requiring integrality off $S \cup \bar{S}$ and requiring $\lambda_w, p_w\bar{\lambda}_w$ to be units for every $w \in S$. On that support put

$$F_p(\lambda_f) = \prod_{w\in S} \chi_w(\lambda_w)^{-1}\chi_w(p_w\bar{\lambda}_w);$$

give it value zero elsewhere. The fixed vacuum in the scaled coordinates gives the integrality condition off $S \cup \bar{S}$. The split calculation expresses the line kernel at $\alpha_0\lambda$ as $l_\beta G_v F_\varpi(\lambda_f)$ times the lower theta kernel with the same P and the real line slice at λ_∞. Here $l_\beta \neq 0$ is the fixed vacuum normalization of the adapted splitting and G_v is the product of the g_w with the unit determinant phase. Thus $|G_v|^2 = Q_S^{-1}$. For the real slice the upper vector is acted on by the real component of the compact argument k of $U(V^+)$ in the pair representation with character $v^{-1}v^{(N+2)/2}$ on both sides. The finite compact action fixes the datum because its determinant effect on the standard rows cancels the splitting-character effect on the units. At each removed plane the rank twist is $v^{-1}v^{N/2}$ on the lower pair.

Unfolding the line and its finite integral. Write $[V]$ for the quotient by rational points of $U(V)$ and K^+ for the fixed adelic maximal compact subgroup of $U(V^+)$, with its full compact Haar measure. For a lower $g \in U(W)(\mathbb{A})$, let $B_\beta(\alpha; h, k)$ denote the line kernel just described, with h in $[V]$ and $k \in K^+$. The line constant term at $(1, g)$ is the fixed multiple $c_{\mathrm{unf}} \neq 0$ of the integral at $z = e + 1$ of

$$\widetilde{\nu}(\lambda)|\lambda|^z \sum_{\xi\in E^\times} B_+(\xi\lambda; h, k)$$

on $E^\times\backslash\mathbb{A}_E^\times \times [V] \times K^+$. Here the line constant-term formula first selects rational isotropic vectors. Its zero orbit on the large side vanishes because the fixed rows kill the added first isotropic column, and a variable block in the zero slice kills that same column. The full stacked matrix then has a kernel. The other vectors are grouped by their line. Witt's theorem gives one orbit of isotropic lines, and unfolding along its stabilizer gives the Levi, lower quotient and compact integrations. In the coordinate λ the actual line Levi acts on its generator by λ^{-1}. Conjugation on its radical, whose coordinates are a column in E^a and a one-dimensional center over the real subfield, has modulus $|\lambda|^{-(a+1)}$. Quotient integration therefore contributes $|\lambda|^{a+1}$. The direct line Weil dilation contributes $\widetilde{\nu}(\lambda)|\lambda|^{-N/2}$, giving $z = e + 1$. For the opposite model its dilation contributes $\widetilde{\nu}(\lambda)|\lambda|^{-(N+2)/2}$, with line coordinate $\bar{\lambda}^{-1}$. The measures here are the fixed quotient measures of the two adapted splittings. For the equal-rank period calculation the same relation holds at (u, g), for $a = N$, with λu in place of λ. The two line-slice determinants in that case give a numerator $|u|^{a+1}$ in the Weil dilation, as required to replace $|\lambda|^{a+1}$ by $|\lambda u|^{a+1}$.

The two line kernels have zero term at the origin by the split support just computed. The mixed slices first take the common zero projection away from the remaining hyperbolic plane. The two scalar line variables are dual in that plane, and the lower theta functional commutes with the line Fourier transform. Adelic Poisson on this line, using self-dual measures, gives

$$\sum_{\xi\in E^\times} B_-(\xi\bar{\lambda}^{-1}; h, k) = |\lambda| \sum_{\xi\in E^\times} B_+(\xi\lambda; h, k).$$

The bilinear Fourier pairing on that line uses conjugation in the dual variable. Under $j_-(\lambda) = \bar{\lambda}^{-1}$ the norm inverts and $\widetilde{\nu}$ is unchanged, since $\bar{\lambda}^{-1}/\lambda^{-1} = \lambda/\bar{\lambda}$. Consequently the Mellin integral at $z = 1 - x$ uses the $\beta = -$ line kernel at exponent x in the j_--variable. The integral at $z = x$ uses $\beta = +$ and $j_+(\lambda) = \lambda$.

Choose the finite idele c with components 1 off $\bar{S}$ and with conjugate component ϖ_w at each $w \in S$. Then $F_\varpi(\lambda_f) = F_1(\lambda_f c)$. In the unfolded idele integral write the line argument as λK_v, with $K_v = \alpha_0 c^{-1}$. One has

$$|K_v| = |\alpha_0|Q_S, \qquad |\widetilde{\nu}(K_v)| = 1.$$

For $x > 1$, unfold the sum over $E^\times$ using the multiplicative quotient measure and split the idele measure in the variable j_β as the chosen real multiplicative Haar times a finite multiplicative Haar ρ_β. This splitting is independent of ν and translation by K_ν preserves the measure. The finite integral is

$$\int_{\mathbb{A}_{E,f}^\times} F_1(\lambda_f)\widetilde{\nu}_f(\lambda_f)|\lambda_f|^x \, d\rho_\beta = \rho_\beta(U_f)L^S(x,\nu), \qquad U_f = \prod_{w<\infty} \mathcal{O}_w^\times.$$

Indeed both components above a selected split prime are units, and their character values cancel F_1 on the unit ideles. At each remaining prime the valuation runs over $j \geq 0$ and its summand is $\widetilde{\nu}_w(\varpi_w)^j q_w^{-jx}$. The normalized coset volume is the same over every valuation tuple, namely $\rho_\beta(U_f) > 0$. Absolute convergence follows by majorization by the Dedekind Euler product for $x > 1$. Here L^S omits $S \cup \bar{S}$.

The real integral is the lower theta integral on the cone vector. It averages the real line slice at λ_∞ over $K^+ \times (E \otimes_{\mathbb{Q}} \mathbb{R})^\times$ against

$$|\lambda_\infty|_E^x \prod_\tau \left(\frac{\lambda_\tau}{|\lambda_\tau|}\right)^{2\mu_\tau}.$$

This fixes the real operator before the finite part of ν is chosen. The mixed-model slice of a Schwartz vector, after acting by k in the compact group, has every lower Schwartz seminorm rapidly decreasing in the line coordinate, uniformly in k. This follows by partial Fourier transformation and restriction in the removed variables. Use the primitive-support theta kernel bound with zero exponent on the height of h. It gives a bound by a fixed Fock moment, uniformly in h for a fixed g, P, ν. On the support of the lower datum Y has full row rank at the selected places, since the matrix formed by stacking P and Y is invertible. Thus the bound applies to this datum. In each complex radial variable the measure is proportional to dr/r, and the weight at zero is r^{2x}. The cone operator and the integral after applying the lower theta functional are therefore absolutely convergent and continuous locally for $x > 0$. These bounds justify Fubini for $x > 1$.

At $x = 1$ use the common Mellin integral on the idele quotient. For $R > 1$ its absolute convergence at R and $1 - R$ follows from the two unfolded formulas. For z in between, the weight $|\lambda|^{\Re z}$ is bounded by $|\lambda|^R + |\lambda|^{1-R}$, giving continuity at the interior points, including $0, 1$ for $R > 1$. For a unitary Hecke character, Tate's continuation allows a pole of the completed L-function at 1 only for the trivial character. The archimedean factors are finite and nonzero there, and omission of finite Euler factors multiplies by polynomials in their local norm powers. Both $\widetilde{\nu}$ and $\widetilde{\nu}^2$ have nonzero angular type, so their partial functions are holomorphic at 1. The trivial character gives the simple pole of the Dedekind zeta function, which persists after the indicated omissions. [29, §§4.4–4.5] If $L^S(1,\nu) = 0$, then for $\sigma > 1$ tending to one the expression

$$\left(\zeta_E^{S\cup\bar{S}}(\sigma)\right)^3 |L^S(\sigma,\nu)|^4 \, |L^S(\sigma,\nu^2)|$$

tends to zero by the simple pole of ζ_E. It is at least one by taking the logarithm of the Euler products and using $3 + 4\cos t + \cos(2t) \geq 0$ term by term. It follows that $L^S(1,\nu) \neq 0$. Thus the factorization at $x = 1$ uses the right limit of the partial Euler product. On either branch the finite and archimedean factors combine to give the complex scalar

$$\gamma = c_{\mathrm{unf}} I_\beta G_\nu \, \widetilde{\nu}(K_\nu)|K_\nu|^x \rho_\beta(U_f) L^S(x,\nu).$$

The factor Q_S^{2x-1} in its squared modulus is the product of Q_S^{-1} from the additive Fourier transform and Q_S^{2x} from the idele translate. For $e \geq 0$ the ratio of the two Π's is $Q_S^{2e+1}|L^S(e+1,\nu)|^2$. For $e < 0$, set $x = -e$ and use $\Pi_{\nu,S}(-j) = \Pi_{\nu,S}(j)^{-1}$. The ratio becomes $\Pi_{\nu,S}(x)/\Pi_{\nu,S}(x-1)$ and has the same formula.

Descent of the exchange scalar. Use the corresponding projected standard column datum with the same P on the other side. Its line constant term is the lower theta integral on the real zero slice times the fixed nonzero vacuum factor κ_0, by the primitive-support small-side constant-term identity. If the upper large integral equals c_+ times the upper small integral on the standard data, then $\gamma I_{\mathrm{large}} = c_+ \kappa_0 I_{\mathrm{small}}$ at the lower

stage. Thus $c_- = c_+\kappa_0/\gamma$. For an upper scalar squared modulus proportional to $\Pi_{\nu,S}(e+1)\mathcal{R}$, the lower squared modulus is proportional to $\Pi_{\nu,S}(e)\mathcal{R}$ with the same $\mathcal{R} = \mathcal{R}_\nu(w)$ from the fixed determinant test above. For the equal-rank period use $a = N$ and the equal-rank constant function identity. Its numerator $|I^+(1)|^2$ has the same factor, up to $C_{\beta,x}$ and the fixed radical volume, as the denominator $\Pi_{\nu,S}(N/2+1)$ when passing down to rank N.

Nonvanishing and the complementary-period norm. For an even-rank space Y with an adapted hyperbolic line write $I_Y(v)$ for its equal-rank theta integral on the signed determinant datum. The split equivariance with zero column rank makes $I_Y(v)$ right invariant under the full unitary group at a selected split place. The derived group is the simply connected, absolutely almost simple special unitary group. Over the selected split place it is the group $\mathrm{SL}_{\dim_E Y}$ with matrix size at least two, so its local points are noncompact. Kneser's strong approximation theorem for simply connected classical groups gives density of its rational points in the adeles away from that place. Equivalently, the rational points times the one local group are dense in the full adelic group. In applying this density to $I_Y(v)(sg)$, the local factor can be commuted past g, since its support stays at that place. Thus $I_Y(v)$ factors through the determinant. In particular it equals its line radical average. For equal lower ranks the one-plane computation has an external scalar u on the line. Its contribution combines with the quotient modulus to give $|\lambda u|^{a+1}$, so the substitution λu removes u from the constant term. The determinant of the ambient line Levi ranges over the norm-one ideles by Hilbert 90, choosing unit lifts at almost all places. The integral $I_Y(v)$ is therefore constant on the adelic group. The same argument applies to the two sign-compatible spaces reached by the local unitary transport after changing the hermitian form: the standard finite datum and the measures are transported, and the remaining Weil factor has modulus one. [2]

We first record the two real preimages used with the one-plane calculation. For the positive branch, the cone slice is read along the null generator b_0 of the added plane on the integrated side. For a compact argument k and real line variable λ, the moved column is $x = k^{-1}b_0\lambda$. Near b_0, the orbit of nonzero isotropic vectors has a smooth local unitary section $h(x)b_0 = x$. Indeed the derivative of the action, $X \mapsto Xb_0$ for skew-adjoint X, surjects onto the real tangent space $\mathrm{Re}\langle b_0, y\rangle = 0$. In a hyperbolic pair this follows by freely choosing the column at b_0, filling its adjoint row, and imposing the stated real condition at their intersection. Choose a smooth nonnegative bump β near b_0 positive there, using an ambient smooth extension of the section from the null manifold on this chart. Extend the lower vector by a Schwartz scalar of value one in the second zero projection, and prescribe the first slices to be $\beta(x)$ times the translates of this extension by $h(x)$, with target weight $n/2 - \mu_\tau$. The real mixed model realizes these slices by partial Fourier inversion. On their support $kh(x)$ multiplies b_0 by λ. The line Levi law thus gives a lower translated vector times $\beta(x)|\lambda|_E^{n/2}$ and angular phase $(\lambda/|\lambda|)^{-2\mu}$. The angular weight of the cone cancels this phase, and translating the lower integrated variable removes the lower unitary element. The remaining real multiplier is the integral of $\beta(x)|\lambda|_E^{z+n/2}$. It is finite since the bump confines every $|\lambda_\tau|$ to a compact subinterval of $(0,\infty)$, and positive by its value on a neighborhood of $(1,1)$. Dividing the upper input by this factor gives any prescribed lower equal-rank integral on its standard datum. This argument uses only the positive branch.

For the two branches in the nonvanishing argument it suffices to have preimages on determinant vectors. Write $n = 2(k_0+q)$ at a lower step and $\mu_\tau = k_0 + m_\tau$ with $m_\tau \geq 0$. Use the Gaussian times the conjugate product of positive-coordinate principal minors of exponents m_τ on the positive block and $m_\tau + 2k_0$ on the negative block. In the notation of the fixed Fock model these minors are taken after multiplying the positive-coordinate matrix by one fixed matrix A_0. Extend A_0 by zero across the complements and by the two marked line pivots $-\alpha_\infty^{-1}(K_W K_V^* + L_W L_V^*)$ in the adapted real frames. For the real compact action each block transforms by the determinant of its compact matrix. The vacuum and splitting character cancel both powers because $n'/2 = k_0 + q(W^+)$ at the upper step. In the mixed slice the equal-sign

[2] For strong approximation use M. Kneser, *Starke Approximation in algebraischen Gruppen I*, J. reine angew. Math. 218 (1965), Satz 2, p. 191.

projection of the added plane has one added λ_τ pivot in each compact sign block, with zero off-diagonal entries relative to the old frame. This is obtained by projecting the two dual isotropic generators onto their equal-sign sums and taking the chosen scalar in the retained line. The two mixed models exchange those generators. Thus each sign minor acquires a factor λ_τ in either model, while the Gaussian acquires the fixed line Gaussian. The angular factor of the cone cancels the conjugate determinant power $\overline{\lambda_\tau}^{2(k_0+m_\tau)}$. With the normalized real scale, the resulting integrand is the lower determinant vector times

$$\exp\left(-\sum_\tau |\lambda_\tau|^2\right)|\lambda|_E^z \prod_\tau |\lambda_\tau|^{2(k_0+m_\tau)}.$$

For $z \geq 1$ its integral is positive and finite. Rescaling gives the desired preimage on each successive plane. The rank-zero vacuum is in this determinant family.

Take J positive definite of rank $2k_0$ and $Y = \mathbb{H} \perp J$. Choose one tame character ν_0 of the prescribed weight $\mu_\tau \geq k_0$. If its $I_{Y,\nu_0}(w)(1)$ vanished for every Schwartz w, every top coefficient of $(\mathbb{H}^{2k_0} \perp Y, J^- \perp Y)$ in the self-dual test would vanish, even for arbitrary large real input. Indeed the real tensor product first gives this for pure tensors in the slice, and the coherent-state integral with all moments gives it for arbitrary slices. For invertible indices outside the negative class use the negative-line rotation above. The coefficient propagation gives zero on all standard rows. Descend through the $2k_0 + 1$ planes, using the determinant preimages and the nonzero Tate factor at each stage. On $(J, 0)$ this would kill the vacuum on the standard finite frame, whose theta integral is the positive volume of the rank-zero quotient. This proves a nonzero period for the chosen character.

Choose w with $I_{Y,\nu_0}(w)(1) \neq 0$, and fix this same real vector for the entire family of characters of weight μ. Apply the descended identity to $(J, J, 0)$, with the rank-zero vacuum as the small vector. Then apply it to $(J, 0, J)$ using the resulting real (J, J)-vector on the small side. In both comparisons the real construction uses the same w at the top. The real cone and zero-slice maps depend on the weights and frames. The top scalar is proportional to $I_{Y,\nu}(w)(1)$, including when that period vanishes, and the descent divides only the scalar by the nonzero one-plane Tate factor. Thus the two real test vectors are fixed and the identities hold for every ν with the prescribed tame data and real weights. Put $\mathcal{R}_\nu = |I_{Y,\nu}(w)(1)|^2/\Pi_{\nu,S}(k_0+1)$. With fixed positive normalization constants B, C, the two exchange scalars satisfy

$$B|c_1|^2 = \Pi_{\nu,S}(k_0)\mathcal{R}_\nu, \qquad C|c_2|^2 = \Pi_{\nu,S}(-k_0)\mathcal{R}_\nu.$$

Use the standard partner tests at the identity frames. Their common middle kernel on the two copies of J has the fixed real vector and the signed determinant datum for ν. Write m_0 for the rank-zero quotient volume. The first comparison integrates this kernel in one variable and gives $m_0 c_1$. The second comparison integrates in the other variable and multiplies by c_2, giving $m_0\xi$, where ξ is the value of the fixed real test on the rank-zero hom space. The finite rank-zero datum has value one. Integrating the two identities over the remaining copies of $[U(J)]$ and using the same Haar gives $m_0 \operatorname{vol}[U(J)]c_1c_2 = m_0 \operatorname{vol}[U(J)]\xi$. Fubini applies to the compact definite quotients and Schwartz tests, so $c_1c_2 = \xi$ for every ν. At ν_0 the two squared-modulus identities show $\xi \neq 0$. Since the two Π's are reciprocal, they give $\mathcal{R}_\nu^2 = BC|\xi|^2$ for every ν. This fixes a positive $\mathcal{R}_\nu$ on the entire family. To pass from the positive J to the actual W_*, use the positive preimage at each added plane. The one-plane factor and the constant function identity multiply numerator and denominator of $\mathcal{R}_\nu$ by the same Π-ratio up to a fixed positive factor. The number of added planes is q_*.

Now use the fixed vector w on $(\mathbb{H} \perp W_*, \mathbb{H} \perp W_*)$ in the top coefficient construction. An arbitrary Schwartz vector on the small pair extends to the required opposite zero slice, since in the real mixed model one can multiply it by a Schwartz scalar of value one there. The $b + 1$ successive removals leave exactly that small vector. At each step the standard column is projected using the original adapted dual frame, and the scalar update is $c_- = c_+\kappa_0/\gamma$. The powers of $\Pi_{\nu,S}$ telescope to the final excess $e = a - N/2$, including steps with $e < 0$ in the two rank-zero tests. The top calculation uses the opposite V_{small}^-; return to V_{small}

by adjoint coordinates, with the finite component at its indicated conjugate half. The conversion of the unitary quotient is measure preserving and the remaining determinant factor is a unitary scalar. The rank twist $\nu^{-1}\upsilon^{N/2}$ on the exterior group converts the two original integrals into the complementary integrals used in the square. It contributes no factor to the squared scalar. The general output is the partner identity with $|c|^2 = C_* \, \Pi_{\nu,S}(e)$, $C_* > 0$ independent of the twist and of the partner tests. It holds for the chosen real vector on the small pair, all Schwartz—Bruhat determinant partners and the fixed measures. A finite quotient can be of greater depth than the tame determinant character.

In triangle z the two exterior ranks are both $r + z + (n+1)(t-z)$. The intermediate ranks in the forward square are $a = r + z + (n+2)(t-z)$ and $b = r + z + n(t-z)$, so the lower endpoint of the plane descent is $e = t - z$. The square character has $\nu^{-1} = (\upsilon^{-h}\chi_0\vartheta)^{n+1}$ for V and the same formula with $\vartheta = 1$ for Z. The finite matrix modulus is $\mathfrak{M}Q_\vartheta$ for V and $\mathfrak{M}^{b_1}$ for Z. In the latter case the tame character factors through the residue fields. The scalar uses conductor exponent one at the selected places in both triangles. The power $1 \le n+1 \le m$ preserves nontriviality there. The two intermediate signs add to the exterior count, and $(n+1)h \ge N/2 - q_*$ at the actual ranks. The equal-intermediate-rank exchange of the reverse square in V uses $e = 0$. Its exchange scalar product is $\Pi_{\nu,S}(0) = 1$. The reversal factors enter through the separate Weil kernel comparison.

The real comparison on the realized form spaces. The complementary integral is an identity at every exterior point. For a fixed small-side real vector v, it pairs v with the constructed large-side vector Tv. The extension to the top zero slice, the fixed real bumps and equal-rank vector, and the real cone operators give a continuous linear choice of T on Schwartz space. This choice depends on the signatures, frames and real weights. It is independent of the finite data and finite part of ν. Thus a pointwise use of the identity may transport the real vector by T, also in the presence of other tensor factors.

For the operator comparison on realized forms, first pair with the two prescribed exterior types. In an oriented path through an intermediate group H, regard each edge as a lift from its exterior group to H, using the transpose with the contragredient type where required by the bilinear pairing. For each edge let π_i be the exterior real type and $\pi_H^{(i)}$ the resulting intermediate type. The signed center and weight calculations give the same intermediate representation on one edge and its contragredient on the other, with the compact positive order reversed in the contragredient factor so that the two chosen lines are paired.

Here are the real hypotheses for using the oscillator quotients on these edges. After removing the determinant centers, every matching adds or removes one positive-labeled central string of length $d = |N' - N|$. All remaining centers are beyond $(d+1)/2$ in absolute value. The signs of the surviving coordinates are precisely the match in the signed parameter calculation. Paul's Theorem 2.7(b) gives the larger discrete series, and Proposition 3.4(ii) gives the reverse first occurrence since that central string is maximal. Negating the parameters or taking the opposite convention uses the contragredient identification of Proposition 2.9. In the forward direction the smaller parameter has no central entry. For odd d its $k_\pi = -1$ and its first occurrence in the indicated tower follows from Ichino's Theorem 6.1(i), with $l = 0$, $t = (d-1)/2$ and both near sets empty. For even d the smaller centered entries are half-integers, $k_\pi = 0$, and Theorem 6.1(ii) applies with $l = 0$, $t = d/2$ and the same empty near sets. In either case removing a hyperbolic plane would make $l < 0$. When the smaller Witt-tower member does not exist, no such check is needed. These are the odd V and the even Z excesses, including $t - 1$ and $t + 1$ in the bordered grid. In the transpose the determinant centers are preserved upon removing a hyperbolic plane by the normalized zero-term formula. Wu's tower constant-term formula in both orientations therefore gives cuspidal intermediate lifts for arbitrary real sections and cuspidal exterior inputs, as in the column calculation. In the realized comparisons the exterior forms are in the cuspidal spaces at the stated levels and real types. For the original bounded row inputs, the constant-term argument uses their integrable discrete-series reproducer and its vanishing real unipotent constant terms to give cuspidality. The core comparison uses the cuspidal multiplicity space at the full deep principal level; it contains the core Poincaré tests in the

corank-one separation. Subsequent triangle inputs remain cuspidal by the tower formula. The Petersson tests separating the square outputs range over the corresponding cuspidal spaces. Their rapid decay and the Schwartz estimates give continuity of the pairings. [13, Theorem 2.7(b) and Propositions 2.9 and 3.4(ii)] [14, Theorem 6.1(i)–(ii)] [19, Proposition 3.1]

For a single edge the smooth oscillator quotient onto the exterior contragredient and the intermediate representation is unique up to scalar. The prescribed joint harmonic has nonzero image in it. In fact the larger discrete series is in the case of Paul's Proposition 3.4(ii). Its lowest compact type has minimal oscillator degree by the calculation in the proof of that proposition. The paired lowest compact type on the smaller side also has minimal degree by the joint harmonic correspondence. In the proof of Howe's Lemma 4.1(c), the joint harmonic of that degree has nonzero image in the maximal isotypic quotient. Its image generates the exterior lowest-type isotypic space under the enveloping algebra of the other member. The same conclusion holds in the irreducible paired quotient. The polynomial fixed by the signed minor construction is the joint highest vector in this harmonic. Pairing its exterior compact factor with the selected conjugate line gives the nonzero intermediate lowest vector. These statements are unchanged by the fixed characters at infinity. [20, Theorem 1, Lemma 4.1(c) and (4.5)] [13, Theorem 2.2(2) and the proof of Proposition 3.4]

To apply the two quotients, fix a product of finite tests and exterior forms on their selected real lines. For a polynomial edge vector Q_i, first integrate against the exterior form. Pairing the lift with any intermediate cuspidal form on the indicated real type gives a continuous bilinear oscillator pairing. After the exterior and intermediate representations are specified, this factors through the single paired quotient by Howe's uniqueness. Equivalently, the lift has the form

$$m_i \otimes u_i(Q_i) \quad \text{in} \quad \mathcal{M}_i \otimes \pi_H^{(i)},$$

where $\mathcal{M}_i$ is the automorphic multiplicity space, m_i contains the finite data and exterior form in the multiplicity space, and u_i is the fixed real quotient contracted against the exterior lowest vector. Normalize $u_i(P_i)$ using the prescribed harmonic P_i. Nonzero images of P_i make this normalization common to all choices of m_i. The Petersson or bilinear contraction through H pairs the two multiplicity vectors and contracts $u_1(Q_1), u_2(Q_2)$ by the invariant dual pairing of the two intermediate types. The latter pairing is nonzero on $u_1(P_1), u_2(P_2)$. Thus, for $Q_1 \otimes Q_2$, the integral through H is its value on $P_1 \otimes P_2$ times the ratio of the real pairings. This ratio depends only on the real vectors and the fixed representations.

Finite residue functions of primitive support are finite linear combinations of products of matrix indicators, using boxes inside the primitive support. On every resulting product both descriptions give the same kernel by the tensor product of the two Weil actions. The primitive-support bound above and the rapid decay of the exterior forms justify the group integrations in either order. The real ratio is the same on each term of the finite expansion. For each fixed finite datum and pair of exterior forms the resulting period is continuous in the Schwartz topology. If the prescribed tensor has nonzero period at one choice of these inputs, the period at that choice divided by its value on the prescribed tensor gives the continuous real functional extending the ratio on polynomial tensors. The identity with that functional extends by density to arbitrary real vectors for every other choice of finite data and exterior forms. If the prescribed tensor gives zero for all choices, polynomial tensor factorization likewise forces zero for all real vectors on that path.

Now take the actual small-path tensor for v and compare Tv on the large side with its actual tensor. The preceding factorization gives κ_* on the large-side form pairing, independent of the twist and of the finite and exterior tests. In the case where the large-side form pairings all vanish take $\kappa_* = 0$. This proves that the actual small-path pairing is κ_*/c times the actual large-path pairing. The resulting identity of operators is on the prescribed form spaces. The bordered endpoint path uses the pointwise comparison with T on its transported vectors.

For a triangle square the intermediate ranks differ by $2(t - z)$. The negative-count equality follows entry by entry. The changes from $(i, n+1), (i+1, n+1)$ to $(i, n+2), (i+1, n)$ simultaneously flip the two

signs exactly when $A \geq i+n+2$, where the old signs are opposite. The stable inequality follows because each exterior node has rank $r+z+(n+1)(t-z)$, only $r+z$ entries can be negative, and $t-z > r+z$. The character product has ν^{-1} equal to $(\upsilon^{-h}\chi_0\vartheta)^{n+1}$ in V, where χ_0 has residue character ω. The same formula with $\vartheta = 1$ holds in Z. The opening character b_i inverts the $(i+1)$-st character and the upper lift uses the $(i+n+2)$-nd. At same-parity ranks, the two roundings add to half the total rank. The weight $(n+1)h$ exceeds the rank of either middle column. Ramification is at the indicated places since $n+1 \leq m$, ω^{n+1} is nontrivial at the chosen places, and ϑ is an m-fold twist. The loss in the scalar thus uses one product of at most $t-z$ ramified ratios at the actual power ϑ^{n+1}.

Set $p_{i,n+1} = p^{+}_{i,n}$, where the new block has size $b = t-z$ and

$$p^{+}\begin{pmatrix} Y & S \\ C & D \end{pmatrix} = \begin{cases} p(-Y)/\#\{S\}, & C = 0,\ D = I_b, \\ 0, & \text{otherwise.} \end{cases}$$

For the two square kernels $L(X)$ and $L'(Y)$, applying the partner rule to the graph of $\left(\begin{smallmatrix} I \\ -W \end{smallmatrix}\right)$ and $(W\ I)$ gives $|\mathrm{GL}_{a_0}(R_*)|L'\left(\begin{smallmatrix} I \\ -W \end{smallmatrix}\right) = (\kappa_*/c)|\mathrm{GL}_{a_1}(R_*)|L(W\ I)$, with a_0, a_1 the smaller and larger intermediate ranks. This follows by summing over their invertible frames. The inverse determinant characters in the partner weights cancel the characters of the kernels. Specialize $W = \left(\begin{smallmatrix} Y \\ 0 \end{smallmatrix}\right)$ after exchanging the common blocks. The determinant-one shear $\left(\begin{smallmatrix} I & S \\ 0 & I \end{smallmatrix}\right)$ on the larger intermediate space leaves its value unchanged. Averaging over S gives exactly p^{+} and the two weights q_n, q_{n+2}. In operator form the identity is

$$q_n a_{i+1,n} d_{i,n} = (\kappa_*/c) q_{n+2} d_{i,n+1} a_{i,n+1} = \frac{q_{n+2}}{\bar{s}_{i,n}} d_{i,n+1} a_{i,n+1},$$

where $s_{i,n} = \overline{c/\kappa_*}$ if $\kappa_* \neq 0$, and is zero otherwise, with the convention $0^{-1} = 0$. Wherever the chain gives nonzero s, $|s_\vartheta/s_1|^{-2} = |c_\vartheta/c_1|^{-2}$.

11 The comparison spaces and their determinants

11.1 Integral comparison forms

Gluing at the chosen split places. Fix j and a table index i. In the prescribed splitting $V_j = V'' \oplus V^\circ$, write $N'' = \dim V''$ and let E'' and J° be the coordinate inclusions of the r marked columns into the respective blocks. At a chosen split place v, write $\mathcal{O}_v$ for its valuation ring and ϖ for the local component of the generator of the chosen modulus $\mathfrak{M}$. Lift the $3m$-row matrix s_i from $\mathcal{O}_v/\varpi^2$ and form the stacked matrix $A_i = [s_i; \varpi I_{N''}]$. Its rows generate a full lattice between $\varpi\mathcal{O}_v^{N''}$ and $\mathcal{O}_v^{N''}$. Choose a basis matrix D_i for that lattice. Both D_i and ϖD_i^{-1} are integral. Let $D_\circ$ be diagonal with ϖ on the J°-coordinates and 1 elsewhere. The two integral matrices

$$\gamma_{i,v} = \begin{pmatrix} \varpi D_i^{-1} & E''(J^\circ)^{\mathrm{t}} \\ 0 & \varpi D_\circ^{-1} \end{pmatrix}, \qquad G_{i,v} = \begin{pmatrix} D_i & -D_i E''(J^\circ)^{\mathrm{t}} \\ 0 & D_\circ \end{pmatrix}$$

have product ϖI. Use $\gamma_{i,v}$ at v, its conjugate inverse transpose at $\bar{v}$, and the identity off the chosen pairs. This gives a finite adelic unitary element γ_i. The inverse at $\bar{v}$ is integral. At v we have $\varpi\gamma_{i,v}^{-1} = G_{i,v}$. These choices use only the untwisted modulus and the fixed splitting.

For $X = [X''; X^\circ]$ modulo ϖ^2, the condition $G_{i,v}X = \varpi[0; J^\circ]$ implies $X^\circ = J^\circ U$ with $U \equiv I_r \bmod \varpi$. The unique $Y \equiv I_r \bmod \varpi$ with $X^\circ Y = J^\circ$ is U^{-1}. The first block of the condition is then equivalent to $D_i X'' Y = D_i E''$, or $A_i X'' Y = A_i E''$, since the two matrices generate the same row lattice. Conversely these two block conditions with any such Y give $G_{i,v}X = \varpi[0; J^\circ]$. The indicator for the glued lattice therefore equals the sum of the product indicators for these Y, with multiplicity one. Chinese remaindering gives the identity modulo $\mathfrak{M}^2$. The upper N'' columns of $G_{i,v}$ also have the following property. A change of basis in $\mathrm{GL}_{N''}(\mathcal{O}_v)$ puts the rows of $s_i \bmod \varpi$ into the designated $3m$ coordinates, since their rank is at most $3m$. In that basis $\varpi G_{i,v}$ vanishes modulo ϖ^2 on all other upper columns. This gives the column bound on the very lattice used for the indicator identity.

One normalization of the large-block vectors. Fix j. In the large block of rank $r + 9m + 1$ use the prescribed $3m$-row residue matrix s_i and the short basis Y_i of the prescribed rank-$(6m + 1)$ definite class. For a free module with an ordered basis, the integral standard-basis theorem of Akin, Buchsbaum and Weyman gives a free basis of the Schur module of any partition consisting of its standard tableaux. Their presentation by exterior columns and straightening allows base change to any coefficient ring. For the ordinary diagrams here take the inner partition empty, transpose the indexing from column lengths to the partition convention, and realize the tableaux by the products of flag minors. We use the resulting basis $b_i(X)$ of products of minors in the initial columns of each matrix X_τ.[3] For a column diagram λ, write f_λ for the product of the leading principal minors with its prescribed multiplicities. Define the Gram matrix D_λ by the coefficient expansion

$$f_\lambda(Y^{\mathrm{t}}X) = \sum_{i,k} (D_\lambda)_{ik}\, b_i(X) b_k(Y).$$

Products and transposes in this formula are taken place by place. Use the bilinear Fischer pairing $[P, Q] = P(\partial)Q(0)$, with $[X^\alpha, X^\beta] = 1_{\alpha=\beta} \prod_a \alpha_a!$. Its matrix $B_\lambda = ([b_i, b_k])_{ik}$ is integral. Put $h_\lambda = [f_\lambda, f_\lambda] > 0$.

The Fischer pairing with f_λ is h_λ times evaluation at the identity on this flag-polynomial space. Indeed, only the row multidegree of f_λ contributes to its pairing or to evaluation at the identity. If a product of basis minors has this multidegree, the number of entries in the first s rows of the ambient matrix is the sum of $\min(s, h)$ over the columns of the diagram, where h is the height of the column. Every column therefore attains its maximum for each s. Taking $s = h$ forces the first h ambient rows in that minor, giving f_λ itself. Linear substitution in the Fischer pairing is adjoint to its transpose, by the chain rule. Hence

$$[f_\lambda(Y^{\mathrm{t}}X), b_k(X)]_X = [f_\lambda(X), b_k(YX)]_X = h_\lambda\, b_k(Y).$$

Row substitution by Y preserves the flag-polynomial space by Cauchy–Binet. Comparing the standard basis coefficients and using symmetry gives $D_\lambda B_\lambda = h_\lambda I$. In particular the polynomial of a Weyl coordinate vector v is $\sum_i (D_\lambda v)_i b_i$, and pairing it with b_k gives $h_\lambda v_k$. This also shows that D_λ is positive definite for the Hermitian pairing $\overline{v}^{\mathrm{t}} D_\lambda w$.

For the large positive block use the complementary diagram. At a place of rank N'', the substitution

$$P(X) \longmapsto \det(X)^c P(X^{-\mathrm{t}} w_0), \qquad c = \sum_{0<h<N''} k_\tau^+(h),$$

with w_0 reversing the columns, sends the basis products of the diagram $\lambda^\vee$, whose multiplicities are $k_\tau^+(N'' - h)$, to signed products of the original minors. Jacobi's identity gives this on each minor, since the complementary columns to the last $N'' - h$ columns are the first h columns. Applying the same substitution in the reverse direction composes to multiplication by $\det(w_0)^c$, so the matrix on the integer bases is unimodular. The integrality here uses determinant straightening over $\mathbb{Z}$. Equivalently, within the flag space one can read the coefficients of any polynomial with coefficients in a subring from the tableau basis by elimination. For the diagonal term order, the leading monomials are distinct and monic, since their exponents record the ordered entries in each tableau row.

On applying this substitution to both variables of the Gram kernel, its argument becomes $w_0^{\mathrm{t}} (Y^{\mathrm{t}}X)^{-\mathrm{t}} w_0$. Jacobi's identity takes each leading principal minor of that argument to the complementary leading minor of $Y^{\mathrm{t}}X$ divided by $\det(Y^{\mathrm{t}}X)$. The determinant factors cancel. Thus there is an integer matrix J with integer inverse for which

$$J^{\mathrm{t}} D_\lambda J = D_{\lambda^\vee}, \qquad h_{\lambda^\vee} D_\lambda^{-1} = J B_{\lambda^\vee} J^{\mathrm{t}}.$$

Here and below the diagram and its norm over all infinite places use the products of the local bases and pairings.

[3] Akin, Buchsbaum and Weyman, *Schur Functors and Schur Complexes*, Theorem II.2.16, *Advances in Mathematics* 44 (1982), pp. 207–278.

Each stacked determinant of height h contains exactly h entries from the active columns and $N'' - h$ added-block entries in every nonzero term. Let

$$d_j^{\mathrm{pol}} = \sum_{\tau} \sum_{0<h<N''} (N'' - h) k_\tau^+(h), \qquad D_j = \sum_{\tau} |\,\mathrm{DetIndex}_\tau\,|.$$

The first number is the common added-block degree and satisfies $d_j^{\mathrm{pol}} \le 31 m D_j$, with $D_j \le 2fb$. It is also the box count of $\lambda^\vee$. Over the ring generated by the entries of the $\tau(Y_i)$, and with the added-block variables as coefficients, each determinant product lies in the standard flag-polynomial space for λ. In the factors attached to standard rows the upper unit rows use the other active columns. Expansion along those unit rows and then the first h variable columns identifies the needed minors. For a factor attached to the short basis, all the active columns are variable, and the same expansion applies. The basis coefficient vector is integral over that ring in the sense above. Therefore multiplying its Weyl coordinate vector by $H = h_{\lambda^\vee}$ clears it over the same ring. The determinant expansion of $f_{\lambda^\vee}$ has at most $d_j^{\mathrm{pol}}!$ terms counted with multiplicity. Each monomial of degree d_j^{pol} has squared Fischer norm at most $d_j^{\mathrm{pol}}!$. The triangle inequality gives $1 \le H \le (d_j^{\mathrm{pol}}!)^3$, so $H \le (d_j^{\mathrm{pol}})^{3 d_j^{\mathrm{pol}}}$, with the empty degree giving $H = 1$. Both the degree and this multiplier are fixed over table choices and twists.

For a factor of height h, the stacked determinant has unit rows in the positive columns from $h + 1$ through $r_+ = \kappa + j - 1$. Subtract multiples of those upper rows from the lower variable rows to clear the entries in these columns. In the remaining determinant the first h columns are the variables of the Weyl polynomial and have zero upper block. The complementary columns have entries from $0, 1, Y_i, z$, with z the specialized added-block variables. The column count in this complement is at most $12m + 1$, since $r_+ - h \le j$ on this diagram. Expanding along the variable columns gives complementary minors as the exterior coordinates. Their squared norm is bounded by the full sum of squared minors of the complementary columns and hence by the product of their squared lengths, by Cauchy–Binet and Hadamard. On $\max |\sigma(z_a)| \le g^4$ the short basis has the same entry bound. Each column has squared length at most g^9, since $r + 11m + 1 \le g$. The Cartan tensor norm is bounded by the product of the exterior norms. Thus the squared Weyl norm of the cleared coordinate vector at any such specialization is at most $H^2 (g^9)^{(12m+1)D_j}$.

By definition the coordinates of the large-block vector are the class contraction of three definite kernels. The first is the degree-zero carrier kernel of ranks $m, 3m$ at unit modulus. The second is the degree-zero kernel of ranks $j, r + 9m + 1$ from the negative block, read via the opposite-space identity. The third is the added-block kernel from rank $12m - j + 1$ to $r + 9m + 1$ with the cleared polynomial. The latter two use $\mathfrak{M}^2 Q_\psi$. The scalar for a standard definite kernel of ranks u, v' and degree d' can be taken to be

$$C_{u,v',d'} = (g^2)^{2fuv'}\, v^{uv'} (g^2 d\, \eta_0)^{d'}.$$

Here v and η_0 are the vacuum and jet parameters and d is the discriminant prime. The normalized vacuum datum at $\mathfrak{M}^2 Q_\psi$ is integral at every place over $\ell \nmid \mathrm{N}(Q_\psi)$ after including $(g^2)^{2fuv'}$: the integral standard-data clearer uses $g^2\,\mathrm{N}(Q_\psi)$ and the omitted norm power is a unit. The jet then uses the scale $g^2 d\, \eta_0$. This is the weighted standard datum calculation with the two stars from the negative block included in the conjugate reading. Coefficients of the residue data are sums of additive roots of unity, or indicators, and are integral.

To apply this bound in the transported gauges, move each finite argument to integral frames at the places above ℓ. A rational change of the representative exists by weak approximation, using the Cayley chart in the unitary group. The short-basis columns in these frames are integral by their construction in the lattice of the corresponding class. Under the change of representative, the positive stacked minors transform by row substitutions on the variable matrix and the associated substitutions on its constant rows. Indeed the right determinants from the change of columns multiply each determinant once. The total exponent of this factor is the number of minors at that place. The two changes of splitting character have product v^{r+9m+1} in the opposite pair, so the determinant scalar is the one in the Weyl action. Thus

transporting the output before reading its coefficients gives the same polynomials on the integral input data. This evaluates the unit gauge on the actual large-block vector. Put $D_{\mathrm{aut}} = \prod_{v\mid\mathfrak{M}} |\mathrm{GL}_{rt}(k_v)|$. In the double class sum use D_{aut}^2 to clear both stabilizer orders. The fixed normalization β_j is the product of this integer, H, the three kernel scalars with the negative one conjugated, and the fixed opposite-space phase. It is nonzero and depends on j.

For the complex norm at $\psi = 1$, a normalized kernel of either degree-zero pair has conjugate modulus $\le e^{x_1}$, where x_1 is the secondary cutoff. For the third apply the weighted Gaussian integral to the homogeneous polynomial after contraction with an arbitrary unit vector in the dual Weyl norm. On the polydisc for its coefficients, the preceding determinant bound and Cauchy–Schwarz bound this contraction by $H(g^9)^{(12m+1)D_j/2}$. Homogeneity extends the estimate to larger balls, using $(1+R)^{d_j^{\mathrm{pol}}} \le e^{d_j^{\mathrm{pol}}R}$. With the actual ranks $\le 40\kappa$ and degree $\le g$, the Gaussian factor and the residue and kernel clearers contribute at most e^{x_1} times the maximum norm of the finite datum. The latter is at most $|\mathrm{GL}_{3m}(\mathcal{O}_E/\mathfrak{M}^2)|$, the number of terms in the translated trace character. All class counts and residue group orders here have logarithms bounded by a fixed power of x. Applying the triangle inequality to the two class sums and taking the supremum in the dual norm yields

$$\log \left\| y_{1,i}(h)^{\sigma} \right\|^2 \le 6x_1 + O_d(x^{99}) + O(fmb \log b).$$

Since $rB_0 = rfa(\log b)^3/(m\kappa)$ and $m^2 \le 4x$, the right side is $o(rB_0)$. This gives the required e^{rB_0} bound at each conjugate, with the same β_j as the finite gauge.

The comparison forms on these lattices. Use the normalized vectors $y_{\psi,i}(h)$ with their one stage multiplier β_j. Their conductor-away gauges and conjugate squared norms satisfy the bounds just proved. For the table entry $\xi_{i,c}$ at a small-block class c, form the Cartan product of $y_{\psi,i}(h)$ and $\xi_{i,c}$. Let $D_{i,h,c,\psi}(\xi)$ be the corresponding point form at $\mathrm{diag}(h, c)\gamma_i$, with the original marked columns and level character. If J_i is the large-block level fixing the stacked rows, set

$$G_\psi \xi = g^{L_j} \sum_{i,c} \sum_{[h]\,\mathrm{at}\ J_i} D_{i,h,c,\psi}(\xi),$$

$$L_j = \sum_{\tau} \left(3m \sum_{a<N''} k_\tau^+(a) + \dim V^\circ \sum_{a<\dim V^\circ} k_\tau^-(a) \right).$$

Here $k_\tau^{\pm}(a)$ count the columns of height a in the two parts of the stage weight. The row-fixing condition for $\varpi I_{N''}$ puts J_i inside the principal level modulo $\mathfrak{M}$, so its classes have trivial stabilizers. At primes of the twist the point form keeps the prescribed determinant character.

The mixed Cartan map used in this product comes from multiplication of the two polynomial modules, after complementing the negative diagram. Multiplication is an integral surjection onto the module for the union of their columns. A standard tableau of the union can be factored by assigning its columns of each height to the two diagrams in the prescribed numbers. Each factor is an integer combination of standard products. The coefficient module of the union is free, so the dual of the multiplication map has an integral left inverse. Evaluations span the dual coordinate spaces, and applying multiplication shows that the dual map sends an evaluation for the union to the tensor of the two evaluations. The Gram kernel of the union is the product of the two Gram kernels, proving that this map is isometric with the normalized Gram forms. In the complemented factor apply the unimodular isometry just proved. Its formula on evaluation vectors gives the contragredient action and the accompanying power $-\sum_{h>0} k_\tau^-(h)$ of the determinant. At the identity the image is the tensor of the highest evaluation for the positive diagram and the reversed highest evaluation for the negative diagram. Place their supporting rows in V'' and V°, respectively, by a permutation of the target coordinates. The orbits under the two block general linear groups span the two coefficient modules on those blocks separately. This proves that their pure tensors lie in the Cartan image. The block extensions preserve their metrics, since the Gram kernels of the extended evaluations restrict

to those on the blocks. Applying the integral left inverse to the pure tensors gives the actual integral product and its asserted norm.

For the gauge bound, transport the coefficient vectors into the local integral frames at h, c. Lift the local change of columns to the compact integral group, using the inverse transpose at the conjugate place. Inserting this matrix between the inverse lattice and h^{-1} replaces h by a compact translate and preserves the input gauge. In the positive part only the designated $3m$ columns of the inverse lattice can have a local denominator ϖ, by the column bound. In the negative part use the conjugate inverse on the dual block, for which at most $\dim V^{\circ}$ columns can have this denominator. Every determinant in a Weyl polynomial is multilinear in the columns. Thus clearing a column in the positive or negative block costs at most its number of minor factors in the corresponding weight. The integral left inverse of the mixed Cartan embedding carries this bound to the coefficient vector of the product. A class value is integral if its transported coordinates belong to the integral Weyl lattice. This condition is imposed on the full space of class values before projection. Since $g \in \mathfrak{M}$, the factor g^{L_j} clears all the resulting losses at the two readings. More precisely, at a place a the input gauge on the small block is the gauge of $\bar{\xi}_{i,c}$ at $a \circ \mathrm{conj}$. The point form is an orbit sum of the transported coefficient vector. Its terms have no extra denominator and level characters are local units. Summing in h, i, c therefore preserves integrality.

The rational Cartan coordinates give the same product norm after applying a coefficient automorphism to the vectors. At a target class a point form contributes at most J_{aut} vectors of that norm, where $J_{\mathrm{aut}} = \prod_{v \mid \mathfrak{M}} |\mathrm{GL}_{rt}(k_v)|$. Indeed reduction of a class stabilizer is injective at the chosen levels by freeness, and $\mathrm{GL}_{r+t}(k_v)$ embeds in $\mathrm{GL}_{rt}(k_v)$ for the ranks under consideration. Each summand comes from a unitary rational transport and a root of unity, and remains an isometry after applying the coefficient embedding to the transported matrices. Apply the triangle inequality to the h-sum and Cauchy–Schwarz to the table index sum. If C_V is the number of target classes, C_J bounds the number of J_i-classes, and I_j counts the pairs (i, c), the result is

$$\|G_1 \xi\|_{\sigma}^2 \le C_V J_{\mathrm{aut}}^2 C_J^2 I_j \, g^{2L_j} e^{rB_0} \sum_{i,c} \|\xi_{i,c}^{\sigma}\|^2.$$

At the trivial twist all these definite levels contain the principal congruence subgroup modulo g^2, with ranks at most $r + t$. The definite class count therefore bounds their logarithmic class counts by x^{99} for large x. The logarithm of I_j is bounded by the two class counts, $3mN'' \log |\mathcal{O}_E/\mathfrak{M}^2|$ for the matrix choices, and $2mfr^8 \log(6m+1)$ for the row choices. These terms and $\log J_{\mathrm{aut}}$ are $O(x^{99})$. The number of boxes for the small negative weight is $O(f(r+t)^2 r^8)$, so the logarithm of its hook number k_j is also $O(x^{99})$. Finally $\sum k_{\tau}^{+} = w_j(\tau) - j - \lfloor (t + \kappa - 1)/2 \rfloor \le 2b$ and $k_{\tau}^{-}(a) \le 2r^8$. It follows that

$$2L_j \log g = O(fmb \log b + x^{99}) = o(rB_0).$$

The last equality uses $m^2 \le 4x$, $a \ge xb$, and $B_0 = fa(\log b)^3/(m\kappa)$. Therefore $k_j \|G_1 \xi\|_{\sigma}^2 \le e^{2rB_0} \|\xi^{\sigma}\|^2$ for sufficiently large scales. The hook is used once, in the normalization of the compressed operator.

11.2 The three measured exchanges

The real matching of the determinant products. Use the ordered positive and negative blocks on the two signed spaces. Let W have signature $(p + a, q)$ and S signature (p, q), with complementary target I_a definite. Write $N = p + a + q$, and put $d_{\tau} = |\mathrm{DetIndex}_{\tau}|$ and $\gamma_{\tau} = q - d_{\tau}$. Order the positive indices of W by the fixed bijection from the stack of p and a rows. In these coordinates let I denote the first p columns and J the remaining a. Write Z_{+} for the positive columns in the I_a-by-W Fock variables, and Y for the positive-positive block of the W-by-S variables. Let Y_{--} be the negative-negative block of the latter. For the constant matrices $\eta_{\tau b}$ of size p by $p + a$, the two polynomials are

$$P_{\eta} = \prod_{\tau, b} \det \begin{pmatrix} \eta_{\tau b} \\ Z_{+,\tau} \end{pmatrix}, \qquad Q_{\eta} = \prod_{\tau} \det(Y_{--,\tau})^{d_{\tau} + N - 2q} \prod_{\tau, b} \det(\eta_{\tau b} \, Y_{\tau}).$$

Let $\eta^0 = (I_p\ 0)$ for all factors, and put $P_0 = P_{\eta^0}$, $Q_0 = Q_{\eta^0}$. Empty determinants and products have value one.

For these two oscillator actions use splitting weights $(\gamma, 0)$. The compact action of a diagonal sign-block pair (x, y) substitutes $y^{\mathrm{t}} Z x^{-\mathrm{t}}$ on equal signs and $y^{-1} Z x$ on unequal signs, with the blocks read separately. Its scalar in a + row block of y is $\det(y_+)^{\gamma_\tau - \#\{x \text{ negative indices}\}}$. For the rank-$a$ compact group of the first oscillator and for the positive-block group acting on the second oscillator this is $\det(y_+)^{-d_\tau}$. It is the splitting character times the vacuum determinant of the opposite-sign variables. The positive-block group acting through x on the first oscillator has scalar one because I_a is definite and the second splitting weight is zero.

Denote the first and second real actions by Ω_A and Ω_B. Integrate in the first Fock space over the circle with normalized measure. At every place let $t(z)$ act on the positive block of W by z^{-2} on I and z^{-1} on J, trivially on the negative block. Put $u_\tau(z) = z^{-1} I_a$. The operator

$$\left(\prod_\tau \det(u_\tau(z))^{d_\tau} \right) \Omega_A(t(z), u(z))$$

substitutes z times each variable outside Z_J and fixes Z_J, with no scalar. Its circle average sends polynomials to their restrictions to Z_J. After that, average the action $\Omega_A(1, U(a))$ over the product of compact groups with mass one. For a polynomial in Z_J, the result transforms under left substitution by the determinant character $\prod_\tau \det(u_\tau)^{d_\tau}$. The identity extends to the full general linear groups by polynomial continuation from the unitary groups. An invertible square $Z_{J,\tau}$ is in the orbit of the identity, so such a polynomial is a multiple of $P_0 = \prod_\tau \det(Z_{J,\tau})^{d_\tau}$. Both compact averaging operators are bounded in the Fock norm by unitarity, and the polynomial vectors are dense there. The composite thus sends each real datum v to $\kappa(v) v_{P_0}$, with a scalar independent of the finite data and all evaluating functionals.

Fix the adelic evaluation point and denote by ℓ_x the continuous first theta functional at its right translate by x. Let ℓ' be a linear functional on the second polynomials. The first functional is invariant under the I_a-compact factor, and $\ell_x(\Omega_A(y, u) w) = \ell_{xy}(w)$ for the real translations in use. Suppose that $\ell_k(v) = \ell'(\Omega_B(1, k) v_{Q_0})$ on the positive compact block, identifying polynomials and their real Fock vectors on the right. Apply the two averages and evaluate with ℓ_k. The $U(a)$ average leaves ℓ_k unchanged. On the other side, the polynomial of $\Omega_B(1, kt(z)) v_{Q_0}$ is

$$z^{a \sum_\tau d_\tau} \left(\prod_\tau \det(k_\tau)^{-d_\tau} \right) Q_{\eta^0 k^{\mathrm{t}}}.$$

Indeed $\det(t_\tau)^{-d_\tau} = z^{(2p+a) d_\tau}$ and $\eta^0 t_\tau^{\mathrm{t}} = z^{-2} \eta^0$, so the circle weight cancels with $\prod_\tau \det(u_\tau(z))^{d_\tau}$. On the first side, substitution in the stacked determinant gives

$$\Omega_A(k, 1) v_{P_0} = \left(\prod_\tau \det(k_\tau)^{-d_\tau} \right) v_{P_{\eta^0 k^{\mathrm{t}}}}.$$

For example, multiply the whole stacked matrix on the right by $k_\tau^{-\mathrm{t}}$ after replacing the upper rows by $\eta^0 k_\tau^{\mathrm{t}}$. Cancelling the common determinant scalar proves the matching for all equal-row powers on each place.

To allow the different rows $\eta_{\tau b}$ in the same product, let $F_\eta(Y) = \prod_{\tau, b} \det(\eta_{\tau b} Y_\tau)$. All these polynomials transform on the right by $\prod_\tau \det(u_\tau)^{d_\tau}$ for $u_\tau \in \mathrm{GL}_p(\mathbb{C})$. The Hermitian Fischer pairing with $\prod_\tau \det(Y_{I,\tau})^{d_\tau}$ is its squared norm times evaluation at $Y_\tau = \binom{I_p}{0}$. In fact one can set the other rows to zero in this pairing, and the stated determinant transformation on the right gives the assertion on the square surviving block. After unitary row translation these pairings recover evaluation at every tuple of orthonormal p-frames. Gram–Schmidt and the right determinant rule recover evaluation at all full-rank frames. Those frames are dense, so the equal-row powers span the space of the F_η.

Every linear relation among the F_η is also a relation among the P_η. For a specialization of $Z_{+,\tau}$ of full row rank, take a matrix A_τ whose p columns are a basis of its kernel and complete it to an invertible matrix

$(A_\tau\ B_\tau)$. Multiplication by this matrix puts the lower-left block of each stacked determinant equal to zero. Thus the stacked determinant is $\det(Z_{+,\tau}B_\tau)\det(\eta_{\tau b}A_\tau)/\det(A_\tau\ B_\tau)$. Its scalar factor is independent of $\eta_{\tau b}$. These specializations are dense, proving the relation as a polynomial identity. The passage from F_η to Q_η is multiplication by the fixed negative minor powers and also preserves all relations. Linearity gives $\kappa(\upsilon)\ell_1(\upsilon_{P_\eta}) = \ell'(Q_\eta)$ for the full family of determinant products.

Determinant partners and the exchange scalar. For the same real spaces let $b' = p+q$, and suppose $a = b' + 2e$ and $N \geq 2q+2$. Let α have weight $\gamma_\tau = q - d_\tau$ and tame character χ_α on $\mathcal{O}_E/\mathfrak{m}$, where $\mathfrak{m}$ is supported on conjugate-disjoint split places. For an arbitrary function F on the rectangular residue matrices use

$$\Phi(X'\ X'') = \chi_\alpha(\det X'')F((X'')^{-1}X'),$$
$$\Psi\begin{pmatrix} Y' \\ Y'' \end{pmatrix} = \chi_\alpha(\det Y')F(-Y''(Y')^{-1}).$$

The expressions are zero where the indicated square block is noninvertible. These are matching data for χ_α^{-1}. To verify this, write the first b' columns of an integral invertible matrix k as $\binom{A}{C}$ and the last a rows of k^{-1} as $(E\ D)$. Jacobi gives $\det D = \det A/\det k$, so the two supports agree on these columns and rows. Where they are nonzero, $EA + DC = 0$ gives $D^{-1}E = -CA^{-1}$. Substitution yields $\Phi(E\ D) = \chi_\alpha(\det k)^{-1}\Psi\binom{A}{C}$, with a primitive frame on each open support. This proof works over the product residue ring including its squared local factors. At the remaining places the datum is the vacuum.

In the normalized complementary-period formula take $\nu = \alpha^{-1}\upsilon^{N/2}$. Its weight at τ is $N/2 - q + d_\tau$, within that formula's stable range. Apply the formula to the fixed real vector υ_{Q_0}. It gives a first real datum υ independent of the tame character and a nonzero scalar c_α for the partner equality, with

$$|c_\alpha|^2 = C\,\Pi_{\nu,S}(e), \qquad \Pi_{\nu,S}(e) = \left(\prod_{w\in S} q_w\right)^{e^2} \prod_{z=1}^{e} |L^S(z,\nu)|^2, \qquad S = \operatorname{supp}\mathfrak{m}.$$

The scalar on the integral over $[U(S)]$ is c_α. Use that integral multiplied by c_α for ℓ' in the real matching. Translation of the original formula gives the required equality on the positive compact block. The transfer of polynomials gives $\kappa(\upsilon)/c_\alpha$ as the multiplier of the first theta period to obtain the second one.

For the resulting P_η, the compact real action of I_a is trivial, because substitution in the a lower rows of the stack contributes exactly the determinant character inverse to its vacuum and splitting phase. At finite integral points of I_a, the multiplier $\alpha(\det k)$ cancels $\chi_\alpha(\det k)^{-1}$ from replacing the matrix argument by $k^{-1}X$. Thus integration over $[U(I_a)]$ is the class sum at full level with weights $1/|\operatorname{Aut}|$, since the real and the standard finite compact factors have volume one. Changing the second splitting character on the first kernel from 1 to β multiplies it by $\beta(\det g)^{-1}$. The same multiplier on the second kernel replaces its first splitting character α by α/β. Hence the exact matched scalar has squared norm $C'/\Pi_{\nu,S}(e)$, with $C' = |\kappa(\upsilon)|^2/C$ fixed across the twists. Nonvanishing in the applications is established by their period identities.

The mass and the three characters in the matching. Write $\chi_\psi = \chi^j\psi$, and $b_\psi = \mu_j^0\psi\,\upsilon^{j+\lfloor(t+\kappa-1)/2\rfloor}$. In the matching construction the first two exchanges on the carrier have complementary rank differences m and $2m$. Their Hecke characters in the norm of the exchange scalar are $\chi_\psi\upsilon^{m/2}$ and $\chi_\psi^{-1}\upsilon^m$, at $m/2$ and m. In the third exchange the ranks are $r+9m+1$ and $3m$, so half the difference is $z_* = (r+6m+1)/2$. The common negative block has rank m, the padded source has rank $r+12m+1$ after including the carrier indices, and this rank exceeds $2m+1$. For the three exchanges the square blocks have the tameness required at $\mathfrak{M}^2Q_\psi$, since χ^j has nontrivial tame part at both chosen sets and ψ is nontrivial at every prime of its

conductor. The real weight check on the third block is $|\operatorname{DetIndex}_\tau| + j + \lfloor (t+\kappa-1)/2 \rfloor = w_j(\tau)$. The product of the carrier character with χ_ψ is $b_\psi v^m$. The character entering the third exchange at z_* is

$$(b_\psi v^m)^{-1} v^{3m+z_*} = \left(\mu_j^0 \psi\, v^{j+\lfloor (t-11m-1)/2 \rfloor}\right)^{-1}.$$

Here $2m + z_* = (11m+\kappa)/2$ and both integers inside the half-rank floors are even. This identifies the third exchange factor with the indicated factor of Γ_j.

The finite change of variables can be made directly. For a square block x'' on the support of the determinant character, write $x' = x''z$. Then the invariant finite datum is $\chi_\psi(\det x'')\, F(z)$, where $F(z) = \sum_{k \in \mathrm{GL}_{3m}(\mathcal{O}_E/\mathfrak{M}^2)} \operatorname{trchar}_i(zk)$. If x'' is singular modulo the local modulus the determinant datum vanishes. Unfold the full-level class sum by representatives for the row-fixing subgroup. The allowed integrally invertible x'' lies in exactly one such coset. On that coset replace x'' by its marked columns and force the other columns to be those of the identity. The invariance of F under the row-fixing subgroup uses $s_i A = s_i$ and cyclicity of the trace. The normalized class measure converts the sum over the cosets to the class sum at the row-fixing level. This leaves the finite indicator of the comparison lattice and the two class variables of its prescribed large-block vector.

The first exchange uses $W = I_m$, $S = 0$, $a = m$, $d_\tau = 0$, and $\alpha = \chi_\psi^{-1}$. The second uses $W = I_m \oplus S_0$ with S_0 of signature $(2m, m)$, $S = I_m^-$, $a = 3m$, $d_\tau = 0$, and $\alpha = \chi_\psi v^m$. The third uses $W = I_{3m} \oplus W_j \oplus I_j^- \oplus I_{3m-j} \oplus I_{6m+1}$, $S = S_0$, $a = r + 9m + 1$, and $\alpha = b_\psi v^m$. Here the original source has $m - j$ negative signs, so the first two nondefinite blocks account for all m negative indices of W in the last exchange. The ambient dimensions are even and at least $2q + 2$ in all three cases. The half-rank differences are $m/2$, m and z_*. The real block weights just checked are exactly $q - d_\tau$ in these spaces. For the first two exchanges the upper-row determinant polynomial on W is the empty product. The corresponding polynomials on S are 1 and $\prod_\tau \det(Y_{--,\tau})^{2m}$. For the third, the negative exponent is $d_\tau + r + 10m + 1$. Its positive factors from the standard rows select the padded minors in $W_j \oplus I_j^- \oplus I_{3m-j}$, and the remaining factors have the indicated rows of Y_i in I_{6m+1}. This is the prescribed polynomial before the row-fixing restriction.

The first class identity is normalized to the constant one on $U(I_m)$, by dividing by the fixed mass of the rank-zero quotient. Conjugating and negating its source yields the constant one on $U(I_m^-)$. The sign change of the oscillator conjugates the two characters and multiplies both by v^m. It sends the finite determinant datum to the new determinant character times a factor of modulus one, since scaling the matrix by the finite norm-minus-one unit scales its determinant by the m-th power of that unit. The fixed phase of the oscillator also has modulus one. Thus the conjugated first normalizer has the same squared norm.

Insert its class expression for the constant function in the integral over $U(I_m^-)$ in the second exchange. For the finite datum take $F = 1$ in the determinant partner formula. The product of the two kernels in this integral is the kernel for the direct sum with the determinant datum on the square block. This follows by separating the rational lattice sum and using the product of the two splitting characters. Apply the second exchange to each class term and split the larger kernel on $I_m \oplus S_0$ again. The untwisted finite function on the I_m block is the standard vacuum, so this term is the full-level class coefficient of the carrier. The resulting class sum on I_{3m} comes with the product of the first two normalizers and the modulus-one phase. On the smaller integral the first character is $\chi_\psi v^m$, so its value is $\chi_\psi(\det x)$ times the fixed carrier. Changing the second character on the class expression to χ_ψ multiplies it by $\chi_\psi(\det x)^{-1}$ and gives the asserted equality.

For the third exchange insert this carrier class sum. Here take $F(z) = \sum_k \operatorname{trchar}_i(zk)$, using the reduction of the rectangular matrix to $\mathcal{O}_E/\mathfrak{M}^2$, with k running through $\mathrm{GL}_{3m}(\mathcal{O}_E/\mathfrak{M}^2)$. On the support of Ψ, right invariance of F makes $F(-Y''(Y')^{-1}) = F(Y'')$. The product kernel therefore has precisely the determinant partner datum on the third space. Taking the class sum and applying the third exchange produces the fixed carrier integral with this sum of translated trace characters. Full level of the carrier and triviality of the tame part of its other splitting character give identical integrals for the translates. Thus the scalar on the row-fixing class sum is the product of the three normalizers, including the modulus-one

phase, divided by $|\mathrm{GL}_{3m}(\mathcal{O}_E/\mathfrak{M}^2)|$. The measure at the smaller row-fixing level is the class sum with the same stabilizer weights, as in the preceding unfolding. Each squared norm contributes the reciprocal of the corresponding factor computed above. It follows that $|\kappa_\psi|^2 = C/\Gamma_j(\psi)$ on the actual comparison vectors, with one twist-independent C.

The Weil phase from the glued lattice uses the source splitting character $v^{\#\{\text{negative signs of } W_j\}}$, which is independent of ψ even on the nonintegral chosen-place point. Its other factor is a power of the finite idele absolute value. Thus the common-adjoint normalization does use an untwisted ω_i.

For the large-block detection calculation put $n_* = 2ft^3$ and $B = 5b$.

11.3 Separation by the large-block products

A first minimum at a half-supremum point. For a moved Hermitian lattice L use the Euclidean norm of its complex coordinates, taking one embedding per conjugate pair. Put

$$\lambda(L) = \min\bigl(1, \inf_{0\neq y\in L} \|y\|\bigr).$$

Suppose a bounded automorphic function U is holomorphic of weight u in an order with the positive indices first. Assume that $\max_{\tau,i} |u_{\tau i}| \le B$, $B \ge 2$, and that every negative entry of u_τ is at least one larger than every positive entry. Let $p \in \mathcal{O}_E$ satisfy $\bar p = -p$ and $0 < \operatorname{Im} \tau p \le C$ at the chosen embeddings, with $C \ge 1$. Suppose U is invariant under the principal level of p. Then every z such that $|U(z)| > \frac{1}{2} \sup |U|$ satisfies $\lambda(L_z) \ge (BC)^{-1/2}$.

Indeed, a shorter nonzero vector has integral Hermitian square. At each embedding the absolute value of that square is bounded by the squared Euclidean norm of the corresponding component, hence is strictly less than one. Its field norm gives isotropy. In the moved frame write v_τ for its real-place coordinates, J for the diagonal signature matrix and $\eta_\tau = \operatorname{Im}(\tau p)\|v_\tau\|^2$. Each η_τ is positive and $B\sum_\tau \eta_\tau < 1$. On the upper half-plane consider

$$M_\tau(s) = I - (\tau p)(s-i) v_\tau v_\tau^* J.$$

Isotropy gives $M_\tau(s+1) = M_\tau(1+i) M_\tau(s)$. The transvection at $1+i$ comes from the rational vector. Its finite component is congruent to I modulo p in the moved frame, since the finite coordinates of the vector are integral. Left rational invariance and the right level therefore make the holomorphic extension of $U(z\,\cdot)$ along $M(s)$ periodic of period one.

To bound this extension, order the coordinates as above and put

$$P_l = 1 - 2\operatorname{Re}((\tau p)(s-i)) \sum_{i'\le l} J_{i'i'} |v_{\tau i'}|^2.$$

The determinant lemma identifies P_l with the leading Gram minor of $M_\tau(s)^* J M_\tau(s)$ divided by the corresponding minor of J. As l increases, the P_l go monotonically from 1 to $P = 1 + (\operatorname{Im} s - 1)\eta_\tau$ through the positive indices and back to 1 through the negative indices. They stay positive. Gram–Schmidt for J thus gives a unitary factor followed by an upper triangular factor with squared diagonal moduli P_l/P_{l-1}. If $\operatorname{Im} s \ge 1$, the logarithms of these ratios have the signs of J and sum to zero. The weight gap bounds their u-weighted sum by $-\log P$. If $0 < \operatorname{Im} s < 1$, the size of the weights bounds the sum by $-2B \log P \le 4B\eta_\tau$, since $\eta_\tau < 1/2$. The modulus of the weight character uses one half of these logarithmic bounds. The extension is consequently bounded by $e^2 \sup |U|$, and tends uniformly to zero as $\operatorname{Im} s \longrightarrow \infty$ by the bound $\prod_\tau P^{-1/2}$ in the first case. In the variable $q = \exp(2\pi i s)$ it extends holomorphically across 0 with value zero. Schwarz's lemma at $s = i$, where $M_\tau(i) = I$, gives $|U(z)| \le e^{2-2\pi} \sup |U|$. This excludes the putative short vector. A nonzero bounded U has such a half-supremum point by the definition of the supremum.

We use two instances. Write $G_0 = g_0 g_1 = \pi_G \bar\pi_G$ for the product of the chosen rational primes, with the prescribed half $\mathfrak{M} = (\bar\pi_G)$, and S for the Hermitian space of signature $(2m, m)$. The fixed carrier ϕ on $U(S)$ has weight $(0; 3m)$. The first two carrier exchanges, at the untwisted stage character, express it as a finite

combination of theta kernels with determinant character on square matrices. The finite support forces an invertible residue matrix in the orientation with S as source as well, since the matrices are square. Apply the primitive-support decay estimate with its short isotropic vector d' in this source and a full frame C' in the fixed target. The corresponding modulation has coefficient $C'^{*}Yd'$ up to the invertible Hermitian and frame factors. It is nonzero on the named support since Y is injective. Thus the additive average vanishes there, and repeated differences of the bounded real translates give every power of the inverse source height. The definite target runs over finitely many classes. The level character cancels the residue determinant character on the full integral group. The real polynomial 1 gives the displayed holomorphic weight. The carrier is nonzero. In its original definition it is the lift of 1 from $U(0,m)$ to $U(S)$, with the unit finite vacuum and the polynomial $\prod_{\tau}\det(X_{\tau,-})^{2m}$. On the compact real source its source matrix coefficient is constant, equal to the positive squared norm of the polynomial vector. In the block-minor formula the source has only negative indices. Its weight there is the positive target rank $2m$ minus the determinant exponent $2m$, and the accompanying conjugate factor has weight zero there. The convergent doubling pairing for $3m > 2m$ sends a zero lift to zero convolution with this coefficient; its convolution with the constant 1 is the positive compact Haar mass times that squared norm. The finite vacuum has positive mass at full level and the source is definite. Thus one can fix a carrier peak z_S, independently of the good form to be tested.

For a good form F on W_j apply the same estimate to $U = \overline{F}$. To see the holomorphic weight, use the injective principal-space lift L at the trivial twist, supplied by the sharp-point pairing on the full auxiliary base. In finite dimension $L^{*}L$ is invertible, so F is the descent of Ly for some y. The highest polynomial of the descent is a product of the joint harmonics of the edge. Its translates on the definite side give the conjugate holomorphic type k on W_j. Its weight has gap at least $300n_{*}$ and $\max|k| \le 5b$. The good level contains the principal level of $G_0\sqrt{-d}$, since $\bar{\pi}_G \mid G_0$. Its cuspidal realization makes F bounded. Thus one may use $B = 5b$ for both forms. For ϕ take $p = \sqrt{-d}$, $C = \sqrt{b}$. For $\overline{F}$ take $p = G_0\sqrt{-d}$, $C = b^{23}$. These imaginary parts are positive on the fixed CM type, and $G_0 \le b^{22}$. The respective first-minimum bounds are

$$\lambda_S(z_S) \ge (5b\sqrt{b})^{-1/2}, \qquad \lambda_{W_j}(z) \ge (5b^{24})^{-1/2} \quad \text{at every peak of } \overline{F}.$$

A sparse theta kernel on the two moved lattices. Let $\widetilde{W_j} = W_j \perp I_j^{-} \perp I_{3m-j}^{+}$, and put $\nu = n_{*} = 2ft^3$. Fix z with $\lambda_{W_j}(z) \ge (5b^{24})^{-1/2}$, and extend it by the identity on the padding. Both standard padding lattices have first minimum one with the cutoff at one. Thus

$$\lambda_S(z_S)\lambda_{\widetilde{W_j}}(z \oplus 1) \ge (5b^{51/4})^{-1} \ge b^{-13}.$$

This estimate uses the same padded point in the theta kernel. Write $s = 3m$, $s' = r + 3m$ for the two ranks and $Q_E = dN_1$. For their moved Hom lattice the lower bound is

$$\lambda_{\mathrm{Hom}} \;\ge\; \frac{\lambda_S\lambda_{\widetilde{W_j}}}{6Q_E(fs)^2ss'}.$$

Here is the transference argument with indefinite signatures. The Hermitian lattice is self-dual at every finite place. The Euclidean dual of its archimedean realization is obtained from the trace dual by the sign matrix and the factor two. The sign matrix is a Euclidean isometry. Since Q_E belongs to the different, the Euclidean dual has minimum at least $2\lambda_S/Q_E$. The Euclidean covering-radius argument preceding classwise evaluation gives an E-basis in the source lattice with lengths at most $6Q_E(fs)^2/\lambda_S$. A nonzero integral Hom matrix takes one basis vector to a nonzero integral target vector. The Hilbert–Schmidt bound on this evaluation gives the displayed inequality, even with the harmless factors ss'. Cutting the first minima at one preserves this lower bound.

The polynomial P on $(S, \widetilde{W_j})$ in the definition of $F_i^{\natural}$ uses the negative m-row minor with exponent $D_\tau + N'' + m$ and the positive $2m$-row minors with the column multiplicities $k_\tau^{+}(h)$. Here $N'' = r + 9m + 1$,

$D_\tau = \Delta' + \sum_h k_\tau^+(h)$, and $\Delta' = h_2 - \lfloor (t+1)/2 \rfloor$. For each minor the rows and columns are distinct and have the same sign, including the rows carried into the padding. Thus P is nonzero and is homogeneous of degree

$$d_P = \sum_\tau \left[m(D_\tau + N'' + m) + 2m \sum_h k_\tau^+(h) \right].$$

We have $D_\tau \le 5b$, $N'' \le t$ and m, f, t, N_1 are bounded by fixed powers of x. The scale inequalities give $100(d_P + 4\nu) \le b^2$ at the sufficiently large test squares. Let v_P be its Fock vector, with the splitting characters v^m and the fixed carrier character, and put $C_P(h,u) = \langle v_P, \Omega(h,u) v_P \rangle$. In particular $C_P(1,1) = \|v_P\|^2 > 0$.

Use the element $c = N_1/\bar{\pi}_G^2 \in E$. It satisfies $c\alpha\, \mathfrak{M}^2 \widehat{\mathcal{O}}_E \subset \widehat{\mathcal{O}}_E$ since $N_1\alpha$ is integral, and $|\tau c| = N_1/G_0$. The real Heisenberg coordinate at τ is scaled by $\sqrt{2\pi \operatorname{Im} \tau \mathfrak{j}}$, where $\mathfrak{j} = -(\sqrt{-d})^{-1}$. The Hom minimum gives

$$b^6 |\tau c| \le \sqrt{2\pi \operatorname{Im} \tau \mathfrak{j}}\, \lambda_{\mathrm{Hom}}.$$

Indeed $G_0 \ge b^{20}$, and $6Q_E(ft)^2 t^2 N_1/\sqrt{2\pi/\sqrt{d}} \le b$ at these scales. The indices of the finite datum below are residue matrices for $\mathfrak{M}^2$. Let $\psi_T(X)$ be their additive trace characters, write Θ_T for their kernels with polynomial P, and set

$$\Psi = \frac{1}{\#\{T\}} \sum_T \overline{\Theta_T(z_S, z \oplus 1)}\, \psi_T.$$

Only the two base points enter these coefficients.

To estimate its kernel, expand each product of two theta values as a sum of matrix coefficients of rational Heisenberg translates. In a Schrödinger polarization this identity follows by expanding both theta sums, writing the difference of their rational position variables as a translation, and applying Poisson summation to the remaining variable with the self-dual measures. The dual rational variable gives the modulation coordinate. The Weyl half-phase on rational pairs is one by the global additive character. Transport by the common adelic Weil element moves the rational vectors into the Hom frame, and its scalar phases cancel between the conjugate pair. The mean over T kills every term except those whose α^{-1}-scaled finite coordinates are integral and zero modulo $\mathfrak{M}^2$. Thus the zero vector contributes $C_P(h,u)$. For all other surviving rational y, cy is in the moved Hom lattice, and the finite vacuum factor has absolute value one. The real Heisenberg coordinates Y of these y have norm at least $L = b^6$. Their pairwise distances are also at least L, by the same argument on a difference.

For clarity, in $e = fss'$ complex Fock variables, let $A = d_P + e$ and $q = 2\nu > e$. For a real Weil operator attached to the matrix σ, the translated coefficient at $Y \neq 0$ satisfies

$$|\langle v_P, \rho(Y)\omega(\sigma) v_P \rangle| \le \|v_P\|^2 \|\sigma\|^{e+2q} [4(A+q)]^q \|Y\|^{-2q}.$$

Indeed the vacuum integral bounds the coefficient by $\|\sigma\|^e \int |v_P(X) v_P(\sigma X + Y)|\, dX$. For $a = (1 + q/A)^{1/2}$, homogeneity and the Gaussian give

$$|v_P(X)| = a^{d_P} |v_P(X/a)| \exp\left(-\frac{q\|X\|^2}{2(A+q)} \right).$$

Use $\|Y\|^2 \le 2\|\sigma\|^2(\|X\|^2 + \|\sigma X + Y\|^2)$, which follows from the triangle inequality and $\|\sigma\| \ge 1$. Bounding the product of the two dilated functions by half the sum of their squares costs $a^{2(d_P+e)} \le \exp(q)$ in the integral, since σ preserves real volume. The Gaussian factor is at most $[4(A+q)\|\sigma\|^2/(\exp(1)\|Y\|^2)]^q$ by $t \exp(-t) \le \exp(-1)$. This proves the bound. In the representation on Hom coordinates $\|\sigma\| \le \|h\|\|u\|$.

Disjoint balls of radius $L/2$ show that at most $(1 + 2R/L)^{2e}$ of the points have norm at most R. Integrating this count against $2qR^{-2q-1}$ for $R \ge L$ gives $\sum \|Y\|^{-2q} \le (5/L)^{2q}$, since $2q \ge 2e+2$ and $3^{2e}\, q/(q-e) \le 5^{2q}$. Consequently at the same two points the entire error is bounded by

$$|\Theta_{\Psi,P}(z_S h, (z \oplus 1)u) - C_P(h,u)| \le b^{-20\nu} \|h\|^{6\nu} \|u\|^{6\nu} C_P(1,1).$$

Here $e + 2q \le 3q$ and $100(d_P + 2q)/L^2 \le b^{-10}$. The finite trigonometric sum and the summable majorant justify the averaging.

The carrier integral in the comparison span. Keep the carrier point z_S fixed. Put $\imath h = h \oplus 1$ on the padded space. For a weight a write m_a for its product of normalized holomorphic principal-minor functions. For the following three weights the two entries apply to every positive and every negative coordinate of S, respectively:

$$k_1 = (-\Delta', 6m + 1), \qquad k_2 = (\Delta', D_\tau + r + 2m), \qquad k_3 = (0, D_\tau + N'' + 2m).$$

At each place the positive minors of P use the whole block of $2m$ positive columns and the negative minor uses the whole block of m negative columns. Linear substitution on the source therefore multiplies each such minor by the determinant of that block of the complex-linear part. In the Fock coefficient this gives

$$C_P(u, 1) = C_P(1, 1)\, m_{k_2}(u)\overline{m_{k_3}(u)}.$$

Indeed the two powers of the positive and negative block determinants from the polynomial are $\sum_h k_\tau^+(h)$ and $D_\tau + N'' + m$. The vacuum for $\widetilde{W_j}$ has m negative coordinates and $r + 2m$ positive ones. Including its splitting weight $m - D_\tau$ gives the displayed weights, up to adding the same integer in all entries of both weights. Such a shift cancels in the product because the determinant of a unitary matrix has absolute value one.

Use the same real Haar measure throughout, and put $B = 5b$, $H = B^{7\nu}$, $J = \int_{U(S)(\mathbf{R})} |m_{k_3}(u)|^2\, du$. Then $0 < J < \infty$ and

$$\int_{U(S)(\mathbf{R})} \|u\|^{6\nu} |m_{k_1}(u)|\, du \le HJ.$$

Here is the estimate at this rank. The weights k_1, k_3 have constant compact blocks, their gaps are at least 300ν by $\Delta' + 6m + 1 \ge 300\nu$ and $D_\tau \ge \Delta'$, and $\max |k_3| \le B$ by the bound on the column counts. The squared k_3-coefficient has integral the reciprocal formal degree at each place. In a Cartan decomposition the modulus of the k_1-coefficient is $|\det(u_{++})|^{-(\Delta' + 6m + 1)}$. The singular values of this block include the hyperbolic cosines in the Cartan decomposition, so $|\det(u_{++})| \ge \max(1, \|u\|/2)$ at that place. Take $a_0 = 10s^3 > 8s^2$. The local decay powers can be taken as a_0 at all places and $a_0 + 6\nu$ at one of largest norm. Both fit the gap. Multiplying the local bounds absorbs $\|u\|^{6\nu}$ at cost $2^{6\nu}$. Use the formal-degree tail integral with $L = 2\lfloor B \rfloor + s$, which bounds the absolute doubled k_3-parameter. It gives a ratio to J of at most

$$2^{6\nu + f a_0} \left[8\{2000(s + L + 1)^6\}^{2s^2} \right]^f \le H.$$

Indeed the bracket is at most $B^{17 s^2}$ for the scales here, $s = 3m \ge 6$, and $2 f s^3 \le \nu$. This proves integrability on the same group on which the carrier will be reproduced.

Define a function on the real group of W_j by

$$K_0(h) = \int_{U(S)(\mathbf{R})} \phi(z_S u) C_P(u, \imath h) m_{k_1}(u)\, du.$$

The coefficient has modulus at most $C_P(1, 1)$ by unitarity, so K_0 is bounded. Its value at one is

$$K_0(1) = \phi(z_S) C_P(1, 1) J \ne 0.$$

In fact $(0, 3m) + k_1 + k_2 = k_3$ since $N'' = r + 9m + 1$. Hence $\phi(z_S\, \cdot) m_{k_1} m_{k_2}$ is holomorphic of weight k_3. For such a function V, averaging $V(c_\alpha u)$ against the inverse compact weight of the ordered circle leaves the lowest term $V(1) m_{k_3}(u)$. Haar invariance and the same circle character of m_{k_3} then give $\int V \overline{m_{k_3}} = V(1) J$. Absolute convergence here follows from the coefficient bound and the preceding integral.

For the chosen z of the sparse-kernel estimate put

$$K_z(h) = \int_{U(S)(\mathbf{R})} \phi(z_S u)\, \Theta_{\Psi,P}\big(z_S u, (z \oplus 1)\imath h\big) m_{k_1}(u)\, du.$$

We next verify that its coefficients use precisely the prescribed finite span. On S use the additional definite space $W^\flat = I_{6m+1}$, whose rank exceeds $2 \dim S$, with characters $v^m, 1$. Fix an inclusion e of the $3m$ source coordinates and take the finite indicator δ_e of its matrix modulo $\mathfrak{M}^2$. At infinity use $Q_0 = \prod_\tau \det X^{\Delta'}_{\tau,e(+),+}$. This vector is nonzero and holomorphic of weight k_1: it is independent of the negative letters and transforms in the positive columns by the indicated determinant power, with the negative vacuum weight $6m + 1$. Fix the actual class representative p of $W^\flat$ and its invertible short basis Y_p. Every real target translate of Q_0 is in the span of the prescribed products $Q_{p,R}$ of determinants with rows of τY_p. Indeed at each embedding its transformed rows are linear combinations of those rows, and multilinearity in each row of each determinant gives exactly the row choices R.

For a rapid function v on $[U(S)]$ of principal level $\mathfrak{M}^2$ write

$$\ell_{p,R}(v) = \int_{[U(S)]} v(x) \Theta_{\delta_e, Q_{p,R}}(x, p)\, dx.$$

If these finitely many functionals vanish on v, the (δ_e, Q_0)-lift of v vanishes everywhere. Indeed the flat target is definite, so its finite class representatives and their real translates exhaust its adelic quotient at this level. The preceding polynomial expansion applies at each representative, and the finite level fixes δ_e.

Recall the sharp-data consequence of the convergent doubling formula. For a zero lift from rank s to rank $M > 2s$ with this finite inclusion indicator, convolution of the input by the real coefficient of its polynomial vanishes identically whenever that coefficient is integrable. Use the primitive-section version of the graph unfolding, since the input here is a rapid function. The marked inclusion gives injective matrices in one slot, so the Heisenberg modulation kills each boundary term. For the definite source the rational isotropic intersections are zero. At the chosen split places two supported integral frames related by a source matrix force it to be integral and to induce the identity modulo $\mathfrak{M}^2$, by the left inverse of the inclusion. The finite self-pairing there is constant on the principal subgroup and has positive mass. At any finite set off those marked pairs, cut the vacuum coefficient down to the integral group in the second argument of the doubling pairing. The first test still has the marked inclusion. For example use s integral hyperbolic pairs in the target there. For the positive standard lattice the pair $(1, \eta), (1, 0) - \xi(1, \eta)$ has Gram matrix $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ with $\eta\bar{\eta} = -1,\ \xi + \bar{\xi} = 1$ in the finite integers of the Weil datum. In its position coordinates the first vacuum is the indicator of integral s-by-s matrices, up to the fixed lattice scale. Against it take the normalized indicator of $\mathrm{GL}_s(\mathcal{O})$ in the same coordinates, retaining the vacuum on the complement. The source acts on the position matrix by right multiplication. Its local Haar Jacobian is one because the determinant of a unitary matrix has norm one to the base field. On this test the first vacuum evaluates to one exactly when the acted matrix is integral. Taking conjugate places together, preservation of the Hermitian form also makes the inverse integral. On the integral group the coefficient agrees with the vacuum self-pairing, including its unitary splitting character on the integral subgroup. This argument uses the hyperbolic polarization only at the added places. Increasing their finite set removes all nonintegral finite translates by dominated convergence in the absolutely convergent doubling formula. Every pairing with the zero lift was zero before this limit. The remaining positive finite mass multiplies the Petersson pairing of any rapid test with the real convolution. Such tests include every smooth compactly supported function on the quotient at the same finite level, so the convolution vanishes weakly there. The real coefficient is integrable and the rapid input is bounded and continuous. Dominated convergence makes the convolution continuous, so it vanishes at every point, including z_S. For the original definite source of the carrier the same argument uses the unit-modulus finite vacuum and $M = 3m > 2m$; the target has m integral hyperbolic pairs there as well. In the doubling formula use the normalized value with $M - \mathrm{Witt}(V) > s$ and the indicated integral hyperbolic pairs for the finite restrictions. [30, Theorem 2]

For Q_0 the real coefficient is $\|v_{Q_0}\|^2 m_{k_1}$. Indeed it has that holomorphic weight, and its value at one normalizes the lowest coefficient by the same circle projection and Hermitian symmetry. Thus vanishing of all $\ell_{p,R}$ implies $\int v(z_S u) m_{k_1}(u)\,du = 0$. Linear algebra on the image of $v \mapsto (\ell_{p,R}(v))$ now writes this latter functional as $\sum_{p,R} a_{p,R}\ell_{p,R}$, with coefficients independent of v.

Apply this identity to $v(x) = \phi(x)\Theta_{\Psi,P}(x,(z\oplus 1)\imath h)$. Its level is $\mathfrak{M}^2$: the second splitting character has trivial tame part and the finite datum factors through that quotient. It is rapid because ϕ is rapid and the theta kernel has at most polynomial growth. The product under $\ell_{p,R}$ is the kernel on $\widetilde{W_j} \perp W^\flat$, with polynomial $P(X_1)Q_{p,R}(X_2)$ and finite function $\Psi(X_1)\delta_e(X_2)$. The primitive additive trace pairing modulo $\mathfrak{M}^2$ gives a finite Fourier expansion of this function, say with coefficients b_T, over exactly the residue matrices in the definition of $F_i^\natural$. Absolute convergence permits termwise integration against the same ϕ. Thus

$$K_z(h) = \sum_{p,R,T} a_{p,R} b_T F^\natural_{(T,p,R)}(z(h,1)) \qquad \text{for every real } h.$$

Only Ψ and hence the b_T may vary with z. The sparse-kernel estimate, $\|\imath h\| = \|h\|$, and $|\phi| \le 2|\phi(z_S)|$ give

$$|K_z(h) - K_0(h)| \le \epsilon_0 \|h\|^{6\nu}, \qquad \epsilon_0 = 2H b^{-20\nu}|K_0(1)|.$$

Indeed $(4H+1)\epsilon_0/|K_0(1)| \le 10(5^{14}/b^6)^\nu < 1$. This proves the approximation at the untwisted square level for every point on the required minimum locus.

On W_j apply the same reproducing statement with the definite small block of rank $t - 9m - 1 > 2r$, the given inclusion datum modulo $\mathfrak{M}^2$, and the negative-letter highest vector of the small weight w. Its translates at a small class are in the span of the negative Weyl basis there, by equivariance of the letter substitution. A table supported at this class with the conjugate of any such coordinate as its entry produces the corresponding functional. To identify the type of the approximants on W_j, apply the third large-block matching identity at twist 1, with the factorization $I_{3m} s_i = s_i$ of the prescribed member matrix. It writes $F_i^\natural$ as κ_1 times a finite class sum of kernels from W_j to the positive definite large block. In these kernels the polynomial uses the positive letters of the untwisted diagram. Its highest vector is holomorphic, and substitution by all definite target translates spans its Weyl module. The sum of its staircase weight and the negative-letter weight w is k, coordinate by coordinate. Thus each $F_i^\natural$ is holomorphic of weight $k-w$. The finite data have level $\mathfrak{M}^2$ with trivial tame part off that modulus. The kernels have moderate growth, so their products with the rapid F give $FF_i^\natural$ rapid at the table level. These conclusions follow by multiplying by κ_1, before its nonvanishing is established.

Vanishing of the given table periods therefore gives the real vanishing integral in the circle isolation argument for each K_z above. Each rotation of a good-form peak retains the same first minimum. Isolation now forces $F = 0$ for the full good space at each j.

Detection at the large-block products. The carrier approximation gives functions K_g in the span of the prescribed $h \mapsto F_i^\natural(g(h,1))$, on the first-minimum locus, with

$$|K_g(h) - K_0(h)| \le \epsilon_0 \|h\|^{6n_*}, \qquad |K_0(1)| > (4B^{7n_*} + 1)\epsilon_0.$$

The function K_0 is bounded and independent of g. Let k be the holomorphic lowest weight of $U = \bar{F}$ and w the weight of the small kernel. By the untwisted third exchange in the preceding argument, each K_g is holomorphic of weight $k - w$. If all the table periods vanish, the preceding small-block reproducing calculation gives

$$\int_{U(W_j)(\mathbf{R})} F(g(h,1)) K_g(h) m_w(h)\,dh = 0.$$

Here m_w is the holomorphic principal-minor function normalized to equal 1 at the identity.

Suppose $F \neq 0$ and choose g_0 with $|U(g_0)| > \frac{1}{2}\sup|U|$. By the first-minimum bound, g_0 lies on the required locus. Every ordered compact diagonal rotation c_α preserves the real lattice norm, so put $K_\alpha =$

$K_{g_0 c_\alpha}$ there. Divide $U(g_0 c_\alpha(h,1))$ by its unit weight character on the rotation. Its circle mean is $U(g_0)m_k(h)$, by the lowest-term projection of the holomorphic principal-minor model. Write $I_k = \int |m_k(h)|^2\, dh > 0$. The weights k, w have dominant compact blocks, signed gaps at least $300n_*$, and $\max|k| \le B$. The real minor integral bound on these weights is

$$\int_{U(W_j)(\mathbf{R})} \|h\|^{6n_*} |m_w(h)|\, dh \le B^{7n_*} I_k,$$

where the integral is over the real unitary points with the same Haar measure as the reproducing identity. Equivalently one may apply the norm integral bound for a function bounded by $C\|h\|^{6n_*}$. In the vanishing identity at $g_0 c_\alpha$ replace K_α by K_{g_0}. Its integral changes by at most $4|U(g_0)|\epsilon_0 B^{7n_*} I_k$, since each approximation differs from K_0 by at most $\epsilon_0\|h\|^{6n_*}$. Now the integrand with K_{g_0} is continuous in α and has an integrable majorant. Average only this fixed function times the rotated translates. The circle mean and real reproducing identity for the weight k applied to $K_{g_0} m_w$ give $\overline{U(g_0)} K_{g_0}(1) I_k$ for that average. Thus $|K_{g_0}(1)| \le 4\epsilon_0 B^{7n_*}$ and $|K_0(1)| \le (4B^{7n_*}+1)\epsilon_0$, a contradiction. The table periods therefore detect every good form on this fixed untwisted span. The integration uses the single fixed K_{g_0}.

11.4 The common adjoint and the marked periods

The common adjoint and the finite tests. The glued indicator identity gives a product of the two block kernels at $g\widetilde{Y}$ for each of its terms. The Fock test splits into the two prescribed polynomials because the mixed Cartan tensor is supported on the disjoint blocks. The adelic point γ_i is the identity at the twist places, so its Weil scalar $\omega_i \ne 0$ uses only the fixed source character and the chosen split places. The representatives $\widetilde{Y}$ are principal modulo $\mathfrak{M}$ there and trivial at the twist places. They fix the good source forms. In the Petersson pairing change g to $g\widetilde{Y}$ in each term. Then they all have the same value, and the multiplicity is $\#\{Y : Y \equiv I_r \bmod \mathfrak{M}\}$.

For a definite group with a full integral finite compact K_0 of volume one, the mass of each class at level J is $[K_0 : J]^{-1}/|\operatorname{Aut}(c)|$, since the translate of J is quotiented by the finite stabilizer. The compact real factor in this class pairing has probability measure. The level character enters the stabilizer action on the coefficient value. Thus the class-sum normalization multiplies the quotient integral by $[K_0 : J]$. The cuspidal convergence of the block periods permits the finite sums to be taken outside the integral.

Write $F_i^\natural$ for the fixed large-block functions and $\mathcal{T}_i(\xi)$ for the small-block theta linear in the i-th table row. The matched large-block identity gives $F_i^\natural = \kappa_\psi$ times the large-block sum before its normalization, with $|\kappa_\psi|^2 = C/\Gamma_j(\psi)$ for a common C. The unfolding just described gives, with the second entry linear,

$$\overline{\kappa_\psi}\langle F, u_\psi^* G_\psi \xi\rangle = C_j \sum_i \overline{\omega_i} \int_{[U(W_j)]} \overline{F F_i^\natural}\, \mathcal{T}_i(\xi),$$

where $C_j \ne 0$ contains g^{L_j}, the conjugated stage multiplier and the congruence multiplicity. This right-hand side is independent of ψ. Use the detection statement for the actual products $F_i^\natural \mathcal{T}_i$. If some $\kappa_\psi = 0$, the matched identity would give $F_i^\natural = 0$ for every i, and detection would force the good source space U_j to be zero. Its dimension is $R > 0$. Similarly, if F is orthogonal to $u_1^* G_1$ of all tables, take a table supported on one i. Its right-hand period vanishes for every small-block table. Detection gives $F = 0$. Thus the common map $A = u_1^* G_1$ is onto, and the scalar is $\gamma_\psi = \overline{\kappa_1/\kappa_\psi}$. It satisfies $|\gamma_\psi|^2 = \Gamma_j(\psi)/\Gamma_j(1)$.

Two applications of the marked period. At each eligible split place w off the level, let $H_w(\xi_w)$ be the Hecke operator at the genuine root from the rectangular vacuum calculation. The comparison forms will be tested on the subspace

$$\operatorname{im} e_0 \cap \ker H_w(\xi_w) \cap L_\psi^\perp. \tag{11.1}$$

Here e_0 is the good tuple idempotent in the target action. In the rank-one space use the trivial tuple idempotent e_{triv} and the descent Θ to $U(1)$. The point forms will initially be tested on the subspace

$$\operatorname{im} e_{\mathrm{triv}} \cap \ker H_w(\xi_w) \cap \ker \Theta. \tag{11.2}$$

In each case testedness means orthogonality to the indicated subspace at every eligible w, in the class-sum pairing.

The marked period for the comparison form integrates over the small block of rank $t - 9m - 1$, with the block of rank $r + 9m + 1$ held fixed. For Y in (11.1), the theta descent is good by equivariance at the chosen places. Adjunction then makes the descent zero. The reflection sums at these places kill Y: their normalized sums are projectors on the determinant character of the rank-one quotient of the corresponding hyperplane stabilizer. Under the finite parabolic induction rule this component can occur only in a Green label with linear factor named by that character. The defining good-label condition excludes that factor for the ω^j on the actual marked block. At a new twist place the character is nontrivial by definition of the conductor. The required rank gap is $2r + 2(r + 9m + 1) + 3 < 6r$, which follows from $r \gg m$. At a chosen place the small-block principal level modulo $\mathfrak{M}^2$ fixes the small kernel. Conjugation of such an element by γ_i changes the identity by a matrix in $\varpi M_{r+t}(\mathcal{O}_v)$, because both $\gamma_{i,v}$ and $\varpi\gamma_{i,v}^{-1}$ are integral. At the conjugate place use the conjugate inverse transpose. At twist places γ_i is the identity and the small-block action keeps the marked rows in the fixed block. These conditions give the same compact stabilizer required by the period, including its determinant character.

For the rank-one point test, keep the first three blocks fixed and expand the flat lift on the last block into its rank-one source classes. For each such class the last-block sum is a marked period of the same form. The compact rank inequality is $2 + 2(12 + 4 + 4) + 3 = 45 < 2z_p$. For Y in (11.2) the label at each chosen place is trivial and its descent is zero. The reflection sum on $\mathrm{GL}_1(k_v)$ is the sum of ω_v^{-j} over $k_v^\times$ and vanishes on this summand. At twist places the rank-one character has the nontrivial tame part of ψ. At chosen places the last block is at the level fixing J_1, which fixes both the flat column datum and the full point form's column $(x_0, J_2, J_3, J_1)^{\mathrm{t}}$. The determinant character is read in the second block and stays fixed. The Cartan contraction against the flat form is equivariant because the flat coefficient type is trivial and the first twelve columns are fixed, as in the point-form construction.

Both applications use the same marked-period implication. Suppose that Y has zero normalized marked descent, the chosen-place reflection sums vanish on Y, and the tame character is nontrivial at every new twist place. Assume the off-level split Hecke equation at $(a_{\mathrm{Frob},r}(w)q_w^{n-r-1})^{-1}$, with $a_{\mathrm{Frob},r}(w)$ the rank-r Frobenius value of the source splitting character. The datum is vacuum off the chosen and new twist pairs. When the fixed block of rank b satisfies $2r + 2b + 3 < n$, the equivariant period over its definite complement vanishes on the observed residue frames. The proof uses the normalized Weil actions on these data and the mixed zero-term formula in the seesaw calculation.

From the marked section to the allowed descent data. Write $A = \mathscr{M} \cup Q$ for the chosen and new twist places of the CM field E. These are split and pairwise distinct even up to conjugation. Let $F = E^+$, $V = I_n$, and let the lower space W have rank r. Use the marked column frame $I_{\mathrm{mk}} = \binom{I_r}{0}$ in each indicated ordering of the rows. Write $Z_v = \binom{Z_{0,v}}{Z_{1,v}}$ in these coordinates. Let χ_v be the residue character of the actual auxiliary pair. The marked finite function, with vacuum at the other finite pairs, is

$$\Psi_0(Z) = \prod_{v\in A} 1_{M_{n,r}(\mathcal{O}_v)}(Z_v) \prod_{v\in\mathscr{M}} 1_{\overline{Z}_v = I_{\mathrm{mk}}} \prod_{v\in Q} \chi_v(\det \overline{Z}_{0,v}),$$

where the residue character is extended by zero at zero. The level K_A fixes the marked columns modulo $\mathscr{M}$ and is block lower triangular modulo Q. Its character on the form is $\lambda(k) = \prod_{v\in Q} \chi_v(\det \overline{k}_{00,v})^{-1}$. The vacuum condition on a datum means the vacuum at every finite pair outside $A \cup \overline{A}$. At a designated split place use the Schrödinger coordinates $M_{n,r}(E_v)$, with the other place of the pair fixed by the unitary relation.

Let y be the algebraic form whose marked descent vanishes, with real Weyl type ρ on L. For a linear functional u on L, use the conjugate scalar coefficient $Y_u(h) = \overline{u(\rho(h_\infty^{-1})y(h_f))}$. For a real Schwartz vector P and $g \in U(W)(\mathbb{A}_F)$, write $D(\Phi, y)$ for the descent of Y_u at g with section $P \otimes \Phi \otimes \mathrm{vac}$ and splitting characters $(a, 1)$. This pairing is linear in Φ and conjugate-linear in y. The defining marked class sum agrees with the conjugate of the adelic descent on the highest harmonic and highest coefficient, multiplied by $\dim L\,[K_0 : K_A]$. Here L is the actual archimedean Weyl type on V, with the highest vector of unit Gram norm. Indeed, on each class the compact average contracts the two L-coordinates by Schur orthogonality with scalar $1/\dim L$. This is the usual compact integral of $\langle u, \rho(k)x\rangle\overline{\langle u, \rho(k)y\rangle}$. The simultaneous finite actions have reciprocal phases, and the class measure is $[K_0 : K_A]^{-1}/|\,\mathrm{Aut}\,|$ with the original volume-one compact factors. Erasing the second splitting character in the descent only removes its character of $\det g$, which is nonzero. Thus the highest harmonic gives zero at every g.

The same holds for every real Schwartz vector with the same marked finite section and for every coefficient u. Identify each coefficient by the Gram form with $u = u_\zeta$, $u_\zeta(v) = \langle \zeta, v\rangle_L$, linear in v. Then $\zeta \mapsto Y_{u_\zeta}$ is linear. For the compact action $R(k)$ it takes $\rho(k)\zeta$ to $R(k)Y_{u_\zeta}$, since the coefficient at hk reads $\rho(k)^{-1}\rho(h_\infty^{-1})y(h_f)$. The bilinear integral is invariant under simultaneous right actions of $U(V)(F_\infty)$ on the oscillator and this coefficient. In the polynomial oscillator it therefore factors through contraction of L with $L^\vee$. At each real place Howe's compact-member theorem gives irreducibility of the multiplicity space of $L^\vee$ for the real W-side Harish–Chandra actions. Write J for the tensor product of those multiplicity spaces. In the marked polynomial the positive minors substitute the transpose of the $U(n)$-matrix and the negative minors its inverse. The first splitting character is $a = v^{r_-}$ with r_- the number of negative columns of W. The negative-column diagram contributes determinant exponent $-\sum_{i=1}^{n-1} c^-_{\tau,i}$ at τ, where $c^-_{\tau,i}$ counts the negative columns of height i. These are the determinant powers of L, and the substitutions make $\bar\zeta \mapsto P_{\bar\zeta}$ an intertwiner from $L^\vee$. It is nonzero on the paired highest line, which is the product of the staircase minors. The compact contraction of this polynomial with ζ at the paired highest vector therefore gives a nonzero vector of J. The same vector generates J. Moving each real W-action to the evaluation point g, and then differentiating, gives zero on the generated polynomial module. The polynomial Fock vectors correspond to finite Hermite sums, dense in real Schwartz space. Continuity of the integral on this space follows from the theta lattice sum over the compact quotient of $U(V)$. This proves $D(\Psi_0, y) = 0$ with all P, u, g.[20, p. 542, (3.9)(a),(d)]

For any $B = (B_v)_{v\in A}$ with $B_v \in \mathrm{GL}_r(E_v)$, embed B as an element of $U(W)(\mathbb{A}_{F,f})$ on those split pairs. The source Schrödinger action on the marked datum replaces $\Psi_0(Z)$ by $\Psi_0(ZB)$ times a nonzero scalar. This follows from change of variables $Z \mapsto ZB$ in the Schrödinger model. The scalar consists of the splitting character and $\prod_v |\det B_v|_v^{n/2}$. Right translation of g by this element then gives $D(\Psi_0(\cdot\, B), y) = 0$. For the integral target elements k at A the same model acts as $\Phi(Z) \mapsto \Phi(k^{-1}Z)$ since a is an unramified splitting power at these places. Haar invariance on $[U(V)]$ gives

$$D(\Phi(k^{-1}\,\cdot), R(k)y) = D(\Phi, y). \tag{11.3}$$

The character projection by K_A, acting on functions by $\overline{\lambda(k)}\Phi(k^{-1}Z)$, thus leaves $D(\Phi, y)$ unchanged. The projection uses compact Haar probability, or the identical average over a finite quotient fixing the function on its compact support.

Fix a centre Z and a product of small matrix boxes C about it. Suppose first that all $Z_{0,v}$ are invertible and the ratios $Z_{1,v}Z_{0,v}^{-1}$ are integral, in the maximal ideal at the chosen places. With $B_v = Z_{0,v}^{-1}$, the function $\Psi_0(\,\cdot\, B)$ is nonzero at Z. Its support is a single K_A-orbit. After multiplying the argument by B on the right, a point of the support is $\binom{X_0}{X_1}$ with $X_0 \in \mathrm{GL}_r(\mathcal{O}_v)$. The integral completion $\begin{pmatrix} X_0 & 0 \\ X_1 & I \end{pmatrix}$ takes I_{mk} to this point, fixes the marked columns modulo the chosen place, and is block lower triangular modulo a twist place. The orbit is compact and open in the product of matrix spaces. The marked function transforms on it by $\Psi_0(kZB) = \overline{\lambda(k)}\Psi_0(ZB)$. Take C contained in this orbit. The character projection of 1_C is supported there and has the same transformation rule. Transitivity makes it a scalar multiple of $\Psi_0(\,\cdot\, B)$. Thus $D(1_C, y) = 0$.

If the ratio condition fails at v, there is an integral row $\ell = (\ell_0, \ell_1)$ with

$$\ell Z_v = 0, \qquad \bar{\ell}_0 \neq 0, \qquad \bar{\ell}_1 = 0 \quad (v \in Q).$$

For singular $Z_{0,v}$, scale a nonzero left null row of that block to have integral entries and a unit entry, and append zeros. Otherwise take a row a of $Z_{1,v}Z_{0,v}^{-1}$ with an entry of maximal modulus at least one at a chosen place or greater than one at a twist place. Append the negative of the corresponding coordinate row in the lower block and divide by the entry of maximal modulus. This gives the three asserted properties.

For every $s = 1 + t\bar{\ell}_0 \in \mathrm{GL}_r(k_v)$, lift the column t integrally and put $k_t = 1 + I_{\mathrm{mk}}\widetilde{t}\,\ell$. The determinant of k_t is a unit because its residue is $\det s = 1 + \bar{\ell}_0 t$. It fixes Z_v and hence every matrix box about Z_v. At a chosen place k_t differs from the lift of $\mathrm{diag}(s, 1)$ by a level element of character one. Writing $\pi_v(s)$ for the chosen action twisted by $\chi_v(\det s)$, we get $R(k_t)y = \chi_v(\det s)^{-1}\pi_v(s)y$. Equation (11.3) makes each term of the weighted reflection sum pair with 1_C to $D(1_C, y)$. The sum kills y, so its nonzero number of terms times $D(1_C, y)$ is zero. At a twist place $\bar{\ell}_1 = 0$, so k_t belongs to the level and acts by $\chi_v(\det s)^{-1}$. The map $t \mapsto 1 + \bar{\ell}_0 t$ takes all nonzero residue values among these s, and the character is nontrivial. Equation (11.3) again gives $D(1_C, y) = 0$. Every locally constant compactly supported function at A is a finite linear combination of such indicators. In the good case one can shrink boxes inside the orbit, using compactness of the support and disjoint refinement of the finitely many boxes. This proves vanishing on all allowed descents to W.

Integral frames and adapted compact subgroups. The field contains a primitive 5040-th root of unity. At places of residue characteristic different from three a primitive cube root has distinct reduction from its conjugate. At the places over three the fourth roots i and $-i$ have distinct reductions. Thus conjugation cannot lie in the inertia subgroup of any nonsplit quadratic local extension of E/F, and such a local extension is unramified. In an unramified quadratic ring every unimodular Hermitian lattice has an orthonormal basis. In its residue form some vector has nonzero norm, by nondegeneracy, polarization and the surjective residue trace. The unit norm map is surjective, first on the residue field and then on the principal units by lifting with the residue trace. Rescale the lifted vector to norm one and repeat on its orthogonal complement. This also extends every integral orthonormal frame. At a split place such a frame gives a matrix inclusion and its left inverse over the valuation ring. Extend its basis by a basis of the kernel of the left inverse and use the inverse transpose at the paired place.

For an isotropic line at a nonsplit finite place choose a primitive integral vector v in it. The unimodular pairing gives an integral partner w with $\langle v, w\rangle = 1$. Replace w by $w - cv$, where $c + \bar{c} = \langle w, w\rangle$, to make w isotropic. The trace onto the base valuation ring is surjective for the unramified extension, also at the dyadic places. The orthogonal complement of the resulting integral hyperbolic plane is unimodular, so the previous basis argument identifies it with the complement for the adapted plane. At a split place the isotropic line gives a line and its annihilating hyperplane in one factor. Integral bases for the saturated flag give transitivity under the integral compact. At infinity the positive and negative parts of an isotropic vector have the same norm, so the product of the two definite unitary groups is transitive on the isotropic lines. Factoring the stabilizer of the line into its Levi and unipotent radical gives the adapted Iwasawa decomposition used below.

Bounded descents up the hyperbolic tower. Let $W_j = W \oplus \mathbb{H}^j$, $a_j = av^j$. Suppose the allowed descents to W_{j-1} vanish. For a fixed allowed section ϕ on $\mathrm{Hom}(W_j, V)$, write

$$F_j(g) = \int_{[U(V)]} \Theta_{W_j,V}(\phi; g, h)\, Y_u(h)\, dh.$$

We show that F_j is bounded. All measures on definite groups are the original quotient measures. The additive means used here are Haar probability on the compact adelic quotients.

Use the adapted Witt splitting of the last plane. Let k be its designated isotropic vector and J_j the form on W_j. The isotropic space of maps $Z = \{d\, k^* J_j \,:\, d \in V\} \subset \mathrm{Hom}(W_j, V)$ acts by modulation of the mixed variable. Its mean on theta, $\Theta_0 = \int_{[Z]} \Theta \circ \rho(z)\, dz$, selects the mixed slice at zero by additive Fourier orthogonality. This is the normalized zero term of the lower-chain calculation, with the same Haar and vacuum scales. For each g, decompose $g = umk_c$ relative to the last-plane parabolic, with k_c in the adapted maximal compact. The zero slice of $\omega(g,h)\phi$ is, up to the nonzero scalar of the plane Levi m_1, the lower Weil action on the zero slice of $\omega(k_c, 1)\phi$. The unipotent part acts trivially on the zero slice. For the Levi, the action $X \mapsto hXg^{-1}$ on the isotropic map block has determinant $\det(h)\,\overline{\det(m_1)}^{\,n}$. The normalized zero-term law multiplies by v of this determinant to power -1 and its absolute norm to power $1/2$. The h-factor changes a_j to a_{j-1} since $\det h$ has norm one. The scalar of m_1 is $v(\det m_1)^n |\det m_1|_E^{n/2}$ for second character 1. In the adapted finite bases the vacuum is a product over the retained space and the split plane. Compact action off $A \cup \overline{A}$ preserves the vacuum line, and the zero slice of the split vacuum is a nonzero multiple of the retained vacuum. Hence the lower section is allowed. The inductive vanishing says that Θ_0 has zero integral against Y_u at every g.

Here is the uniform estimate for the difference $\Theta - \Theta_0$. For a point g let $H(g) = \max(1, \lambda_1(g)^{-1})$, with $\lambda_1(g)$ the first Euclidean minimum of the moved integral lattice in normalized archimedean coordinates. We first recall the polynomial bound for Θ in this height and give the reduction needed for its uniform use. All moved lattices have one covolume by the determinant norm identity for an adelic unitary matrix. Those with a fixed positive lower bound on λ_1 have uniformly bounded Euclidean $\mathbb{Z}$-bases and basis inverses. One proof chooses a shortest primitive vector v, bounded above by the volume pigeonhole bound. The orthogonal quotient lattice has covolume divided by $\|v\|$ and its shortest vector has length at least $\sqrt{3}\|v\|/2$, by rounding a lift by multiples of v. Induction gives a bounded basis of the quotient and bounded lifts. The bounded covolume and basis then bound the inverse. In these bases the matrices of an integral basis of $\mathcal{O}_E$ and the Hermitian Gram entries take finitely many values. They are integral and bounded at infinity. For a fixed table of these values, the map between the two bases is a rational E-linear isometry. It sends one lattice to the other, so the moved finite matrices differ modulo the fixed integral compact. The real coordinates and inverse coordinates are bounded in this identification. This proves cocompactness of the bounded-height part modulo rational points. In particular the definite group has a compact quotient and admits, at every h, an E-basis $d_1, \dots, d_n$ of integral vectors of uniformly bounded moved lengths.

Two moved integral vectors of length less than one have Hermitian product zero. Indeed its value is an algebraic integer of absolute value less than one at every embedding. At each real place the span of these short vectors is thus totally isotropic. For the diagonal signature matrix J, an orthonormal basis (v_i) of this span and the vectors (Jv_i) are mutually Euclidean orthogonal isotropic frames dual for the Hermitian form. The real unitary dilation with factors $2, 1/2$ on these frames and 1 on the complement doubles the short lengths and shrinks no length by more than two. Archimedean right translation by its inverse changes the minimum to at least $\min(2\lambda_1, 1/2)$. At most $\log_2 H(g) + O(1)$ such steps bring the point to the bounded-height region. Thus a rational representative lies in a fixed compact set times archimedean matrices whose norms and inverse norms are $O(H(g))$.

For any fixed finite Schwartz factor ϕ_f, the theta distribution in a rational Schrödinger polarization is a lattice sum and is bounded by a fixed real Schwartz seminorm. The finite compact representatives give finitely many translates of ϕ_f since it has an open stabilizer. One can take the real seminorm p_s to be the sum of L^2-norms of products of at most s infinitesimal Heisenberg coordinate operators. For s large this controls the polynomially weighted supremum norm used in the lattice sum, by Euclidean Sobolev embedding. Weyl covariance moves each coordinate operator through the real Weil action by the corresponding inverse symplectic matrix. Unitarity then bounds p_s of its translate by a fixed power

of the matrix and inverse norms times p_s of the vector. The matrix on the space of maps is $X \mapsto hXg^{-1}$. Rational invariance, the previous reduction, and compactness for the definite variable therefore give

$$\left|\Theta(\omega(g,h)(f_\infty \otimes \phi_f))\right| \le CH(g)^B p_s(f_\infty)$$

uniformly in h, g, with B, s, C depending on the fixed finite factor and the spaces. The constant works on every real vector with this finite factor.

When $H(g) > 2$, a shortest integral vector is isotropic by the same algebraic-integer argument. Rational Witt extension takes it to the fixed k in the last plane. Replace g by this rational translate, so that k has moved length $H(g)^{-1}$ and integral finite coordinates. The rank-one maps $d_i k^* J_j$ form an E-basis of Z. Their moved coordinates are integral at finite places and $O(H(g)^{-1})$ at infinity, by multiplying the coordinates of d_i by the Hermitian dual coordinates of k.

Take $\phi = f_\infty \otimes \phi_f$; finite sums of such tensors suffice. Choose a positive integer M such that finite Heisenberg translations by $M\widehat{\mathcal{O}}_E$ in each normalized matrix coordinate fix ϕ_f. This is possible by smoothness of the finite oscillator. Represent each coefficient in $\mathbb{A}_E/E$ by an adele with finite part in $M\widehat{\mathcal{O}}_E$ and bounded archimedean part. Additive approximation first clears the finite congruences, after which reduction by $M\mathcal{O}_E$ bounds the archimedean part. For $z \in [Z]$ use these representatives in the displayed basis of Z. After applying Weil covariance to move $\rho(z)$ across $\omega(g,h)$, its finite part fixes ϕ_f, and it is a real Heisenberg translation of size $O(H(g)^{-1})$. This uses the same normalized product coordinates at every place.

For an integer $N \ge 1$, rational invariance and the isotropy of Z give

$$(\Theta - \Theta_0)(\omega(g,h)\phi) = \int_{[Z]^N} \Theta\left(\prod_{i=1}^{N}(1-\rho(z_i))\,\omega(g,h)\phi\right) d\mathbf{z}.$$

Indeed each nonempty subset of the factors averages the translate by a Haar-uniform point of $[Z]$, and the alternating coefficients of the nonempty subsets sum to -1. In this identity rational changes of any representative have no effect, since the shifts on Z commute and rational shifts fix theta. For a small real Heisenberg shift v, $\rho(v) - 1 = \int_0^1 \rho(tv)\, d\rho(v)\, dt$. Each factor costs one order of p_s and gains $O(\|v\|)$. Applying the theta bound with the same finite factor gives $O(H(g)^{B-N} p_{s+N}(f_\infty))$. Choose $N \ge B$. For $H(g) \le 2$, compact reduction bounds the kernel and the lower zero-slice integral is still zero. Rational invariance of the ordinary descent permits the chosen translate when $H(g) > 2$. Now integrating the difference proves boundedness of F_j, since Y_u is bounded on the definite quotient of finite measure.[4]

The Hecke equation in the tower. Use the original split place w off $A \cup \overline{A}$ and let $q = q_w$. Set $m = r+2j$, $e = n - m$ and use the conjugate root

$$X = \overline{(a_{\mathrm{Frob},r}(w) q^{n-r-1})^{-1}}, \qquad |X| = q^{1-n+r/2}$$

on the conjugate scalar coefficient. The coefficient map is conjugate-linear and commutes with the finite right actions, so $H_w(X)Y_u = 0$. The datum is the vacuum at this pair of places, and both relevant splitting characters are unramified. Use the Hecke polynomial with the K_w-double-coset coefficient $(-X)^i q^{i(i-1)/2}$, Haar measure giving K_w mass one. Put

$$\alpha = a_{j,\mathrm{Frob},m}(\overline{w}), \qquad \beta = 1_{\mathrm{Frob},n}(\overline{w}), \qquad X' = q^e \alpha X/\beta.$$

The character defining $a_{j,\mathrm{Frob},m}(v)$ is $z \mapsto a_j(z/\bar{z})v(z)^{-m}$, evaluated at the uniformizer and multiplied by $q_v^{m/2}$. The β convention is the same in rank n with trivial norm-one character.

For the bilinear descent integral, move the whole bi-compact double coset through the quotient integral. The translating elements are inverted with no measure factor because the split group is unimodular.

[4]The real seminorm argument also uses the normalized zero term in the lower-chain calculation.

Inverse conjugate transpose then identifies the reversed double coset with the positive double coset at $\bar{w}$, where the compact again has mass one. Evaluate its action on the vacuum at $\bar{w}$ in matrix coordinates $Z \in M_{n,m}(E_{\bar{w}})$. For $U \subset k_{\bar{w}}^{n}$, choose an integral γ_U of determinant valuation $\dim U$ with kernel U modulo the uniformizer and elementary divisors only $1, \varpi$. The target double-coset representatives are γ_U^{t}. The Schrödinger multiplier on the vacuum is $\alpha^{\dim U} 1_{Z^{\mathrm{t}}\gamma_U^{-1}\ \mathrm{integral}}$, from the determinant character and $|\det \gamma_U|^{-m/2}$. The indicator is one exactly when Z is integral and $U \subset \ker \overline{Z}^{\mathrm{t}}$. The source action of $H_w(X')$ in these opposite coordinates uses γ_U^{-1} on the right, with $U \subset k_{\bar{w}}^{m}$. Its vacuum multiplier is $\beta^{\dim U} 1_{Z\gamma_U^{-1}\ \mathrm{integral}}$, with the condition that Z is integral and $U \subset \ker \overline{Z}$. These representatives come from the lattices between $\varpi\mathcal{O}^a$ and $\mathcal{O}^a$, with transpose interchanging the left and right cosets. For a kernel of dimension d, summing the coefficients by the subspace q-binomial identity gives $\prod_{i<d}(1 - q^i \alpha X)$, or the same expression with $\beta X'$ in place of αX on the source. The two kernel dimensions differ by e for every integral Z. Since $\beta X' = q^e \alpha X$, the product on the target splits at e. The identity on the vacuum and source equivariance now give the equation on F_j

$$0 = \left(\prod_{i<e} (1 - q^i \alpha X) \right) H_w(X') F_j.$$

The spherical property follows from the compact action on the vacuum at w. The transfer is unchanged for arbitrary data at other places.

At this stage $1 \le j \le b_{\mathrm{fix}}$ and $n > 2r + 2b_{\mathrm{fix}} + 3$. We have $e + m/2 = n - r/2 - j \le n - r/2 - 1$ and

$$m + 1 + e/2 = (n + r + 2j)/2 + 1 \le n - r/2 - 1.$$

Thus $q^e|\alpha X| = q^{1-j} \le 1$, so every factor in the product is nonzero. Also $q^{m+1}|X'| = q^{(2r+2j+4-n)/2} \le 1$. For a bounded spherical function f on rank m, the nonconstant terms of $H_w(X')$ have total absolute coefficient at most

$$\prod_{i<m} (1 + q^i |X'|) - 1 \le \exp\left(|X'| \sum_{i<m} q^i \right) - 1 < 1.$$

Here $\sum_{i<m} q^i \le q^m$, so the exponent is at most $1/q \le 1/2$. The coefficient of degree zero is the identity. Taking sup norms in $H_w(X')f = 0$ forces $f = 0$. It applies to F_j by the preceding estimate. Induction therefore gives vanishing for all allowed data on $W \oplus \mathbb{H}^{b_{\mathrm{fix}}}$ with character $a v^{b_{\mathrm{fix}}}$.

Extraction of the complementary period. Write $V = V_1 \oplus V_2$ in the actual block ordering, with $V_1 = I_b$, $b = b_{\mathrm{fix}}$, and $H = U(V_2)$. We first use the split smaller group $\mathbb{H}^b$. Let f be continuous on $[U(V)]$, right invariant under the integral compact K_{off} supported off $A \cup \overline{A}$, with every allowed descent to $\mathbb{H}^b$ zero for the character v^b. For the standard frame $x_0 = \binom{I_b}{0}$ and an arbitrary $h_1 \in U(V_1)(\mathbb{A}_F)$, put $X(h) = h^{-1} x_0 h_1$. In the split mixed model the central element with upper-right block $B - B^*$ multiplies the summand for a rational frame x by $\psi_E(-\jmath \operatorname{tr}(x^* x B))$. This follows by substituting $-2\jmath B$ in the radical action, whose scalar at the mixed variable x is $\psi_E(\frac{1}{2} \operatorname{tr}(z x^* x))$. The Fourier coefficient of the descent at $\jmath I_b$ is the mean with the multiplier $\psi_E(\jmath \operatorname{tr} B)$ on $M_b(\mathbb{A}_E)/M_b(E)$, with probability Haar. Additive orthogonality selects exactly the frames of Gram matrix I_b. Rational Witt extension makes these frames a single orbit with stabilizer $H(F)$. Unfolding gives integration over $H(F)\backslash U(V)(\mathbb{A}_F)$ with its quotient measure. At a fixed value of the source Levi the Schwartz sums are absolutely integrable. In fact a compact set of representatives of $[U(V)]$ gives finitely many finite supports and uniformly bounded real Schwartz seminorms, and the rational sum on each support is absolutely and uniformly convergent. This justifies Fourier integration and unfolding term by term.

Let κ be the scaling idele of the Weil model. Its finite component is α_f. At τ its real component is $(2\pi \operatorname{Im} \tau(\jmath))^{-1/2}$. Put $\Delta = \jmath \kappa \bar{\kappa}$. In the split mixed model use the finite vacuum and choose the real Schwartz vector to be $\exp(-\sum_{\tau,i,j} |X_{ij,\tau}|^2)$ in the variable $x = \kappa X$. At the finite places the two isotropic

half lattices in these coordinates are dual under $(B, X) \mapsto \psi_E(\Delta_f \operatorname{tr}(B^{\mathrm{t}} X))$. Indeed $\Delta_f = \jmath \alpha_f \bar{\alpha}_f$ has the inverse-different valuations. The normalized vacuum is fixed by Heisenberg shifts in each of these half lattices. To check the phase here, its finite matrix coefficient on the integral phase-space lattice is the unit phase $\psi_f((c_0 \jmath/2)\langle v, v\rangle)$ in the standard finite frame, with $c_0 = (1 - \jmath^{-1})/2$. Each half lattice is isotropic, so its phase is one, including at dyadic places. Invariance under the modulating half forces support on integral X, by self-duality, and invariance under the translating half makes the function constant there. The finite vector is nonzero, so the constant is nonzero. This determines every allowed off-level finite vacuum component and allows the same choice at the marked pairs. Thus the combined datum at κX is, up to a common nonzero constant,

$$w(X) = \exp\left(- \sum_{\tau,i,j} |X_{ij,\tau}|^2 \right) 1_{\widehat{\mathcal{O}}_E^{n \times b}}(X_f).$$

A Heisenberg shift with mixed coordinates $yk = 0$ and $yl = \kappa \bar{u}$ in the adapted target frames leaves the mixed argument unchanged. The mixed action multiplies the datum by the character with argument $(\jmath/2) \operatorname{tr}((\overline{\kappa \bar{u}})^{\mathrm{t}}(x + x))$. At $x = \kappa X$ this is $\jmath \kappa \bar{\kappa} \sum u_{ij} X_{ij}$. Taking u supported at infinity and the marked pairs leaves every off-level vacuum intact. The determinant of h contributes no character in the split model since $v^b(\det h) v(\det h)^{-b} = 1$. Using the Levi matrix κh_1 on the b-plane gives $X(h)$ above and contributes the fixed nonzero Levi scalar. Consequently vanishing of these Fourier coefficients gives

$$\int_{H(F) \backslash U(V)(\mathbb{A}_F)} \psi_E\big(\Delta \sum u_{ij} X(h)_{ij}\big)\, w(X(h)) f(h)\, dh = 0 \qquad \text{for all such } u.$$

Disintegrate the last displayed integral through $H(\mathbb{A}_F) \backslash U(V)(\mathbb{A}_F)$, with the measure induced by the original Haar measures on H and $U(V)$. The function to be tested there is $P(h) = \int_{[H]} f(th)\, dt$. It is continuous by compactness of $[H]$. The frame map identifies the cosets with their frames $X(h)$. By integral transitivity its integral support is compact modulo $H(\mathbb{A}_F)$. The function P there depends only on the coordinates of X at infinity and the marked pairs. Indeed if these coordinates agree, an element of K_{off} carries one full integral frame into the other. Left H and right compact invariance then give equality of P. The induced function on the compact image of the observed coordinates is continuous. Push the positive measure $w(X(h))\, d\dot{h}$ to this image. The additive characters tested above separate its points. Their linear span is a self-conjugate algebra containing the constant function. By Stone–Weierstrass the induced function times this measure is zero. The image of x_0 is in its support, since x_0 occurs at $\operatorname{diag}(h_1, 1)$, integrality is open there, and the Gaussian is positive. Thus $\int_{[H]} f(\operatorname{diag}(h_1, t))\, dt = 0$.

To apply this to the original Y_u, fix an allowed small-block datum Φ_2, its character b_2 on the source determinant, and the actual g, h_1 of the marked table. Its finite translate is $\gamma = x A_{\mathrm{pt}}$, where x is the divided translation and A_{pt} the integral point lift of the prescribed residue frames. Both are supported on $A \cup \overline{A}$, so $Y_u(h\gamma)$ has the same off-level invariance and its allowed top descents still vanish. This follows by right translation of the datum in the integral, leaving the off-level components fixed. Choose a datum Φ_1 on the first block with fixed finite vacuum and $\Theta^{a, b_2^{-1}}_{W, V_1}(\Phi_1; g, h_1) \neq 0$. One can choose it directly in a rational Schrödinger polarization. The finite Weil translate of the vacuum is a nonzero locally constant compactly supported function. A rational vector meets its nonzero set by finite additive approximation. The rational vectors on its finite support are discrete at infinity because their denominators are bounded. Choose a real smooth compactly supported section after the real Weil translation that isolates the chosen rational vector, and apply the inverse real Weil translation. The theta sum at the given point is nonzero.

The function

$$f(h) = \Theta^{a,1}_{W,V}(\Phi_1 \otimes \Phi_2; g, h) Y_u(h\gamma)$$

has zero allowed descents to $\mathbb{H}^b$ by the top vanishing for $W \oplus \mathbb{H}^b$. In the orthogonal sum on the source the two kernels multiply, and the product character is $a v^b$. Its right K_{off}-invariance follows from vacuum invariance and the support of γ. On the block $V_1 \oplus V_2$ the same tensor product on the target gives the two

theta factors with b_2^{-1}, b_2, whose product character is 1. The fixed first factor is nonzero, so the extracted period gives, upon taking complex conjugates,

$$\int_{[H]} \overline{\Theta^{a,b_2}_{W,V_2}(\Phi_2; g, t)}\, \overline{Y_u(\mathrm{diag}(h_1, t)\gamma)}\, dt = 0.$$

All these integrals use the given class measures. Expanding the Cartan contraction in the actual Weyl coordinates gives a finite sum of these periods. The invariant Gram form and Cartan equivariance cancel the real compact action on the sum, so its integral is $[K_{0,H} : K]^{-1} \sum_c f_c / |\operatorname{Aut}(c)|$ when the resulting finite contraction is fixed by the actual chosen level K of H. The original class sum therefore vanishes with its specified weights. In the comparison application the principal level at $\mathscr{M}^2$ fixes the kernel and $\gamma_{i,v}^{-1}(k_v - 1)\gamma_{i,v}$ is in $\varpi_v M_n(\mathcal{O}_v)$ at every chosen v. Here $\gamma_{i,v}$ is integral by the conjugate side of the divided-lattice condition and $\varpi_v \gamma_{i,v}^{-1}$ is integral by its chosen side. At the conjugate place use the inverse transpose. At a chosen place $A_{\mathrm{pt}} I_{\mathrm{mk}}$ is the specified column in the table. The conjugated complementary element fixes it and hence conjugation by A_{pt} fixes I_{mk}. At twist places $x = 1$ and $\overline{A}_{\mathrm{pt}}^{\mathsf{t}} I_{e'} = I_{\mathrm{mk}}$ for the prescribed rows e' of the fixed block. The complementary motion fixes those rows, so its conjugate by $\overline{A}_{\mathrm{pt}}$ fixes I_{mk} under transpose. The corresponding determinant in the level character is one. Thus the conjugation uses the original divided lattice and table coordinates. The rank gap is $2r + 2(r + 9m + 1) + 3 < 6r$, since $r + 1 = \kappa + m$ and $\kappa > 100m + 9$. In the point application the rank is $r = 1$, the first $12 + 4 + 4 = 20$ positions are fixed and $45 < 2z_p$. At the chosen places the level fixes J_1 and therefore the full column $(x_0, J_2, J_3, J_1)^{\mathsf{t}}$. The determinant character is read from the second block. When the first three factors are fixed, the last-block contraction is on the same flat character and the same chosen-level classes as the point test. Its real compact type is trivial, so the Cartan equivariance used in the class sum is just invariance on this last block. In both cases the finite translate is $\gamma = xA_{\mathrm{pt}}$, the specified divided translation followed by the prescribed integral point lift. The rank-one test uses the trivial divided translation. This proves the vanishing of the comparison adjoint tables and the testedness of the rank-one form on their respective false spaces.

Apply this implication to the same tables. For a form Y in $\operatorname{im} e_0 \cap \ker H \cap L_\psi^\perp$ the marked period vanishes in each $\mathcal{T}_i(G_\psi^* Y)$. Applying the common adjoint formula at $G_\psi^* Y$ then gives $AG_\psi^* Y = 0$. Hence $G_\psi A^* F$ is orthogonal to that false-lift space for each F. The tested forms and their integrality thus refer to the same G_ψ that compresses the single-lift determinant.

11.5 The compressed determinant

Lemma 11.1 (Mixed Gram inequality)**.** *Suppose the Hermitian pairing is represented by a rational positive definite matrix. For a field automorphism σ of* $\mathbf{C}$ *put $\sigma^{\#} = \mathrm{conj} \circ \sigma \circ \mathrm{conj}$. If $X = (x_i)$ and $Y = (y_i)$ are two u-tuples of vectors, then*

$$\left|\sigma \det(\langle x_i, y_j \rangle)\right|^2 \le \det\left(\langle x_i^{\sigma^{\#}}, x_j^{\sigma^{\#}} \rangle\right) \det\left(\langle y_i^{\sigma}, y_j^{\sigma} \rangle\right).$$

In particular $\sigma(\langle x, y \rangle) = \langle x^{\sigma^{\#}}, y^{\sigma} \rangle$.

Proof. The rational Gram entries and linearity in the second variable give the one-vector identity. Apply Cauchy–Schwarz to the two wedge products in the complex exterior power, with the induced Gram pairing equal to the determinant of pairings of the factors. □

The compressed determinant. Set $K_\psi = k_j P_{L_\psi} G_\psi G_\psi^*|_{L_\psi}$ and $d_j(\psi) = \det K_\psi$. The matched comparison gives $u_\psi^* G_\psi = \gamma_\psi A$, with A surjective and $\gamma_\psi \neq 0$, for every trial ψ. Thus u_ψ is an isomorphism from U_j onto L_ψ. Taking determinants of $u_\psi^* K_\psi u_\psi$ on U_j gives

$$\lambda_j(\psi) d_j(\psi) = k_j^R |\gamma_\psi|^{2R} \det(AA^*),$$
$$\frac{\lambda_j(1)}{\lambda_j(\psi)} = \frac{d_j(\psi)}{d_j(1)} \left(\frac{\Gamma_j(1)}{\Gamma_j(\psi)}\right)^R.$$

In particular $d_j(\psi) \neq 0$. The identity $G_\psi^* u_\psi = \overline{\gamma_\psi} A^*$ shows injectivity of $G_\psi^*|_{L_\psi}$. It also gives testedness of $G_\psi G_\psi^* z$ for every $z = u_\psi F$ in L_ψ, since the marked argument applies to $G_\psi A^* F$. These two facts will allow a basis of such tests to be chosen by maximal minors at the coefficient reading.

For the archimedean estimate, the projection to L_1 is contractive on tested forms for each coefficientwise norm. Indeed the good idempotent e_0 is self-adjoint for this norm because the finite group acts unitarily. At a split place off the level write H for the test Hecke polynomial operator. Its adjoint for the coefficientwise metric is the Hecke polynomial at the conjugate place with its parameter conjugated in that metric. It commutes with H and with e_0. Consequently H is normal for every such metric and the orthogonal complement of $P = \operatorname{im} e_0 \cap \ker H$ is the same vector space $\ker e_0 + \operatorname{im} H$ for all of them. The lift image L_1 has the good type selected by e_0. It lies in $\ker H$ by the split vacuum calculation on that genuine lift. A tested form, orthogonal in the original metric to $P \cap L_1^\perp$, therefore splits as its L_1-projection plus a vector in this common orthogonal complement. The two terms are orthogonal in the coefficientwise metric as well.

Let $z \in L_1$ be a nonzero eigenvector of K_1, with eigenvalue μ, and put $\xi = G_1^* z \neq 0$. Write $\sigma^\# = \mathrm{conj} \circ \sigma \circ \mathrm{conj}$, and choose ξ' with $(\xi')^{\sigma^\#} = \xi^\sigma$. Both the form metric and the table metric have rational Gram entries. With the inner product linear in its second entry, applying σ to the pairing of $G_1 \xi'$ with $\mu z = k_j P_{L_1} G_1 \xi$ gives

$$\sigma(\mu) \|\xi^\sigma\|^2 = k_j \big\langle (G_1 \xi')^{\sigma^\#}, (P_{L_1} G_1 \xi)^\sigma \big\rangle.$$

The vector $G_1 \xi$ is tested by the adjoint identity above. Using the contractivity and the two bounds for G_1 shows that the right side has modulus at most $e^{2rB_0} \|\xi^\sigma\|^2$. Thus $|\sigma(\mu)| \leq e^{2rB_0}$ and $|\sigma(d_j(1))| \leq e^{2RrB_0}$. The one-vector mixed Gram identity justifies the coefficient pairing used here. Since the bound holds for every field automorphism of $\mathbf{C}$, $d_j(1)$ is algebraic. Indeed any transcendental element can be sent to a transcendental element of arbitrarily large modulus by extending each to a transcendence basis and extending the induced isomorphism to $\mathbf{C}$.

12 Integral rank-one tests

The signed contribution of the rank-one tests. Fix one rank-one auxiliary pair p at stage j. Write F_ψ for its nonzero trivial-character lift, with the compact jet and vacuum scales, and h_1 for the class number of its chosen-level $U(1)$. Let P_ℓ be the sum over σ of the local asks at $u = w(\sigma)$, including the factor $(1 - \chi_d(\ell))/2$. Then

$$\sum_\ell 2P_\ell \log \ell = \sum_\sigma \mathrm{Ask}(w(\sigma), \sigma). \tag{12.1}$$

Indeed only prime powers contribute. For inert ℓ only the odd powers survive and for the ramified prime all powers survive. The supports are finite at $\ell \leq x_1$ by the minimum weight.

For each such j, p, fix the ramified pair q_2, the scalar $V_{jp} \neq 0$ of the rank-one test, and the flat pair of the same rank used in its untwisted norm comparison. Set

$$S_0 = \sum_{\sigma \in \Phi} S(w_0(\sigma)), \qquad a_{jp} = I_{jp} / \|F_1\|^2, \qquad A_{jp} = a_{jp} V_{jp} \overline{V_{jp}},$$

where $I_{jp} \neq 0$ is the integral algebraic part at the trivial twist of the flat pair. Choose V_{jp} as in the character selection. Its conjugates have modulus at most $\exp(B_0)$. The flat-part comparison gives the bound $\exp(B_0 - S_0)$ for the conjugates of a_{jp}.

For every trial ψ and each ι with $\ell \nmid \mathrm{N}\, Q_\psi$, the construction gives a tested form Ξ_a for each $a = j_\iota$ and $a = j_\iota \circ \mathrm{conj}$, with

$$|a(V_{jp})|_\ell \ell^{-P_\ell} \leq \left| a\left(h_1^2 \langle \Xi_a, F_\psi \rangle / \rho_{q_2}(\psi) \right) \right|_\ell, \qquad \ell \nmid \mathrm{N}\, Q_\psi. \tag{12.2}$$

The test at the conjugate embedding uses the same ψ, and each Ξ_a is integral at all embeddings above ℓ.

The flat algebraic parts and the untwisted multiplier. For a flat pair of rank $n = 2z \le t$, the diagram is empty and the coefficient space is one dimensional with rational Gram matrix 1. At the trivial twist the denominator clearer is g^{2fn}. Set $D_{\mathrm{flat}} = D_{\mathrm{aut}} g^{4fn}$, so its integral algebraic part is $I = D_{\mathrm{flat}} \|F_1\|^2$. This is integral at every coefficient place. Indeed the vacuum clearers make $g^{2fn} F_1$ integral by the normalized indicator calculation, and D_{aut} clears the stabilizer weights of the final class sum. The same argument off Q_ψ applies to the ramified flat block with its conductor power a unit.

For the same flat pair, expand the unramified indicator on matrices modulo $\mathfrak{M}$ into finitely many moves of the normalized vacuum. The coefficients have conjugate moduli at most the number g^{fn} of residue matrices. At order zero, the conjugate Gaussian integral in the weighted jet bound, with $L = 1$, gives a bound of

$$C_n = g^{fn} \exp\bigl(C_d (N_1 n + 1)^2 + fn \log 4\bigr)$$

for every conjugate point kernel. Indeed the Weyl polynomial is the constant 1; the Gaussian majorant has real dimension $2fn$ and the factor from its integration is 4^{fn}. The level determinant scalars have modulus one under the coefficient embeddings. The class counts of $U(1)$ and $U(n)$ are at most $e^{x^{30}}$ at their chosen levels. The level comparison for this count uses the principal level modulo $\bar{\pi}_G \pi$, which is contained in both actual levels. Thus after summing first in the source group and then in the target group, the mixed pairing between the σ and conj σ conj readings gives

$$|\sigma(\|F_1\|^2)| \le e^{3x^{30}} C_n^2.$$

This follows by Cauchy–Schwarz with inverse stabilizer orders as positive rational weights. Here $\log D_{\mathrm{aut}} \le f(rt)^2 \log g$. The bounds $f, N_1 = O(x^{11} \log x)$, $n = O(x^5)$ and $\log g = O(x^{\varepsilon})$ give $\log C_n = o(x^{33})$ and $\log D_{\mathrm{flat}} = o(x^{33})$. Hence $|\sigma(I)| \le e^{x^{34}}$ for all the flat blocks on this datum.

To compare the untwisted norm of the main rank-one pair to the flat pair of the same rank, write $m_\tau = c_p(\tau) - z_p$. The conjugate CM jet transport uses the same scalar and the same relabelling of places for the two pairs. The transported conjugate Hecke character uses the same chosen level and is unramified elsewhere. Apply the rank-one norm computation at the transported character. The scale of each degree is $(g/\Omega)^{m_\tau}$ in the squared norm, from the actual $\bar{\pi}_G$ factor and the conjugate CM period. Its row hook is $m_\tau!$. The character and vacuum phases cancel in the mixed conjugate pairing. For σ fixing K this gives

$$\log |\sigma(\|F_1\|^2)| = c_{\sigma,n} + \sum_\tau S(m_\tau) + \log d_p + \log |L^S(z_p, \nu_{\sigma,p}^{-1})|,$$

where d_p is the dimension of its coefficient module and $c_{\sigma,n}$ is common to the two untwisted pairs. The transported character is unitary on the rank-one torus. Absolute convergence at $z_p \ge 2$ bounds the last logarithm in modulus by $[E : \mathbf{Q}] \log \zeta(2) \le 4f$. Since $d_p \ge 1$ and the flat dimension is 1, subtracting the formula for the main norm from that for the flat algebraic part I_{jp} gives

$$\log \bigl|\sigma(I_{jp}/\|F_1\|^2)\bigr| \le \log D_{\mathrm{flat}} + 8f - \sum_\tau S(m_\tau).$$

An automorphism of $\mathbf{C}$ restricting to conjugation on K can be composed with complex conjugation on the output without changing any modulus. The same bound therefore holds for all conjugates. As $g/\Omega \ge 1$ and $m_\tau \ge w_0(\tau)$, $S(m_\tau) \ge S(w_0(\tau))$. Finally $\log D_{\mathrm{flat}} + 8f \le B_0$. This gives the required bound $e^{B_0 - S_0}$ for the fixed multiplier.

Curve coefficients and the twelve-dimensional indicator. Fix the rank-one pair p at stage j, with degrees $n_\sigma = c_p(\sigma) - z_p$. The rank-one inequalities give $w_j(\sigma) - t \le n_\sigma \le w_j(\sigma)$. Since $j \ge 1$, we have $w_j - w_0 = 2jr^8 > t$. Rounding and $j \le m$ give $w(\sigma) \le w_0(\sigma) \le n_\sigma \le w(\sigma) + x^{42}$ and $n_\sigma \le 2b$ for large x. We will find a nonzero t_p of conjugate modulus at most $\exp(fmb/\kappa)$ such that for every prime ℓ and reading $a_\ell : \mathbf{C} \to \mathbf{C}_\ell$ there is a rank-12 residue indicator at $(\bar{\pi}_G)$ and a Taylor monomial of total degree n_σ in each row σ. The coefficient at scale $\bar{\pi}_G \eta_0$, multiplied by the vacuum scale to the twelfth power and by $\prod_\sigma n_\sigma!$, has absolute value at least $\ell^{-P_\ell} |a_\ell(t_p)|_\ell$.

Packing the Frobenius directions. Put $H = a/x$, $L = a/x_1$, and $\Delta = x^{42}$. Thus $w_\sigma = H$ on the thick set $\mathcal{J}$, and $w_\sigma = L$ elsewhere, with $w_\sigma \le n_\sigma \le w_\sigma + \Delta$ and $n_\sigma \le 2b$. Write $c_* = \lceil 60b/\kappa \rceil$, $n_\sigma^- = \max(n_\sigma - c_*, 0)$, and $D_\sigma = (\prod_{h=1}^{c_*}(n_\sigma^- + h))^2$. Set $\omega_\ell = (1 - \chi_d(\ell))/2$ and

$$P_{\ell,\sigma} = \omega_\ell \sum_{\substack{e \ge 1 \\ e \text{ odd or } \ell = d}} \left(w_\sigma/\ell^e - w_{[-\ell^e]\sigma} \right)_+, \qquad P_\ell = \sum_\sigma P_{\ell,\sigma}.$$

For a nonsplit ℓ, the same live powers have

$$\omega_\ell \sum_{\substack{e \ge 1 \\ e \text{ odd or } \ell = d}} \ell^{-e} = \frac{1}{\ell - 1} - \rho_\ell, \qquad \rho_\ell = \begin{cases} 1/(\ell^2 - 1) & \ell \text{ inert}, \\ 1/(2(\ell - 1)) & \ell = d. \end{cases}$$

First let $\ell \le \kappa$ be inert. It is prime to N_1 and $\ell \ge 13$, since $2, 3, 5, 7, 11$ all split. Put $u = [-\ell]$ and let ϕ be reduction modulo 2^{s+1} on the place group. If $\phi(u)^v = 1$ for $0 < v \le 2s$, then $2^{s+1} \mid \ell^{2v} - 1$. For any odd z and positive v, factoring the difference of squares and then successively squaring gives

$$v_2(z^{2v} - 1) = v_2(z - 1) + v_2(z + 1) + v_2(v).$$

In fact the odd part of v leaves the valuation unchanged, and each doubling after z^2 adds one. The smaller of the first two valuations is one, so $2^s \mid v(\ell - 1)$ or $2^s \mid v(\ell + 1)$. This is impossible, since $2^s = m^{22} \ge x^{11}$ and $2s(\kappa + 1) = O(x^5 \log x)$.

Join σ to $u^i\sigma$ whenever both are thick and $0 < i < s$ is odd. These directed edges have distinct endpoints and are disjoint. Indeed a thick triple of the form $\sigma, u^i\sigma, u^j\sigma$, with $0 < i < j \le 2s$, satisfies the three-term identity in the construction of the thick set. It would force $\phi(u)^i = 1$. Two outgoing edges give such a triple. Two distinct incoming edges of exponents $i < j$ give one by starting at the source of the longer edge, with exponents $j - i$, j; equal exponents give the same source. Two consecutive edges give one with exponents $i, i + j$.

On an edge out of σ, choose $i_\sigma = i$ and $y_\sigma = \lfloor n_\sigma/\ell^i \rfloor$. Elsewhere set $y_\sigma = 0$ and $i_\sigma = 0$. The outgoing cost $\ell^i y_\sigma$ is at most n_σ. Each target is thick and has only its single incoming cost, which is at most $(H + \Delta)/\ell \le H \le n_{u^i\sigma}$. For two active sources σ, τ, the equality $u^e\tau = u^{i_\sigma}\sigma$ at a positive odd $e < i_\sigma$ would give the same excluded triple, now with exponents $i_\sigma - e, i_\sigma$ from σ. At $e = i_\sigma$, the equality forces $\tau = \sigma$. Thus the capacity and lowest-term conditions both hold. At the ramified prime $\ell = d$, take $i_\sigma = y_\sigma = 0$ for every σ. No live edge there takes thick to thick, since all thick places have class -1 modulo M, whereas $d \equiv M/2 - 1 \pmod M$ and $4 \mid M$ give $d^2 \equiv 1 \pmod M$ and $-d^e \ne 1 \pmod M$ for every $e \ge 1$.

For a live $e < s$ that is outside a thick-to-thick edge, the loss in replacing n_σ/ℓ^e by $(w_\sigma/\ell^e - w_{[-\ell^e]\sigma})_+$ is at most $2L$. For a thin source this follows from $n_\sigma \le L + \Delta$. For a thick source the target is thin, and the loss is at most $L + \Delta/\ell^e$. Here $\Delta \le L$. At the single edge of a source the bound $n_\sigma/\ell^e \le y_\sigma + 1$ suffices. The tail $e \ge s$ is at most $2b \sum_{e \ge s} 2^{-e} = 4b/2^s$. The factor ω_ℓ is one in the inert case and one half at d, so in both cases

$$\left(\frac{1}{\ell - 1} - \rho_\ell \right) n_\sigma \le P_{\ell,\sigma} + y_\sigma + 2sL + 1 + 4b/2^s.$$

In the weighted scales $b_2 \le L + N_1 \le 2L$, $s = O(\log x)$, and $L = O(bx/x_1)$, with $x_1 \ge x^{10^6}$. For $\ell \le \kappa \le 3x^5$ we have $c_*/\ell \ge (20/3)b/x^{10}$. This also covers $\ell = d$ for large x. Since $4b/2^s \le 4b/x^{11}$, it follows uniformly that $2sL + 1 + 4b/2^s + 4b_2 \le \lfloor c_*/\ell \rfloor$. The product of c_* consecutive positive integers in the definition of D_σ is divisible by $c_*!$. Summing over the places gives the needed small-prime bound

$$\sum_\sigma \left(\left(\frac{1}{\ell - 1} - \rho_\ell \right) n_\sigma + 4b_2 - y_\sigma \right) \le P_\ell + \sum_\sigma v_\ell(D_\sigma). \tag{12.3}$$

For an inert $\ell > \kappa$, put $i_\sigma = 1$ and $y_\sigma = \min(\lfloor n_\sigma^-/\ell \rfloor, n_{[-\ell]\sigma})$. The partner map is a permutation, and $\lfloor n_\sigma^-/\ell \rfloor \le 2b/\ell < c_*$. At a target with $n_\sigma < c_*$ the outgoing cost is zero and the incoming cost is capped at

n_σ. At any target with $n_\sigma \ge c_*$ its outgoing cost is at most $n_\sigma - c_*$ and its incoming cost at most c_*. This gives the capacity in both cases. The least allowed odd exponent is one, and partners are distinct at that exponent. Put $J_\sigma = \prod_{h=1}^{c_*}(n_\sigma^- + h)$. If $n_\sigma < c_*$, $n_\sigma!$ divides $J_\sigma = c_*!$. Otherwise $n_\sigma! = n_\sigma^-! J_\sigma$. In either case, writing $q_\sigma = \lfloor n_\sigma^-/\ell \rfloor$, Legendre's formula gives

$$v_\ell(n_\sigma!) \le q_\sigma + v_\ell(q_\sigma!) + v_\ell(J_\sigma).$$

Here $q_\sigma \le c_*$, and $c_*! \mid J_\sigma$. Also $n_\sigma^- \le w_\sigma$ since $c_* \ge \Delta$. Therefore

$$q_\sigma - y_\sigma \le (w_\sigma/\ell - w_{[-\ell]\sigma})_+ \le P_{\ell,\sigma}.$$

Using $D_\sigma = J_\sigma^2$ yields

$$v_\ell(n!) - \sum_\sigma y_\sigma \le P_\ell + \sum_\sigma v_\ell(D_\sigma), \qquad n! = \prod_\sigma n_\sigma!. \tag{12.4}$$

The frame of length g in the rank-one gauge. Set $\nu = N_1$. We use the lattice $\mathfrak{a}_0 = E \cap \alpha\widehat{\mathcal{O}}_E$ and the period factor s_1 from the normalized source gauge. The list of powers of ζ has frame constant ν. For every rational prime $q \mid \nu$, use the list w_{jkh} from that gauge, before repetition. Put $A_0 = (d+1)/4$, choose $0 \le z_0 < \nu$ as in its construction, and write $t_1 \bar{t}_1 = A_0 \nu e_0$, with $e_0 \in \mathbf{Z}$. Denote $\nu_q' = \nu/q^{v_q(\nu)}$. The frame constant of the unrepeated list is

$$C_q = \nu + \nu_q'(1 + 2A_0 z_0) + 2q(\nu_q')^2 A_0 e_0.$$

This is a positive integer by the same diagonal frame sum. It is 0 modulo each prime divisor of ν other than q and ν_q' modulo q. Thus $C = \sum_{q\mid\nu} C_q$ is prime to ν, where the sum runs over primes. It is of polynomial size in x, since $z_0 < \nu$, $e_0 = O_d(\nu)$ and ν is of polynomial size. Put $R = 3|n|$, where $|n| = \sum_\sigma n_\sigma$. We have $g \ge b^{20}$ and $R \le 6fb$, so $g - R(C+\nu) \ge (\nu-1)C$. The integer on the left is $a'\nu + b'C$ with $a', b' \ge 0$, by first choosing $0 \le b' < \nu$ in its residue class after multiplication by C^{-1}. Repeat each vector of the powers-of-ζ list $R + a'$ times, and each vector of the other lists $R + b'$ times. This gives

$$\sum_i c_i \Phi_\sigma(w_i)\overline{\Phi_\tau(w_i)} = g\, 1_{\sigma=\tau}, \qquad c_i \ge 3|n|.$$

The list is in $\mathfrak{a}_0$, so for any reading $a_\ell : \mathbf{C} \to \mathbf{C}_\ell$ all coefficients $\mu_{i,\sigma} = a_\ell(s_1\overline{\Phi_\sigma(w_i)})$ are integral. There is an index i at which all are units. If the period factor is a unit, use $w_i = 1$. Otherwise $\ell \mid \nu$ and the readings of $\overline{\Phi_\sigma(w_i)}$ see the side where $\alpha = \Phi_\nu'(\zeta)^{-1}$. In the block $q = \ell$, take the vector with $(j,k,h) = (0,1,0)$. This is $t_1 \nu_\ell'(1-\zeta_\ell)/\nu$. The factor t_1 is a unit on that side. Its valuation is therefore $1/(\ell-1) - v_\ell(\nu) = v_\ell(\alpha)$ for every σ, and the period factor cancels it. In particular s_1 and its conjugate are units at every nonsplit ℓ. At these primes g is also a unit, since both of its prime factors split.

Invariant derivatives in the marked elliptic chart. Work in the marked normalized parameter $Y_0(z) = z + O(z^2)$ of the elliptic germ calculation. Let $\mathcal{O}$ be the valuation ring of $\mathbf{C}_\ell$. Translation by the marked point identifies this chart with the formal group at the origin of the good elliptic model. Its integral group law expresses $Y_0(z+w)$ as a power series with integral coefficients in $Y_0(z), Y_0(w)$. For nonsplit ℓ, the same elliptic germ calculation gives $T \in \mathcal{O}[[Y_0]]$ with

$$|\varpi|_\ell = \ell^{-\rho_\ell}, \qquad S(X) = T(\varpi X)/\varpi \in \mathcal{O}[[X]], \qquad \overline{S} - \overline{S}^{\,q_0} = X,$$

where $q_0 = \ell^2$ for inert ℓ and $q_0 = \ell$ at d. The construction of this local parameter uses only $\ell \ge 13$ and the indicated splitting type, so applies to every nonsplit ℓ.

For split ℓ the same parameter conclusion holds with $q_0 = \ell$ and $\rho_\ell = 1/(\ell-1)$. To verify this case, choose $a_0 \in \mathcal{O}_K$ of valuation one at the prime of K defined by the reading and valuation zero at the

conjugate prime. The CM isogeny $[a_0]$ extends to the good elliptic model. Its degree has ℓ-valuation one by the norm of a_0. Its reduced differential is zero, so its inseparable degree is exactly ℓ. In an integral formal parameter its reduction starts in degree ℓ. Average $Y_0(\zeta z)$ with coefficients $\zeta^{-1}/(\ell-1)$ over the $(\ell-1)$-st roots of unity in $\mathbf{Z}_\ell$, giving an integral parameter T. The integral action for these scalars follows by limits of integer multiplications. Write $Z(T)$ for the inverse in the differential coordinate. It has only degrees 1 mod $\ell-1$, and Z' is integral since the invariant differential is integral in T. The coefficient of $[a_0]$ at degree ℓ, in the T coordinate, is $(a_0-a_0^\ell)Z_\ell$, where $Z_\ell = [T^\ell]Z$. Thus $v_\ell(Z_\ell) = -1$. Choose $\varpi^{\ell-1} = -1/Z_\ell$. In degree $1+h(\ell-1)$ for $h \geq 2$, the coefficient of $\varpi^{-1}Z(\varpi X)$ has valuation at least $h - v_\ell(1+h(\ell-1)) \geq 1$. The last inequality follows from $1+h(\ell-1) < \ell^h$. Thus its reduction is $X - X^\ell$, and its integral compositional inverse is the desired S.

Let A_m be the maximum of $|m![z^m]J(Y_0(z))|_\ell$ over $J \in \mathcal{O}[[Y_0]]$. The maximum is attained on one of the finitely many monomials Y_0^h with $0 \leq h \leq m$. The integral addition law gives $A_{m+n} \leq A_m A_n$: express $J(Y_0(z+w))$ in the two parameters and take the coefficient of $z^m w^n$ times $m!n!$. In the differential coordinate this is $(m+n)![z^{m+n}]J$. Put $B_m = A_m/\ell^{(\rho_\ell - 1/(\ell-1))m}$. The reduced equation for S gives $\overline{S} = \sum_{j\geq 0} X^{q_0^j}$, by iteration starting at zero constant term. Therefore at $m = q_0^j$ the coefficient of T in degree m has absolute value $|\varpi|_\ell^{1-m}$. Legendre's formula gives $B_{q_0^j} \geq |\varpi|_\ell$. If some $B_m < 1$ with $m > 0$, submultiplicativity and division of every index by m would force $B_n \to 0$. Thus for every m an integral series J_m in the marked parameter satisfies

$$|m![z^m]J_m(Y_0(z))|_\ell \geq \ell^{(\rho_\ell - 1/(\ell-1))m}.$$

Twelve factors on the marked frame. Write $Z = (Z_\sigma)_\sigma$, $\mu_i(Z) = \sum_\sigma \mu_{i,\sigma} Z_\sigma$, and $\Theta = \prod_i \theta(\mu_i(Z))^{c_i}$, with coefficients read by a_ℓ and $\theta(0) = 1$. Let $\mathcal{R}$ be the image of the integral formal series ring in the variables $Y_0(\mu_i(Z))$. For each multi-index a, choose $H_a = \prod_i Y_0(\mu_i)^{h_i(a)}$ maximizing $N(a) = |[Z^a](H_a\Theta)|_\ell$ among monomials of total degree at most $|a|$. The same bound applies to the coefficient of $J\Theta$ for every $J \in \mathcal{R}$ by the ultrametric inequality. If a product of k such factors has a coefficient at a, the coefficient formula bounds it by some $\prod_{j=1}^k N(a_j)$ with $\sum_j a_j = a$. The same bound holds for any factor obtained by rescaling Z by a scalar in $\mathbf{Z}_\ell$, because this multiplies each coefficient by a scalar of norm at most one.

Put $L_\theta = -\log\theta$. The marked isogeny formula implies $\exp(t^2 L_\theta(z) - L_\theta(tz))$ is an integral unit in $Y_0(z)$ for $t \in \mathbf{Z}_\ell$. Use nonzero integers and coefficientwise continuity in the marked parameter. Choose $t_1, \ldots, t_4 \in \mathbf{Z}_\ell$ with $\sum_h t_h^2 = -2$, by the four-squares theorem modulo each power of ℓ and compactness. It follows that $U_\theta = \Theta(Z)^2 \prod_h \Theta(t_h Z)$ is an integral unit in $\mathcal{R}$ with constant coefficient one. Its inverse belongs to $\mathcal{R}$ by expanding the inverse of an integral formal series with constant coefficient one. Now $\Theta^{-1} = (U_\theta^{-1}\Theta)\prod_h \Theta(t_h Z)$, so each coefficient of Θ^{-1} is bounded by a five-fold product of maxima N over a splitting of its degree.

For a residual degree r', apply the single-variable derivative bound in order $|r'|$ to the form with unit coefficients. It gives an element of $\mathcal{R}$ with coefficient at r' at least $\ell^{(\rho_\ell - 1/(\ell-1))|r'| + v_\ell(r'!)}$ in absolute value, by the multinomial formula. Here $v_\ell(r'!) = \sum_\sigma v_\ell(r'_\sigma!)$. Writing the element as a product with Θ times Θ^{-1} gives this bound on the product of six maxima at indices summing to r'.

At an inert prime, let y_σ, i_σ be the packing above and $\sigma^* = [-\ell^{i_\sigma}]\sigma$. Put

$$A = \sum_\sigma \ell^{i_\sigma} y_\sigma e_\sigma, \qquad B = \sum_\sigma y_\sigma e_{\sigma^*}, \qquad b' = A + B \leq n.$$

Use $\theta(\varpi X) = \exp(\lambda \Xi(X))V_0(X)$ from the marked germ with no torsion translation. Here $V_0(0) = 1$, $|\lambda|_\ell = \ell^{1-1/(\ell-1)}$, and $\overline{\Xi} = \lambda_0 \overline{S}^{\ell+1}$ with $\lambda_0 \neq 0$. Thus $F(Z) = \Theta(\varpi Z) = \exp(\lambda\Psi(Z))V(Z)$, where $\Psi = \sum_i c_i \Xi(\mu_i)$ and $V = \prod_i V_0(\mu_i)^{c_i}$ are integral and $V(0) = 1$. Write $r_{i,\tau} = \overline{\mu_{i,\tau}}$ in the residue field. As $\overline{S} - \overline{S}^{\ell^2} = X$, the coefficient at $Z_\tau^{\ell^e}$ of $\partial_\sigma \overline{\Psi}$, for odd $e \geq 1$, is $\lambda_0 \sum_i c_i r_{i,\tau}^{\ell^e} r_{i,\sigma}$. These are the only possible monomials. Odd Frobenius on the reduction of $\overline{\Phi_\tau(w_i)}$ acts as the embedding $\Phi_{[-\ell^e]\tau}$ on w_i. Frobenius undoes the

conjugation on K and sends the cyclotomic generator to its indicated power. This reading applies to the frame vectors, since $\ell \nmid dN_1$ and they are integral at ℓ. The frame identity gives

$$\sum_i c_i r_{i,\tau}^{\ell^e} r_{i,\sigma} = \overline{a_\ell(s_1)}^{\,\ell^e+1} \overline{g}\, 1_{[-\ell^e]\tau=\sigma}.$$

Thus every indicated coefficient is nonzero, even when individual multiplicities vanish in the residue field.

For each partial derivative, the new integral factor in $\lambda^{-|B|}F^{-1}\partial^B F$ comes from multiplying by the corresponding partial of Ψ, modulo the maximal ideal. Indeed differentiation of any integral factor, or of V, has integral coefficients and comes with the factor λ^{-1}. Induction gives the reduction $\prod_\sigma (\partial_{\sigma^*}\overline{\Psi})^{y_\sigma}$. Give the active directions weight one and the inactive ones weight $\ell^{\max i_\sigma}$ when there is an active direction. In each of the indicated partials the chosen monomial is the unique one of lowest weight. For active directions this uses the packing's collision property and the bijection at a fixed exponent. An inactive direction contributes only exponents at least ℓ. The product has a unit coefficient at A. If all $y_\sigma = 0$, the product is the constant one and the same conclusion holds.

Take a largest Cauchy summand at A and write its indices as $\gamma + \beta = A$ in the product $F^{-1}\partial^B F$. The coefficient of $\partial^B F$ at β is the coefficient of F at $\beta + B$ times $B!\binom{\beta+B}{B}$. The binomial is integral. Undoing the scaling by ϖ gives

$$|[Z^\gamma]\Theta^{-1}|_\ell\, |[Z^{\beta+B}]\Theta|_\ell \geq \ell^{(1-1/(\ell-1))|B|+\rho_\ell|b'|+v_\ell(B!)}.$$

This yields six maxima at indices summing to b'. At d use $A = B = b' = 0$ and six degree-zero factors. The corresponding inequality then has right side one.

For $\ell \leq \kappa$ inert or $\ell = d$, take the other six maxima from the residual degree $r' = n - b'$. Every entry of B and r' is at most $2b < \ell^{2b_2}$, as $b_2 \geq a/x_1$ on the weighted scales. Legendre's formula bounds the combined two factorial losses by $4fb_2$. Adding the two six-factor exponents gives at least $\rho_\ell|n| + \sum_\sigma y_\sigma - 4fb_2$, which by (12.3) and $v_\ell(n!) \leq |n|/(\ell-1)$ is at least $v_\ell(n!) - P_\ell - \sum_\sigma v_\ell(D_\sigma)$.

For inert $\ell > \kappa$, all $i_\sigma = 1$, so $\rho_\ell|b'| = |B|/(\ell-1)$. The first six maxima have exponent at least $\sum_\sigma y_\sigma$. For the residual degree, use the elements of $\mathcal{R}$ with linear terms $a_\ell(\bar{s}_1 s_1 g)Z_\tau$ obtained by summing $c_i a_\ell(\bar{s}_1 \Phi_\tau(w_i))Y_0(\mu_i)$ over i. Their coefficients as combinations of the generators are integral by the same period bound at the conjugate reading. The common linear scalar is a unit. The product of their powers with exponents r'_τ, multiplied by Θ, has a unit coefficient in degree r'. Use this maximum and pad with degree-zero maxima. Equation (12.4) gives the same total exponent. Finally for split ℓ, $P_\ell = 0$ and $\rho_\ell = 1/(\ell-1)$, so the residual construction alone in degree n, with six more degree-zero maxima, supplies exponent $v_\ell(n!)$. In every case there are twelve arrays $a_1, \dots, a_{12}$ with sum n and

$$\prod_{j=1}^{12} N(a_j) \geq \ell^{v_\ell(n!) - P_\ell - \sum_\sigma v_\ell(D_\sigma)}. \tag{12.5}$$

Recovery of the rank-12 indicator. Use the arrays a_b of (12.5), where here b indexes the twelve columns. Each $Y_0^{h_i(a_b)}\theta^{c_i}$ is a degree-c_i elliptic section, since $Y_0\theta^3$ is a section of degree three and $3h_i(a_b) \leq 3|n| \leq c_i$. Set $\delta = \pi_G$, so $\delta\bar{\delta} = g$. With $\eta = (2\pi/\sqrt{d})^{1/2}$, the product on the standard rank-12 space is

$$\prod_{b=1}^{12} \prod_i (Y_0^{h_i(a_b)}\theta^{c_i}) \left(\frac{\Omega_0}{\bar{\delta}\eta} \sum_\tau \overline{\Phi_\tau(w_i)} W_{\tau b} \right).$$

For a period $\bar{\delta} r_\infty(\gamma)$, $\gamma \in \mathfrak{a}_0$, the argument in each column is $\Omega_0 \operatorname{Tr}_{E/K}(\gamma \bar{w}_i) \in \Lambda$, by the trace-dual identity for $\mathfrak{a}_0$. The frame identity and $|\delta|^2 = g$ make the sum of the pulled back Hermitian forms the standard Fock form. The elliptic canonical multipliers then multiply to the standard multiplier on $L' = \bar{\delta} L^{12}$, where $L = r_\infty(\mathfrak{a}_0)$. This is the lattice and multiplier calculation of the finite vacuum projection, now with $\bar{\delta}$ as the lattice scale.

The rational finite translation formula identifies the translates of the vacuum on this space with translates of the standard finite vector. In fact $L'^{\vee} = \delta^{-1}L^{12}$, by self-duality of L, and $L'^{\vee}/L^{12}$ indexes a basis of translates of the vacuum. The characters of L^{12}/L' on these vectors are distinct by the Riemann pairing, and the dimension for the polarization on L' is g^{12f}. Hence the product is the entire function of a vector U in the finite part of $\mathrm{Hom}(E, E^{12})$ with its standard Weil data. Every finite value $U(\mathrm{finVec}\, x)$ is integral at each reading. For a rational γ this follows by translating the elliptic sections to the corresponding torsion points and using the normalized torsion height bound on each $Y_0^h \theta^{c_i} = (Y_0\theta^3)^h \theta^{c_i - 3h}$. The same bound at arbitrary finite arguments follows by the local constancy of the finite translation and approximation by rational γ. The Heisenberg phases are roots of unity.

The Taylor scale $\bar{\delta}\eta_0 = \bar{\delta} s_1 \eta/\Omega_0$ turns the complex linear forms in the product into precisely the forms read as μ_i above. The twelve groups of variables are disjoint. Thus if L_α denotes coefficient extraction at $\alpha_{\sigma b} = (a_b)_\sigma$ in this Taylor scale, then $\sum_b \alpha_{\sigma b} = n_\sigma$ and

$$\left|a_\ell(L_\alpha(U))\right|_\ell = \prod_{b=1}^{12} N(a_b).$$

To read a coefficient of a residue indicator, use Schur orthogonality for the finite group of shifts $\mathcal{A} = L'^{\vee}/L'$, of order g^{24f}. For each class choose a representative $r_\infty(\gamma)$ with $\gamma \in \delta^{-1}\mathfrak{a}_0^{12}$, and put $x = \alpha_{\mathrm{W}}^{-1}\gamma$ at the finite adeles, where α_{W} is the idele defining $\mathfrak{a}_0$. Then δx is integral. Write $R_x = \varrho(\mathrm{finVec}\, x)$ for the induced shift. The shifts act projectively, since L' acts by its scalar multiplier. The vacuum translates give one-dimensional eigenspaces for L^{12}/L', and the shifts permute these eigenspaces transitively. Thus their action is irreducible. Averaging the operator $u \mapsto L_\alpha(u)U$ under conjugation yields a scalar operator; its scalar is its trace $L_\alpha(U)$ divided by the dimension g^{12f}. The Heisenberg shifts at x and $-x$ are inverse operators. Applied to the scaled vacuum $F_{\mathrm{vac}} = v^{12}\,\mathrm{vac}$, the average gives

$$|\mathcal{A}| L_\alpha(U) F_{\mathrm{vac}} = \sum_{x \in \mathcal{A}} g^{12f} L_\alpha(R_x F_{\mathrm{vac}}) R_{-x} U.$$

Fix a finite argument y_0 with $u_0 = F_{\mathrm{vac}}(\mathrm{finVec}\, y_0) \neq 0$. Such a point exists because the standard vacuum has a nonzero finite factor and its Gaussian has nonzero value at zero. Evaluation at y_0 in the last identity gives the coefficient comparison with $g^{24f} u_0 L_\alpha(U)$. All its weights $g^{12f}(R_{-x}U)(\mathrm{finVec}\, y_0)$ are integral.

For $\mathfrak{M} = (\bar{\delta})$, every $R_x\,\mathrm{vac}$ with δx integral is a sum of standard residue indicators at $\mathfrak{M}$, with root-of-unity coefficients. In the Schrödinger coordinates at these split places, the component of x conjugate to the residue coordinates multiplies the vacuum by the additive character pairing $\bar{x}$ with the residue coordinates. It factors through $\mathfrak{M}$, since $\bar{x}$ times $\bar{\delta}$ is integral. The components of the shift outside this part are integral and act on the vacuum by their semicharacter. Expanding the residue character as its values times the point indicators gives the assertion, in the normalized vacuum with no averaging denominator. This is the finite translation formula in the same standard model as the entire product. Let F_y be the vacuum multiplied by the indicator of the residue fibre y modulo $\mathfrak{M}$ in these coordinates. The ultrametric inequality supplies a residue y for which the rank-12 coefficient satisfies

$$\left|a_\ell(v^{12} L_\alpha(F_y))\right|_\ell \geq \left|a_\ell(g^{24f} u_0 L_\alpha(U))\right|_\ell.$$

Choose, before any prime or reading,

$$t_p = g^{24f} u_0 \prod_\sigma D_\sigma.$$

This is nonzero, and the fixed integral vacuum bound gives $\log|\gamma(u_0)| \leq O_d(N_1^2)$ for every complex automorphism γ. The other factors contribute at most $24f \log g + 2f c_* \log(3b)$, since $c_* = o(b)$ and $n_\sigma \leq 2b$. Their sum is at most fmb/κ for large x, using $\log b = o(m)$, $\log g = O(\log b)$, and the polynomial bound on N_1. Multiplying the indicator coefficient by $n!$ and dividing by t_p, equation (12.5) gives the required lower bound ℓ^{-P_ℓ}. The monomial and the residue can be chosen individually at each reading. The stage, pair, degree, modulus, character at the depth point, and t_p are fixed.

Integral point forms for the rank-one tests. Use the partition $2z_p = 12 + 4 + 4 + (2z_p - 20)$ of the rank-one target. The three flat blocks are taken at the twist ψ, the trivial twist, and the trivial twist, in that order. They have ranks $4, 4, 2z_p - 20$. Since $m < z_p \le r$, all three even ranks lie in $[4, t]$ and the last rank is positive. Let α be any exponent array on the first block with $\sum_b \alpha_{\tau b} = m_\tau$ at each place τ. Write $X^{\tilde{\alpha}}$ for its extension by degree zero on the other blocks. The rank-one harmonic model has one row of length m_τ in the jet polynomial at τ. Let N_λ be the integral tableau basis in rank $2z_p$ for this one-row diagram. Put $c_\lambda = [X^{\tilde{\alpha}}]N_\lambda$. Its entries are integers. The highest polynomial at each τ is a single variable to the power m_τ. Thus the Fischer identity for this Weyl module has hook $\prod_\tau m_\tau!$. The Fischer pairing of $X^{\tilde{\alpha}}$ with N_λ is $\big(\prod_{\tau,b} \alpha_{\tau b}!\big) c_\lambda$. The polynomial of the Weyl vector c is uniquely determined by its Fischer pairings with the basis. It follows that its harmonic is

$$P_c = \frac{\prod_\tau m_\tau!}{\prod_{\tau,b} \alpha_{\tau b}!} X^{\tilde{\alpha}}.$$

For any matrix fixing the first twelve columns of the identity, the transposed row substitution fixes $X^{\tilde{\alpha}}$. By transpose covariance of the Fischer pairing the Weyl action fixes c. In this one-row pair all columns come from the positive sign, so the negative determinant exponent is zero. The coefficient identity and the fixed-vector identity are polynomial identities over $\mathbb{Q}$. They hold for the same frames and levels at both coefficient readings, for whichever monomial and indicator are chosen in each reading.

Fix the rank-12 residue column x_0 modulo $\mathfrak{M}$. In the full point form at $\mathrm{diag}(1, h_2, h_3, h_1)$, use the column $(x_0, J_2, J_3, J_1)^{\mathrm{t}}$ at the chosen places, where J_a is the standard column of the indicated flat block. It is primitive there because J_1 is a unit column. At a twist place read the marked row of the second block for the determinant character. The prescribed column at the chosen places and the prescribed row at the twist places can each be completed to an invertible residue matrix. Lift these matrices by Chinese remaindering to one compact finite adelic element. The class delta at the transported point pairs with a form by evaluation with no mass factor, since its stabilizer sum cancels the stabilizer divisor in the class pairing. Transport through this compact element preserves the integral gauge. Before this compact transport the point acts trivially on the first twelve columns, so its gauge on c is one or less at every embedding over ℓ.

Reduce all four finite data modulo $\mathfrak{M}^2 Q_\psi$. At a chosen place the four indicator constraints are exactly the four consecutive blocks of the full column constraint. At a twist place the second block alone contributes ψ_w of the marked row. By Chinese remaindering their product is the datum of the full point form. The real test is P_c on the first block and the constant polynomial on the other three blocks. Choose the splitting character of the first block such that its product with the three flat splitting characters is the full target splitting character. The oscillator representation on the direct sum is the tensor product in these characters, and the theta sum separates into the four independent lattice sums. The result is the product of the four kernels with coefficient $\prod_\tau m_\tau! / \prod_{\tau,b} \alpha_{\tau b}!$. This proves the finite and harmonic identity before taking any class means.

Average the point forms over the flat class variables using the three lifts of the constant function, each multiplied by its denominator clearer and the coefficient $\mathrm{N}(Q)^{1-d'}$ of its flat rank d'. As in the finite class formula, multiplication by D_{aut} clears the stabilizer orders in each class sum. Each flat coefficient is integral away from its conductor after including the two powers of its denominator clearer, since the class average uses an integral point kernel and the flat coefficient pairing has Gram matrix 1. The factor $\mathrm{N}(Q)^{1-d'}$ is a unit at ℓ. Thus the resulting Ξ is integral simultaneously at all readings over ℓ. The scalar averages are exactly $A/(h_1 v^{d'})$, where A is the flat integral algebraic part, v is the rank-one vacuum scale and h_1 is the chosen-level class count of the rank-one source. The three ranks sum to $2z_p - 12$. Substitution gives the normalized point test with the multiplier $h_1^3 v^{2z_p - 12}$. The marked-period verification supplies its testedness.

Character selection for the normalized rank-one test. Keep the same pair and any trial stage character ψ. Write h_1 for the number of classes of $U(1)$ at its chosen level. The depth calculation gives a fixed nonzero t_p and, for each embedding a above ℓ, a rank-12 indicator and monomial Z^α of degrees $m_\tau = c_p(\tau) - z_p$. Let $s_\alpha(g)$ be its kernel at the finite argument g, with the vacuum factor v^{12} and with the coefficient normalization $\eta^{\sum_\tau m_\tau} \prod_\tau m_\tau! / \prod_i \alpha_i!$. Here η is the scaled jet parameter. At $g = 1$ this is the depth coefficient in the asserted normalization, since an unscaled derivative multiplies a Taylor coefficient by $\prod_i \alpha_i!$. Thus $|a(s_\alpha(1)/t_p)|_\ell \geq \ell^{-P_\ell}$.

Choose z inverse to the representative of the identity class and use Fourier inversion on the finite torus quotient. The class quotient of $U(1)$ is a finite abelian group C. Each character of C lifts to a weight-zero character, tame on the chosen places and unramified elsewhere. For each character χ, form the average

$$B_\chi = \frac{1}{h_1} \sum_{c \in C} \chi(\det(g_c z))\, s_\alpha(g_c z)$$

on the chosen representatives g_c. Orthogonality gives $\sum_\chi B_\chi = s_\alpha(1)$. The ultrametric inequality selects χ with $|a(B_\chi/t_p)|_\ell \geq \ell^{-P_\ell}$, including when $\ell \mid h_1$. The division by h_1 is in Fourier inversion and costs no power in this lower bound. In the oscillator kernel the source splitting factor is $b(\det g)^{-1}$. Replacing b by b/χ on the rank-12 block supplies precisely the character in the sum.

Let $\mathcal{X}$ be the finite group of the permitted weight-zero characters on the chosen places. Write A_ϕ and B'_ν for the two fixed untwisted algebraic parts of the flat pairs of ranks 4 and $2z_p - 20$, with characters $\phi, \nu \in \mathcal{X}$. Let $A_2(\psi)$ be the analogous algebraic part of the fixed ramified rank-4 pair, so $\rho_{q_2}(\psi) = A_2(\psi)/A_2(1)$. Use the same normalized point form above for $\chi = \phi\nu$. The four-block pairing with the unscaled full kernel $K_\psi(g)$ gives

$$h_1^3 v^{2z_p-12} \langle \Xi, K_\psi(g) \rangle = \frac{\prod_\tau m_\tau!}{\prod_i \alpha_i!} A_\phi B'_\nu A_2(\psi)\, K_{12,b/\chi,\alpha}(g),$$

where $K_{12,b/\chi,\alpha}$ is unscaled. Testedness comes from the rank-one marked period.

Now $F_\psi = v^{2z_p} \eta^{\sum_\tau m_\tau} \sum_{c \in C} K_\psi(g_c z)$, since class translation preserves the unweighted sum. Substitution in the pairing identity cancels v^{2z_p-12}, leaves h_1^{-1}, and gives exactly

$$\frac{h_1^2 \langle \Xi, F_\psi \rangle}{\rho_{q_2}(\psi)} = A_\phi B'_\nu A_2(1) B_\chi.$$

Choose V_{jp} before any embedding by

$$V_{jp}^{|\mathcal{X}|} = \prod_{\nu \in \mathcal{X}} t_p A_2(1) A_\nu B'_\nu.$$

Every factor is nonzero. For fixed χ, the products $t_p A_2(1) A_{\chi/\nu} B'_\nu$ have that same total product. Their largest absolute value at a is at least $|a(V_{jp})|_\ell$. Use the corresponding two auxiliary characters in the point test. Each complex conjugate of V_{jp} is bounded by the product of the uniform archimedean bounds for the three flat parts and t_p. The flat parts contribute $3x^{34}$ and t_p contributes fmb/κ, whose sum is at most B_0 for large x. This proves (12.2) on either reading of the same ψ. The characters and monomial may vary with the reading, while V_{jp}, ψ and the ramified pair remain fixed.

Multiply the two inequalities (12.2). The product of the pairings, with the second conjugated, equals $\|F_\psi\|^2 \langle \Xi_{j_\iota}, P_{\mathrm{CF}_\psi} \Xi_{j_\iota \circ \mathrm{conj}} \rangle$. Write $J_{jp}(\psi) = h_1 \langle \Xi_{j_\iota}, P_{\mathrm{CF}_\psi} \Xi_{j_\iota \circ \mathrm{conj}} \rangle$. It is nonzero by the two nonzero tests. The ratios $\rho_{q_2}(\psi)$ are real. Hence

$$\left| j_\iota \left(V_{jp} \overline{V_{jp}} \frac{\rho_{q_2}(\psi)^2}{h_1^3 \|F_\psi\|^2} \right) \right|_\ell \leq \ell^{2P_\ell} |j_\iota(J_{jp}(\psi))|_\ell.$$

Use the bound with $J_{jp}(\psi)$ and multiply by $h_1^3 I_{jp}$, by the unit power of the conductor norm and by $\rho_{q_2}(\psi)^{-2}$. The first factor is integral and the last has local logarithm at most $2v_\ell$ whenever the trial character satisfies

the ramified inverse-ratio estimate with allowance v_ℓ. By the Rallis normalization of the rank-one algebraic part, $\rho_p(\psi)^{-1}$ is $\|F_1\|^2\|F_\psi\|^{-2}$ times that conductor power. We obtain

$$\log|j_\iota(A_{jp}/\rho_p(\psi))|_\ell \le \log|j_\iota(J_{jp}(\psi))|_\ell + 2P_\ell \log \ell + 2v_\ell. \tag{12.6}$$

The factorization of $\Gamma_j(\psi)/\Gamma_j(1)$ has $\kappa/2$ such rank-one factors and $5m$ ramified factors, all squared, and a conductor norm power of absolute value one at each reading away from its conductor. The rank-one factors occupy $7m/2 < z \le 7m/2 + \kappa/2$ in the unsigned list. The other factors come from the first $7m/2$ places in that list and the two signed lists of lengths m and $m/2$. Put

$$K_\Gamma = \prod_{j,p} A_{jp}^2, \qquad Q_\Gamma(\vartheta) = \prod_{j=1}^{m} \frac{\Gamma_j(1)}{\Gamma_j(\vartheta^j)},$$

with p ranging over the sharp pairs in K_Γ.

13 Determinants of the realized chains

The realized reversals in the source triangle. We use the reversal on the same extended source tests that occur in the first commutation. Let the two original ranks be n_0, n_1 and let D be their finite matrix datum over $R_* = \mathcal{O}_E/\mathfrak{m}$. Set $u = \overline{\alpha_f \varpi_\mathfrak{m}}/(\alpha_f \varpi_\mathfrak{m})$, with $\varpi_\mathfrak{m}$ the prescribed finite generator. On the two finite half-spaces the normalized adjoint of a matrix map exchanges its two polarizations. Conjugation first passes to the opposite source form, with a finite multiplication by a unit of norm minus one and the conjugated polynomial. On returning by the adjoint map the same unit occurs as the change of Fourier variable. The identity $c_0^{-1} = -\bar{c}_0$ for the fixed unit c_0 of norm minus one shows that the resulting primitive character is the original $\psi_\mathfrak{m}$. At a split pair the scaled vacuum support on the second half has now been moved by u; the two factors in u are the conjugate and original scales of this support. Off the modulus the factors from the fixed Weil scale are included in u in the same way.

For a point indicator at A the adjoint coordinates carry its finite translation to the modulation $\psi_\mathfrak{m}(\mathrm{tr}(YA))$, in the transposed variables. This follows from the Heisenberg pairing of those two polarizations. Use the same gauge as in the diagonal line reversal. The point indicator is normalized by the measure $1/|R_*|^{n_0 n_1}$ of its coset in the standard matrix box. The character function has absolute value one on that box. Since the normalized adjoint map is unitary, the factor relating the indicator kernel to the unnormalized Fourier character kernel has squared modulus $|R_*|^{-n_0 n_1}$. The global change of theta functional has modulus one by the rational Heisenberg normalization. At infinity the conjugation and adjoint carry a matrix coordinate to its signed transpose, so a polynomial is changed to the signed transpose $P^\natural$ of its coefficientwise conjugate. In particular vacuum moduli introduce no further twist-dependent factor. Linearity in the functions on the residue quotient gives

$$\overline{K_{D,P}(g,h)} = \tau_\vartheta K_{\widehat{D},P^\natural}(h, gu), \qquad |\tau_\vartheta|^2 = |R_*|^{-n_0 n_1}.$$

Here the two determinant characters exchange positions, with the same signs as the two original groups. This formula uses the full matrix dimension for each edge.

The real entry of a reversed edge. On a reversed edge take the matrix entry of the smooth one-edge oscillator quotient on the two selected real lines. For any polynomial Q write $\ell_\infty(Q)$ for that entry, with the determinant characters and real measures of the edge. A theta pairing with the fixed exterior source and target types factors through the same quotient by Howe's multiplicity one, as in the square comparison. The joint harmonic P_{edge} has nonzero entry by the minimal-degree calculation there. Thus replacement by Q multiplies the pairing by $\ell_\infty(Q)/\ell_\infty(P_{\mathrm{edge}})$. The finite part of the pairing is the same in this calculation.

Equivalently, when Q is represented on the selected line by convolution by α_∞ on the other group, the ratio is

$$\int \alpha_\infty(x)\,\overline{e'(x)}/e'(1)\,dx,$$

where e' is the prescribed lowest-vector coefficient. The opposite matrix entry conjugates this scalar.

For the actual reversal use the transported harmonic $P^{\natural}$ on the reversed edge. The signed transpose and coefficient conjugation interchange the mixed-sign Laplacians of the two real actions. They exchange the paired compact types, with the reversed compact order on the contragredient argument. They send the nonzero joint highest line to the corresponding joint highest line in those orders. Thus $P^{\natural}$ has nonzero matrix entry. Using the chosen generator of that same line changes the reversal by a fixed nonzero real-place normalization, independent of the finite test, points and twist.

For the adjoint companion in V, the reversed square has equal intermediate ranks and the same weight and character product, so the complementary integral has $e = 0$. Its finite identity uses the *same* p^+. More precisely, for an additive Fourier variable $\left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$, the transform of $\overline{p^+}$ is $1_{C=0}\widehat{\bar{p}}(-A)\psi(\operatorname{tr} D)$. This is the trace expansion of the block product and orthogonality of the average over S. The corner sum over D now equals the matrix Gauss sum times the value at $D = I_b$. To see its modulus, work at a field factor of the squarefree ring R_*. The determinant covariance forces a function of D to vanish at singular D, by a stabilizer scaling a complementary line with any determinant. The same argument shows that its additive Fourier transform is zero at singular dual matrices. At invertible dual matrices the values differ by unitary determinant characters. Finite Plancherel applied to the function $\chi(\det D)$ on invertible matrices gives $|G_b|^2 = |R_*|^{b^2}$, after multiplying the field factors. The free block adds a factor $\#\{S\} = |R_*|^{bN_0}$, where N_0 is the rank at length n. Thus the reversed comparison uses $|G_b \#\{S\}|^{-2} = |R_*|^{-b(b+2N_0)}$. The translations from the two reversals move through the intervening upward lift up to its unitary determinant phase. Taking adjoints of the reversed edges gives a factor of squared modulus $|R_*|^{-N_0N_1}$ for the first reversal and $|R_*|^{N_1N_2}$ for the inverse second reversal. Thus the scalar $t_{i,n}$ in the companion has squared modulus a fixed multiple of $|R_*|^{N_1(N_2-N_0)-t(t+2N_0)} = |R_*|^{t^2}$, since $N_1 = N_0 + t$ and $N_2 = N_0 + 2t$ in V. The real type convolutions are fixed in the twist. In detail, one first writes $\langle d_{i,n}a^*_{i,n+1}x, y\rangle$ as a multiple of the pairing of $a^*_{i,n+1}x$ with the reversed lift of $R(u^{-1})y$. Move the upper adjoint to the other argument, commute the two lifts by the reversed rule, move the central translation through the lower lift, then reverse the remaining left edge back. This gives $t_{i,n}\langle d_{i,n+1}x, a_{i+1,n}y\rangle$. The scalar introduces no ramified ratio. The first commutation is available in both triangles, while the adjoint companion uses the squarefree ring of V.

The Gram recurrence. The square law just proved is

$$q_n a_{i+1,n} d_{i,n} = \frac{q_{n+2}}{\overline{s_{i,n}}}\, d_{i,n+1} a_{i,n+1},$$

with the reciprocal of zero read as zero. Its adjoint companion is $\langle d_{i,n}a^*_{i,n+1}z, y\rangle = t_{i,n}\langle d_{i,n+1}z, a_{i+1,n}y\rangle$. Here $s_{i,n}$ costs one square and $t_{i,n}$ costs zero. The first descent is adjoint to the first lift on the base forms under the inclusions, and the composite $d_{i,0}a_{i,0}e_i = e_{i+1}B_i$ is injective there.

Induction on n gives scalars $s_i^{(n)}, t_i^{(n)}$, products of the first n successive square factors in row i, such that

$$M_{i+1,n} d_{i,0} a_{i,0} = \frac{q_{n+1}}{q_1 q_0 \overline{s_i^{(n)}}}\, d_{i,n} M_{i,n+1},$$
$$\langle z, a_{i,n} M_{i,n} e_i y\rangle = t_i^{(n)} \langle d_{i,n} z, M_{i+1,n} e_{i+1} y\rangle.$$

For the first formula multiply by the next weighted lift $q_n a_{i+1,n}$ and use the square law. For the second move $a_{i,n+1}$ to the first variable, apply the current identity and then its adjoint companion, using $q_n a_{i+1,n} M_{i+1,n} = M_{i+1,n+1}$. For $n = 0$ the first formula follows from $M_{i,1} = q_0 a_{i,0}$ and the second is the first-arrow adjoint identity. The first formula also proves injectivity of the weighted chains on the base by induction, since

applying the last descent to a zero image gives a zero image of $B_i y$ under the next shorter chain. For any nonzero base form y, the left side of the first identity evaluated at $e_i y$ is therefore nonzero. Hence $s_i^{(n)} \neq 0$ for all the factors used in the norm comparison. For the second identity take $z = a_{i,n} M_{i,n} e_i y$ for a nonzero base vector y. The left side is then positive by chain injectivity, so $t_i^{(n)} \neq 0$ on the same range. This establishes the nonvanishing of the realized comparison factors used here.

The first-arrow Gram form is $q_0 \langle a_{i,0} e_i x, a_{i,0} e_i y \rangle$ on the good base forms. Write S_i for its operator.

The determinant of the chain. Pairing the two images under $M_{i,n+1}$ and using both induction formulas reduces their pairing to that for $M_{i+1,n}$ on $B_i x, y$, multiplied by $q_n q_1 q_0 \cdot t_i^{(n)} s_i^{(n)} / q_{n+1}$. Thus the full Gram form on W for the weighted chain in row zero is that of $S_0 \cdots S_{m-1}$, times $\left(q_0 q_1^m / q_m\right) \prod_{i=0}^{m-1} t_i^{(m-i-1)} s_i^{(m-i-1)}$. The cost of the latter product of square factors is at most $\sum_i (m-i-1) \leq m^2$ squares. This recurrence takes determinants of the product without a choice of a common eigenbasis.

The first-arrow calculation gives the precise determinant ratio

$$\det S_{j-1} = \left(\frac{C_{j-1}}{A_j \Delta_\vartheta} \right)^R \lambda_j(\vartheta^j).$$

The conversion from the weighted triangle norm to the class-sum norm multiplies its Gram form by the final weight dimension and by the last restriction level index. The following count identifies that index with the factors used in the single-lift comparison.

The terminal level in the chain norm. For the fixed chain let J_m be its last restriction level. At a new place with residue field of size q, the block sizes of its parabolic are $r, t, \ldots, t$ with m blocks of size t. Each strictly upper block is arbitrary, so its unipotent order is q^{e_m}, where $e_m = mrt + \binom{m}{2} t^2$. Writing $G_u(q) = |\mathrm{GL}_u(\mathbf{F}_q)|$, the local index is

$$\frac{G_{r+mt}(q)}{G_r(q) G_t(q)^m q^{e_m}}.$$

For the single auxiliary lift the local target index is $G_{r+t}(q)/(G_r(q) G_t(q) q^{rt})$. Its section has q^{rt} free residues after fixing the marked submatrix, and the source projection contributes $G_r(q)$ before squaring. The normalization divides by $G_r(q)$. Thus the local factor of the actual auxiliary Gram scalar Δ_ϑ is $G_{r+t}(q)/G_t(q)$. Meanwhile the triangle weight q_n acquires $G_{r+nt}(q)$ at this place. The same level modulus occurs for the m powers, by the definition of the trial m-fold twist. Multiplying gives

$$[K_{\mathrm{full}} : J_m] = C \, \mathrm{N}(Q_\vartheta)^{-e_m} \Delta_\vartheta^m \frac{q_m}{q_0 q_1^m},$$

with C fixed by the untwisted chosen places. Hence the weight product cancels on the realized chain, leaving a conductor power of absolute value one at each reading away from Q_ϑ. The remaining passage to the terminal class-sum norm uses the dimension of its compact coefficient type and the actual realization of the weighted arrows. The square and adjoint recurrences use the same descent datum at a node in both applications. Use the common Rallis operator and Hecke transfer in the chain calculation above.

Taking absolute values thus gives

$$Z_\vartheta = |c_\vartheta|^R \prod_{j=1}^{m} \lambda_j(\vartheta^j), \qquad Z_\vartheta \neq 0.$$

The square-factor product controls c_ϑ up to a fixed constant and a conductor norm power. At any reading ι with $\ell \nmid \mathrm{N}\, Q_\vartheta$ the latter is a unit. Suppose the ramified inverse ratios at all charged powers of this trial character have local logarithm at most v_ℓ. There are at most t squared ratios per square and at most m^2 squares. Dividing the squared bound by two gives

$$\log |j_\iota(|c_1/c_\vartheta|)|_\ell \leq m^2 t v_\ell. \tag{13.1}$$

Fix the product of the untwisted compressed determinants

$$D_0 = \prod_{j=1}^{m} d_j(1), \qquad \log|\sigma D_0| \le 2RrmB_0 \quad \text{for every } \sigma.$$

14 Endpoint periods and a completion of the rows

Endpoint period maps. Write $\mathcal{H}_0$ for the full principal-level space on Y_0 and $\mathcal{H}_c$ for the prescribed core forms on $Z_{0,0} = Y_0 \perp \langle 1 \rangle$. Let $\mathscr{F}$ be the level space on the definite end of the core column. Write $S : \mathscr{F} \to \mathcal{H}_c$ for its fixed untwisted descent. It is onto by the transpose calculation below.

For each real translate a_∞ let $r_{a_\infty} : \mathcal{H}_c \to \mathcal{H}_0$ represent the bordered restriction pairing by Riesz duality. Write K for the chosen principal level on Y_0 and K^{deep} for the level of modulus $\bar{\pi}_G^{b_1}$ there. The restriction identity against a test of the twisted bottom level uses the same r_{a_∞}. Indeed, for $k \in K$ choose v supported at the chosen places with reduction inverse to k at depth b_1. Then v fixes the test, since it is principal modulo the chosen places and is trivial on the twist places. The element kv lies in K^{deep}, so its block image fixes the core form. By simultaneous finite translation in the period, replacing the test by its k-translate does not change the period. Averaging over K gives the asserted pairing against r_{a_∞}, independent of the twist. Put $Q_{a_\infty} = r_{a_\infty} S$.

Write T_ϑ for the final-level chain on $\mathcal{H}_0$, with restriction Θ_ϑ to W. At Y_m let K_{full} be the full integral terminal level. Write $K_\vartheta^{\mathrm{pr}}$ for its subgroup with identity reduction at the chosen and twist places, in the specified split coordinates. Write $K_\vartheta^{\mathrm{res}}$ for the subgroup with identity reduction at the chosen places and upper block triangular reduction at the places of Q_ϑ, with block sizes $r, t, \dots, t$ and m blocks of size t. The latter is the restriction level used in the chain norm. Its tame character on an element with diagonal blocks $k_{ll,w}$ is $\prod_{w|Q_\vartheta} \prod_{l=1}^{m} \vartheta_w^l(\det k_{ll,w})$, trivial on $K_\vartheta^{\mathrm{pr}}$. The initial path identity, exchanges and normalization of the last path give

$$[K_{\mathrm{full}} : K_\vartheta^{\mathrm{pr}}] \int \Delta_{h,v}\, T_\vartheta(Q_{a_\infty}\phi) = \frac{\alpha_{a_\infty}}{c'_\vartheta} \mathcal{E}_\vartheta(\Delta_{h,v}, \phi). \tag{14.1}$$

The integral takes the automorphic realization of the algebraic form at level $K_\vartheta^{\mathrm{pr}}$, with $\Delta_{h,v}$ the dual-weight class delta test at h of value v. The determinant uses the class sum at $K_\vartheta^{\mathrm{res}}$. The two readings evaluate the same lift at a given adelic argument. The normalized path calculations below give the same c' for all real translates, with a scalar α_{a_∞} that may depend on the translate.

The first path and its level weights. The weighted lift with the inclusion datum and the descent on the first column of a grid use the same kernel. Their transpose identity follows by Fubini in the bilinear pairing of a form and its conjugate type, with the factor q_j on each arrow. Each endpoint image after a nonempty sequence is fixed by the kernel of reduction in its triangle ring. The inclusion datum is fixed there on either side, since both determinant characters are trivial on this kernel. In Z take the actual untwisted ring of depth b_1. The core descent maps its terminal principal-level space to the bottom principal-level space. If a bottom vector annihilates the image, test the transpose identity with the conjugate of its upward image at the terminal level. The resulting pairing is the squared norm of the upward image. The first commutation and the injectivity of the first arrow, applied inductively to the core triangle, make the upward image injective, so the bottom vector is zero. This proves surjectivity of the required descent S on the full core level.

For the initial path put $R = \mathcal{O}_E/(\mathfrak{b}Q_\vartheta)$. In the grid let $I_{z,a,b}$ be the index of the residue subgroup at node (a, b). This subgroup is trivial on the first column of V. In Z and off the first column of V it is the kernel of reduction to $\mathcal{O}_E/\mathfrak{b}$. Thus $I_{1,a,b}$ and $I_{0,a,b}$ for $a > 0$ are independent of the twist. On the first column of V, the ratio $q_j/I_{0,0,j}$ is independent of the twist by reduction of matrices from $\mathfrak{b}Q_\vartheta$ to $\mathfrak{M}Q_\vartheta$. At a truncated discrete valuation ring the kernel consists of all matrices congruent to the identity at the

smaller depth. Its size is the corresponding residue cardinality to the power N^2 at a node of rank N. On the first column of Z the two orders are equal.

Expand the two first-column descents by pulling back their respective inclusion data to R, and then summing point indicators. The coefficient is either zero or one. The matrices on the twisted edges have left inverses because they reduce to the inclusion at every maximal ideal and hence split over each local factor of R. The sum has at most $|R|^{(2n_m)^2(2m)}$ terms. Let Φ be this sum as a linear functional on functions of tuples. If g_p acts at node p and a matrix on edge $p \to q$ changes to $g_q^{-1} X g_p$, the character of the edge is evaluated at $\det(g_q)/\det(g_p)$. Write $\eta(g)$ for the product of these unitary factors over the two legs. Initially Φ has the required covariance for the residue subgroup at every node because reductions to the triangle rings act trivially there.

Choose one representative t_x for each tuple of cosets of these subgroups at the $2m$ nonterminal nodes. At the bottom node of Z, use the block extension of the chosen bottom representative in V. Keep the former coset index in the sum so its multiplicity is retained. At the terminal nodes take the identity. For a function F of the tuple write $A_g F(X) = F(g \cdot X)$, and put $\Lambda_0(F) = \sum_x \eta(t_x)\Phi(A_{t_x^{-1}}F)$. To check covariance, permute the cosets by an admissible family g. The error $t_{gx}^{-1} g t_x$ lies in the level at each node. At the modified bottom node of Z this holds because the error in V is the identity and both lifts are the same block extension there. The character of this error cancels against the changed coefficient $\eta(t_{gx})$. Thus $\Lambda_0(A_g F) = \eta(g)\Lambda_0(F)$. The coefficient in each term is a root of unity. This second sum adds at most $2m$ tuple powers, giving order $4m$ in the compact model.

In the actual bilinear period at the bottom, simultaneous translation in the two groups by k and $k \perp 1$ preserves the integral for every real translate. Interior translations on each leg cancel, leaving exactly $\eta(t_x)^{-1}$ on the term read at t_x^{-1}. Hence the coset sum contributes the product of the $2m$ vertex indices. This converts the weight product from the expansion into $\prod_{j<m}(q_j/I_{0,0,j})$, a fixed constant. At the terminal principal level $K_\vartheta^{\mathrm{pr}}$ the finite index is q_m. The stage periods put $I_{0,0,m}$ on the test form before using the two grid paths. Converting this last factor to q_m adds $q_m/I_{0,0,m}$. The initial identity therefore holds with $C_0 = \prod_{j=0}^{m}(q_j/I_{0,0,j})$, with all the q_j here taken on V. Insert the twist-free bordered Riesz map r_{a_∞} and the core descent S, and transpose the V-leg. This gives the left side of (14.1) with $Q_{a_\infty} = r_{a_\infty} S$ and the right side $C_0 \mathscr{P}_0$. The datum Λ_0 is independent of a_∞.

Pointwise transport through the off-diagonal squares. At a corner (a, b) the old path between the exterior nodes $(a, b+1)$ and $(a+1, b)$ goes through (a, b). The new path goes through $(a+1, b+1)$. Thus moving the descent past the lift uses the complementary identity from the smaller to the larger intermediate space. Apply it to the kernel integrated over that intermediate group at every pair of exterior arguments. For an off-diagonal square with corner (a, b), the two intermediate ranks differ by $2t$. At offset zero the negative count in either grid depends only on b and is $\lfloor (m - b + z)/2 \rfloor$. Indeed, the only negative entries are the stem positions with $b < A \le m - 1 + z$ and $A + b$ odd. This verifies the signed sum condition. The exterior rank is at least $2(r + z + t)$, and each exterior node has at most $r + z$ negative signs. Since $t > r + z$, the rank is at least twice the negative count plus two. The total rank is at most $4n_m$ and the weight $w_{a+1} + w_{b+1}$ exceeds $2n_m$. The floor and ceiling in V add to half the total rank because the two ranks have the same parity. The two products in Z have integral halves. In the square rule this gives $\nu^{-1} = \mu_{a+1}^0 \mu_{b+1}^0 \vartheta^l$, where $l = b + 1$ for V and $l = a + 1$ for Z. The residue exponent of ω at chosen places is $a + b + 2 \le 2m$, and the twist is nontrivial at its places by the m-fold condition. Use the pointwise complementary-period identity with $e = t$ on the finite functions of modulus $\mathfrak{M}^{b_1} Q_\vartheta$. Its arbitrary residue depth covers both grids, including the ordinary V-square of the diagonal move.

Let $\mathcal{V}_0(a_\infty)$ be the real test on the initial path. Move the right translate on the core form into the adjacent real Weil vector by covariance of its theta kernel. This uses the same real action at every twist, since ϑ has weight zero. At each square let $T_\square$ be the small-to-large real map from the pointwise identity. The zero-slice extension uses fixed Schwartz factors of value one. The top construction uses the fixed compact chart and bump and tensors with its fixed determinant vector. The real cone maps have positive

exponents along this descent to $e = t$. Their estimates give continuity in the Schwartz topology. Thus $T_{\square}$ depends only on the real data of the square. The finite Tate factors remain in the normalized comparison scalar $c_{\square}(\vartheta)$, whose squared modulus is the fixed multiple of $\Pi_{v,S}(t)$ in the complementary identity.

Use $T_{\square} \otimes 1$ on the whole real test, with the other edges as tensor factors. First take a finite sum of decomposable real tests. At fixed exterior arguments and fixed finite tuple, the spectator factors come out of the integral over the square intermediate group. The finite tuple map below then extends linearly over the partner graph with the same $T_{\square}$ on every term. The primitive-support estimate in the split model bounds the absolute intermediate integral by a Schwartz seminorm times a power of the exterior heights. It applies to the complete column and row stacks of the partner data. To pair at an exterior node, integrate the prefix of the leg from its initial form. The one-edge first-occurrence calculation applies in the traversed orientation to each pure Schwartz factor. The output is a smooth cusp form of the prescribed type. The theta kernel's derivative bounds and the zero-constant-term rapid decay estimate bound its rapid-decay seminorms by Schwartz and input cusp seminorms at the fixed finite level. These estimates justify Fubini at the intermediate group and the adjacent exterior node. Rapid decay also bounds the paired border integral. A short vector in the smaller lattice is still short in the larger lattice of the block embedding, and a fixed translate changes the height bound by a constant. For each fixed finite tuple this gives continuous multilinear bounds on pure tests and hence continuous linear periods on the Schwartz space of the direct sum. Density of finite sums of pure tests gives the same identities for arbitrary real tests. The estimates may depend on the fixed finite tuple. The real map and the scalar in the pointwise identity are common to all tuples.

The finite move on the actual tuple sums uses the following direct description. For matrices $(X, Y) = (X_k, Y_k)$ in the partner graph give the pair coefficient $\chi(\det k)/|\mathrm{GL}_b(R)|$, with b in this formula the smaller intermediate rank, and give all other pairs coefficient zero. Completions of a pair differ by a determinant-one block shear, so the coefficient is well defined. For fixed primitive Y, sum over its orbit Yg with weight $\chi(\det g)^{-1}$. The two partner sums then give the pointwise square rule with this graph coefficient. Indeed each of the $|\mathrm{GL}_b(R)|$ replacements gives the same product of weights and kernel values. If a block of Y has a left inverse, the opposite block of X does as well. To see this, reduce at each maximal ideal and use $\ker X = \operatorname{im} Y$: a vector in the kernel of the opposite block would give a vector of $\operatorname{im} Y$ with zero projection on the original primitive block. That projection is injective on $\operatorname{im} Y$. Full column rank in all residue fields gives a left inverse over R.

Use one representative for each right $\mathrm{GL}_b(R)$-orbit of Y. Retaining this representative and multiplying its graph coefficient by $|\mathrm{GL}_b(R)|$ gives the same functional as keeping the average. This follows by averaging over the admissible change of basis at the removed node, extended by one on the extra line in a diagonal move. The action is free on the primitive frames. Under exterior changes g, new-row change f and old-column change v, the completion changes from k to $g^{-1}k\,\mathrm{diag}(v, f)$. The character of its determinant is exactly the quotient of the new and old edge characters. Consequently the new functional has the required covariance and root-of-unity coefficients. The two new matrix variables add at most two to the tuple order. The finite split matching at depth b_1 uses this graph over R.

At stage j move the last vertical descent past the j lifts in each grid. Apply the square rule with all other edges fixed and with the complementary pairing as the linear functional. These $2j$ moves raise the order by $4j$ and use $2j$ square products. What remains is the pair of descents to the diagonal (j, j). There is an isometry $Z_{j,j} = V_{j,j} \perp L_j$, with line sign $(-1)^j$, by the negative counts and the actual diagonal basis identifications. In the off-diagonal position the groups $Z_{j,j+1}$ and $V_{j+1,j}$ are identified by a relabeling, since they have equal negative counts and the same even rank $N_j = r + (2j + 1)t$. Use this common group for the seesaw. At the old diagonal the V-rank is odd: adding the line preserves the ceiling of half the rank. Thus the kernel on the Z-edge leaving that diagonal splits into the corresponding untwisted V-kernel and the kernel from L_j of line character one. At the new diagonal the floor increases by one. The new twisted V-kernel is multiplied by the kernel to L_{j+1} of line character v. At the common even-rank group

the twisted and untwisted characters differ by ϑ^{j+1}. This twist enters the line reversal only through a unitary determinant factor at the line argument.

The first seesaw splits the old edge matrix into the columns for $V_{j,j}$ and the column for its signed line.

The normalized reversal on the diagonal line. Fix $j < m$. Use the identification of $Z_{j,j+1}$ with $S = V_{j+1,j}$ on both kernels. Write $N = N_j$ for its rank, s_i for its signs, and $\epsilon = (-1)^{j+1}$. The old and new lines are $L_{-\epsilon}$ and L_{ϵ}, respectively. Use $R = \mathcal{O}_E/\mathfrak{m}_\vartheta$ with $\mathfrak{m}_\vartheta = \mathfrak{b}Q_\vartheta$, and write λ for its finite idele generator, equal to one off the support A of the modulus. This support is disjoint from its conjugate. Put $u = \overline{\lambda\alpha_{\mathrm{W}}}/(\lambda\alpha_{\mathrm{W}})$, $k = (1, u^j)$ and b, a' equal to the horizontal and vertical edge characters on S. Denote the modulus kernel by $K^{U,V}_{a,b}[\Psi, P](g, h)$, with character multiplier $a(\det h)b(\det g)^{-1}$ on the underlying Weil action on $\mathrm{Hom}(U, V)$. There are a primitive additive character ψ of R and an $f_j = f_j(\vartheta) \neq 0$ such that, for each $g \in U(S)(\mathbb{A})$, column $v \in R^N$ and polynomial P,

$$K^{L_{-\epsilon},S}_{1,b}[\mathbf{1}_v, P](k, g) = f_j \sum_{y \in R^{1\times N}} \psi(\mathrm{tr}(yv))\, K^{S,L_\epsilon}_{a',v}[\mathbf{1}_y, P^\vee](g, (1, u^{j+1})), \qquad |f_j|^2 = |R|^{-N}. \tag{14.2}$$

Here v is the restriction of the fixed splitting character to the norm-one ideles. The map $P \mapsto P^\vee$ is complex linear. In the i-th column variable it substitutes $-s_i\epsilon$ times the corresponding row variable, at each infinite place. Thus the substitution is independent of ϑ. This is a complete-kernel identity. The conjugate reversal of a chain edge uses coefficient conjugation and a highest-line comparison on realized form spaces.

Let c_{neg} be the fixed integral finite unit of norm -1 in the standard signed line, and put $c_0 = c_{\mathrm{neg}}^{\epsilon}$. The old line point is $x_- = (1, c_0)x_+$, where x_+ is the standard point of L_ϵ. Set $\ell_0 = (1, \bar\alpha_{\mathrm{W}}/\alpha_{\mathrm{W}})$ and $\xi = (1, -c_0\alpha_{\mathrm{W}}/\bar\alpha_{\mathrm{W}})$. Then $x_+\ell_0 = \bar\xi x_-$. In pointed coordinates let D_S and D_- be the standard Hermitian matrices of S and $L_{-\epsilon}$. They give the prescribed signs at infinity and are identity matrices at finite places. The adjoint identification from new row coordinates to old column coordinates is

$$Y \longmapsto \xi D_S Y^* D_-.$$

It identifies rational vectors for the new line pointed at $x_+\ell_0$ with the old rational vectors. Indeed, for an old rational matrix M the adjoint $J_-^{-1}M^*J_S$ is rational, where J_-, J_S are the rational Hermitian matrices. If $X = x_S^{-1}Mx_-$ is its old column coordinate, the new row coordinate of the adjoint is $(x_+\ell_0)^{-1}x_-D_-X^*D_S$. The displayed identification sends this row to X.

In the coordinates scaled by α_{W} on the finite model, this map is $Y \mapsto -c_0Y^*$. The self-dual additive character there is $\psi_0(z) = \psi_f(\jmath\alpha_{\mathrm{W}}\bar\alpha_{\mathrm{W}}z)$, where $\bar\jmath = -\jmath$ is the fixed Weil datum. Thus $\psi_0(-\bar z) = \psi_0(z)$. The identity $\bar c_0c_0 = -1$ gives preservation of the finite Heisenberg bicharacter, including its factor $1/2$, and of the twisted integral-box vacuum. At infinity the maps are $z \mapsto r\bar z$ on coordinates of old sign $r = -s_i\epsilon$ and new sign $-r$. They preserve the Gaussian and transport the bicharacter with sign $-r$ to that with sign r. With $\zeta_r(z) = \bar z$ for $r = 1$ and $\zeta_r(z) = z$ for $r = -1$, one has $\zeta_r(r\bar z) = r\zeta_{-r}(z)$. This gives precisely $P^\vee$ in the same Fock normalization.

For the action $X \mapsto gXh^{-1}$ of the old pair, put $T = g \otimes h^{-\mathsf{T}}$, where h lies in the old line group. The action of the exchanged pair in the new coordinates has matrix $T^{-\mathsf{T}}$ after the coordinate permutation. The explicit Weil action in the vacuum model is

$$\omega(T)F = \Delta(T) \int F(X)\rho(TX)e_0\, dX,$$

where ρ is the Heisenberg action and e_0 the vacuum. The unweighted integrals intertwine by the change of coordinates, which preserves the vacuum Haar measure. To compare the multipliers, transposition preserves the finite elementary divisors and the real principal minors. In the real factor $\Delta_r(T) = \overline{\det T_{--}}$, Jacobi's complementary-minor identity gives $\overline{\Delta_r(T)} = \det(T)\Delta_{-r}(T)$. At the finite places $\Delta_f(T) = \chi_f(\det D)$,

where $D\widehat{\mathcal{O}}_E^N = \widehat{\mathcal{O}}_E^N \cap T^{-1}\widehat{\mathcal{O}}_E^N$ and $\chi_f = \mathrm{Y}_f|\cdot|_f^{-1/2}$, with Y_f the finite part of the full idele splitting character. One may use TD for the corresponding lattice of T^{-1}. As $|\det T|_f = 1$, it follows that $\Delta_f(T^{-1}) = \mathrm{Y}_f(\det T)\Delta_f(T)$. Together with $\Delta_r(T^{-1}) = \overline{\Delta_r(T)}$, these relations give

$$\frac{\Delta_{\mathrm{old}}(T)}{\Delta_{\mathrm{new}}(T^{-\mathsf{T}})} = v(\det T)^{-1} = v(\det h)^N v(\det g)^{-1}.$$

Consequently the old characters $1, b$ become $b_* = b^{-1}v^N$, v on the exchanged kernel.

The theta functional is invariant under rational Heisenberg operators. In a rational Schrödinger polarization it is unique up to scalar. Rational modulations force an invariant distribution to be supported on the rational lattice, since that lattice is its own annihilator. At a lattice point the modulation factors minus one give local coordinates at infinity, so the distribution has no derivative terms there. Rational translations make all point masses have the same coefficient. The adjoint preserves the vacuum and gives an isomorphism of the rational quotients. Both theta functionals have mean square one on translates of their vacua in Haar probability measure on these quotients. Their proportionality factor has modulus one. Moving the new line point from $x_+\ell_0$ to x_+ puts ℓ_0 into the line argument and contributes the unitary character $b_*(\det \ell_0)$ in the identity for the new kernel in terms of the old. This follows by multiplying the underlying Weil operators and keeping the explicit determinant character on the moved point. [5]

To compute the factor with the prescribed finite datum, start with the Heisenberg translate of the old vacuum by $-\lambda^{-1}\tilde{v}$ at A, in the scaled finite coordinates, for an integral lift $\tilde{v}$ of v. The split Schrödinger function of this translate is $\mathbf{1}_{\mathcal{O}_A^N}(X - \lambda^{-1}\tilde{v})$. Acting by $d = (1, \bar{\lambda}/\lambda)$ in the old line changes its argument to $\mathbf{1}_{\mathcal{O}_A^N}(\lambda^{-1}(X - \tilde{v}))$, the specified residue indicator. On the split integral model the action on functions is pullback by the inverse matrix times the square root of its absolute determinant, with the unitary splitting and edge characters. Here the inverse matrix on the position coordinates is $dI_N = \lambda^{-1}I_N$ at A, so the scalar C of this action satisfies

$$|C|^2 = \prod_{w\in A} |\lambda_w|_w^{-N} = |R|^N.$$

This computation uses the characteristic function of the residue class on the original vacuum.

On the new side the transported translation is supported on $\bar{A}$, with coordinate $\overline{\lambda^{-1}(-\bar{c}_0)\tilde{v}^{\mathsf{T}}}$ there. In the split Schrödinger model on A, translation at $\bar{A}$ multiplies by the character obtained by pairing with the new position variable. Thus the new named function is $\psi(\mathrm{tr}(yv))$ on the integral box, with $\psi(t) = \psi_0(\lambda^{-1}(-\bar{c}_0)\tilde{t})$ on components in A. The self-dual lattice pairing and the unit $-\bar{c}_0$ make ψ primitive, including at the prime powers of $\mathfrak{b}$. This verifies the sign and the orientation of the finite Fourier row. Since $\ell_0 d = (1, u)$, the new line point is exactly $k\ell_0 d = (1, u^{j+1})$. The adjusted kernel in terms of the old kernel at k has multiplier C times the theta and point phases. On inversion it has squared modulus $|R|^{-N}$. Finally changing b_* to a' on the first character slot of the new kernel places the factor $(b_*/a')(\det(1, u^{j+1}))$ on the right of (14.2). Its modulus is one and it is independent of g, v, P. Expand the character function on the residue box in the row indicators to obtain the unnormalized sum in (14.2).

The signed real adjoint, including the fixed change of polarization, acts continuously on Schwartz space. Hence the identity extends by density at every fixed g and finite column. It acts on the line factor of a real test with identity on all other factors. The finite Fourier row is unchanged. With C_j the normalized complementary scalar for the ordinary diagonal V-square, the combined diagonal denominator is C_j/f_j. Its squared inverse-line factor is $|R|^{N_j}$. Relative to the trivial twist this factor contributes $\mathrm{N}(Q_\vartheta)^{N_j}$, a unit at the coefficient place selected for that twist. The determinant uses the squared modulus, in which the unitary phases cancel.

The diagonal square and stage product. Use the tensor product of the Weil kernels in the first seesaw. It separates the Z-edge into the ordinary V-edge and the old line kernel on the common exterior group

[5] For the rational summation model and the adélic oscillator action see Weil, *Sur certains groupes d'opérateurs unitaires*, Acta Math. 111 (1964), §§18–19 and §41.

S, using the same product measures. At the two exterior arguments, reverse the line by (14.2) and apply $T_{\square}$ to the ordinary V-square. This square moves from $V_{j,j}$ to $V_{j+1,j+1}$, of rank $2t$ greater. The line kernel is a spectator for the square, since it depends on the exterior S-argument and the indicated pointed line argument. The marked old Z-edge has a left inverse, so its restriction to the non-line columns does also. The partner graph therefore preserves the needed primitive block. In the second seesaw tensor the new row for L_{j+1} with the new V-kernel. For pure tests this is the direct sum identity of the kernel and for general tests it is the corresponding continuous Schwartz tensor identification.

In the split presentation the line factor multiplies the exterior prefix on S. The theta line kernel with its point fixed is of moderate growth, with the same Schwartz seminorm bounds for its derivatives. Its product with the rapid prefix is rapid. The intermediate integral is still bounded by the primitive-support estimate for the ordinary square. Thus the two exterior integrations and the square can be grouped as above, even with the line factor. All real operations in this move are on the kernel before pairing with the exterior forms.

The denominator of the combined move is the product of the $2j$ off-diagonal square scalars and the normalized diagonal denominator C_j/f_j of the line comparison. The tidy order goes up by $4j+2$. At every stage the real maps act on the same tensor for every finite tuple and both input forms. The finite graph and Fourier coefficients are those of the normalized move and are independent of the real translate. Multiplication yields one real tensor $\mathcal{V}_m(a_\infty)$ on the two last paths, one last finite datum of order $4m+2m^2$, and one nowhere-zero denominator c'_ϑ costing at most $2m^2$ squares. The actual number of Hecke products in this construction is $\sum_{j<m}(2j+1) = m^2$. Keep the larger budget in the endpoint estimate. All modulus powers in these relative scalars are powers of $\mathrm{N}\, Q_\vartheta$ times fixed constants, hence they have absolute value one after division by the untwisted value and evaluation away from the conductor.

The last path and the single multiplier on the completion. The final paths on both grids run from $(0,m)$ to (m,m) by increasing the first index. Every edge is a lift from the smaller to the larger definite space. In the coordinates (a,b) the two lists of nodes are $V_{j,m}$ and $Z_{j,m}$. The compact coefficient models for these oriented edges use the negatives of the grid parameters.

For the terminal pairing, compare the doubled V-row parameter with the negative of the doubled Z-row parameter. Let $\mathcal{B}_d(c)$ denote the length-d string $c+d-1, c+d-3, \dots, c-d+1$. At each real place the entries are

	V	reflected Z
above the center	$\mathcal{B}_t(2w_l)$	$\mathcal{B}_{t-1}(2w_l), \quad 2w_i - h - 1$
central stem	$\mathcal{B}_{\kappa-1}(0)$	$\mathcal{B}_{\kappa-1}(-1)$
below the center	$\mathcal{B}_t(-2w_l), \quad h-2w_l$	$\mathcal{B}_{t+1}(-2w_l)$

Here $1 \le l \le m$ and $1 \le i \le m+1$. The two blocks at $2w_l$ sit between $2w_{l+1}-h-1$ and $2w_l-h-1$ and alternate with a V-entry at each end. The central reflected entries lie one below the respective V-entries. The last central reflected entry is above $h-2w_1$. After $h-2w_l$ the two blocks at $-2w_l$ alternate with a reflected entry at each end. These separations use $w_{l+1}-w_l = h$, $h > t+\kappa+2$ and $2w_1 - h > \kappa$. In decreasing order the two full lists therefore strictly alternate, beginning and ending with the reflected list.

An integral compact terminal pairing. Use the two definite spaces $V_{m,m}$ and $Z_{m,m}$ with their fixed block embedding. At a CM-type place let n be the rank on the V-side. Write $\lambda = (\lambda_1, \dots, \lambda_{n+1})$ for the highest weight whose doubled parameter is the negative of the terminal Z row parameter, and $\nu = (\nu_1, \dots, \nu_n)$ for the highest weight with the terminal V row parameter. The two actual algebraic weights are λ and $-w_0\nu = (-\nu_n, \dots, -\nu_1)$. For descending rank d entries the i-th doubled parameter is $2\mu_i + d + 1 - 2i$. The two successive doubled gaps are thus $2(\lambda_i - \nu_i)+1$ and $2(\nu_i - \lambda_{i+1})+1$. Both are positive, so $\lambda_i \ge \nu_i \ge \lambda_{i+1}$.

Use the Weyl lattice of integral flag-minor coordinates and its Gram form from the Fischer calculation. For an arbitrary highest weight μ of rank d, keep the exponent μ_d of the determinant separately and put $\kappa_h = \mu_h - \mu_{h+1}$ for $1 \le h < d$. With X_i the columns of X, there is an integral coordinate map

$$I_\mu : \operatorname{ev}_X \longmapsto \bigotimes_{h=1}^{d-1} (X_1 \wedge \cdots \wedge X_h)^{\otimes \kappa_h}.$$

Here ev_X lists the standard flag-minor products of this shape at X. Each coordinate of the displayed tensor is an integer linear combination of those products by integral flag-minor straightening. Applying Cauchy–Binet to each factor makes the pairing of two image evaluation vectors the product of the same principal minors of $X^* Y$ that defines their Gram pairing. The evaluation vectors span because the standard products are independent polynomials. Thus I_μ is an isometry for the Weyl norm, with the ordinary exterior tensor norm on its target. Its action on the representation includes the single external factor $\det(g)^{\mu_d}$. [6]

In the Z-side tensor let e_L be the extra line basis vector. For each height h shorten $\gamma_h = \lambda_h - \nu_h$ columns by taking the $\bigwedge^{h-1} V \wedge e_L$ component and dropping e_L. Keep the remaining $\kappa'_h = \nu_h - \lambda_{h+1}$ columns by projecting onto $\bigwedge^h V$. These operations are integral contractions. After the tensor factors are reordered and a full-height factor $\bigwedge^n V$ is written in its volume basis, the resulting map D has the tensor target of I_ν. The number of factors of height $h < n$ is $\kappa'_h + \gamma_{h+1} = \nu_h - \nu_{h+1}$. For $\operatorname{diag}(g, t)$ the full-height factors give the determinant exponent $\lambda_{n+1} + \kappa'_n = \nu_n$; the added line has exponent

$$s_\tau = \lambda_{n+1} + \sum_{h=1}^{n} \gamma_h = \sum_{i=1}^{n+1} \lambda_i - \sum_{i=1}^{n} \nu_i.$$

On the V-side use the signed complementary wedge map $C_h : \bigwedge^h V \to \bigwedge^{n-h} V$. Its matrix in the exterior bases is a signed permutation and $C_h \bigwedge^h g = \det(g)(\bigwedge^{n-h} g^{-\mathsf{T}}) C_h$. Applied on every factor of $I_{-w_0 \nu}$, with the factors reordered, this gives C into the tensor target of I_ν. The V-side determinant exponent after this operation is $-\nu_1 + (\nu_1 - \nu_n) = -\nu_n$.

Take the ordinary bilinear tensor contraction of $C I_{-w_0 \nu} v_-$ and $D I_\lambda v_+$. Its matrix is integral, and its character under the simultaneous block action is exactly t^{s_τ}, since the tensor actions $g^{-\mathsf{T}}$ and g cancel. This identity holds in characteristic zero and hence in the coefficient readings used in the transported integral gauges. The contraction is nonzero. On the V-side evaluate at the reversed basis flag, whose complementary factors give the standard highest tensor of I_ν up to sign. On the Z-side evaluate at $X_i = e_i + e_L$ for $1 \le i \le n$ and $X_{n+1} = e_L$. Each retained factor has standard highest wedge coefficient one. Each shortened factor has coefficient one at $e_1 \wedge \cdots \wedge e_{h-1}$, coming from the line entry of its last column. The contraction of these two tensors is therefore a sign. Tensor over the CM-type places, using the specified reindexing of the block. This gives a single integer matrix B_{end} with

$$\|B_{\mathrm{end}}\|_{\mathrm{bil}} \le 1, \qquad \sum_\tau |s_\tau| \le C_{\mathrm{W}} + f\lceil 20b \rceil.$$

The first inequality is in the two unscaled Weyl Gram norms, by the isometries above and the norm-one exterior projections. The second follows from the bounds for the actual Z-side diagram and determinant exponents. These are the norm and character needed at $h \perp u^m$.

Take any real coefficient covectors L_-, L_+ in the terminal compact models, with the translate already in the real datum. At the finite points h and $h \perp u^m$, write their algebraic values as v_-, v_+. The compact integral is

$$\int L_-(\rho_-(k^{-1}) v_-)\, L_+(\rho_+((k \perp 1)^{-1}) v_+)\, dk = \beta(L_-, L_+)\, B_{\mathrm{end}}(v_-, v_+).$$

[6] For integral straightening apply Akin, Buchsbaum and Weyman, *Schur Functors and Schur Complexes*, Adv. Math. 44 (1982), Theorem II.2.16, to a straight shape (transpose the diagram in their row convention). Map each exterior factor of height h to the minor on the first h columns of the variable array. The tableau relations are the flag-minor relations, and the basis images are the standard flag-minor products. Cauchy–Binet supplies the Gram kernel as above.

Indeed the invariant space for the common compact block is one-dimensional. In the Schur character on $n+1$ letters, the boxes with letter $n+1$ form a horizontal strip. Removing it gives exactly one branching term for each interlacing shape on n letters, with the remaining tableau arbitrary. A common determinant shift treats arbitrary integral highest weights and gives the indicated added-line character t^{s_τ}. Pairing with the weight $-w_0\nu$ takes the term of weight ν with multiplicity one at every real place. Transport to the conjugate Weyl models takes the invariant contraction to the same B_{end}. Thus β is a fixed bilinear functional of the real covectors alone. The original finite measure is $I_{0,m,m}^{-1}$ times the measure giving the grid terminal level mass one. Apply this rescaling to the tuple-averaged integrand, whose right invariance and free classes give the finite class sum. This is the factor $I_{0,m,m}^{-1}$ in the last-path cancellation. A fixed nonzero scalar multiple of B_{end} changes $\beta(L_-, L_+)$ by the inverse scale at every twist.

The two definite CM jet comparisons. Write $z = 0, 1$ for the V, Z legs. The rank at step j is $N_{z,j} = r + z + j(t+z) + m(t-z)$ for $0 \le j \le m$. All signs are positive and the adelic frame is the standard one, including at $V_{0,m} = Y_m$. At each embedding in the CM type Φ, let l_j be the negative of the grid parameter at this node. Write $l_{j,i}$ in decreasing order when assigning the highest weight $\lambda_{j,i} = (l_{j,i} - N_{z,j} - 1 + 2i)/2$. The finite matrices and the Weyl bases keep their original indices. The decreasing order specifies the highest flags.

For one edge put $n = N_{z,j}$, $n' = N_{z,j+1}$ and $\Delta = n' - n = t + z$. Its splitting characters are

$$a = \gamma_{z,j,\vartheta}\upsilon^{\lfloor (n+1-z)/2 \rfloor}, \qquad b = \gamma_{z,j,\vartheta}\upsilon^{\lfloor (n'+1-z)/2 \rfloor}, \qquad \gamma_{z,j,\vartheta} = \begin{cases} \mu_{j+1}^0 & z = 0, \\ \mu_{j+1}^0\vartheta^{j+1} & z = 1. \end{cases}$$

Here μ_{j+1}^0 has real weight $-w_{j+1}$ and tame part ω^{j+1}, and υ is restricted to norm-one ideles. Write a_τ, b_τ for the weights of a, b at τ. They are independent of ϑ.

The new positive block lies above the old grid entries. The preceding positive centres are at most $2w_j$, separated from $2w_{j+1}$ by at least $2h$. The small central base entries have modulus at most κ, the remaining base entries $2(h_2 - w_i)$ are negative, and the negative blocks are centred below zero. The inequalities for w_1, h, t, κ put these entries below the first positive block. After negating the parameter the new block is thus the last Δ entries. In the decreasing orders,

$$l_{j+1,i} - l_{j,i} = \Delta + 2(a_\tau - b_\tau) \quad (i \le n), \qquad l_{j+1,n+k} = 2a_\tau + \Delta - n + 1 - 2k \quad (1 \le k \le \Delta).$$

For $z = 0$, m is even and t is odd. The old entries shift by $+1$ for even j and -1 for odd j. The parity $n \equiv j+1 \bmod 2$ gives this shift from the ceilings in a_τ, b_τ. The new negated block has centre $-2w_{j+1} + (n \bmod 2)$. For $z = 1$, the old entries are unchanged, n, n' are even, and the new negated centre is $-2w_{j+1}$. The highest weights therefore satisfy $\lambda_{j+1,i} = \lambda_{j,i} + a_\tau - b_\tau$ for $i \le n$ and $\lambda_{j+1,i} = a_\tau$ for $i > n$.

Set $e_{z,j,\tau} = \lambda_{j,n}$. The diagram $\kappa_{z,\tau}$ has column multiplicity $\lambda_{j,i} - \lambda_{j,i+1}$ at $i < n$, multiplicity $e_{z,j,\tau} - b_\tau$ at n, and zero at $n < i < n'$. The multiplicity at n is nonnegative by the decrease of the target weight. Successive arrows thus use truncations of one diagram, with the determinant exponents $e_{z,j,\tau}$ read separately at their ranks. In this rectangle the polynomial degree is $D_{z,j} = \sum_\tau \sum_{h<n'} h\kappa_{z,\tau}(h)$. It includes the columns of full source height n. The bound on the grid parameters gives row lengths $O(b)$, so $D_{z,j} = O(f n_* b)$, with $f = |\Phi|$ and $n_* = 2n_m$. Since $f = O(x^{11}\log x)$, $n_* = O(x^{11/2})$, $r \asymp x^5$ and $b_1 \ge b$, this is at most $C_{\mathrm{W}} = b_1 r^4$ for large x.

On the transported tuple S use its component $S_e = S((z,j,m),(z,j+1,m))$. Put $R_\vartheta = \mathcal{O}_E/\mathfrak{m}_\vartheta$, where $\mathfrak{m}_\vartheta = \mathfrak{b}Q_\vartheta$. The raw finite datum is $F_e = \mathrm{name}_{\mathfrak{m}_\vartheta}(\mathbf{1}_{S_e})v_0$ on $\mathrm{Hom}(E^n, E^{n'})$. Naming means multiplication by that residue indicator in the specified split Schrödinger coordinates, with the vacuum elsewhere. The analytic arrow uses F_e, and the algebraic arrow uses $F_e^\circ = \upsilon^{nn'}F_e$. Here $\upsilon = c_s \ne 0$ is the rank-one vacuum multiplier of the same datum. The Hom space has rank nn' over E, so this is the product normalization $c_s^{nn'}v_0$ on its vacuum. Use the CM Taylor scale $c = \mathrm{per} = (s_0/u_0)\eta/\Omega_0 \ne 0$. At the actual finite source and target points g, h put

$$J_e(g,h) = \mathcal{J}_c\big(\omega_{a,b,f}(g,h)F_e^\circ\big), \qquad (\mathcal{J}_c(F))_{i',i} = S_{\kappa_z}(\partial)_{i',i}\, T_c(F)(0).$$

Here the polynomial entries of S_{κ_z} in the Weyl bases are integral. The finite action is $\omega_{a,b,f}(g,h) = a(\det h)b(\det g)^{-1}\omega_f(G(g,h))$, where $G(g,h)$ acts by $X \mapsto hXg^{-1}$. The Taylor coefficients and factorial differential pairing are those of the definite Fock calculation.

For the compact conjugate Weyl actions ρ, ρ' set $P_{\ell,\xi}(Z) = \ell(S_{\kappa_z}(Z)\xi)$. The real action substitutes $y_\tau^{\mathsf{T}} Z_\tau x_\tau^{-\mathsf{T}}$ and contributes $\det(y_\tau)^{a_\tau} \det(x_\tau)^{-b_\tau}$. Right substitution by $x_\tau^{-\mathsf{T}}$ in the rectangular Schur matrix contributes $\kappa_{z,\tau}(n)$ powers of $\det(x_\tau)^{-1}$. Thus the transformed coefficient is $P_{\ell\circ\rho'(y^{-1}),\,\rho(x)\xi}$. Homogeneity and the factorial pairing give, for every finite F,

$$c^{D_{z,j}}\,\Theta\big(\omega_{a,b}((x,g),(y,h))(\mathsf{v}(P_{\ell,\xi})\otimes F)\big) = \ell\big(\rho'(y^{-1})\mathcal{J}_c(\omega_{a,b,f}(g,h)F)\rho(x)\xi\big).$$

Adelic invariance of Θ then gives the rational covariance $J_e(\gamma g, \delta h)\rho(\gamma) = \rho'(\delta)J_e(g,h)$ by testing every ℓ, ξ.

Take the two highest evaluation vectors u, u' on the decreasing flags and put $\ell' = \langle u', \cdot\rangle_{\mathrm{W}}$. The Gram minors of a permutation flag are one, so the highest vectors have norm one. The shifted source entries $l_{j,i} - 2b_\tau + n'$ equal $l_{j+1,i} - 2a_\tau + n$ and exceed Δ, since $\lambda_{j+1,i} \geq a_\tau$. The new shifted string is $\Delta + 1 - 2k$. This is the positive definite harmonic match. At τ, the product H_τ of leading minors in the ordered rectangle, with the specified column multiplicities, is joint highest. The opposite-sign Laplacians are vacuous. The lower row and upper column unipotents preserve the minors, and the diagonal characters are the row and column lengths with the splitting weights. These triangular groups have an open orbit, obtained by elimination when the successive leading minors are nonzero and by diagonal rescaling. Hence semi-invariants of the same character on this orbit are proportional. The prescribed joint harmonic is the product of $u_\tau H_\tau$, with $u_\tau \neq 0$ fixed.

If P_τ is the permutation matrix of the ordered source flag, the rectangular Schur action takes the evaluation at P_τ to that at $Z_\tau P_\tau$ after multiplication by $\det(P_\tau)^{\kappa_{z,\tau}(n)}$. The target Gram covector extracts the leading minors in its ordered rows. Thus the actual harmonic is $\eta_{z,j} P_{\ell',u}$, where $\eta_{z,j} = \prod_\tau (u_\tau \det(P_\tau)^{\kappa_{z,\tau}(n)})$.

Write $I_{z,j,m} = [K_{z,j}^0 \,:\, K_{z,j}]$ for the source index. In the real source variable the algebraic form has coefficient $\langle u, \rho(x^{-1})\phi(g)\rangle_{\mathrm{W}}$. Its product with the kernel averages the conjugates of $u\langle u, \cdot\rangle_{\mathrm{W}}$ over compact probability Haar. The average has trace one and commutes with the irreducible action, so it is the identity divided by the source Weyl dimension $d_{z,j}$. After this average, the finite integrand is proportional to $J_e(g,h)\phi(g)$. The old measure gives $K_{z,j}^0$ mass one, while the algebraic class mean gives $K_{z,j}$ mass one. The ratio is $I_{z,j,m}^{-1}$ on any common refinement fixing the integrand at this tuple. On that refinement both measures divide by the same rational stabilizer orders. Thus the analytic lift at (y,h) equals

$$\frac{\beta_{z,j}}{I_{z,j,m}\, v^{nn'}}\, \ell'\big(\rho'(y^{-1})(\mathscr{L}_e\phi)(h)\big), \qquad \beta_{z,j} = \frac{\eta_{z,j}}{d_{z,j}\, c^{D_{z,j}}},$$

where $\mathscr{L}_e$ integrates $J_e(g,h)\phi(g)$ in the source measure with $K_{z,j}$ of mass one. For raw finite datum on both sides the multiplier is $\beta_{z,j}/I_{z,j,m}$. In either case $\beta_{z,j}$ is fixed by the real coefficient models and the CM scale, uniformly for all twists, tuples and class arguments.

Compact averaging of the transported tensor. On the final paths the input on each leg has its prescribed compact type, with its fixed coefficient covector. For a real place τ, the polynomial Fock space of an edge is the rectangular polynomial space with the determinant characters of the action above. Its Cauchy decomposition has one pair of Schur modules for each partition with at most n rows. This follows by applying Schur–Weyl decomposition to the two tensor powers and taking the diagonal symmetric-group invariants. With the compact-action conventions of $P_{\ell,\xi}$ above, the entering weight forces the partition

$$(\lambda_{j,1} - b_\tau, \dots, \lambda_{j,n} - b_\tau).$$

Its last entry is nonnegative by the target weight formula. The outgoing coefficient weight is the same partition padded with Δ zeros and shifted by a_τ. It is exactly λ_{j+1}. Thus the entering type forces the outgoing type at every step, with no choice of intermediate compact type. This component of the Schwartz

oscillator is the span of the vectors $\mathsf{v}(P_{\ell,\xi})$ after taking all real places. Indeed the compact character projector is continuous on Schwartz space and the polynomial Fock vectors are dense there. The projected polynomial component is finite-dimensional.

For the projected edge vector $\mathsf{v}(P_{\ell,\xi})$, let L be the covector of the entering form, with algebraic value $\phi(g)$. Schur averaging in the real source variable integrates

$$\rho(x)\xi\, L \circ \rho(x^{-1}) \quad \text{to} \quad \frac{L(\xi)}{d_{z,j}}\,\mathrm{id}.$$

The resulting automorphic coefficient at (y,h) is

$$\frac{L(\xi)}{I_{z,j,m}\, v^{nn'}\, d_{z,j}\, c^{D_{z,j}}}\; \ell\big(\rho'(y^{-1})(\mathscr{L}_e\phi)(h)\big).$$

The arithmetic operator $\mathscr{L}_e$ is the very same finite jet integral as for the joint harmonic. The outgoing real covector is ℓ. The finite measure and scale factors are precisely those of the normalized harmonic comparison, with the real factor $L(\xi)/\eta_{z,j}$ relative to its $\beta_{z,j}$.

For a pure projected tensor on both legs, write $\ell_{z,j}$ for the successive real covectors and $\xi_{z,j}$ for the entering vectors of its edge factors. After retaining the normalized edge constants, the remaining real coefficient of the period is

$$\beta(\ell_{0,m},\ell_{1,m})\prod_{z=0}^{1}\prod_{j=0}^{m-1}\frac{\ell_{z,j}(\xi_{z,j})}{\eta_{z,j}}.$$

Here $\beta(\ell_{0,m},\ell_{1,m})$ is the compact terminal coefficient of B_{end} just computed. This is a multilinear contraction on the fixed finite-dimensional edge components. Composing it with the tensor product of the compact projections gives one continuous linear functional $\beta(\mathcal{V})$ on the real Schwartz tests of the last paths. The period identity extends to arbitrary tensors by Schwartz density and continuity of the integrals. In particular for $\mathcal{V}_m(a_\infty)$ the functional is the same for each term of the finite tuple sum, for both input forms, and for every ϑ. The finite operator in every term is the composition of the prescribed jet matrices and the integer pairing B_{end}, at its indicated block point.

Cancellation on the last path. The reduction of $K^0_{z,j}$ to $\mathrm{GL}_{N_{z,j}}(R_\vartheta)$ is surjective at the chosen split places. At $V_{0,m}$ the reduction subgroup defining the level is trivial. At every other node it is the kernel of reduction modulo $\mathfrak{b}$. Thus every target index and every Z-index is independent of ϑ. At the start of V, $K_{0,0}$ is contained in the principal delta level of modulus $\mathfrak{M}Q_\vartheta$, since $\mathfrak{b}=\mathfrak{M}^{b_1}$. At the start of Z the level is the core principal level of $\mathfrak{b}$. The two initial forms have these levels.

Write the arrow factor as

$$C_{z,j}\frac{I_{z,j+1,m}}{I_{z,j,m}}, \qquad C_{z,j}=\frac{\beta_{z,j}}{v^{N_{z,j}N_{z,j+1}}I_{z,j+1,m}(1)}.$$

Along a leg these factors multiply to $\prod_j C_{z,j}$ times $I_{z,m,m}/I_{z,0,m}$. The compact terminal pairing at $h\perp F^m_\vartheta$ contributes $I^{-1}_{0,m,m}$ and the contraction $\beta(\mathcal{V}_m(a_\infty))$ accounts for the full real tensor. Combining this with the initial stage factor $I_{0,0,m}$ gives the common multiplier

$$\alpha_{a_\infty}=C_0\,\beta(\mathcal{V}_m(a_\infty))\,\frac{I_{1,m,m}}{I_{1,0,m}}\prod_{z,j}C_{z,j}.$$

This proves (14.1) with the same c'_ϑ from the pointwise path moves.

14.1 Corank-one periods at the specified level

Proposition 14.1. *For every nonzero $y \in \mathcal{H}_0$, there is a core form $u \in \mathcal{H}_c$ and a real translate a_∞ with*

$$\int_{[U(Y_0)]} \overline{y(h)}\, u(\iota(h)a_\infty)\, dh \neq 0.$$

Here ι fixes the added positive line and the integral is the Petersson pairing defining r_{a_∞}.

Proof. Write the input as $y = t(w)$, with t the nonzero row multiplicity map and w its prescribed real coefficient vector. We first construct a reproducing period for a vector w_* in the same smooth row representation. Use the core representation π and the row representation σ at each real place. He's branching theorem for $U(p,q)$ restricted to the subgroup $U(p-1,q)$ fixing a positive line gives one Hilbert copy of σ if the strictly decreasing merge of their signed Harish–Chandra entries meets the following condition. Consecutive entries from the same group have opposite signs, and consecutive entries from different groups have the same sign. This is the discrete part of the restriction and its multiplicity is one. [35, Definition 1 and Theorem 0.1]

The core entries are $-\lambda_i - 1/2$, for $0 \le i \le r$, and the row entries are λ_i, for $0 \le i < r$. The shared signs are positive for $i < \kappa - 1$, then alternate from positive. The last core entry at $i = r$ is positive because m is even. From largest to smallest the merge first contains the reflected core tail, from indices r to $\kappa - 1$, alternating in sign with positive endpoints. The middle contains only positive signs and alternates the row entries $\kappa/2 - 1 - i$ with the reflected core entries, starting on the row side and ending on the core side. The row tail starts positive and alternates. The separation of the tails from the middle follows from $b_1, b_2 \gg r^C$ for each fixed C. This checks the adjacency at both interfaces. Let j be the resulting isometric intertwiner into π, and take the initial unit vector v_0 on the specified ambient lowest compact-type line.

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

$$\int_{U(Y_0)_\infty} \overline{k_\infty(\iota h)} f(xh)\, dh_\infty = c_\infty f(x), \qquad c_\infty \neq 0, \tag{14.3}$$

$$\int_{U(Y_0)_\infty} |k_\infty(g\iota h)|\, dh_\infty \le e^Q |c_\infty| \prod_{v|\infty} |g_v z_0|_v^{-(2n+5)}. \tag{14.4}$$

The first variation along the row Lie algebra has an L^1-bound with constant $e^Q|c_\infty|$. Moving that variation to the coefficient in (14.3) gives the row Lipschitz constant $e^Q\|f\|_\infty$. The required uniform range is $Q = Cdn^4 \log b$ with an absolute C. The vector before translation must also have sufficient matrix-coefficient decay to form a cuspidal Poincaré series. The lower bound on the Hilbert projection and the transversal and first-variation bounds use the same translated vector. The reproducing identity follows from branching and Schur orthogonality.

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

For $a, c \in \mathcal{D}$ let $e_{a,c}(h) = d_\sigma \overline{\langle a, \sigma(h)c\rangle}$, with d_σ the formal degree for the chosen Haar measure. Schur orthogonality gives $e_{a,c} * e_{a',c'} = \langle c, a'\rangle e_{a,c'}$ and $e_{a,c}^* = e_{c,a}$. On any unitary representation τ of H, integrate τ against these L^1 functions. If $\tau(e_{u_0,u_0})$ has nonzero image, choose a unit vector z in its image. The matrix-unit identities make $a \mapsto \tau(e_{a,u_0})z$ an isometry on $\mathcal{D}$. Left translation of e_{a,u_0} by $g \in H$ gives $e_{\sigma(g)a,u_0}$, so

the isometry extends to a copy of σ in τ. Applied to the orthogonal complement of $j\sigma$, uniqueness shows that $\pi(e_{u_0,u_0})$ vanishes there. Factoring each $e_{a,c}$ through this projection gives the same vanishing for all matrix units. On $j\sigma$ they act by the usual rank-one operators.

Taking the coefficient against v gives

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

The integral extends continuously to all a, c by $k_v \in L^1(H)$. With $c = j^*v/\|j^*v\|$, this is the reproducing eigenvector with eigenvalue $\|j^*v\|^2/d_\sigma$. The vector is smooth since j^* is a bounded H-intertwiner and v is smooth. Multiplying the identities over the real places gives (14.3) in the prescribed measures.

The smooth tensor in the row multiplicity space is algebraic. Its map t extends to the Hilbert tensor, as follows. Freeze all but one real factor. Point evaluation is continuous on the remaining smooth Fréchet space by the defining Sobolev condition on multiplicity maps. Every such cuspidal value is a bounded function by rapid decay for finite-level smooth cusp forms of moderate growth, finite under the center. See Miller–Schmid [18, Theorems A and 2.23–2.24]. The uniform boundedness principle applied to all point evaluations therefore gives continuity of this one-variable map into bounded functions and hence into $L^2([U(Y_0)])$, since the quotient has finite volume. The pullback of the L^2 inner product is invariant on the compact-finite module. Such a form is a nonnegative scalar multiple of the model form: on each finite-dimensional compact isotypic space compare the forms by the Riesz operator, then invariance makes the combined operator an endomorphism of the irreducible Harish–Chandra module. Density of compact-finite vectors in the smooth topology extends the identity to the smooth factor. On pure tensors, applying the identity in each factor in turn gives

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

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

First apply the integrated Schur operator on the Hilbert space of the row representation and use the automorphic intertwiner. The pointwise integral against the smooth bounded row form is continuous by the L^1-bound on the kernel. By Fubini against compactly supported continuous functions on the quotient it agrees with the Hilbert-space integral. The resulting equality of continuous functions holds at each x.

The quantitative reproducing kernel. Work at one real place with $G = U(p, q)$ and $H = U(p-1, q)$, and write $N = p + q = n$. Use the adapted frame with signature matrix $J = \mathrm{diag}(\epsilon_1, \dots, \epsilon_N)$, in which the subgroup fixes $z = e_{i_0}$ with positive sign. Let λ and μ be the ambient and row Harish–Chandra parameters in ambient decreasing order. Their differences within each parameter are integers. Opposite-sign gaps in both lists are at least r^8, since the tail spacing is $2r^8$ and the thick and thin rounded weights separating the tails from the middle exceed every fixed power of r. Thus the opposite-sign gaps exceed $1000N^3$ at the chosen large scale. The entries are $O(b)$, and we may use $B = b^2$ and $\log(8NB) \le 6\log b$ in bounds.

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

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

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

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

$$\|d\pi(X_1) \cdots d\pi(X_j) v_0\| \le L^j \prod_{i=1}^{j} \|X_i\|_F \qquad (j \le N + 2). \tag{14.5}$$

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

$$\int_G \|g^{-1} z\|^2 |F(g)|^2 \, dg \le T^{8N} \int_G |F(g)|^2 \, dg = T^{8N} / d_\pi. \tag{14.6}$$

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

For the tensor splitting, use the θ-stable Borel in the ordered frame, with inducing weight $\lambda - \rho_G$. The tensor identity for cohomological induction applies to the dual standard representation, whose Borel filtration has one-dimensional factors of weights $-e_i$. In the Dolbeault description the isomorphism of complexes with the tensor differential is the one in Prudhon's Proposition 3.6. Equation (14) places the filtration by Borel constituents in that complex. In the regular good range, induction from the Borel has cohomology only in degree $\dim(\mathfrak{n} \cap \mathfrak{k})$, with the discrete series of shifted parameter in that degree. [36, Proposition 3.6 and (14)] [37, Theorems 5.35, 5.99 and 8.2, and Chapter XI]

A shift $\lambda - e_i$ is singular only when $\lambda_i - \lambda_{i+1} = 1$. Its adjacent root $\alpha = e_i - e_{i+1}$ is compact by the gap bound. Induce in stages through the compact rank-two Levi on these coordinates. The outer half-sum pairs to zero with $\alpha^\vee$, because the outer roots are stable under the corresponding simple reflection. The root pairing of the shifted parameter is zero, so subtracting the inner half-sum gives degree -1 on the compact flag line $\mathbf{P}^1$. Both cohomology groups of $\mathcal{O}(-1)$ vanish. Thus the singular factor contributes zero. Every other shift is in the original chamber and is regular. The long exact sequences of the filtered complex then give a filtration by the discrete series π_i of these regular parameters, all in the same cohomological degree. Their ordinary infinitesimal characters are distinct because each tuple is strictly decreasing and the tuples differ. The center therefore splits the filtration. The passage from this Harish–Chandra module to smooth vectors uses the Casselman–Wallach equivalence for admissible modules of finite length, which applies to the finite-dimensional tensor product here. [38, Theorems 1.1 and 10.6]

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

$$[x, p_i y] = s_i \langle R_i x, R_i y \rangle, \qquad s_i \in \{1, -1\}. \tag{14.7}$$

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

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

This is the specialization of the antisymmetrizer identity in [39, Proposition 2.7(4), Remarks 2.8(1) and 2.11] under $T(u) \mapsto I + \mathcal{E}/u$. Group the product by the row of its last factor, after fixing the column sequence consisting of all columns except a followed by the final column. This gives the stated sign and Capelli shifts. The eigenvalue of $Q_G(u)$ also follows by evaluation on a highest weight vector: the diagonal contribution in column l is $u - l + 1 + \lambda_l - \rho_{G,l}$. The largest displaced column in each other summand gives a raising operator acting on the highest weight vector after scalar factors. Keep every original λ_k in Q_G. Evaluating the identity on the surviving eigenspaces and taking $u_i = -\lambda_i + (N-1)/2$ gives

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

The denominators have absolute value at least 1. By (14.5), each signed minor applied to v_0 has norm at most $(N-1)!(B+N+L)^{N-1} \leq T^{4N}$. In particular, $|s_i| \|R_i x_l(v_0)\|^2 \leq T^{4N}$ by (14.7) and the (l,l) minor.

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

$$P\big((p_i x_{i_0}(y))_{i_0}\big) = \gamma_i P y, \qquad \gamma_i = \frac{\prod_j (\mu_j - \lambda_i + 1/2)}{\prod_{k \neq i}(\lambda_k - \lambda_i)}, \qquad |\gamma_i| \leq T^N. \tag{14.8}$$

Put $\xi = j u_0$. The $L^1(H)$ matrix unit $e = d_\sigma \overline{\langle u_0, \sigma(\cdot) u_0 \rangle}$ from the preceding argument acts on $\pi|_H$ as the orthogonal projection onto $\mathbb{C}\xi$. It acts on $\pi_i|_H$ as an orthogonal projection Q_i, by the self-adjoint idempotent identities. Choose a smooth vector η with $\langle \eta, \xi \rangle = 1$ and put $a = x_{i_0}(\eta)$. Since H fixes this coordinate, integrating the equality in (14.7) on $\Pi(h)y$ against $e(h)$ gives, for smooth tuples y,

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

With $y = a$, (14.8) yields $s_i \|b_i\|^2 = \gamma_i$. This averaging takes place in the unitary Hilbert spaces of the summands. For each l we obtain

$$(g^{-1})_{l i_0} F(g) = \sum_i s_i \langle b_i, \pi_i(g) R_i x_l(v_0) \rangle.$$

Schur orthogonality bounds the square integral of each summand by T^{5N}/d_{π_i}. Harish–Chandra's formal-degree formula on the same group and for the same Haar measure is proportional to $\prod_{a<b} |\lambda_a - \lambda_b|$, with a common factor for all these parameters. For the compact central factor, Haar integration of the squared coefficient contributes the same factor in every representation, since the central character is unitary. Decreasing one coordinate by 1 changes only $N-1$ factors and each ratio is at most 2, since every difference in a surviving shift is a positive integer and changes by at most one. Thus $d_\pi / d_{\pi_i} \leq 2^{N-1} \leq T^N$. [10, (3.10), (3.12) and Theorem (8.1)] Cauchy–Schwarz over i, followed by summation over l, introduces at most N^3 in the bound T^{6N}/d_π. This proves (14.6).

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

$$\operatorname{vol} D(R) \leq (1 + 2R^2/\delta)^{2r_D^2} \operatorname{vol} B_D(\delta). \tag{14.9}$$

Indeed $g^{-1} = J g^* J$ gives $\|g^{-1}\|_F = \|g\|_F$. A maximal δ/R-separated subset of $D(R)$ gives a cover by its left translates of $B_D(\delta)$ and the Euclidean packing bound in real dimension $2r_D^2$ gives the number of points. With $\delta = 1/4$ and dyadic shells, (14.9) gives

$$\int_D \|g\|_F^{-u} \, dg \leq 2\, 33^{2r_D^2} \operatorname{vol} B_D(1/4), \qquad u \geq 4r_D^2 + 1. \tag{14.10}$$

Let $Y = 2T^{8N}$. By (14.6), the region $S = \{g : \|g^{-1} z\|^2 \leq Y\}$ carries square integral at least $1/(2d_\pi)$. This region is left H-invariant. A vector u of Hermitian norm one is the image of z under a matrix b_u with

$\|b_u\|_F^2 \le N + 2\|u\|^2$. To construct it, take $u = u_+ + u_-$ in the positive and negative subspaces, with lengths a_+, a_- satisfying $a_+^2 - a_-^2 = 1$. First map z to the direction of u_+ by an element of the compact group, then act in the span of the two directions by the block with diagonal a_+ and off-diagonal a_-. The compact map alone suffices when $u_- = 0$.

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

$$c = \|P\pi(g_*)v_0\|^2/d_\sigma \ge \frac{\operatorname{vol} B_H(1/4)}{4 d_\pi \operatorname{vol} G(R_G)}. \tag{14.11}$$

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

For $\delta = (4L)^{-1}$, the logarithm of $g \in B_G(\delta)$ belongs to the real Lie algebra and has Frobenius norm at most 2δ. This follows from the convergent logarithm series and $g^{-1} = Jg^*J$. Integrating the derived orbit along the exponential and using (14.5) shows that $|\langle v_0, \pi(g)v_0\rangle| \ge 1/2$ on this ball. Schur on G gives $d_\pi \operatorname{vol} B_G(\delta) \le 4$. As $R_G \le T^{O(N)}$, (14.9) and (14.11) yield

$$\|g_*\|_F \le T^{O(N)}, \qquad c \ge \operatorname{vol} B_H(1/4) \exp(-C_0 N^3 \log T) \tag{14.12}$$

for an absolute C_0.

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

$$\|P_V \tau(g)|_V\|_{\mathrm{op}} \le \exp(100M^3)\|g\|_F^{-40M^2}. \tag{14.13}$$

Use the polar decomposition $g = k \exp X$, with X Hermitian and supported in the off-sign blocks. The compact factor is isometric and preserves V. The unitary group is connected reductive with finite center on its derived group. Sun's bottom-layer formula applies to the compression of $\exp X$ for the upper triangular θ-stable Borel and the Cartan character of weight $\beta - \rho_G$. The bottom type is the lowest compact type, with highest weight $\beta + \rho_{\mathrm{nc}} - \rho_K$, and occurs with multiplicity one. The weight is dominant for the induced compact ordering because the entries of β are decreasing with successive gaps at least 1. [11, Theorem 6.3 and (25)–(26)]

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

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

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

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

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

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

The sum of these variances with weights $\beta_i - \beta_{i+1} + 1$ has coefficient $\beta_a - \beta_b + b - a$ at $|Z_{ab}|^2$. Every nonzero entry has opposite signs, so the sum is at least $1000M^3\|Z\|_F^2/2$. For $t \ge 1$ the compression norm is therefore at most $\exp(-125M^2 t/2)$. For $t < 1$ it is at most 1 by unitarity. The inequality $\|g\|_F \le \sqrt{M}e^t$ proves (14.13). This argument also supplies the row integrability used above by (14.10), with $M = N - 1$.

Transversal integrals and first variation. Let $v = \pi(g_*)v_0$, $k(g) = \langle v, \pi(g)v\rangle$, $G_0 = \|g_*\|_F$, and $D_0 = 40N^2$. Conjugating (14.13) by g_* gives

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

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

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

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

For X in the Lie algebra of H with $\|X\|_F \le 1$, let $A_h(t) = k(\exp(-tX)h)$. The bound for the second derived word in (14.5) and $\|\operatorname{Ad}(g_*^{-1})X\|_F \le G_0^2$ give $\sup_t |A_h''(t)| \le L^2 G_0^4$. For $|t| \le 1$ we also have $|A_h(t)| \le M_0 e^{D_0} \|h\|_F^{-D_0}$, since the operator norm of $\exp(tX)$ is at most e. Taylor's formula at $t = \|h\|_F^{-D_0/2}$ bounds

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

Both integrals are bounded using (14.10). Dividing by the same c of (14.12), we get the stated transversal and first-variation bounds with a single $\exp(C_1 N^3 \log T)$. For the product over real places they give one $Q \le C d n^3 \log b$, with $d = [F : \mathbb{Q}]$. The weaker exponent $C d n^4 \log b$ in the depth argument also suffices. Using the reproducing identity at $x \exp(tX)$, translation puts the variation on the kernel as $k(\exp(-tX)h) - k(h)$. Its L^1-norm is at most $|t|$ times the first-variation bound, by integration along that curve and left invariance of Haar measure. This gives the stated Lipschitz bound for f and supplies the transversal estimate for the Poincaré seed.

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

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

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

The covolume uses one complex embedding over each real place of F, and quotients use the orthogonal quotient norms. Write $\ell(D)$ for the minimum size of a nonzero sublattice and $h(D)$ for the maximum size of a nonzero torsion-free quotient. Put $\ell = \ell(L)$. For any nonzero integral ideal $\mathfrak{b}$ and $0 \ne y \in \mathfrak{b}L$, write $L \cap Ey = \mathfrak{a}y$. The quotient $L/(L \cap Ey)$ is torsion-free and hence projective over $\mathcal{O}_E$. Splitting this quotient and multiplying the summands by $\mathfrak{b}$ gives $\mathfrak{b}L \cap Ey = \mathfrak{b}\mathfrak{a}y$. Thus $1 \in \mathfrak{b}\mathfrak{a}$ and $\mathrm{N}\mathfrak{a} \le 1/\mathrm{N}\mathfrak{b}$. At each complex place the line map $c \mapsto cy_v$ multiplies area by $\|y_v\|^2$. The covolume of $\mathfrak{a}$ in the chosen complex Lebesgue measure is $2^{-d}|\Delta_E|^{1/2}\mathrm{N}\mathfrak{a}$. Substituting in the rank-one size formula gives $\operatorname{sz}(\mathfrak{a}y)^{2d} = \mathrm{N}\mathfrak{a} \prod_v \|y_v\|^2$, hence

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

In particular, failure of the product lower bound with $s^d\sqrt{\mathrm{N}\mathfrak{b}}$ forces $\ell < s$.

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

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

Indeed, the trace pairing of $\mathcal{L}$ with the reverse homomorphisms is $\operatorname{Tr}_{E/\mathbb{Q}}\operatorname{tr}(BT)$. Its trace dual is $\mathfrak{D}_E^{-1}\operatorname{Hom}_{\mathcal{O}_E}(D_2, D_1)$, as one sees in integral bases at the finite places. It follows that the Euclidean dual elements are $(2B_v^*)_v$ for B in this trace dual. In the real Hilbert–Schmidt product they pair with T by $\sum_v 2\operatorname{Re}\operatorname{tr}(B_v T_v)$, which is the field trace. The forward homomorphism lattice has real dimension $2dab$, and the local reverse matrix trace pairing is perfect over the integers before applying the field trace.

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

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

The dj nonzero singular values then give $\big(\sum_v \|B_v\|_{\mathrm{HS}}^2\big)^{1/2} \geq \sqrt{d}\,l/(u\operatorname{rd}(E))$ by the arithmetic–geometric mean inequality. Banaszczyk's transference bound $\mu(\mathcal{L})\lambda_1(\mathcal{L}^\vee) \leq \dim_{\mathbb{R}}(\mathcal{L}_{\mathbb{R}})/2 = dab$ proves the stated bound. For either rank zero the homomorphism space has radius zero. See Banaszczyk [40, Theorem (2.2)].

The short isotropic step. Suppose $\ell < 1/2$. The Arakelov Harder–Narasimhan filtration of L has saturated steps and strictly increasing successive sizes. The normalized degree divided by rank equals $-\log \mathrm{sz}$. At a step R_i the largest quotient size is its last successive size, and the least size of L/R_i is the next successive size. These assertions follow by comparing the degree polygon with the chords ending or starting at the vertex for R_i. They also give that the first successive size is ℓ. See Bost–Chen [41, §§1.2–1.3].

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

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

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

The last bound follows from $h \leq \rho/2$ and $0 < \rho < 1$.

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

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

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

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

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

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

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

Let Y be a real logarithm in the radical. Apply the first radius bound to $Y/(2g)$. Write $Y = 2gY_0 + A + 2gZ_1$ with Y_0 rational and integral on L, Z_1 central and $\|A\|$ at most $2g$ times the first radius. Apply the second radius bound to $(2gZ_1 + [A, 2gY_0]/2)/g$. This gives a rational center Z_0 integral on L and a real central error Z' with $\|Z'\|$ at most g times the second radius. The two-step Baker–Campbell–Hausdorff identity gives

$$\exp Y = \exp(2gY_0 + gZ_0)\exp A\exp Z', \qquad \|A\| + \|Z'\| \leq \frac{r^2\sqrt{d}}{4}\,g\,\mathrm{rd}(E)\rho.$$

The first exponential is congruent to the identity modulo g on L. Indeed its finite sum is $1 + 2gY_0 + gZ_0 + 2g^2Y_0^2$, and the same argument applies to its inverse. This gives the principal level on L even at primes over 2.

For $n \in N_W(\mathbb{A}_F)$, first move its finite component into $x_f K_Y((g))x_f^{-1}$ by left multiplication by a rational element of N_W. The polynomial logarithm identifies its finite-adelic points with those of its Lie algebra, where the approximation follows from the density of F in its finite adeles. Use the preceding real factorization on the remaining real component. Its rational factor has finite components in the same conjugate principal subgroup. Expressed in the frames transported by x, the outcome is

$$nx = \gamma x\Big(\big(\exp A_v \exp Z'_v\big)_v, k_f\Big), \qquad k_f \in K_Y((g)),$$

and the sum of the two real norms over the places is at most $r^2 d\, g\,\mathrm{rd}(E)\rho/4$ by Cauchy–Schwarz. For any cuspidal row f invariant under this finite level and with real Lipschitz constant L_f in the stated norms, average $f(x) - f(nx)$ over $N_W(F)\backslash N_W(\mathbb{A}_F)$ with probability Haar. Its second term averages to zero by cuspidality. If $\ell < s \leq 1/2$, the variation estimate gives

$$|f(x)| \leq L_f r^2 d\, g\,\mathrm{rd}(E)(2s)^{1/q}.$$

Combined with the line calculation, this is the short-lattice implication at the prescribed bottom level used in the Poincaré argument.

The Poincaré test at the core level. Set $A = r^2 d\, g\,\mathrm{rd}(E)$, $\mathfrak{m} = \mathfrak{M}^{b_1}$, and $s = \frac{1}{2}(2e^Q A)^{-m}$, where $Q \leq Cdn^4 \log b$ is the uniform real loss just proved for the same translated lowest-type vector. The finite condition is

$$e^Q A(2s)^{1/m} \leq \tfrac{1}{2}, \qquad 2e^{Q+(6n+9)d} \leq R_{\mathrm{c}}, \qquad R_{\mathrm{c}} = s^d\sqrt{\mathrm{Nm}}. \tag{14.14}$$

The first inequality is equality. For the second, $\log \mathrm{N}\mathfrak{m} = d b_1 \log g$, whereas

$$\log 2 + Q + (6n+9)d + d\log(1/s) = \log 2 + Q + (6n+9)d + d\big(\log 2 + m\log(2e^{Q}A)\big).$$

This is $O(d^2 m r^4 \log b)$. The datum has $b_1 \asymp b$, $\log g \asymp \log b$, and d, m, r bounded by fixed powers of the size variable. Hence the last display is less than $\frac{1}{2} d b_1 \log g$ at the chosen scale, as required. Use $R = R_{\mathrm{c}}$ below.

The short-lattice estimate of Section 14.1 says that if some nonzero y of the transported rational row lattice $\mathfrak{a}\Lambda_Y(x)$, for a nonzero integral ideal $\mathfrak{a}$, satisfies

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

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

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

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

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

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

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

Here is the justification on the ambient smooth representation. At a real place the translate in the first coefficient vector changes the constant in (14.13). The bound with exponent $40n^2$ then holds when the second vector is either in the lowest type or its chosen translate. For a fixed second K_∞-finite vector an exponent $20n^2$ suffices. Indeed the lowest type generates the irreducible Harish–Chandra module under the derived action. For a derivative of order $j > 0$ of a translated lowest-type coefficient, use the values on a finite stencil in real exponential coordinates, at scale $h \le 1$, exact on polynomials of degree at most $2j-1$. Such a stencil is obtained by interpolation on a $2j$-point grid in each variable. The absolute value of the derivative is bounded by $C_j(h^{-j}\|g\|_F^{-40n^2} + h^j)$. The first term uses the power decay, uniform for each right shift in the stencil. The remainder uses derivatives of total order $2j$, bounded independently of g by unitarity and smoothness of the lowest-type vector. Taking $h = \|g\|_F^{-20n^2/j}$ proves the assertion for every fixed derived word and hence for each fixed compact-finite vector.

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

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

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

of the rational term. The same bound for every fixed derived word allows differentiation term by term on the compact-finite module.

For a compact-finite second vector, unfold the constant term of this series along a proper rational parabolic. The quotient of its unipotent radical by the rational points is compact, and the absolute bound above permits unfolding. Each resulting integral over the adelic unipotent radical vanishes by the real discrete-series constant-term calculation. The connected compact center reduces this calculation to the derived group. Both coefficient vectors in the untranslated coefficient are compact-finite, so the coefficient is center-finite and belongs to the Harish–Chandra Schwartz space. Its constant term vanishes on every proper unipotent radical. This includes the real radical in the unfolded adelic term. Translation of the first vector moves the translate into the argument of the matrix coefficient and leaves that theorem applicable. The level is K_C and the right compact and derived actions on the Poincaré series agree with those on the second vector.

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

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

In $\int_{[U(Y_0)]} \overline{P_x(\imath h)} f(h)\, dh$, unfold using right cosets by $\imath U(Y_0)(F)$. For each coset represented by γ the finite part of the integral over $U(Y_0)(\mathbb{A}_F)$ has support either empty or a left translate of $K' = \imath^{-1}(K_C)$ in $U(Y_0)(\mathbb{A}_{F,f})$. The orthogonal lattice sum gives the row congruence modulo $\mathfrak{m}$. For a unitary automorphism of a self-dual lattice, the inverse is the adjoint. Taking the adjoint of the inverse congruence therefore gives the congruence modulo $\overline{\mathfrak{m}}$ as well. Each prime of E over $g_0 g_1$ belongs to one of the two chosen split halves, and $(g_0 g_1)$ is squarefree in $\mathcal{O}_E$. Thus $\mathfrak{m} \cap \overline{\mathfrak{m}} \subset (g_0 g_1)$ and K' lies in the bottom principal subgroup. For the stabilizer coset, left translation by x, K'-invariance and (14.3) give the term $\mathrm{vol}(K') c_\infty f(x)$.

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

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

By (14.4), the absolute tail is at most

$$\|f\|_\infty e^{Q} |c_\infty| \,\mathrm{vol}(K') \sum_{0 \neq y \in \mathfrak{m}\Lambda_C(\imath x)} \prod_{v \mid \infty} \big((1 + |\imath(x_v)^{-1} y|_v)/3\big)^{-(2n+5)}.$$

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

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

This bound also justifies absolute unfolding. The last two inequalities of (14.14) make the tail at most $\|f\|_\infty |c_\infty| \operatorname{vol}(K')/2$, strictly smaller than the absolute value of the stabilizer term. The resulting period is nonzero.

Let v_* be the tensor of the translated lowest-type vectors in the diagonal seed. Use compact approximate identities of finite Peter–Weyl support for the product of the adapted maximal compact groups. Their averages of v_* converge in the smooth topology. Averaging the corresponding right translates of P_x commutes with the locally normal sum and gives the series of these compact-finite second vectors. By continuity of evaluation and bounded continuity of P_x, the resulting pointwise identity is $P_x = T_{\mathrm{c}}(v_*)$.

For an arbitrary nonzero w in the smooth row tensor, the pairing against $\overline{P_x}$ defines a continuous functional on the row Hilbert tensor. Indeed t is bounded into L^2, and $P_x \circ \iota$ is bounded on the finite-volume row quotient. The tensor of the local irreducible unitary representations is irreducible for the product real group, so the span of the real translates of w is Hilbert-dense. The functional is nonzero on w_*, hence nonzero on some real translate of w. Right invariance of the quotient measure moves the inverse translate to P_x. The resulting ambient vector is still a translate of the fixed lowest-type vector v_0 at each real place. Evaluating T_{c} on the tensor of those vectors gives $u \in \mathcal{H}_{\mathrm{c}}$. The translates used in the seed and in moving the row variable combine to give a_∞, with the same finite level throughout. The nonzero period proves the claimed separation. □

The endpoint scalar at a specified coefficient place. Set $\zeta_\vartheta = c'_\vartheta / c'_1$ for this common path. Write $q_\vartheta = \#(\mathcal{O}_E/Q_\vartheta)$ for the norm of the added conductor. At a coefficient reading ι with $\ell \nmid q_\vartheta$, each square gives a product of at most t squared inverse ramified ratios and a power of the conductor norm. The latter is an ℓ-unit. Suppose the inverse ratios at the charged powers of the trial character have local logarithm at most v_ℓ. There are at most $2m^2$ squares over the two paths, so multiplication in (14.1) gives

$$\log |j_\iota(|\zeta_\vartheta|^{-2})|_\ell \le 4m^2 t v_\ell. \tag{14.15}$$

The residue indicators in the arithmetic jet. For $A \in \mathrm{Mat}_{n',n}(R_\vartheta)$, put $[A,B] = \sum_{i,j} A_{ij}B_{ij}$. Let ψ_ϑ be the primitive additive character at the chosen generator λ_ϑ of $\mathfrak{m}_\vartheta$, and let U_ϑ be the support of the modulus. It meets none of its conjugate places. For a lift $\widetilde{B}$ of B on U_ϑ, let $x_B = \overline{\mathbf{1}_{U_\vartheta} \lambda_\vartheta^{-1} \widetilde{B}}$. Character orthogonality gives the exact identity

$$F_e^\circ = \frac{1}{|R_\vartheta|^{nn'}} \sum_B \psi_\vartheta(-[S_e, B])\, T(x_B) v_{\mathrm{int}}, \qquad v_{\mathrm{int}} = v^{nn'} v_0.$$

Here T is normalized finite Heisenberg translation. The Fourier mode $\psi_\vartheta([X,B])$ acts by $T(x_B)$ in the conjugate coordinates. Its quadratic normalization has phase one since the support is on one side of each conjugate pair. After every complex automorphism the absolute coefficient sum in the displayed expansion is one.

For k, k' in the standard compact groups, the finite Weil action on the indicator datum is

$$\omega_{a,b,f}(k,k') F_e^\circ = \chi_{e,\vartheta}\left(\frac{\det \widetilde{k'}}{\det \widetilde{k}}\right) v^{nn'}\, \mathrm{name}_{\mathfrak{m}_\vartheta}(\mathbf{1}_{\widetilde{k'} S_e \widetilde{k}^{-1}}) v_0.$$

The tildes on k, k' now denote reductions. Indeed the transformed indicator is evaluated at $(\widetilde{k'})^{-1} X \widetilde{k}$, and the splitting characters give the displayed determinant quotient. The components of $\chi_{e,\vartheta}$ are ω_w^{j+1} in V and $(\omega_w \vartheta_w)^{j+1}$ in Z. The indicator index thus changes by the inverse of the action defining tidiness. The product of the determinant factors is cancelled by the inverse character from tidiness. This gives joint invariance in the class variables with the fixed input levels, using $(k, k \perp 1)$ at the terminal block. The arithmetic readings preserve this exact identity on the combined finite sum.

Transport of the integral gauge. Put $J = g^{b_1}$, $q_\vartheta = \mathrm{N} Q_\vartheta$. For an embedding $\iota : \mathbf{C} \to \mathbf{C}_\ell$ with $\ell \nmid q_\vartheta$, use the Weyl gauge at a finite point g_0 given by the maximum absolute value after transporting by g_0^{-1} at the readings $\iota \circ \Phi_\tau \circ \mathrm{conj}$. The transport uses the determinant exponents $e_{z,j,\tau}$. Use the ι-integral span Std_J of $T(x) v_{\mathrm{int}}$ with Jx integral above ℓ. Since $q_\vartheta J \in \mathfrak{m}_\vartheta$, $J x_B$ is integral above ℓ. Also $|R_\vartheta|$ divides $|\mathrm{N}_{E/\mathbf{Q}}(q_\vartheta J)| = (q_\vartheta J)^{2f}$. The Fourier identity shows that $J^{2fnn'} F_e^\circ \in \mathrm{Std}_J$, using the fact that the phases are roots of unity.

Make g, h integral at the places over ℓ by independent rational unitary changes. Weak approximation follows by the Cayley chart on skew-Hermitian matrices and approximation in its coordinates. Rational covariance transports the two gauges to these new points, where integral Weyl matrices and their integral inverses preserve the maximum norm. The finite action there preserves Std_J. Indeed, the move of the vacuum is the finite Weil multiplier divided by its squared complex modulus times the sum indexed by the intersection lattice quotient. Its multiplier and index are ℓ-units, since the standard lattice and its inverse image agree at the primes above ℓ. For the multiplier one uses the product of the local determinant absolute values of the intersection frame. The algebraic character valuation formula gives units for both determinant characters at the integral points. Conjugation of the translations moves x by $X \mapsto h X g^{-1}$, preserving the integral condition on Jx over ℓ.

The normalized CM product calculation applies with $\delta = J$ and the multiplier $c_0 = 1$ of per. The prime-to-ℓ frame in that calculation expresses the normalized data by elliptic products at y_i with $n_0 J y_i \in \Lambda$, where n_0 is prime to ℓ. The coefficients $\iota((s_0/u_0)\overline{\Phi_\tau(w_i)})$ are integral, and the frame-length denominator is a unit. The derivative algebra of the marked elliptic model has integral constant terms and is preserved by $d\,\partial_z$. The isogeny identity then makes the derivatives of the translated factors integral after rescaling by $\overline{n_0 J} d$. The chain rule on the products bounds the factorial jet paired with a homogeneous integral polynomial of degree $D_{z,j}$ by $|\iota(Jd)|_\ell^{-D_{z,j}}$. The integral Schur entries have exactly this degree when nonzero. Combining the entrywise estimate with the ultrametric inequality and the rational changes, the matrix

$$J^{2fnn'} (Jd)^{D_{z,j}} J_e(g, h)$$

preserves the integral gauges between the two class points. The fixed integer $\nu = J^{2fn_*^2} (Jd)^{C_{\mathrm{W}}}$ works for all edges and all ι above $\ell \nmid q_\vartheta$.

Complex readings of the same arithmetic edge. The normalized Taylor series of the finite translates have algebraic coefficients by the marked CM products and finite recovery in the CM Taylor transport. The finite Weil moves are combinations of translations with algebraic multipliers, and the extra determinant characters are algebraic as well. Thus the integral Schur polynomials give algebraic jets on the same data. The scalar $\beta_{z,j}$ is used in the complex automorphic identity before these arithmetic readings.

For σ fixing $\mathbf{Q}(\sqrt{-d})$, use the semilinear Taylor map L_σ on the standard Hom space. It permutes the places by $\Phi_\tau \mapsto \sigma \circ \Phi_\tau$, sends $T(x)$ to $T(b_\sigma x)$, and changes per to t_σ, with $|t_\sigma|^2 = \Omega^{-1}$. Here $b_\sigma \bar{b}_\sigma = \chi_{\mathrm{cyc}}(\sigma)$ for the finite CM Tate multiplier b_σ. Factor $b_\sigma = z_\sigma a_\sigma$ with $z_\sigma \bar{z}_\sigma = 1$ and a_σ an integral unit equal to one on U_ϑ. At $\overline{U_\vartheta}$ take $a_\sigma = \chi_{\mathrm{cyc}}(\sigma)$, and complete the factorization on the other places with integral units of the specified norm. The vector $a_\sigma x_B$ is then the translating vector for $\chi_{\mathrm{cyc}}(\sigma) B$, modulo the integral lattice. Its phase is paired with $\sigma \psi_\vartheta(-[S_e, B]) = \psi_\vartheta(-[S_e, \chi_{\mathrm{cyc}}(\sigma) B])$. The change of index is bijective in the same residue ring. If $L_\sigma v_0 = c_\sigma \omega_f(z_\sigma) v_0$, the result is

$$L_\sigma F_e^\circ = \sigma(v^{nn'}) c_\sigma \omega_f(z_\sigma)\, \mathrm{name}_{\mathfrak{m}_\vartheta}(\mathbf{1}_{S_e}) v_0.$$

Transport leaves the finite move G at the same argument. Its multiplier is $\sigma(\Delta_f(G))/\Delta_f(G)$. Indeed its vacuum pairing is $\Delta_f(G)/|\Delta_f(G)|^2$, and the squared modulus is a rational lattice index. Comparing the vacuum pairings gives the claimed multiplier on the vacuum, and translation covariance gives it on the span. The factor $\sigma(a(\det h)/b(\det g))$ is retained. Both of these factors have modulus one by the principal-ideal norm formulas. For σ not fixing the quadratic field, compose on its output with complex conjugation for the norm estimate, since the Weyl Gram matrices are real.

The moved vacuum bound now gives an entire majorant $Me^{\|Z\|^2/2}$ at the trivial twist, where $M = \exp(C_d(N_1 nn' + 1)^2)$. The absolute coefficient sum of the translated vacua is one. Write $D_0 = fnn'$ and let $r_{\tau,i}$ be the row lengths in the rectangle. For unit input and dual test vectors in the conjugate Weyl norms, the exterior tensor model bounds the coefficient polynomial by the product of the singular values to these row powers. For one singular value s and row length r_0, the r_0-th term of $e^{s^2/2}$ yields $s^{r_0} \le (r_0!)^{1/2} 2^{r_0/2} e^{s^2/4}$. The scale t_σ contributes $\Omega^{-r_0/2}$ for the row. The factorial pairing is Gaussian integration against $e^{-\|Z\|^2}\pi^{-D_0}\,dZ$. The product of the two bounds leaves $e^{-\|Z\|^2/4}$, whose integral against $\pi^{-D_0}\,dZ$ is 4^{D_0}. Using $\log 2 \le 1$ gives the operator bound

$$\|\sigma J_e(g,h)\|_{\mathrm{W}} \le M 4^{D_0} \exp\left(\frac{1}{2}\sum_{\tau,i}\left[\log(r_{\tau,i}!) + (1-\log\Omega) r_{\tau,i}\right]\right).$$

This includes degree zero. The scale in the squared bound is Ω^{-1} per box at the same period normalization, which gives the one-step loss for the endpoint class sum.

Finite class sums and an integer that clears the endpoint period. In the arithmetic realization of $\mathcal{E}_\vartheta$, the last path has two definite chains of m arrows. Write $K_{z,j}$ for the level at step j on side $z \in \{-,+\}$. Put $J = g^{b_1}$ and keep $\mathfrak{m}_\vartheta = \mathfrak{b}Q_\vartheta$ and $q_\vartheta = \#(\mathcal{O}_E/Q_\vartheta)$. The residue ring for every arrow is $\mathcal{O}_E/\mathfrak{m}_\vartheta$, and $q_\vartheta J \in \mathfrak{m}_\vartheta$. Write $n_* = 2n_m$ for the rank bound on both chains and $C_{\mathrm{W}} = b_1 r^4$ for their bound on the total number of polynomial Weyl boxes. Each terminal level is read at its prescribed definite group. The $+$ terminal point above a $-$ point h is the block point $\mathrm{diag}(h, F_\vartheta^m)$, in the given order of its columns, where

$$F_\vartheta = \frac{\overline{\lambda_\vartheta \alpha_{\mathrm{W}}}}{\lambda_\vartheta \alpha_{\mathrm{W}}}.$$

Here λ_ϑ is the finite idele generating the modulus, and α_{W} is the fixed Weil idele. The integer branching matrix pairs the two terminal vectors. Its character on the added line has exponents $s_\tau = e_\tau + \sum_h \gamma_\tau(h)$, where e_τ is the determinant exponent and $\gamma_\tau(h)$ counts the shortened columns of height h. There are at most as many shortened columns as boxes, so $\sum_\tau |s_\tau| \le C_{\mathrm{W}} + f\lceil 20b \rceil$. Its operator norm as a bilinear form for the two Weyl norms is at most $(C_{\mathrm{W}}^{C_{\mathrm{W}}})^3$.

The last-path datum is a linear combination of at most

$$|\mathcal{O}_E/\mathfrak{m}_\vartheta|^{n_*^2(4m+2m^2)}$$

evaluations at edge tuples, with every nonzero coefficient a root of unity. Expand the two kernels on each tuple before summing over classes. An intermediate level change on one side moves through the two adjacent jets. The change in the datum is the product of its edge characters, which cancels the reciprocal factor from the Weil action. At the input vertices use level invariance of the two forms. At the paired terminal vertex move the two arguments by k and $\mathrm{diag}(k,1)$, so the last-coordinate character is one. The resulting sum over edge tuples is separately right invariant in the class variable at each integration. This is the joint class-sum argument applied to two branches with a terminal contraction. Every endpoint level lies in the chosen principal level, so the free-level calculation gives class mass one. The endpoint period is the sum over $2m+1$ class variables of the combined bracket. Inside the bracket each summand is one root of unity times the branching pairing of two products of m jet matrices, applied to the original endpoint values. The count uses both input levels and a single terminal level.

Use the common integer

$$\nu = J^{2fn_*^2}(Jd)^{C_{\mathrm{W}}}.$$

The transported edge estimate shows that ν times each edge matrix preserves the integral Weyl gauges when $\ell \nmid q_\vartheta$.

The factor JN_1 clears both F_ϑ and its inverse at every finite place over $\ell \nmid q_\vartheta$. Indeed, $q_\vartheta J/\lambda_\vartheta$, $N_1\overline{\alpha_{\mathrm{W}}}$, $\overline{\lambda_\vartheta}$ and α_{W}^{-1} are integral finite adeles. Their product is $q_\vartheta JN_1F_\vartheta$. Conjugation gives the same result for F_ϑ^{-1}, and q_ϑ is a unit at the coefficient place. In transported Weyl coordinates the integer branching form is evaluated on the two integral vectors obtained after the m steps on each side. Undoing the transport multiplies the pairing by the character of the extra line at F_ϑ^m. The common integer

$$N_{\mathrm{end}} = (JN_1)^{m\sum_\tau |s_\tau|}\nu^{2m}$$

therefore makes each summand of the combined bracket integral. Apply the ultrametric inequality to conclude that $N_{\mathrm{end}}\mathcal{E}_\vartheta$ is integral when both input forms are integral in their transported Weyl gauges. The integer depends only on the datum. From $b_1 = O(b)$, $\log J = O(b\log b)$, $n_m = O(x^{11/2})$ and $f = O(x^{11}\log x)$, its logarithm is $O_d(x^{30}b^2)$, hence at most $x^{99}b^2$ eventually.

At the trivial twist the residue cardinality is J^f. The definite class count with $q \le n_*$ and $M = J$ bounds the number of classes at each endpoint level by $\exp(x^{36}(1+\log J))$, since the level contains the principal subgroup modulo J.

The one-step bound on the same arithmetic edge gives logarithmic loss at most

$$C_d(N_1n_*^2+1)^2 + fn_*^2\log 4 + \tfrac{1}{2}C_{\mathrm{W}}(1+\log(C_{\mathrm{W}}+1)+|\log\Omega|).$$

Here $\log(h!) \le h\log(h+1)$ for each row of length h, and $|\log\Omega| = O_d(1+\log x)$ by the field and CM normalization. Applying this bound to the $2m$ jets, the norm bound to the branching form and the triangle inequality to the class and tuple sums gives a total log loss $O_d(x^{60}b)$. If the two input Weyl norms are bounded by B, B' on these conjugate values, the endpoint bound is $|\sigma\mathcal{E}_1| \le \exp(x^{200}b^2)BB'$.

Short endpoint forms. At the definite source $\mathscr{F}$ use the classwise evaluation construction with $q \le n_*$ and $D = C_{\mathrm{W}}$. Its frames, including the ones for the determinant exponents of this source weight, have entry bounds $P_{\mathrm{frame}} = 6dN_1(fn_*)^2$. The resulting delta forms span $\mathscr{F}$ and are integral at all coefficient embeddings in the endpoint level.

The squared Weyl norm of an evaluation vector is the product of the leading Gram minors of its matrix arguments, each raised to its column multiplicity. This is the highest-weight Weyl kernel evaluated on the matrix and its conjugate. Hadamard's inequality bounds the squared norm including the determinant factors by $(n_*^2P_{\mathrm{frame}})^{2(C_{\mathrm{W}}+fn_*)}$, uniformly on the conjugates, because C_{W} bounds the total number of boxes over all embeddings. This is at most $\exp(x^{99}b^2)$ for large x.

The integral duality used by the principal tests. At $Y_m = V_{0,m}$ the row lift and the delta test have opposite doubled parameters. At each place their algebraic highest weights are μ and $\mu^\vee = -w_0\mu$, with w_0 reversing the entries. The test multiplicity at height h is the row multiplicity at height $N-h$, where $N = \operatorname{rank} Y_m$. Apply the Jacobi and Fischer calculation to these diagrams, in the original Weyl coordinates. Write A for the integral isometry from the row to the test diagram, obtained by the complementary-minor substitution. Its inverse is integral. Since the two decreasing flags are reverse to each other, its sign can be fixed to send the highest evaluation of the row to that of the test. For a matrix g, the substitution replaces the row action by $g^{-\mathsf{T}}$ times the determinant raised to the number of columns. The determinant exponents complete the two actions to exact contragredients, since $\mu_1 - \mu_N$ is that column count and $\mu_N^\vee = -\mu_1$.

With $G^\vee$ the test Gram matrix, the invariant bilinear pairing therefore has integral matrix $B = G^\vee A$. Its dual norm is the row Weyl norm by the isometry. Let $F^\vee$ be the integral Fischer matrix of the test shape and $h^\vee$ the squared Fischer norm of its highest flag-minor product. The common calculation gives $F^\vee G^\vee = h^\vee I$,

so $B' = A^{-1}F^{\vee}$ satisfies $B'B = h^{\vee}I$. The grid parameter bound also applies to the reflected weight, whose number of polynomial boxes $D^{\vee}$ is $O(f n_* b) \le C_{\mathrm{W}}$. The determinant expansion and factorial norm bound give $1 \le h^{\vee} \le (D^{\vee}!)^3$, including the empty shape. Thus $\log h^{\vee} \le 3C_{\mathrm{W}} \log C_{\mathrm{W}}$.

Use the delta form of the opposite weight at class h in the principal test level $K_{\vartheta}^{\mathrm{pr}}$, with value v. Write $\Delta^{\flat}$ for its finite algebraic form. Its highest coefficient at (x, g) is $\Delta^{\flat}(g)^{\mathsf{T}} B \rho(x) u$, since $Au = u^{\vee}$ and B is invariant. It pairs with the row-lift coefficient by compact Schur orthogonality. As in the definite edge calculation, averaging the Gram-unit highest projector divides by the row weight dimension d_w. The index in (14.1) gives the class-sum measure on the principal test level, and freeness gives mass one for that delta's class. The result is $v^{\mathsf{T}} B (T_{\vartheta} Q_{a_\infty} \phi)(h)/d_w$. The determinant uses the same pointwise row lift at $K_{\vartheta}^{\mathrm{res}}$. Its extraction uses the integral inverse identity even at primes dividing $h^{\vee}$.

One completion of the selected rows. Use the integral pairing and principal delta tests just established.

If $Q_{a_\infty} \neq 0$, then $\alpha_{a_\infty} \neq 0$. Indeed the untwisted chain on $\mathcal{H}_0$ is injective by the row triangle, and the terminal delta tests separate its values. For each such translate choose s_{a_∞} with $s_{a_\infty} \alpha_{a_\infty} d_w / c_1' = h^{\vee} N_{\mathrm{end}}$ and write $f_{a_\infty,\phi} = s_{a_\infty} Q_{a_\infty} \phi$ for a short source form ϕ.

The normalizing factors s_{a_∞} are independent of the twist. In the nonvanishing case the pairing of $\zeta_{\vartheta} T_{\vartheta} f_{a_\infty,\phi}$ with a dual vector is $h^{\vee} N_{\mathrm{end}} \mathcal{E}_{\vartheta}(\Delta_{h,v}, \phi)$. At each coefficient place use a representative h integral above its rational prime, by weak approximation. For each row of B', the corresponding delta test is integral in its Weyl gauge. The endpoint clearing shows that the pairing is $h^{\vee}$ times an integral scalar. By $B'B = h^{\vee}I$, cancellation proves coordinate integrality of $\zeta_{\vartheta} T_{\vartheta} f_{a_\infty,\phi}$, including at primes dividing $h^{\vee}$.

For an automorphism σ and a dual vector v, take the delta test with coefficient $\sigma^{-1}v$. The conjugate delta has Weyl norm at most $\|v\|_{\mathrm{W}}$ on the entire class, by unitarity of the rational action at every embedding. The pairing matrix is integral and fixed by σ. The trivial-twist endpoint bound and norm duality therefore give

$$\big\|\sigma(T_1 f_{a_\infty,\phi})(h)\big\|_{\mathrm{W}}^2 \le \exp\big(2\log h^{\vee} + (2x^{200} + 3x^{99})b^2\big) \le \exp(x_1^2 b^2).$$

The determinant and norm costs are uniform in a_∞. Apply the projection transfer with the identity translate to these rows. Write P_W for the orthogonal projector onto W. The transfer is by the fixed nonzero multiple $|\mathcal{G}|P_W$, and it preserves the local integrality after multiplying by the same ζ_{ϑ}. The resulting rows have untwisted squared class-sum norms at most e^{B_z}, where $B_z = x_1^2 b^2 + x^{99}$.

Let $L_y \subset W$ be the span of the original selected rows. If a vector in W is perpendicular to both L_y and every projected row just constructed, it is perpendicular to $r_{a_\infty} \mathcal{H}_{\mathrm{c}}$ for every a_∞, because the short source forms span and S is onto. The corank-one period then forces the vector to be zero. Consequently the projected short images span W/L_y.

14.2 Allowances at the coefficient primes

The allowances in this subsection are determined by the selected square and its fixed datum. For each rational prime ℓ use $v_\ell(\ell) = 1$ and set

$$c_\ell = v_\ell(G) + 2^{-\ell}, \qquad G = \prod_{v \text{ chosen}} |\mathrm{GL}_{rt}(\mathbf{F}_{q_v})|, \qquad v_\ell = 9 c_\ell \log \ell.$$

The last symbol, used without an argument, is the allowance for a ramified ratio. Write $B_0 = \mathcal{B}/(m\kappa)$. We will obtain allowances $\beta_{j,\ell} \ge 0$ for the two local determinant tests and bound their sum at each stage by B_0. The sum of the ramified allowances is at most x^{999}.

Special values for the critical parity. Let $L/\mathbf{Q}$ be abelian CM of degree n_L, with maximal real subfield F_L, and let N_L be divisible by the conductors of its primitive Dirichlet characters. For $k \ge 2$ put

$$H_{\ell,k}(L) = H^2_{\text{ét}}(\mathcal{O}_L[1/\ell], \mathbf{Z}_\ell(k)), \qquad w_k(L) = |H^0(L, \mathbf{Q}/\mathbf{Z}(k))|.$$

Use the corresponding notation for F_L. The factor $\Lambda_k(L)$ is the product of $L(1-k,\xi)$ over primitive associates of the characters of L with $\xi(-1)=(-1)^k$. By Artin factorization it equals $\zeta_{F_L}(1-k)$ for even k, and $L(1-k,\chi_{L/F_L})$ for odd k.

For even k, Huber–Kings at odd primes and Kolster at 2 give

$$\log_\ell |H_{\ell,k}(F_L)| = v_\ell(\Lambda_k(L)) + v_\ell(w_k(F_L)).$$

The real regulator in this parity has rank zero. The torsion in degree one is the ℓ-part of $w_k(F_L)$ by the coefficient sequence. The Euler factors removed above ℓ are ℓ-units at $1-k$. [42, Theorems 1.3.1 and 1.4.1] [43, Appendix A, Theorem A.1 and Corollary A.3]

For odd $k \geq 3$, Kolster's higher relative class number formula for abelian CM fields gives

$$\log_\ell \frac{|H_{\ell,k}(L)|}{|H_{\ell,k}(F_L)|} = v_\ell(\Lambda_k(L)) + v_\ell(w_k(L)) + v_\ell(Q_k) - ([F_L:\mathbf{Q}]+1)v_\ell(2), \qquad Q_k \in \{1,2\}.$$

These are the ordinary cohomology groups of the respective rings of ℓ-integers. [44][7]

Let $P_k = \prod_{p \leq k+1} p$, and let f_ξ be the primitive conductor of ξ. The special-value formula is

$$L(1-k,\xi) = -\frac{f_\xi^{k-1}}{k} \sum_{a=1}^{f_\xi} \xi(a) B_k(a/f_\xi).$$

It follows by decomposing the Dirichlet series into Hurwitz zeta functions [46, 25.11.14 and 25.11.36]. Expanding the Bernoulli polynomial and using von Staudt–Clausen on B_j for $2 \leq j \leq k$ shows that $kP_k f_\xi L(1-k,\xi)$ is an algebraic integer. The same expansion and $|B_j| \leq j!$ give $|B_k(u)| \leq 3k!$ on $[0,1]$, so $|L(1-k,\xi)| \leq 3(kf_\xi)^k$. The parity list is Galois-stable and the factors are nonzero by the functional equation. Thus

$$M_k(L) = \Lambda_k(L) \prod_{\xi(-1)=(-1)^k} kP_k f_\xi$$

is a nonzero rational integer, and $P_k \leq (2k)^k$ gives

$$\log |M_k(L)| \leq n_L\big(\log 3 + \log(kN_L) + k\log(2k) + k\log(kN_L)\big).$$

In particular $v_\ell(\Lambda_k(L)) \leq v_\ell(M_k(L))$.

If $\ell^a \mid w_k(L)$, the cyclotomic image in $(\mathbf{Z}/\ell^a\mathbf{Z})^\times$ has index at most n_L and exponent dividing k. Its order is at most k at odd ℓ and at most $2k$ at 2, by the cyclic structure of the odd prime-power units and the two cyclic factors at 2. It follows that $\ell^a \leq 2n_L k$ at odd ℓ and $2^a \leq 4n_L k$ at 2. Each odd prime in $w_k(L)$ contributes at most $\log(2n_L k)$ to its logarithm, and there are at most $n_L k$ of them. In particular $\log w_k(L) \ll n_L k \log(2n_L k)$.

A common integer for the Tate-twisted cocycles. Put $N_{\mathrm{cyc}} = dN_1$, $F = E^+$, and $n_E = [E:\mathbb{Q}]$. The weighted field $E = K(\mu_{N_1})$ lies in $\mathbb{Q}(\mu_{N_{\mathrm{cyc}}})$. Since $12 \mid N_1$, we have $E = F(i) = F(\zeta_3)$. The discriminant of E/F divides both (4) and (3) by the orders generated by i and by ζ_3. Hence E/F is unramified at every finite place. The restriction and corestriction maps between its rings of ℓ-integers satisfy $\mathrm{cor}\,\mathrm{res} = 2$ and $\mathrm{res}\,\mathrm{cor} = 1 + c$, where c is complex conjugation on E.

The cocycles to be bounded have coefficients in $(O/\ell^v O)(-e)$, with $0 \leq e < t$. They are unramified outside ℓ. At $e = 0$ they kill inertia above ℓ, so their length for $O = \mathbb{Z}_\ell$ is at most $v_\ell(h_E)$ by ordinary class field theory. At $e > 0$ they vanish on the kernels of the cyclotomic character in the decomposition groups above ℓ, and they obey $z(cgc^{-1}) = -z(g)$.

[7] See also Kolster–Sands, *Annihilation of motivic cohomology groups in cyclic 2-extensions*, §5 and Theorem 6.1, pp. 181–183. Apply their absolute abelian formula with the cyclic extension L/F_L of degree two. Their factor $2^{[F_L:\mathbb{Q}]+1}/Q_k(L)$ and their relative order $h_k(L)/h_k(F_L)$ give the displayed valuation after moving the power of two to the other side.

Take $e > 0$ and $O = \mathbb{Z}_\ell$. Embed the coefficients into $B = (\mathbb{Q}_\ell/\mathbb{Z}_\ell)(-e)$. At $v \mid \ell$ the local cocycle factors through the cyclotomic image of the decomposition group, which is open because $E_v/\mathbb{Q}_\ell$ is finite. At odd ℓ this image has a procyclic subgroup inside $1 + \ell\mathbb{Z}_\ell$ of index prime to ℓ. At 2 it is itself procyclic because $i \in E$. In either case a topological generator of the procyclic subgroup acts on B with its action minus one a nonzero scalar. That scalar acts surjectively on the divisible module B, so H^1 of the subgroup vanishes. The prime-to-ℓ quotient at odd ℓ has vanishing H^1 on the invariant module as well. Thus the global cohomology class belongs to the strict localization kernel at the places over ℓ.

Conjugation acts on coefficients by $(-1)^e$, so the cohomology class of an odd cocycle has eigenvalue $\epsilon = (-1)^{e+1}$. Every coboundary as a function on G_E is even, since $\chi(cgc^{-1}) = \chi(g)$. An odd cocycle that is a coboundary in B is therefore killed by two. The map to the strict localization kernel is injective at odd ℓ, and has kernel of order at most two at $\ell = 2$. For the latter assertion the image of the coboundary map is B/B^{G_E}, again a divisible group whose two-torsion has order two, since B^{G_E} is finite cyclic.

Set $k = e + 1$ and

$$H_{L,\ell,k} = H^2_{\text{ét}}(\mathcal{O}_L[1/\ell], \mathbb{Z}_\ell(k)) \qquad (L = E, F).$$

At finite coefficients the cohomology over the ring of ℓ-integers agrees with Galois cohomology unramified outside ℓ. Poitou–Tate duality pairs the finite localization kernels in degrees one and two. Filtered direct limits commute with these kernels since the set of places above ℓ is finite. The finite cohomology groups in degree one give the Mittag–Leffler condition for identifying inverse-limit cohomology in degree two. Consequently the strict localization kernel is the equivariant Pontryagin dual of the wild kernel $A \subset H_{E,\ell,k}$. There are only complex infinite places. The ℓ-adic groups $H_{L,\ell,k}$ are finite for $k \geq 2$.[47, VI, §8] [43, Propositions 6.12–6.13] The signed parts have the same orders on the two sides, including at 2. Indeed

$$|(A^\vee)^\epsilon| = |A/(1 - \epsilon c)A| = |A^\epsilon| \leq |H^\epsilon_{E,\ell,k}|.$$

Here superscripts are the kernels of $c - \epsilon$. This yields the required bound on the raw finite cocycles from the signed part of $H_{E,\ell,k}$, with the one possible factor of two just proved. [48, Theorem I.4.10(a) and Proposition II.2.9]

Use the critical special values with $L = E$, $F_L = F$ and $N_L = N_{\text{cyc}}$, so that $\Lambda_k = \Lambda_k(E)$. This compares the ordinary groups $H_{L,\ell,k}$ calculated with the same coefficients as the wild kernel.

At odd ℓ, $H^+_{E,\ell,k} \simeq H_{F,\ell,k}$ by transfer, and $H_{E,\ell,k}$ is the sum of its two sign parts. Thus the even-k formula controls the invariant part and the odd-k ratio controls the anti-invariant part, both by $v_\ell(\Lambda_k) + v_\ell(w_k(E))$, using $w_k(F) \mid w_k(E)$.

For the coefficient prime 2, abbreviate $H_L = H_{L,2,k}$. If k is even then $2H^+_E \subseteq \operatorname{res} H_F$, so $|H^+_E| \leq |H_F||H_E[2]|$. If k is odd, then $k \geq 3$ and $(1 + c)H_E \supseteq 2 \operatorname{res} H_F$. The kernel of restriction is killed by 2, and the kernel of doubling on its image is contained in $H_E[2]$. Consequently

$$|H^-_E| = \frac{|H_E|}{|(1 + c)H_E|} \leq \frac{|H_E|}{|H_F|}|H_F[2]||H_E[2]|.$$

The coefficient sequence embeds $H_L/2H_L$ into $H^2_{\text{ét}}(\mathcal{O}_L[1/2], \mu_2)$. By the Kummer sequence the latter has order at most

$$h_L \, 2^{\#\{v \mid 2 \text{ in } L\} + r_1(L)}.$$

Indeed, the class group of $\mathcal{O}_L[1/2]$ is a quotient of the ordinary class group, and the 2-torsion of the Brauer group has local invariants only at the real and dyadic places. Since H_L is finite, $|H_L[2]| = |H_L/2H_L|$. A common upper bound for the dyadic error, including the possible factor of two on passing from raw cocycles to classes, is therefore

$$C_2 = 1 + 2n_E + \lceil \log_2 h_E \rceil + \lceil \log_2 h_F \rceil.$$

The bound at 2 now uses the critical part of the wild kernel and applies to the weighted field.

For E the integer $M_k = M_k(E)$ from the special value calculation satisfies $v_\ell(\Lambda_k) \le v_\ell(M_k)$ and

$$\log |M_k| \ll n_E k \log(2kN_{\mathrm{cyc}}), \qquad \log w_k(E) \ll n_E k \log(2n_E k).$$

It follows that the single positive integer

$$X = h_E\, 2^{C_2} \prod_{k=2}^{t} |M_k| w_k(E)$$

bounds every finite cocycle length by $v_\ell(X)$ for $O = \mathbb{Z}_\ell$. The factor 2^{C_2} is shared because the same bound is used separately for each e. For a finite coefficient extension $O/\mathbb{Z}_\ell$, the defining cocycle conditions commute with finite free base change. In a $\mathbb{Z}_\ell$-basis of O, each of the coefficient components obeys the identical conditions at the same modulus ℓ^ν. Ordinary O-length is therefore the ramification index times $\mathbb{Z}_\ell$-length. With $v_\ell(\ell) = 1$, the normalized bound remains $v_\ell(X)$.

Finally E, F have root discriminants at most N_{cyc} by the cyclotomic different and the discriminant tower formula. Minkowski's bound and the count of ideals of each possible norm give $\log h_L \ll n_E^2 \log(2N_{\mathrm{cyc}})$ for $L = E, F$. Summing the displayed estimates gives

$$\log X \ll n_E^2 \log(2N_{\mathrm{cyc}}) + n_E t^2 \log(2n_E t N_{\mathrm{cyc}}).$$

The local length and its sum over primes. Put $h_2 = r^8$, take b_1 to be the datum's integral rounding of a/x, and set $g = g_0 g_1$ and

$$A = (2b_1)! \left(\prod_{\max(0,\, b_1 - 30h_2^3) < u \le b_1 + 30h_2^3} u \right)^{h_2^3}.$$

The normalized local budget and the finite length used for the Selmer tests are

$$b_\ell = h_2^{27}\big(c_\ell + v_\ell(X)\big) + 4f\, 10^6 \frac{m}{\kappa}\big(b_1 v_\ell(g) + v_\ell(A)\big), \qquad M_\ell = \lfloor r b_\ell \rfloor + 1.$$

The construction can use $\beta_{j,\ell} = b_\ell \log \ell$ at every stage. The main block of the local construction will use coefficient modulus ℓ^{M_ℓ}. Its first-order Hecke test will use a larger depth ν, with the Tate cocycle bound at that depth provided by $v_\ell(X)$. The local lattice estimates apply at both depths.

For a positive integer C, the sum of $v_\ell(C)\log \ell$ over a finite set of primes is at most $\log C$. Also $\sum_\ell 2^{-\ell} \log \ell \le \sum_{n \ge 1} n 2^{-n} = 2$. Each chosen g_i has f places of norm g_i, and $|\mathrm{GL}_{rt}(\mathbf{F}_{g_i})| \le g_i^{(rt)^2}$. Hence $\log G \le f(rt)^2 \log g$. On the scales of D, this is $O(x^{31+\varepsilon} \log x)$. The sum of the allowances $v_\ell = 9c_\ell \log \ell$ is at most $9(\log G + 2) \le x^{999}$ on every finite set of primes for large x.

The bounds on X, f, N_{cyc}, r and t imply that $h_2^{27}(\log G + 2 + \log X)$ is polynomial in x. It is $o(B_0)$ because $a \ge xb$ and $b = \exp(x^\varepsilon)$. For the other term, the interval product in A has at most $60h_2^3$ terms. Since $b_1 = O(b)$, h_2 is polynomial in x and $\log g = O(\log b)$, we have $\log A = O(b \log b)$. Consequently its contribution together with that of g is

$$O\Big(f \frac{m}{\kappa} b \log b\Big) = O\left(\frac{B_0}{(\log b)^2}\right),$$

using $a \ge xb$ and $m^2 \le 4x$. This proves the claimed finite-prime budget for large x, uniformly over the coefficient isomorphisms and over all stages.

15 The untwisted determinant and the choice of a prime

Keep the n rows $y = (y_1, \dots, y_n)$ of (8.40), their translates and their chosen short preimages. Write L_y for their span in the R-dimensional space W, and put $u = R - n$. The endpoint construction gives a spanning family in W/L_y. Choose u of its rows $z = (z_1, \dots, z_u)$ mapping to a basis of the quotient. This fixes $w = (y, z)$ before any finite coefficient place is chosen. Each untwisted completion row has squared norm at most $\exp B_z$ at every complex embedding, where

$$B_z = x_1^2 b^2 + x^{99}.$$

If $u = 0$, the tuple z is empty.

For a trial character ϑ write

$$H_\vartheta = \det \mathrm{Gram}(\Theta_\vartheta w).$$

At the trivial character the chain is injective on W, so $H_1 > 0$. At each field automorphism σ of $\mathbf{C}$, successive orthogonal projection of the added rows bounds the squared volume of the coefficientwise frame $(\Theta_1 w)^\sigma$ by $\exp(V + uB_z)$. Here $V = nh_{\mathrm{sh}}$ uses the short-row height from (7.9) on the original n rows. Let c denote complex conjugation and put $\sigma^\# = c\sigma c$. The class weights and the Weyl pairing matrix are rational. Thus σH_1 is the mixed Gram determinant for the two frames at $\sigma^\#$ and σ. Applying Cauchy–Schwarz on the exterior power to those frames gives $\log|\sigma H_1| \le V + uB_z$.

Let $d_j(1)$ be the untwisted compressed determinant at stage j. There are $\kappa/2$ sharp pairs p at each of the m stages. Using their fixed multipliers $A_{jp} = a_{jp} V_{jp} \overline{V_{jp}}$, set

$$D_0 = \prod_{j=1}^{m} d_j(1), \qquad K_\Gamma = \prod_{j,p} A_{jp}^2, \qquad Y = H_1 D_0 K_\Gamma^R.$$

All factors are nonzero. The archimedean estimates for the compressed determinant and the sharp pair give $\log|\sigma d_j(1)| \le 2RrB_0$, $\log|\sigma V_{jp}| \le B_0$, and $\log|\sigma a_{jp}| \le B_0 - S_0$. The bound for V_{jp} also applies to its complex conjugate, since σc runs through the same automorphisms. Counting the factors yields

$$\log|\sigma Y| \le V + uB_z + 2RrmB_0 + Rm\kappa(3B_0 - S_0). \tag{15.1}$$

In particular, Y is algebraic. Indeed, if it were transcendental, extend $\{Y\}$ to a transcendence basis of $\mathbf{C}$ over $\overline{\mathbf{Q}}$. Replacing Y by hY and fixing the other elements gives an automorphism of the corresponding rational function field for each integer $h \ge 1$. It extends to an automorphism of its algebraic closure $\mathbf{C}$, contradicting the uniform upper bound in (15.1).

The fixed row preimages determine F_ℓ in (7.10). The numerical allowances $\beta_{j,\ell}, v_\ell, P_\ell$ were also fixed on the untwisted datum. For every prime ℓ define the nonnegative real number

$$\begin{aligned} A_\ell = F_\ell &+ R(r+\kappa)\sum_{j=1}^{m} \beta_{j,\ell} + 2Rm\kappa P_\ell \log \ell \\ &+ \left[4um^2 t + R(2m\kappa + 10m^2 + m^2 t)\right] v_\ell. \end{aligned} \tag{15.2}$$

We now choose the only place where a fitting estimate can be needed.

Lemma 15.1. *Let $Y \in \overline{\mathbf{Q}}^\times$, let f_Y be its monic minimal polynomial over $\mathbf{Q}$, and write $e = \deg f_Y$. Put $q_\ell = e^{-1}\log|f_Y(0)|_\ell$, with $|\ell|_\ell = \ell^{-1}$. If $|f_Y(0)| \ge 1$, then $\sum_\ell q_\ell \le 0$. If $|f_Y(0)| < 1$, one can choose a prime ℓ_* and a field isomorphism $\iota : \mathbf{C} \longrightarrow \overline{\mathbf{Q}}_{\ell_*}$ from Y and the allowances A_ℓ alone such that*

$$\log|\iota(Y)|_{\ell_*} \le A_{\ell_*} \quad \Longrightarrow \quad \sum_\ell q_\ell \le \sum_{\ell:\, q_\ell > 0} A_\ell.$$

Proof. Write $|f_Y(0)| = a_Y/b_Y$ in lowest terms. By prime factorization, $\sum_\ell q_\ell = (\log b_Y - \log a_Y)/e$, and $q_\ell > 0$ precisely at primes dividing b_Y. If $a_Y < b_Y$, choose a prime $\ell_* \mid b_Y$ maximizing $q_\ell - A_\ell$ on this finite nonempty set, breaking ties by the least prime. Among the ℓ_*-adic conjugates of Y, choose one of largest absolute value. Its logarithm is at least q_{ℓ_*}, the mean of their logarithms. This choice defines an embedding of $\mathbf{Q}(Y)$. Extend it to algebraic closures of $\mathbf{Q}$ and then along a bijection of transcendence bases. Such a bijection exists because the two algebraically closed fields $\mathbf{C}$ and $\overline{\mathbf{Q}}_{\ell_*}$ have characteristic zero and cardinality $2^{\aleph_0}$, while $\overline{\mathbf{Q}}$ is countable. The resulting ι sends Y to the chosen root. The displayed hypothesis gives $q_{\ell_*} - A_{\ell_*} \le 0$, hence the same inequality at every prime dividing b_Y. The remaining q_ℓ are nonpositive. □

In the second case fix this ℓ_* and write $j_\iota = \iota$, $j_\iota^{\#} = j_\iota c$. The next construction works for any such prescribed reading. Its first-order estimate is taken at j_ι. The two integral tests use j_ι and $j_\iota^{\#}$ on the same rows and with the same auxiliary character.

16 Fitting at the selected coefficient place

Fix the square datum, the untwisted basis and its completion, the comparison scalars, and the stage and sharp test constructions. Their lifts at a trial character use these fixed data. If the norm test requires a local fit, fix its chosen coefficient place $\iota = (\ell, j_\iota)$, where $j_\iota : \mathbf{C} \simeq \overline{\mathbf{Q}}_\ell$ and $|\ell|_\ell = \ell^{-1}$. The paired reading is $j_\iota^{\#} = j_\iota \circ \mathrm{conj}$. We construct one finite-order norm-one character θ of added conductor prime to ℓ. Each $\vartheta = \theta^j$, $1 \le j \le m$, will have the same added support. Its inverse avatar will be trivial on the full decomposition groups over ℓ. Use the allowances fixed in the coefficient budget. In this section v denotes the Hecke character v_E of E, and $v_{U(1)}$ denotes its norm-one restriction. The required outputs are the following.

1. Write $\vartheta = \theta^j$ and $Z_{j,\vartheta}$ for the image of the single lift on good forms on $U(W_j)$. Let $g_{j,\vartheta}$ be the comparison map from the tables, and k_j its positive integral hook factor. On $Z_{j,\vartheta}$ the specific compressed Gram operator
$$K_{j,\vartheta} = k_j \, \mathrm{pr}_{Z_{j,\vartheta}} \, g_{j,\vartheta} g_{j,\vartheta}^* \big|_{Z_{j,\vartheta}}$$
satisfies
$$|j_\iota(\det K_{j,\vartheta})|_\ell \le \exp(Rr\beta_{j,\ell}).$$
Here $\dim Z_{j,\vartheta} = R$, and the unitary source rank is r.

2. The $\kappa/2$ auxiliary lines at stage j use the lower members $p = p_j(z)$ of the unsigned character pairs of lift rank $2z$, where $7m/2 < z \le 7m/2 + \kappa/2$. They use one fixed character as the lift rank varies. Let $Z_{p,\vartheta}$ be the line of the lift of the constant function from $U(1)$, and let h' be the number of $U(1)$-classes at the chosen level. Use the two integral test forms Ξ, Ξ' from the jet test, at j_ι and $j_\iota \circ \mathrm{conj}$ with the same p, ϑ and the same fixed height witness for the test. Their projected pairing satisfies
$$\left| j_\iota\Big(h' \langle \Xi, \mathrm{pr}_{Z_{p,\vartheta}} \Xi' \rangle \Big) \right|_\ell \le \exp(\beta_{j,\ell}).$$
Both forms are integral at every embedding above ℓ. Their Hecke tests use the tuple idempotent labelled $X - 1$ at the chosen places and the kernel of the descent map from $U(2z)$ to $U(1)$.

3. For the ramified pairs and signs used in this comparison, write $\rho_{p,s}(\vartheta)$ for the twisted over untwisted Hecke-value ratio. The local fit gives $\log |j_\iota(\rho_{p,s}(\vartheta)^{-1})|_\ell \le v_\ell$. The list consists of the auxiliary flat pair q_2 with sign $+1$ in the jet test, the lower members of the unsigned stage-j pairs up to $7m/2$ with sign $+1$, the lower members of the two signed families up to m and $m/2$ with signs $+1$ and -1 respectively, and the lower members with sign $+1$ in the scalar square calculations for the chain and the completion. Each of the latter calculations uses one product over the lower points of its specified ramified pair, at its own specified stage.

The table map and the jet test are the ones in the single-lift proof. The square calculations use the grid-pair and triangle-pair characters on the fixed datum.

The two determinant tests come from one diagonal Hecke estimate at j_ι. Both integral inputs are constructed at the indicated tests, using the same θ in the two readings.

16.1 The Galois rows and their local conditions

Let η_f be the finite part of an algebraic Hecke character of E with infinity exponents m_σ, so on $E^\times$ it is $\prod_\sigma \sigma(x)^{-m_\sigma}$. For the selected coefficient reading j_ι, put

$$\eta_\ell(\mathrm{Art}_E^{\mathrm{geom}}\, x) = j_\iota(\eta_f(x)) \prod_{w\mid\ell} \prod_{\sigma\mapsto w} (j_\iota\sigma)(x_w)^{m_\sigma}.$$

Here $\sigma \mapsto w$ means that $j_\iota\sigma$ induces w, and its value on x_w uses its continuous local extension. The right side kills principal ideles. Choose an open subgroup of the finite integral units on which η_f is trivial. The quotient by this subgroup and by $E^\times$ is a finite ray class group. A power of each representative of the quotient is therefore principal up to that subgroup, which shows that the values of η_f are algebraic and lie in a single number field. The corrected character has values in a finite ℓ-adic field. Its values on integral ideles are units, since the smooth character has finite image on the integral units and each local algebraic factor there is a unit. The finite ray quotient shows that all values of the corrected character are units. It is continuous and takes values in a profinite group, so global reciprocity gives the continuous character displayed above.

Write $\eta_w^\iota(x) = j_\iota(\eta_f(\mathrm{idele}_w\, x))$. For $w \nmid \ell$ the avatar on the local Weil group is η_w^ι via geometric reciprocity. Its inertia has finite image, and is trivial if η_f is unramified there. For $w \mid \ell$ write $L = E_w$, $f_w = [k_w : \mathbf{F}_\ell]$, and $\lambda(x) = \prod_{\sigma\mapsto w}(j_\iota\sigma)(x)^{m_\sigma}$. Then $\eta_\ell \circ \mathrm{Art}_w^{\mathrm{geom}} = \eta_w^\iota\lambda$. Remove the finite character δ of G_L corresponding to the restriction of η_w^ι on units, extended to have value 1 at a chosen uniformizer ϖ. Put $\xi = \eta_\ell/\delta$. With Conrad's arithmetic Artin map the algebraic character of ξ on units is λ^{-1}. Enlarge the coefficient field to contain all the embeddings of L. Proposition B.4(i) of Conrad gives crystallinity of ξ. His convention counts the cyclotomic character as having weight 1. Proposition B.3 applied to λ^{-1} gives m_σ in that convention, hence the jump at $j_\iota\sigma$ is $-m_\sigma$ here. Proposition B.4(ii), for the covariant crystalline module, gives the scalar

$$\xi(\mathrm{Art}_w^{\mathrm{ar}}\, \varpi)^{-1}\lambda(\varpi)^{-1} = \eta_w^\iota(\varpi)$$

for φ^{f_w}. On the potentially crystalline periods the operator for a Weil element s of geometric degree b is the linear operator induced by $\varphi^{bf_w}s$. For ξ it has scalar $\eta_w^\iota(\varpi)^b$ because s fixes the crystalline periods. For δ, periods may be formed over the finite extension killing δ. They have jump zero and on their span the operator is $\delta(s)$, since $\varphi^{bf_w}s$ acts trivially on the unramified period scalars. The extension cut out by δ is totally ramified, since δ has value 1 on the symbol of ϖ, so these period scalars have unramified subfield L_0. Tensoring gives $\eta_w^\iota(x)$ when s has symbol $\mathrm{Art}_w^{\mathrm{geom}}(x)$, with $\mathrm{ord}_w\, x = b$. This gives every degree-zero inertia operator and the degree-one geometric Frobenius operator at every inducing embedding. The crystalline module is the full rank-one module when the smooth local character is unramified. With a ramified smooth part, the monodromy is zero after the finite twist, and the jumps are the same. This calculation applies over each completion of E, including those of residue characteristic two.[8]

For a norm-one character b with weights c_τ, write $b^{\mathrm{bc}}(x) = b(x/\bar{x})$. Its algebraic character

$$\eta_{b;k,h} = b^{\mathrm{bc}}\upsilon^{-k}|\cdot|^{h-k/2}$$

has exponents $h + c_\tau - k$ and $h - c_\tau$ at τ and $\bar{\tau}$. Indeed $b^{\mathrm{bc}}\upsilon^{-k}$ has angular exponent $2c_\tau - k$ and the norm factor gives radial exponent $2h - k$. Its smooth finite part is unramified at a split pair when b is trivial on the integral norm-one units of that pair, since υ and the norm are unramified. For an upper auxiliary pair

[8] Conrad, *Lifting global representations with local properties*, Appendix B, Propositions B.3 and B.4.

of ranks r', n, the character $a = v_{U(1)}^{r_-}$ has weight r_-, the number of negative source signs. Its normalizing avatar is that of $\eta_{a^{-1};-r',0}$. Its jumps at τ, $\bar{\tau}$ are $r_- - r'$ and $-r_-$, and its smooth parameter is unramified at every finite place. The lower block uses $\eta_{(\mu_j^0)^{-1};0,(r'-t')/2}$, with $t' = n - r'$. It subtracts $(r' - t')/2 + \epsilon_\sigma w_j$ from the source jump when μ_j^0 has weight $-w_j$. The same formula with the actual point weight applies for the rank-one source. Its finite character on a chosen-place inertia operator is the inverse of the tame part of μ_j^0 through reciprocity. A trial stage character ϑ, specialized to θ^j in the fit, uses the weight-zero avatar of $(\vartheta^{-1})^{\mathrm{bc}}$, with conductor prime to ℓ. Its inertial action at an added place is the character labelling the inverse-twist determinant slot on the marked block. Conjugation inverts its norm-one parameter, so it has the same value at the inverse-conjugate inertia element. The cyclotomic character is the avatar of $|\cdot|$, with geometric Frobenius value q_w^{-1}, jump -1, and trivial inertia on the degree-zero operators at $w \mid \ell$. Its inverse powers give the normalized string and its first power gives the jump $-u - 1$ for the Tate dual $V^\vee(1)$.

For the conjugate reading $j_\iota^\# = j_\iota \circ \mathrm{conj}$, write the ℓ-adic embedding $j_\iota\sigma$ as $j_\iota^\#\bar{\sigma}$. The same formula uses $m_{\bar{\sigma}}$ and the smooth value read by $j_\iota^\#$. Both readings use the same stage character and the same conjugate pair of local tests.

The catalogue of normalized rows. For stage j the source of the auxiliary lift is W_j, of rank r, and its target is the standard definite space of rank $r + t$, $t = 5r$. The common Hecke system is taken on the *good* principal-level forms on W_m. The rigid-label transfer compares the traces of the common Hecke algebra on the good spaces of W_j and W_m. For any of the finitely many systems in their sum, interpolate a polynomial in common Hecke operators that is 1 on this system and 0 on the others. Its trace gives equality of the dimensions of these parts. The nonzero source of each genuine packet in the stage-j lift therefore has its Hecke system on W_m. For each nonzero isotypic part fix a nonzero component of one good tuple idempotent. The group-algebra action at the chosen places commutes with every Hecke operator used here. Thus taking that component preserves the Hecke system. Its source attachment is taken on the standard model of W_m. The label compatibility needed is annihilation by the polynomials for the actual squarefree product type on the indicated tame inertia, and the Hodge jumps of the actual Harish–Chandra parameter at ℓ. For the rank-one sharp pair one starts with the constant on $U(1)$, the trivial source representation and source label $X - 1$. Its target has rank $2z$, with $t' = 2z - 1$. Its normalized row is a character obtained from the source by the same rank normalization. The trivial source character itself would have jump zero and could not supply the row conditions.

The Galois representation of a good lower packet. Take the nonzero good component on W_m selected for the lower Hecke system. In its cuspidal decomposition take an irreducible constituent π on which this component projects nontrivially. The projection preserves its archimedean type, its Hecke values and its finite label idempotents. At a split place over each of the two distinct primes g_i, the primitive Green type forces the local component to be depth-zero supercuspidal, as in the good trace transfer. The definite archimedean factors are cohomological with the regular parameters of the prescribed Weyl modules. The odd rank $r = 3^{a+1}$ and the two distinct supercuspidal split places give the hypotheses of Harris and Labesse, Theorem 2.2.2. For a cuspidal unitary representation of odd rank with these split factors the theorem gives cuspidal conjugate-self-dual base change. It matches the given factor at a quasi-split split place and gives unramified base change at a quasi-split spherical place with unramified quadratic extension. It matches the infinitesimal characters at infinity. Thus the regular cohomological factors on the definite group give a cohomological Π of $\mathrm{GL}_r(\mathbf{A}_E)$. At a split place v of F, with $w, \overline{w}$ above it, use the projection onto the w-matrix as the isomorphism of $U(W_m)(F_v)$ with $\mathrm{GL}_r(E_w)$. The theorem gives $\Pi_w = \pi_v$ in that frame and $\Pi_{\overline{w}} = \pi_v^\vee$ in the conjugate frame. In particular this local identification holds at a specified split place over the coefficient prime.[9]

[9] Harris–Labesse, Theorem 2.2.2, p. 675.

For each ℓ and the chosen coefficient embedding j_ι, Chenevier and Harris, Theorem 3.2.3, give the continuous semisimple rank-r representation R of G_E attached to the cohomological conjugate-self-dual cuspidal Π. Clause (a′) compares its Frobenius-semisimplified Weil–Deligne parameter at $w \nmid \ell$ to $\mathrm{rec}(\Pi_w \otimes |\det|_w^{(1-r)/2})$. The comparison has the same Weil action and monodromy dominated by the automorphic monodromy, so zero automorphic monodromy implies zero Galois monodromy. Clause (b) gives de Rhamness and the labelled jumps specified by the cohomological weight. At a place over ℓ with a maximal-compact fixed vector, clause (c) gives crystallinity and computes the polynomial of the first linear crystalline Frobenius from that same Tate-normalized local parameter. At a split spherical w, let $\alpha_1, \dots, \alpha_r$ be the normalized Satake parameters of π_v, with the uniformizer going to geometric Frobenius. With the compact subgroup of volume one, the Hecke eigenvalue of the j-th minuscule double coset is $q_w^{j(r-j)/2} e_j(\alpha_1, \dots, \alpha_r)$. In the monic Hecke polynomial the extra factor is $q_w^{j(j-1)/2}$, so its roots are $q_w^{(r-1)/2}\alpha_i$. This is exactly the shift $\Pi_w \otimes |\det|_w^{(1-r)/2}$ used in Theorem 3.2.3. At $w \nmid \ell$, part (a′) gives trivial inertia, zero monodromy and this characteristic polynomial for geometric Frobenius. At $w \mid \ell$, part (c) gives crystallinity and the same characteristic polynomial for $\varphi^{[k_w:\mathbf{F}_\ell]}$. For a geometric Frobenius lift s, this is the first linear power in the crystalline Weil operator. Indeed, on crystalline periods the diagonal action of s is trivial. The power of φ is then fixed by the action of s^{-1} as $x \mapsto x^{q_w}$ on the residue field. Over the coefficient field the Frobenius permutes the components of the unramified base field, so the polynomial holds in every component. Cayley–Hamilton gives annihilation on the comparison module for every inducing embedding and every such lift. Degree-one places of E are split over F, so the exact equalities at almost all of them give the required Frobenius limits. Chebotarev and Brauer–Nesbitt identify R with any semisimple attachment to this lower Hecke system.[10]

Fix a real place of F and write $\lambda_1 \geq \cdots \geq \lambda_r$ for the highest weight of the Weyl module used in the definite algebraic forms. The scalar realization sends an algebraic form f to $L(\tau_\lambda(\overline{g_\infty}^{-1})f(g_f))$. Right translation acts on the functionals L through $(\tau_\lambda \circ c)^\vee$. Since the group is compact unitary, this has highest weight λ at the chosen embedding Φ. The numbers $h_i = \lambda_i + (r+1)/2 - i$ are strictly decreasing integers. The archimedean base-change parameter on $\mathbf{C}^\times$ has characters $(z/\overline{z})^{h_i}$. Thus its two infinitesimal characters are the weights h_i and $-h_i$, with the latter in reverse order in the dominant chamber. The coefficient system for cohomology with Π_∞ has the opposite duality, with highest weights $\mu_{\Phi,i} = -\lambda_{r+1-i}$ and $\mu_{\overline{\Phi},i} = \lambda_i$. By the formula $\mu_{\sigma,i} + r - i$ in Chenevier and Harris, the labelled jump multisets are $\{i - 1 - \lambda_i\}_i$ and $\{\lambda_i + r - i\}_i$, respectively. These are the jumps of the lower form before the row twist. The convention gives jump -1 to the cyclotomic character.[11]

The lower principal level imposes its reductions only at split places. At a nonsplit finite place it is the full stabilizer of the standard self-dual lattice. The unit conjugate difference in $\mathcal{O}_E$ shows that E/F is unramified at every finite place. After the unramified quadratic extension the lattice's integral unitary model becomes GL_r over the integers. Reductive descent gives a hyperspecial subgroup of the quasi-split p-adic group. Its spherical unramified base change is included in Harris and Labesse, Theorem 2.2.2. Thus parts (a′) and (c) also give unramifiedness or crystallinity there, as needed for the lower row off the split level. The constant lower packet in rank one has the trivial representation with polynomial $X - 1$ and jump zero before applying its own row twist.

The tame labels of these lower packets. At a chosen split w, write the squarefree label as $c_w = \prod_a Q_a$ and put $d_a = \deg Q_a$. Its finite Green representation is the parabolic induction of the cuspidal Green representations τ_a attached to the factors. The label idempotent has nonzero image on the selected π_v. By finite Frobenius reciprocity, π_v contains the inflation of $\boxtimes_a \tau_a$ to the parahoric for the blocks of size d_a. Moy and Prasad, Theorem 6.11(2), place π_v in the normalized induction of depth-zero supercuspidals π_a with these finite cuspidal types. Distinct factors give distinct types in each fixed dimension, so the inducing

[10] Chenevier–Harris, *Construction of automorphic Galois representations, II*, geometric reciprocity convention and Theorem 3.2.3.

[11] Chenevier–Harris, *Construction of automorphic Galois representations, II*, equation (1.6), p. 61 of the published version.

supercuspidals are pairwise inequivalent up to unramified twist. Their singleton segments are unlinked, and Zelevinsky, Theorem 4.2, gives irreducibility of the induction. The local parameter of $\Pi_w = \pi_v$ is consequently the direct sum of the parameters of the π_a, with zero monodromy.[12]

Write $\omega_D = \exp(2\pi i/D)$, and use the faithful character T of the algebraic residue field from the label. For a root x of Q_a, set $D_a = q_w^{d_a} - 1$ and take b_x with $T(x) = \omega_{D_a}^{b_x}$. The prescribed Green representation τ_a uses the regular character $\theta_x(y) = T(y)^{b_x}$ of $\mathbf{F}_{q_w^{d_a}}^{\times}$. In particular its character on the regular elliptic element y in this torus is $(-1)^{d_a-1} \sum_{j=0}^{d_a-1} \theta_x(y^{q_w^j})$. The depth-zero local Langlands formula for π_a induces the Weil character from the unramified extension of degree d_a, with this restriction on units through geometric reciprocity. Both the rectifier and the rank normalization are unramified.

With $D = q_w^{r!} - 1$, an inertia element s of tame exponent b has $T(\mathrm{res}(s(z)/z)) = \omega_D^b$ for a D-th root z of a uniformizer. Put $z_a = z^{D/D_a}$. Under the geometric Artin map, a unit u_a representing the image of s in abelianized inertia of the unramified extension has $\mathrm{res}(u_a) = \mathrm{res}(s(z_a)/z_a)$. Thus $\theta_x(\mathrm{res}(u_a)) = \omega_{D_a}^{b_x b} = T(x)^b$. The Frobenius conjugates give the roots $j_\iota(T(x^{q_w^j})^b)$ for $0 \le j < d_a$. Taking all factors gives the polynomial for c_w at s.[13]

Away from ℓ, part (a′) of Chenevier–Harris transfers this finite inertia action to R, since monodromy is zero and Frobenius semisimplification preserves inertia. The polynomial is the characteristic polynomial of $R(s)$, so it also supplies the determinant of $R(s)$. At a chosen place over ℓ, apply Theorem 1.1 of Barnet-Lamb, Gee, Geraghty and Taylor, *Local-global compatibility for $l = p$, II*. For a regular algebraic, essentially conjugate-self-dual cuspidal representation of Shin-regular weight it identifies the full Frobenius-semisimplified Weil–Deligne parameter at every place above the coefficient prime with the Tate-normalized Langlands parameter. The regular algebraic Π has odd rank, so its weight is Shin-regular. Both representations have the same semisimple attachment by their almost-everywhere spherical polynomials. The theorem identifies the Frobenius-semisimplified Weil–Deligne representation at the place over ℓ, in particular its zero monodromy and its action of inertia. Potential semistability with zero monodromy gives potential crystallinity. The theorem holds for any coefficient prime. In this datum the level primes are odd, so at $\ell = 2$ the chosen-place types are already covered by the prime-to-ℓ clause. Every inertia Weil pair here has Frobenius exponent zero, so its linear operator is the inertia action on the potentially crystalline periods. Cayley–Hamilton gives the required polynomial on this comparison module as well. The conjugate places follow from the dual local matching and the same geometric reciprocity. The polarization $R^c \simeq R^\vee \otimes \varepsilon_\ell^{1-r}$ follows by comparing the normalized spherical polynomials at conjugate places and using Chebotarev. On a geometric Frobenius, ε_ℓ^{1-r} gives q_w^{r-1}, which takes the inverse of each normalized root to the normalization of the inverse Satake parameter.[14]

These results supply the unshifted labelled lower column. The stage twist and the later row bounds apply to this column as before. The same lower representation gives the genuine block-plus-string representation wherever the theta polynomial has that factorization. The general attachment and its local clauses also apply to the other packets in the upper catalogue.

Galois attachment on the definite catalogue. Keep the field $E = K(\mu_{N_1})$ at the selected square and write $F = E^+$. For each rank $a \ge 1$ in the catalogue, use the group of the standard positive definite Hermitian form over E/F. In particular, this applies to the upper groups of ranks $6r$ and $2z$, at the base, twisted and deeper levels of the coefficient test. The weight includes the determinant weight of the form at that level. The field contains μ_{5040}, and $E = F\mathbf{Q}(i)$ and $[F : \mathbf{Q}] \ge 2$. The latter inequality follows already from $\mathbf{Q}(\zeta_5 + \zeta_5^{-1}) \subset F$. At a finite prime of residue characteristic p, choose a root of unity in E of order 3 or 5 prime to p. Its reduction still has that order, so complex conjugation cannot act trivially on the residue

[12] Green, §§7–8, Moy–Prasad, Theorem 6.11, p. 118, and Zelevinsky, Theorem 4.2, pp. 184–185.

[13] Bushnell–Henniart, *The essentially tame local Langlands correspondence, I*, §2.2 and Corollary 3.3, Part III, Proposition 21, and Adrian–Roe, *Rectifiers and the local Langlands correspondence*, Proposition 3.4.

[14] Barnet-Lamb–Gee–Geraghty–Taylor, Theorem 1.1, with Shin-regularity on the preceding page.

field if it stabilizes the prime. This proves that E/F is unramified at all finite places. The standard self-dual lattice gives a reductive model by descent from GL_a over the unramified quadratic ring, so its stabilizer is hyperspecial at a nonsplit finite place. The base, twisted and deeper levels of the upper catalogue keep this stabilizer at every nonsplit place. The chosen g_i, the new level places Q and the Taylor–Wiles primes are all split. Thus when a nonsplit completion is read at ℓ, its actual upper packet has the spherical vector used in the crystalline comparison.

At a split $v = w\overline{w}$, use the projection i_w onto the w-matrix. In that frame the standard compact is $\mathrm{GL}_a(\mathcal{O}_{E_w})$.

Let f be a nonzero algebraic eigenform in the catalogue, with monic polynomials P_w at almost all places of prime norm. Fix the coefficient isomorphism $j_\iota : \mathbf{C} \longrightarrow \overline{\mathbf{Q}}_\ell$ extending the reading under consideration. Realize f in scalar automorphic forms by taking matrix coefficients against the dual of its weight module. The definite rational group is anisotropic and the quotient is compact. Project onto an irreducible automorphic constituent π on which the realization is nonzero. Projection preserves the weight and all the Hecke equations on f. For a specified local test one may first take the local translate, mean or finite idempotent occurring in that test, and project the resulting nonzero vector. The uniqueness below gives the same attachment for all such choices with this Hecke system.

Labesse's Corollary 5.3 applies since the archimedean factors of π are discrete series on compact groups and satisfy the sign condition $(*)$. The local matches are equivalently those of Thorne's Theorem 4.2, for the very group of the standard form in his equation (4.1). Write the base change as $\Pi = \boxplus_i \Pi_i$. Each Π_i is discrete and individually conjugate self-dual. The base change has $\Pi_w = \pi_v \circ i_w^{-1}$ at every split place, and at a hyperspecial spherical place its local representation is the spherical base change. At infinity its infinitesimal character is the base change of that of π. [15]

By the residual spectrum theorem of Mœglin and Waldspurger, Π_i is the Speh representation attached to a unitary cuspidal σ_i of rank a_i and a positive integer d_i. Its exponents are $\nu_{ik} = (d_i - 1)/2 - k$ for $0 \le k < d_i$. Here $\sum_i a_i d_i = a$. The Speh datum is unique, so $\sigma_i^c \simeq \sigma_i^\vee$ for each i. In fact conjugation and duality replace σ_i by $(\sigma_i^c)^\vee$ and reverse the same string of exponents. Write

$$t_{ik} = \nu_{ik} + \frac{a_i - a}{2}, \qquad \pi_{ik} = \sigma_i \otimes |\det|^{t_{ik}}.$$

At a complex embedding τ, let h range over the infinitesimal entries of $\sigma_i|\det|^{\nu_{ik}}$. Together over i, k these are the entries of Π. They are distinct and lie in $(a-1)/2 + \mathbf{Z}$ at each embedding. Thus the entries $h + (a_i - a)/2$ of π_{ik} are distinct and lie in $(a_i - 1)/2 + \mathbf{Z}$. They give a dominant integral weight after subtracting the Harish–Chandra shift in rank a_i. This is regular algebraicity. The representation π_{ik} is cuspidal and polarized by

$$|\cdot|_F^{2t_{ik}} \delta_{E/F}^{a_i \bmod 2}.$$

Indeed the absolute value on E is the pullback by norm and the quadratic character is trivial on norms and has sign -1 at each real place. In particular the parity correction here is a twist of the automorphic cuspidal representation by a power of the norm. The Galois attachment is taken only after the twist has made the constituent regular algebraic.

Theorem 2.1.1 of Barnet-Lamb, Gee, Geraghty and Taylor attaches a continuous semisimple representation to each regular algebraic polarized cuspidal π_{ik}. Its local parameter off ℓ has the rank-a_i Tate shift $(1 - a_i)/2$. At every place over ℓ it is de Rham and its jumps for coefficient weight $(b_1, \dots, b_{a_i})$ are $b_j + a_i - j$. At a place over ℓ where the automorphic component is unramified it is crystalline. We form the continuous semisimple sum

$$R = \bigoplus_{i,k} r_{\ell, j_\iota^{-1}}(\pi_{ik}).$$

[15] Labesse, *Changement de base CM et séries discrètes*, Corollary 5.3, and Thorne, *A p-adic approach to the existence of level-raising congruences*, §4, Theorem 4.2. Labesse's local spherical transfer uses the hyperspecial test and the split transfer uses the indicated matrix projection.

The total local shift for this summand is $t_{ik} + (1 - a_i)/2 = v_{ik} + (1 - a)/2$. At a hyperspecial spherical place the base change is unramified. Each cuspidal seed there is unramified as well. Indeed the spherical representation is a subquotient of the normalized induction from the translated seeds. Taking maximal compact invariants is exact, and by the Iwasawa decomposition the induced fixed space is the tensor of the fixed spaces of the seeds. At $w \nmid \ell$ the local comparison gives unramifiedness. At $w \mid \ell$ the spherical clause of the same theorem gives crystallinity, also when E_w is ramified over $\mathbf{Q}_\ell$. Both statements hold for every ℓ, including 2. [16]

To calculate the labelled jumps at τ, a rank-a_i block with infinitesimal entry H has jump $(a_i - 1)/2 - H$. This is the formula in Theorem 2.1.1(3) with the coefficient weight dual to the infinitesimal weight. Substituting $H = h + (a_i - a)/2$ gives $(a - 1)/2 - h$. At the chosen CM embedding Φ over a real place, the scalar realization acts on the functionals by $(\tau_\lambda \circ c)^\vee$, of highest weight λ in that frame. Hence the entries there are $h_j = \lambda_j + (a+1)/2 - j$, and at the conjugate embedding they are their negatives in reverse order. The resulting lists are

$$\{j - 1 - \lambda_j \,:\, 1 \le j \le a\}, \qquad \{\lambda_j + a - j \,:\, 1 \le j \le a\}.$$

This is the labelled jump formula on the actual weight of f. The cyclotomic character in this convention has jump -1.

At a split spherical w, with normalized Satake roots α_j of $\pi_v \circ i_w^{-1}$, the rank-a monic Hecke polynomial has roots $q_w^{(a-1)/2}\alpha_j$, counted with multiplicity. Indeed the j-th minuscule Hecke eigenvalue is $q_w^{j(a-j)/2} e_j(\alpha)$, and its coefficient in degree $a - j$ carries the factor $(-1)^j q_w^{j(j-1)/2}$. The total shift of the direct sum gives exactly this polynomial at all spherical split places away from ℓ. A place of E of prime norm is split over F, since E/F is unramified. Thus outside a finite set of places w of prime norm it equals $j_\iota P_w$ on every geometric Frobenius lift. This implies coefficientwise approximation on the Frobenius filter. More explicitly, on the triples (w, A, s) with A above w and s geometric, in every nonempty open subset of G_E the equality holds on every such triple outside the same finite set of w. As s approaches any $g \in G_E$ on this filter, the coefficients tend to those of $\det(X - R(g))$. If a second continuous semisimple representation has the corresponding approximation property for each error $\eta > 0$, use a nonempty open subregion furnished by its Frobenius filter, and then a place of prime norm outside both finite exceptional sets by Chebotarev. Its trace differs from that of R by at most η on a dense set, and thus everywhere by continuity. Letting η tend to zero and applying Brauer–Nesbitt proves uniqueness.

Thorne's Corollary 4.3 attaches a continuous semisimple representation to this same π, at any coefficient prime. It is unramified at every place $w \nmid \ell$ over a spherical factor, and it is de Rham at every $w \mid \ell$. Its almost-everywhere spherical polynomials agree with those of R. Thus uniqueness identifies the two attachments. For every split w, including a specified $w \mid \ell$, clause (3) gives

$$\mathrm{WD}(R|_{G_{E_w}})^{F\text{-ss}} \;\simeq\; j_\iota \mathrm{rec}_{E_w}\left((\pi_v \circ i_w^{-1}) \otimes |\det|_w^{(1-a)/2}\right).$$

Here local reciprocity sends a uniformizer to geometric Frobenius. In particular, if f has Hecke polynomial Q at a specified spherical split w, projection preserves Q and the characteristic polynomial on the local parameter is exactly $j_\iota Q$. Frobenius semisimplification keeps the inertia action and the monodromy operator. For the characteristic polynomial at any Weil element, write it as inertia times a power of Frobenius. Some power of Frobenius centralizes the finite inertia image, so the unipotent part of Frobenius centralizes that image as well. Removing the commuting unipotent factor leaves the characteristic polynomial unchanged, with every multiplicity. [17]

For a spherical $w \nmid \ell$, this is the polynomial of $R(s)$ for any geometric Frobenius lift s. For a spherical $w \mid \ell$, put $f_w = [k_w : \mathbf{F}_\ell]$. The comparison uses the first linear power φ^{f_w} on the crystalline periods.

[16] Mœglin–Waldspurger, *Le spectre résiduel de* GL(n), main theorem and §1.11, Anastassiades–Thorne, proof of Corollary 3.4, and Barnet-Lamb–Gee–Geraghty–Taylor, *Potential automorphy and change of weight*, Theorem 2.1.1(2)–(4), p. 538.

[17] Thorne, *A p-adic approach to the existence of level-raising congruences*, equations (4.1)–(4.2), Theorem 4.2 and Corollary 4.3, pp. 19–21.

For a lift s of geometric Frobenius the paired operator there is $\varphi^{f_w}(s \otimes R(s))$, where $s \otimes R(s)$ fixes the crystalline periods. The components of the unramified base field are permuted by φ, so the characteristic polynomial on each component is $j_t Q$. Cayley–Hamilton gives the required polynomial on the comparison space for every prolongation A and every compatible inducing embedding. More generally, at a split place where the local parameter has $N = 0$, the same argument reads its characteristic polynomials on all Weil elements. In residue characteristic ℓ the representation is then potentially crystalline, and the operator of a geometric-degree-k element on the potentially crystalline periods is $\varphi^{f_w k}(s \otimes R(s))$. At degree zero it is the finite inertia action. Away from ℓ, $N = 0$ gives finite inertia on the original representation and the Weil action agrees with that restriction. These conclusions use the local parameter of the same split component that retains the test.

At almost all split places the conjugate Satake roots are their inverses, since the conjugate matrix projection is inverse transpose. The cyclotomic value at a geometric Frobenius is q_w^{-1}. The two normalized sets of roots therefore give $R^c \simeq R^\vee \otimes \varepsilon_\ell^{1-a}$ by uniqueness. For the determinant let ω_π be the central character of π on the norm-one ideles. The algebraic character

$$x \longmapsto \omega_\pi((x/x^c)I_a)\,|x|_E^{-\binom{a}{2}}$$

has determinant avatar $\det R$ by comparing it with the product of the normalized roots at almost all split places. At a split w its value on a unit equals the scalar by which that central matrix acts. On a spherical uniformizer it equals $(-1)^a Q(0)$. The infinity exponents are $\sum_j \lambda_j - \binom{a}{2}$ at the chosen embedding and $-\sum_j \lambda_j - \binom{a}{2}$ at the conjugate embedding. The separate rank-one avatar calculation applies to this character and to the row twists, with the same j_t. Thus the spherical attachment applies to the whole catalogue before any geometric theta or local label identities are used.

Sorting the local exponents on the same packet. Let $F = E_w$ with residue size q and uniformizer ϖ, normalized by $|\varpi| = q^{-1}$. Let π be the smooth unitary local action at this split place. Write I for the upper pro-p Iwahori and P_I for its averaging projector. Fix a nonzero I-fixed vector x, and let V_x be the span of its local translates. We use the finite-dimensional space V_x^I on the actual forms, with all the other levels and the weight unchanged. Suppose that diagonal units act on x by $\prod_i \tau_i(t_i)$. Suppose that each permutation σ with $\tau_{\sigma(i)} = \tau_i$ for every i acts on x by a scalar. We will find characters $\theta_i : F^\times \to \mathbb{C}^\times$ with $\theta_i|_{\mathcal{O}_F^\times} = \tau_i$ and a functional L nonzero on x with

$$L\big(\pi(t)P_I v\big) = \xi(t)L(v), \qquad \xi(t) = \prod_{i=1}^{n} \theta_i(t_i)|t_i|^{n-i}$$

for every diagonal t with $\operatorname{ord}(t_1) \geq \cdots \geq \operatorname{ord}(t_n)$. We will also find a nonzero pro-Iwahori eigenvector in V_x whose torus character comes from a permutation of the same θ_i, with $|\theta_i(\varpi)|$ in nonincreasing order. The unit values and the uniformizer values are permuted together.

The dominant operators and a rank-one identity. For a dominant diagonal t put $A_t = P_I \pi(t) P_I$. These operators multiply as $A_t A_u = A_{tu}$. To see this, write the pro-Iwahori by elimination as upper unipotents over $\mathcal{O}_F$, diagonal units congruent to 1, and lower unipotents with entries in $\varpi\mathcal{O}_F$, in that order. In a term $P_I \pi(tiu) P_I$ of the middle average, move the upper and diagonal parts of i left across t, and the lower part right across u. They are still in I, by dominance, so each term equals A_{tu}. In particular, these operators commute. They preserve V_x^I, as do their complex adjoints $P_I \pi(t^{-1}) P_I$.

For adjacent positions $i, j = i+1$, write s for their transposition, also viewed as its permutation matrix, and $x_+(c) = 1 + cE_{ij}$, $x_-(c) = 1 + cE_{ji}$. If L is any I-invariant functional, the cosets of $I \cap s^{-1} I s$ in I have representatives $x_+(c)$ indexed by $\bar{c} \in k_F$. Indeed this intersection adds the single condition that the ij-entry reduce to zero. If z is fixed by $x_+(\mathcal{O}_F)$, it follows that $L(\pi(s)P_I z) = L(\pi(s)z)$. If z is fixed by $x_-(\mathcal{O}_F)$, take 0 for the zero coset and units u for the other residues. Direct multiplication of the two-by-two blocks gives

$$s x_+(u) = x_+(u^{-1})\, d(u)\, x_-(u^{-1}), \qquad d(u) = \operatorname{diag}_{i,j}(-u^{-1}, u).$$

The remaining diagonal entries of $d(u)$ equal 1. Averaging gives the identity

$$qL(\pi(s)P_I z) = L(\pi(s)z) + \sum_{\bar{u}\neq 0} L(\pi(d(u))z). \tag{16.1}$$

The same calculation applies to the vector-valued map P_I.

Every permutation operator $P_I\pi(\sigma^{-1})$ is injective on I-fixed vectors. Indeed if $P_I\pi(\sigma^{-1})v = 0$ and σ has adjacent inversion $i < j$, the vector $z = \pi(\sigma^{-1})v$ is fixed by $x_-(\mathcal{O}_F)$. Diagonal units normalize I, so $P_I\pi(d(u))z = 0$. Equation (16.1) gives $P_I\pi(s)z = 0$, removing that inversion. Iteration gives $P_I v = v = 0$. Now the rotation ω with $\omega e_1 = \varpi e_n$, $\omega e_a = e_{a-1}$ for $a > 1$, normalizes I. For every entry that crosses the end of the ordered basis, conjugation supplies or removes one power of ϖ, exactly the change between the lower and upper conditions. For $1 \le a < n$, the diagonal w_a with ϖ in the first a positions and 1 elsewhere is a permutation matrix times ω^a, since those are the positions that wrap in ω^a. Thus A_{w_a} is injective on I-fixed vectors. Their product A_b, with $b = \prod_{a<n} w_a$, is injective and b is strictly dominant.

An exponent nonzero on the original vector. Choose a common eigenvector of the commuting A_t on the τ-space of V_x^I, which contains x. Its common eigenvalues give a multiplicative function ξ. Injectivity of A_b gives $\xi(b) \neq 0$. Every torus element becomes dominant after multiplication by a large power of b; division by $\xi(b)^k$ therefore extends ξ to the full torus as a character. Removing the powers of $|t_i|$ in the displayed convention gives the θ_i. Their unit characters are the prescribed ones because A_t on I-fixed vectors is just $\pi(t)$ for diagonal units.

A common eigencharacter of commuting operators in finite dimension over $\mathbb{C}$ occurs also on the dual with the transpose action, and conversely. Indeed joint generalized eigenspace decomposition for finitely many generators of the operator algebra gives the same characters after transposition. In each nonzero joint block, commuting nilpotents have a common kernel in the space and in its dual. Thus on V_x^I there is a nonzero dual eigenfunctional with character ξ. Write it as $\langle h, \cdot\rangle$ with $0 \neq h \in V_x^I$, using linearity in the second argument. Since the adjoints preserve V_x^I, we have $A_t^* h = \overline{\xi(t)}h$ there and on the ambient representation. Therefore $L(v) = \langle h, P_I v\rangle$ satisfies the exponent law and $L(h) > 0$.

If $L(x) = 0$, it vanishes on all $\pi(\sigma)x$ for finite permutations σ. In fact moving the diagonal units past σ shows that a nonzero value would require σ to preserve the unit characters of x. Its scalar action would then give zero as well. For any vector y fixed by the lower unipotents over $\varpi\mathcal{O}_F$ and the diagonal units congruent to 1, and any dominant t, the same Gauss factorization gives $L(\pi(t)y) = L(\pi(t)P_I y) = \xi(t)L(y)$. In particular this applies to a finite permutation of x, which is fixed by the principal congruence subgroup.

Write any $g \in \mathrm{GL}_n(F)$ as $ut\sigma k$, with u upper unitriangular, t diagonal and $k \in I$. An upper triangular basis of the lattice $g\mathcal{O}_F^n$ gives the Iwasawa factorization. Reduction of its compact factor gives the triangular–permutation–unitriangular factorization by elimination over k_F; the last factor lifts to I. For large m, $b^m u b^{-m} \in I$ and $b^m t$ is dominant. The preceding evaluation gives $L(\pi(b^m g)x) = \xi(b^m t)L(\pi(\sigma)x) = 0$. This vanishing holds eventually on every finite sum of translates of x, in particular on h. It contradicts $L(\pi(b^m)h) = \xi(b^m)L(h) \neq 0$.

A strict adjacent inversion. Let L now denote an exponent nonzero at $\pi(\sigma^{-1})x$ for a current permutation σ. Use θ_i for the characters in the current order. Suppose $|\theta_i(\varpi)| < |\theta_j(\varpi)|$ for $j = i + 1$, and take $s = (i\ j)$. The exchanged character is

$$\xi_s(t) = \xi(t \circ s)|t_i/t_j|,$$

because the modulus exponents at adjacent positions differ by one. Put $S = L\pi(s)P_I$ and $w = w_i$.

First consider a dominant t with gap $d = \mathrm{ord}(t_i) - \mathrm{ord}(t_j) \le 1$. The vector $\pi(t \circ s)\pi(s)P_I v$ is fixed by the lower unipotents with entries in $\varpi\mathcal{O}_F$ and the diagonal units congruent to 1. For a lower elementary matrix this follows by conjugating back through $t \circ s$ and s. Its entry still has order at least one except possibly for the lower ji entry, which becomes an upper entry of order at least zero. The vector $\pi(s)P_I v$ has the same fixed groups. Apply the dominant evaluation above to both vectors, after multiplying $t \circ s$ by a sufficiently large power of b, and divide by $\xi(b)^m$. This gives

$$L\pi(s)\pi(t)P_I v = \xi(t \circ s)S(v). \tag{16.2}$$

For gap zero, $\pi(t)P_I v$ is also fixed by $x_+(\mathcal{O}_F)$, so the upper-root case of the averaging formula gives $SA_t = \xi_s(t)S$ at that gap. For $t = w$, $\pi(w)P_I v$ is fixed by $x_-(\mathcal{O}_F)$. Equations (16.1) and (16.2) give

$$qS(A_w v) = \xi(w \circ s)S(v) + \xi(w)L(v)\sum_{\bar{u}\neq 0}\xi(d(u)). \tag{16.3}$$

Here diagonal units normalize I, so $L\pi(d(u)) = \xi(d(u))L$. On $\ker L$, this is $SA_w = \xi_s(w)S$. For a positive gap d of a general dominant t, the diagonal t/w^d is dominant with gap zero. Multiplicativity and invariance of $\ker L$ therefore give $SA_t = \xi_s(t)S$ on $\ker L$ for every dominant t. Since $A_w - \xi(w)$ maps into $\ker L$ and commutes with the dominant operators,

$$L' = S(A_w - \xi(w))$$

is an exponent for ξ_s.

We show $L'\big(\pi((\sigma s)^{-1})x\big) \neq 0$. Put $y = \pi((\sigma s)^{-1})x$, so $L(\pi(s)y) \neq 0$. The functional L reads the unit characters of $\pi(\sigma^{-1})x$ in the current order, since it is nonzero there. If the two current unit characters differ, those of y disagree with the character of L and $L(y) = 0$. The I-invariance of x gives either the whole upper or the whole lower adjacent integral root group fixing y, depending on the order of $(\sigma s)i$ and $(\sigma s)j$. By the two cases of the averaging formula, $S(y)$ is $L(\pi(s)y)$ or $q^{-1}L(\pi(s)y)$, and is nonzero. If $L'(y) = 0$, the eigenrelation on $\ker L$ implies $\xi_s(w) = \xi(w)$, hence $\theta_j(\varpi) = \theta_i(\varpi)$.

If the two unit characters agree, the transposition $\sigma s\sigma^{-1}$ preserves the unit type of x. Thus y is a nonzero scalar multiple of $\pi(\sigma^{-1})x$. The two orders of the adjacent roots show that both $x_+(\mathcal{O}_F)$ and $x_-(\mathcal{O}_F)$ fix y. Now $S(y) = L(\pi(s)y) \neq 0$. The two-by-two factorization gives $L(\pi(d(u))y) = L(\pi(s)y)$ for each nonzero residue, using invariance of L on the left and of y on the right. If $L'(y) = 0$, (16.3) says

$$q\,\xi(w)S(y) = \xi(w \circ s)S(y) + (q-1)\xi(w)S(y).$$

It follows that $\xi(w) = q\,\xi_s(w)$, hence $\theta_i(\varpi) = q\,\theta_j(\varpi)$. Both possible conclusions contradict the strict inversion. Thus the adjacent exchange preserves the required nonzero value at the permuted vector.

Swapping the strict adjacent inversions in turn gives the sorted character. Its exponent is nonzero on the I-average of a permutation of x, which belongs to V_x^I. The finite-dimensional dual argument gives a nonzero common eigenvector of the same character in V_x^I. This is the sorted vector on the original packet. The original functional L still reads the ordered coordinates on x. For example the first-line operator in the local ties is $q^{n-1}A_{w_1}$, so that $L(Uv) = \theta_1(\varpi)L(v)$ on I-fixed vectors.

The two applications. For the determinant character in the local ties, all the unit characters agree and every permutation acts by its determinant character. For a nontrivial first-entry character, the character-preserving permutations fix the first index and lie in the line parahoric, where they act by a scalar. In particular the result applies to a nonzero vector on either space, including a putative difference of the two sides of the Frobenius trace identity.

For a finite product of determinant types at a split place, let $P_{\mathrm{fin}} \subseteq \mathrm{GL}_n(k_F)$ be the parabolic with the designated consecutive blocks of sizes m_a, and $\lambda = \prod_a \chi_a \circ \det$ its Levi character, inflated across the unipotent radical. At this stage the χ_a are distinct. On a principal-level vector, the nonzero central character projection for the finite representation induced from λ means that its finite cyclic span contains an irreducible constituent of that induction. Indeed the complex character sum acts by the corresponding character multiplicity times the common group order divided by the constituent dimension. Finite Frobenius reciprocity gives a nonzero vector in the same cyclic span transforming by λ under P_{fin}. Its lift is I-fixed, has exactly m_a occurrences of χ_a on the diagonal units, and every permutation preserving these unit characters preserves the blocks and acts by a scalar. Apply the sorting argument to that vector. This gives the sorted eigencharacter, including the required multiplicities of each old determinant label. All the operations are local, so their vectors retain the Hecke system away from w.

The spherical vector for a sorted unramified exponent. Keep the local cyclic representation of the actual sorted vector at w. In this paragraph index the coordinates by $0, \dots, n-1$. Put $K = \mathrm{GL}_n(\mathcal{O}_w)$, write I for the upper pro-p Iwahori, and write π for the local action. Let L be a nonzero exponent on the cyclic representation with character

$$\xi(t) = \prod_{i=0}^{n-1} \theta_i(t_i)|t_i|^{n-1-i}, \qquad \lambda_i = \theta_i(\varpi), \qquad |\lambda_0| \geq \cdots \geq |\lambda_{n-1}|.$$

All the θ_i are unramified. If one starts from the sorted pro-Iwahori eigenvector, take its joint generalized-character block for the commuting dominant operators on the finite space of I-fixed vectors in its own cyclic span. The ideal generated by the operators minus their eigenvalues is nilpotent on this block. The quotient by its image is nonzero, so a nonzero covector on the quotient, extended by zero on the other blocks, supplies the dual eigenfunctional with character ξ. Its extension by P_I is the exponent L. Invariance under I and the trivial diagonal unit character make L invariant under the upper Iwahori in K.

We will use the dominant evaluation also for a few diagonals outside the dominant chamber. For a vector y fixed by the lower unipotents congruent to 1 and the diagonal units congruent to 1, Gauss factorization of I gives $L(\pi(t)y) = \xi(t)L(y)$ for dominant t. Indeed the lower and diagonal factors in the average fix y, and dominance carries the upper factor into I after conjugation by t. If both y and $\pi(a)y$ have these fixed groups, multiply the diagonal a by a sufficiently large power b^k of a strictly dominant diagonal b. Applying the evaluation with $b^k a$ on y and with b^k on $\pi(a)y$ proves

$$L(\pi(a)y) = \xi(a)L(y). \tag{16.4}$$

This uses only the nonzero character value $\xi(b)^k$.

Descent of vanishing through the parahoric flag. For $0 \leq m \leq n$, let P_m be the inverse image in K of the subgroup preserving the first m initial coordinate subspaces, of dimensions $1, \dots, m$. Its reduction is block upper triangular for singleton blocks at the first m coordinates and the remaining full block. Thus $P_0 = K$ and P_n is the upper Iwahori. Suppose L is zero on the K-fixed vectors. We prove by induction on m that it is zero on the P_m-fixed vectors.

For $m < n$, let $C_m \subset P_m$ stabilize the line $k_w e_m$ in the reduction, and let $R_m \subset P_m$ be the subgroup whose m-th row is supported on the m-th column in the reduction. Then $C_m \subset P_{m+1}$. Fix a C_m-fixed vector x. For $i \geq m$, put $s_i = (m\, i)$, with $s_m = 1$, and $L_i = L\pi(s_i)$.

The orbit of e_m under the reduction of P_m consists of the columns $c \in k_w^n$ having a nonzero coordinate at or after m. For such c define $v_c = \pi(p)x$ using a lift $p \in P_m$ with $pe_m = c$ in the reduction. The choice of lift is immaterial by the line stabilizer. If the last nonzero coordinate of c is i, an upper triangular matrix sending e_i to c gives $L(v_c) = L_i(x)$. For example, take the matrix obtained by replacing column i of the identity by c; its diagonal entry there is nonzero. Set

$$a_i = (q-1)q^i L_i(x).$$

There are exactly $(q-1)q^i$ columns with last nonzero coordinate i, and $\sum_c v_c$ is P_m-fixed. The outer induction hypothesis therefore yields $\sum_{i \geq m} a_i = 0$.

Now set $z = \sum_{c_m \neq 0} v_c$. It is fixed by R_m because every such matrix scales the m-th coordinate of the column by a nonzero scalar. For $i \geq m$, applying s_i exchanges the condition $c_m \neq 0$ for $c_i \neq 0$. If the last nonzero coordinate is i, there is no additional restriction. At every last coordinate $k > i$, the fraction of columns with $c_i \neq 0$ is $1 - q^{-1}$. Thus

$$L_i(z) = a_i + (1 - q^{-1}) \sum_{k>i} a_k, \qquad L_m(z) = a_m/q. \tag{16.5}$$

Let d_m be diagonal with ϖ^{-1} at m and 1 elsewhere. Conjugation $d_m^{-1} C_m d_m$ is contained in R_m. It divides the off-diagonal m-th column by ϖ, multiplies the off-diagonal m-th row by ϖ, and preserves the earlier flag conditions. Thus $\pi(d_m)z$ is C_m-fixed. We have

$$L_i(\pi(d_m)z) = \lambda_i^{-1} q^{n-1-i} L_i(z). \tag{16.6}$$

For this identity move d_m past s_i, replacing it by d_i. The vector $\pi(s_i)z$ is fixed by the principal congruence subgroup. The lower root groups congruent to 1 also fix $\pi(d_i s_i)z$. Conjugating back through d_i can turn a lower root of order one into one of order zero only if its column is i and its row is greater than i. Conjugating further by s_i then puts it in column m below row i, where R_m imposes no zero condition. All other conjugated roots still belong to the principal congruence subgroup. The diagonal units congruent to 1 also fix both vectors. Equation (16.4) gives the stated shift.

The inner induction increases j from m to n. Its assertion is that, for every C_m-fixed x, vanishing of $L_i(x)$ for all $i \geq j$ implies $L_m(x) = 0$. The case $j = m$ is immediate. To pass from j to $j+1$, assume $L_i(x) = 0$ for $i > j$ and put $C = (q-1)q^{n-1}$. Form the C_m-fixed vector

$$x' = \lambda_j \pi(d_m)z - Cx.$$

For $i \geq j$ the tail in (16.5) vanishes and $q^{n-1-i}a_i = CL_i(x)$. The shift gives $L_i(x') = C(\lambda_j/\lambda_i - 1)L_i(x) = 0$. The inner induction hypothesis gives $L_m(x') = 0$. But $L_m(z) = a_m/q$, so this last equality is

$$C\left(\frac{\lambda_j}{q\lambda_m} - 1\right) L_m(x) = 0.$$

The coefficient is nonzero by $|\lambda_j| \leq |\lambda_m|$ and $q > 1$. This includes equality of moduli and repeated roots. The inner induction at $j = n$ gives the vanishing on every C_m-fixed vector, hence on P_{m+1}-fixed vectors. The outer induction now gives vanishing at the P_n-average of every vector. Invariance under P_n forces $L = 0$, a contradiction.

The mean and its Hecke polynomial. On the finite-dimensional I-fixed subspace of this cyclic representation let h represent L in the invariant Hermitian form, linear in its second argument. Since $L = LP_I$, it is represented by the same I-fixed h on the cyclic space. Its spherical mean $g = P_K h$ represents the restriction of L to the K-fixed subspace and is nonzero by the preceding descent. The compact average stays in the cyclic space. In particular this is a vector in the packet supplying the sorted exponent.

For $0 \leq d \leq n$, let T_d be the unnormalized spherical Hecke operator of the diagonal with d copies of ϖ and $n-d$ copies of 1. It sums the lattice translates indexed by $\varpi\mathcal{O}_w^n \subseteq \Lambda \subseteq \mathcal{O}_w^n$ of residue codimension d. Partition the corresponding subspaces of k_w^n by the pivot set S of the echelon basis with last nonzero coordinates as pivots. Write A for the complement of S, of cardinality d. At a pivot $s \in S$ the basis vector is e_s plus a linear combination of the e_i with $i \in A$ and $i < s$. These coefficients vary freely. The number of subspaces with this pivot set is

$$q^{\#\{(i,s)\,:\, i\in A,\ s\in S,\ i<s\}} = q^{\sum_{i\in A}(n-1-i)-\binom{d}{2}}.$$

Lift that basis by an upper unitriangular matrix u. A representative for the lattice is ua_A, where a_A has ϖ precisely on the coordinates A. Indeed the pivot columns lift the echelon vectors and the remaining columns are ϖe_i. Since L is upper-Iwahori invariant, it reads this translate as $L(\pi(a_A)v)$ for v spherical. The vectors v and $\pi(a_A)v$ satisfy the fixed group conditions in (16.4). In the second case, conjugation back by a_A changes an entry in a lower elementary matrix by at most one power of ϖ, so the conjugated matrix is still integral. Consequently

$$L(\pi(a_A)v) = \xi(a_A)L(v) = q^{-\sum_{i\in A}(n-1-i)} \prod_{i\in A} \lambda_i \, L(v).$$

Multiplying by the cell counts and summing gives

$$L(T_d v) = q^{-\binom{d}{2}} e_d(\lambda_0, \ldots, \lambda_{n-1}) L(v),$$

where e_d is the elementary symmetric polynomial.

The spherical operators are normal on the K-fixed subspace. In the invariant Hermitian form the adjoint comes from the inverse double coset. Transposition preserves the Haar measure giving K volume

one and every K-double coset by the elementary divisor decomposition. As an anti-involution it reverses convolution, so the spherical operators commute. Put $a_d = q^{-\binom{d}{2}} e_d(\lambda_0, \dots, \lambda_{n-1})$. The functional equation gives $T_d^* g = \overline{a_d} g$, and normality makes $\|(T_d - a_d)g\| = \|(T_d^* - \overline{a_d})g\| = 0$. The usual Satake parameters are $q^{-(n-1)/2}\lambda_i$, where $q^{1/2}$ is positive. The factor $q^{d(n-d)/2}$ for the unnormalized sum on their degree-d elementary symmetric function gives precisely a_d. The monic Hecke polynomial in this convention is therefore

$$\sum_{d=0}^{n} (-1)^d q^{\binom{d}{2}} a_d X^{n-d} = \prod_{i=0}^{n-1} (X - \lambda_i).$$

The formula counts repeated roots with their multiplicities. The local translates are supported on the split pair $\{w, \bar{w}\}$, so g has the same Hecke system outside that pair. Thus on this vector an unramified exponent can be read from its spherical polynomial.

The split Iwahori polynomial on the tested channel. Take a sorted eigenvector from the local exponent calculation, and project it nontrivially to an irreducible constituent π_w. It has the same Hecke system away from the split pair. With I the upper pro-p Iwahori and U the upper unipotent group, let j_U denote the projection to the unnormalized U-coinvariants. For a dominant diagonal d and an I-fixed vector y,

$$j_U(P_I \pi_w(d) y) = d\, j_U(y).$$

In fact factor each element of I into an upper unipotent, a diagonal principal unit and a lower unipotent in that order. The upper factor disappears in the coinvariants, and conjugation by d^{-1} takes the other two factors into I. The image of the sorted eigenvector is nonzero. If it vanished, write the vector as a finite sum of differences $(\pi_w(u) - 1)z$ for $u \in U$. Averaging over a compact subgroup of U containing these u's kills it. A sufficiently large power of a strictly dominant diagonal conjugates that compact subgroup into I. Averaging after translation by this power then kills the vector, contrary to its nonzero eigenvalue for the dominant operator.

In the normalized Jacquet module, the diagonal character attached to this vector has entries

$$\alpha_i = \xi_i \,|\cdot|^{(N-1)/2}.$$

Indeed the averaging eigencharacter is $\prod_i \xi_i(d_i)|d_i|^{N-i}$ and normalization divides it by $\delta_B(d)^{1/2} = \prod_i |d_i|^{(N+1-2i)/2}$. The identity on dominant diagonals and their inverses determines the action on the entire diagonal torus. The principal-unit invariants in the Jacquet module are finite dimensional by admissibility. A character that occurs on a vector there also occurs as a quotient, by taking its simultaneous generalized eigenspace and dividing by the maximal ideal for that character. Averaging over diagonal principal units extends this quotient to the Jacquet module. Normalized Frobenius reciprocity gives $\pi_w \hookrightarrow i_B^{\mathrm{GL}_N}(\alpha)$. Sorting by decreasing $|\xi_i(\varpi)|$ means that the real exponents of the α_i are nondecreasing, since $|\varpi| = q^{-1}$. Thus $\pi_w^\vee$ is the Langlands quotient of the decreasing normalized induction of the α_i^{-1}. Within a group of equal exponents the inducing unitary characters give the tempered induction. The local parameter of π_w is therefore the sum of the characters α_i through geometric reciprocity and has zero monodromy. The Tate normalization multiplies each summand by $|\cdot|^{(1-N)/2}$, giving exactly ξ_i. [18]

For a Weil element s of nonnegative geometric degree, let x be the inverse image of its abelianization under geometric Artin reciprocity. The split comparison gives the characteristic and annihilating polynomial $\prod_i (X - j_\iota(\xi_i(x)))$ on the actual operator for s. In residue characteristic ℓ this is the linear operator on the potentially crystalline periods. At the separator, which lies away from ℓ, this gives finite inertia on the original representation. A Frobenius lift and the selected inertia element commute there. In the Frobenius-semisimplified parameter this follows from the character decomposition, and the omitted

[18] Rodier, *Représentations de* $\mathrm{GL}(n, k)$ *où* k *est un corps* p*-adique*, Bourbaki exposé 587, §§4.3–4.4, Theorem 3, and Thorne, Notation for normalized induction and geometric reciprocity.

unipotent factor commutes with inertia. For the mixed operator at $s = \phi\tau^{b}$ use both values from each same ξ_i in this polynomial. Sorting permutes the unit and uniformizer values together, so the polynomial is the one attached to the original exponent functional in the local tie. The projection to π_w leaves the global semisimple attachment unchanged by uniqueness.

The product label and its inertial polynomial. Let $c = \prod_a Q_a$ be the squarefree Green factor of the type, with distinct monic irreducible Q_a and nonzero constant term. It may be 1. With x_{η_b} the names of the determinant characters, write

$$C(X) = c(X) \prod_{b} (X - x_{\eta_b})^{t_b}.$$

The distinct factors have disjoint roots before the indicated repetitions. On a nonzero vector of the type projector, at least one irreducible constituent has a nonzero image of the same projector. The restriction there to the finite quotient $\mathrm{GL}_N(k_w)$ contains the type given by the parabolic product of the cuspidal Green representations of the Q_a and the determinant representations $\eta_b \circ \det$ of the indicated ranks.

Write this local representation π as a Langlands quotient of the standard induction of essentially square-integrable segment representations Δ_j. Taking fixed vectors under the kernel of reduction is exact over **C**. By Iwasawa decomposition the fixed space of the standard induction is the finite parabolic induction of the block fixed spaces of the Δ_j. In fact a function in the induction is determined on the maximal compact, and after imposing reduction-kernel invariance it is a function on the finite group with the stated parabolic equivariance. The modulus of the parabolic is one on its integral subgroup. Since π has the required finite label, it occurs already in this finite induction.

Each Δ_j is an unramified twist of a segment of length m_j on a supercuspidal ρ_j. It is a subquotient of induction of the corresponding unramified twists of ρ_j. Applying the same calculation shows that any nonzero fixed space of Δ_j requires a nonzero reduction-kernel fixed space of ρ_j. At depth zero this is the finite cuspidal $\bar{\rho}_j$ of its maximal compact type. Indeed the depth-zero compact induction is from the maximal compact times the center. In the Cartan decomposition, a noncentral diagonal has a gap between consecutive valuations. The conjugated reduction kernel has determinant valuation zero, so its intersection with the compact times the center lies in the compact. The reduction of this intersection contains the unipotent radical across the gap. Cuspidality kills the corresponding fixed values in the compact induction. The central double coset gives precisely $\bar{\rho}_j$. Consequently every constituent of the finite fixed space of Δ_j has the single finite cuspidal support $\bar{\rho}_j$, with multiplicity m_j.

The distinct Green factors contribute distinct cuspidal supports, each once. The determinant block of rank t_b has support η_b in rank one, repeated t_b times. The separability of $c \prod_b (X - x_{\eta_b})$ in the actual label makes all these supports disjoint. The finite parabolic induction ring factors by cuspidal support. In particular the multiplicity of the given product is the product of the multiplicities on these separate supports. For every Green factor, its occurrence forces its sole segment to have length one. For a determinant block, its occurrence forces $\eta_b \circ \det$ to occur in the induction of the finite fixed spaces on that support. Here the supercuspidal is of rank one, so Δ_j is a determinant twist of the local Steinberg of rank m_j. Its fixed space is the same twist of the finite Steinberg. This follows by applying exact fixed vectors to the Steinberg quotient of the induction of the trivial Borel character by the larger parabolic inductions and using the Iwasawa calculation for each term. By Frobenius reciprocity after twisting by $\eta_b^{-1} \circ \det$, the trivial character can occur there only when each Steinberg factor is trivial. The finite Steinberg is nontrivial in rank greater than one. Consequently each segment on this support has length one as well. In the local correspondence the standard Langlands quotient therefore has zero monodromy. The Green reciprocity calculation for the lower packet applies separately to each supercuspidal singleton of the standard data. If its residue label has root x of degree e, the corresponding Weil character is induced from the unramified extension of degree e, with the regular character T^{a} on residue units, where $T(x) = \exp(2\pi i a/(q^{e} - 1))$. The unramified rectifier and the rank twist do not affect inertia. For a tame symbol of exponent a_s at an

unramified degree divisible by every e, the Frobenius-conjugate inertial values are $T(x^{q^j})^{a_s}$ for $0 \le j < e$. For a determinant slot the same calculation in degree one gives the value of η_b on the Artin unit, repeated t_b times. Thus the inertial multiset has exactly the prescribed multiplicities. The calculation applies to every irreducible admissible local factor containing the label.[19]

For the values on the chosen inertia element, keep the faithful residue character T and the chosen prolongation. Take $f_0 \ge 1$ divisible by all the degrees of the Q_a, and a $(q^{f_0}-1)$-st root z of a uniformizer. Write b for the tame exponent of s, so

$$T\big(\overline{s(z)/z}\big) = \exp(2\pi i b/(q^{f_0}-1)).$$

The split comparison thus gives the characteristic polynomial

$$\prod_{x\,:\,C(x)=0} \big(X - j_\iota(T(x)^b)\big)$$

on the actual inertia operator, counting roots of C with multiplicity. It also gives the determinant, the finite inertial type and the same labelled jumps, using the degree-zero operator on potentially crystalline periods at $w \mid \ell$. The finite inertia action is diagonalizable in characteristic zero, so its distinct-root polynomial also annihilates it.

The labels of the fitted catalogue. At each chosen split place the lower packet has the prescribed squarefree label of degree r'. The character-corrected action on the upper block scales its roots by the name of the inverse prescribed tame character. The remaining t' slots have the trivial determinant character. Inducing from the parabolic and applying the finite central projector gives the character tests of these product labels on the cyclic spans. At every added place $w \in Q$ there are only determinant slots, of ranks r' and t'. The marked determinant reads the inverse fitted character and the complementary determinant reads the trivial one. The added conductor is disjoint from the chosen places. Every prescribed power stays nontrivial there, so the fitted character gives the full tame part on the marked block and the two determinant names are distinct. The shifted squarefree label at the old places omits 1 by the good tuple condition, so its roots are disjoint from the trivial slots. The projections commute with the other finite levels, so these statements apply to the whole catalogue at each allowed auxiliary level.

At a chosen place with $q = g_i$, each label value at inertia has order dividing $q^a - 1$ for some $1 \le a \le r$, since a root of the squarefree factor has irreducible degree at most $r' \le r$ and the determinant slots are defined over the base residue field. The finite-image inertia matrices away from ℓ are semisimple, and $q^a - 1$ divides the factor $|\mathrm{GL}_{rt}(\mathbf{F}_q)|$ of G. Hence $N_x(s)^G = I$. For a chosen tame generator, its exponent is prime to $q^{r!} - 1$, taking $f_0 = r!$ in the tame symbol. Thus every nontrivial label root still gives a nontrivial eigenvalue. This is used for the pure row, where there are no complementary trivial slots. At a chosen place above ℓ the degree-zero Weil operators have the same eigenvalue orders. Their eigenvalues lie in unramified coefficient extensions of degree at most r, since here $q = \ell$. Taking the product of the minimal polynomials of the labelled values, with multiplicity, gives the common annihilator of degree at most $(r + t)r$ used for the crystalline test. The normalizing avatar of the upper model has unramified finite parameter at these split places. It shifts the labelled jumps by its own algebraic weights and preserves the inertial Weil eigenvalues. The additional character of the lower row scales its inertia labels by the value of its tame avatar, with the corresponding shift of its labelled jumps. These assertions transport to the conjugate chosen place by sending each inertial element s to $cs^{-1}c$, where c is complex conjugation. For the unshifted upper attachment the polarizing exponent is $1 - n$. The upper normalizing avatar η adds jumps $-u, -r_-$ at a conjugate pair with $u + r_- = r'$, and $\eta^c\eta = \chi^{r'}$ by its algebraic character formula.

[19] See Zelevinsky, *Induced representations of reductive p-adic groups II*, §§6,9 for the segment classification, Moy–Prasad, Theorem 6.11, p. 118 for the depth-zero support, Schneider–Zink, §5, Proposition 6 for the finite generalized Steinberg fixed spaces, and Shotton, §6.5, Proposition 6.9 for the factorization of finite induction by cuspidal support. For depth-zero local Langlands and the geometric reciprocity convention, see Bushnell–Henniart, Part I, §2.2 and Corollary 3.3, Part III, Proposition 21.

Thus the normalized packet satisfies $N_x^{c} \simeq N_x^{\vee} \otimes \chi^{1-t'}$. Away from ℓ the cyclotomic factor is trivial on inertia. Above ℓ it is crystalline and its degree-zero Weil action is trivial. Thus on inverse conjugation the inertial operator is transposed, at the conjugate inducing embedding in the period interpretation. The same labelled polynomial and root orders hold there, and the labelled jumps transport by the polarizing weights.

For $w \in Q$, the squarefree Green factor is 1 and the product-label calculation uses $f_0 = 1$. Put $a(s) = \psi(s)$, the inverse fitted avatar, on inertia at w. Local reciprocity identifies it with the character labelling the marked determinant block. Since $w \nmid \ell$, the normalizing avatar is trivial on inertia. The normalized packet matrix $N_x(s)$ has finite order and is therefore diagonalizable. The two determinant slots give the exact relation

$$(N_x(s) - a(s)I)(N_x(s) - I) = 0.$$

If $a(s) = 1$, the matrix is the identity by semisimplicity. Its determinant is $a(s)^{r'}$. The inverse fitted avatar has the same value under inverse conjugation, so the same quadratic relation holds on $Q \cup \bar{Q}$ for the same normalized packet, including nonlift packets.

The labelled weights on the stage and sharp targets. At an infinite place of the full stage pair, write $u = r - r_{-}$. Use the positions $p = 0, \dots, n-1$ in descending order on the definite target and $q = 0, \dots, r-1$ on the source. The target highest weight in these positions is

$$\lambda'_p = \begin{cases} \varsigma_q + w_j - u + p - \lfloor t/2 \rfloor, & p = q < u, \\ 0, & u \le p < u + t, \\ \varsigma_q + w_j - u + p - \lfloor t/2 \rfloor, & p = q + t,\ q \ge u. \end{cases}$$

To see this from the diagram, the positive columns give its adjacent differences up to position u, and the negative columns give the differences at the opposite end in reverse order. In the interiors these differences are $\varsigma_{q-1} - \varsigma_q - 1$. The positive boundary length is $\varsigma_{u-1} + w_j - \lfloor t/2 \rfloor - 1$, and the first negative weight after the zeros is $\varsigma_u + w_j + \lceil t/2 \rceil$ when the negative branch is present. On the selected square $h_2 > \lceil t/2 \rceil$. The active entries satisfy $\varsigma_q + w_j \ge h_2$ on the positive branch and $\varsigma_q + w_j \le -h_2$ on the negative branch. At the final stage $u = r$, the positive column formula uses the next entry of the same row-parameter function, $\varsigma_r = h_2 - w_{m+1}$. Hence $\varsigma_r + w_m = -h_2$ there and the positive boundary formula also holds at that stage. The cutoffs in the two column formulas are therefore active exactly at their boundary differences. Starting at the zero block and summing the differences gives the displayed highest weight, with the determinant exponent cancelling the negative columns on the zero block. Equivalently its half Harish–Chandra parameter at an active position is $\varsigma_q + w_j - u + r/2$, since t is odd. The zero positions have the ordinary consecutive half-sum entries $(n-1)/2 - p$.

The unshifted target jumps at $\tau, \bar{\tau}$ are $p - \lambda'_p, \lambda'_p + n - 1 - p$. The upper normalizing avatar adds $-u, -r_{-}$, respectively. At the active positions this gives $\lfloor t/2 \rfloor - (\varsigma_q + w_j)$ and $\lfloor t/2 \rfloor + (\varsigma_q + w_j)$. At the zero positions $p = u + i$ the two entries are $i, t - 1 - i$, giving the same multiset $0, \dots, t-1$ at each embedding. The active entries agree with the normalized row on the definite lower packet. Its unshifted jump is $(r-1)/2 - \epsilon_\sigma \varsigma_q$, and the lower normalizer subtracts $(r-t)/2 + \epsilon_\sigma w_j$.

For a sharp point of rank $n = 2z$ with character of weight $-c_p(\tau)$, the target has one row of length $c_p(\tau) - z$ and determinant exponent zero. The source has weight zero and $r_{-} = 0$. The upper shift $(-1, 0)$ then gives active jumps $z - 1 - c_p(\tau)$ and $z - 1 + c_p(\tau)$ and the complementary multiset $0, \dots, 2z-2$ at either embedding. The weight-zero twist ϑ gives the same weights. The upper weights just computed are the weights of the ambient spaces as well, so the general attachment reads them on the base, twisted and deeper packets. The decomposition as representations applies to genuine lifts by their Frobenius identity.

In monic form, if sr_v multiplies all roots by v, the identity becomes

$$\mathrm{sr}_{A_w^{-1}}\, Q_w(X) = \prod_{i=0}^{t'-1} (X - q^i)\ \mathrm{sr}_{q^{t'}/B_w}\, P_w(X), \qquad \frac{q^{t'}}{B_w} = q^{(t'-r')/2} \mu^{\mathrm{bc}}(\varpi_w)^{-1}.$$

Here A_w^{-1} is the value of the upper normalizing character $\eta_{a^{-1};-r',0}$ from the rank-one calculation. The lower scalar is the value of $\eta_{\mu^{-1};0,(r'-t')/2}$. Both exponents are integral in the formula defining an algebraic character. Consequently the untwisted lower normalization is fixed before fitting, and its change at the trial stage character is the inverse weight-zero avatar of ϑ. The string is the polynomial of the inverse cyclotomic powers at a geometric Frobenius. Comparison on the Frobenius filter gives the trace identity for the same lower Hecke system, transferred from W_m to W_j as above. The root identity holds at each prescribed off-level split place on the given vector, independently of the approximation used in the trace comparison.

For a pair with source rank r', target rank $n = r' + t'$ and trial character ϑ, the upper representation is normalized by the avatar with geometric Frobenius value $j_\iota(A_w)^{-1}$. The lower factor of the spherical transfer gives one fixed twist T of the source attachment and the inverse avatar ψ of ϑ. Taking the trace coefficient and using Frobenius density gives, for every genuine lift with that source,

$$\operatorname{tr} N_\vartheta(g) = \psi(g) \operatorname{tr} T(g) + \sum_{0 \le a < t'} \chi(g)^{-a}. \tag{16.7}$$

This needs only equality by limits of the Hecke coefficients at geometric Frobenii. A separate local assertion at spherical places states that the Hecke polynomial kills Frobenius. The latter will serve the Hecke tests below. Equality (16.7) for all powers of g determines the characteristic polynomial by Newton's identities over the coefficient field. The trace-nondegenerate attachments in use have semisimple group-algebra image. Indeed its radical pairs to zero with the image under the trace, since every product in the radical is nilpotent. Thus Brauer–Nesbitt gives the actual representation over $\overline{\mathbf{Q}}_\ell$ as the direct sum of $T\psi$ and the displayed characters. One may choose the integral lattice in this representation to be that direct sum. The compact congruence argument makes this choice on its genuine lift packets. Its other packets carry their stable integral lattices in the same local class.

Let $\epsilon_\sigma = 1$ on the chosen CM type and -1 on its conjugate. For the full range the normalized jumps are

$$u_q(\sigma) = \lfloor t/2 \rfloor - \epsilon_\sigma \big(\varsigma_q(\sigma) + w_j(\sigma) \big), \qquad 0 \le q < r.$$

The shift of the lower attachment is $(r - t)/2 + \epsilon_\sigma w_j(\sigma)$. For $q + 1 < \kappa$ we have $\varsigma_q = \kappa/2 - 1 - q$, so $|\epsilon_\sigma u_q + w_j| \le t$. Its jump is above the $2h_2^2$ small window in absolute value. On the remaining m coordinates write $i = q + 2 - \kappa$, $1 \le i \le m$. Then $\varsigma_q + w_j = (j - i + 1/2)h$, with $h = 2h_2$. The modulus of the shifted jump is between $h_2 - t/2$ and $(2m - 1)h_2 + t/2$. The scale inequalities put it in $[h_2/2, h_2^2]$, and $m\kappa \le rm$ gives the required bound on the number of small jumps. At a conjugate pair of embeddings the two jumps in each coordinate sum to $2\lfloor t/2 \rfloor$, hence at most r are negative. For a sharp point $p = (\nu, z)$ the jump is $z - 1 - \epsilon_\sigma c_p(\sigma)$, with $-c_p$ the weight of ν. Since $m < z \le r$ and $w_j + r - t \le c_p \le w_j$, the high-window inequality follows from $|w_j - c_p + \epsilon_\sigma(z - 1)| \le t - r + z - 1$. There is no small jump and at most one negative jump at the two embeddings. These computations also identify the Tate-dual jumps as $-u - 1$ in this convention.

Integral rows and scalar estimates. A continuous representation of the compact group G_E into $\mathrm{GL}_d(\overline{\mathbf{Q}}_\ell)$ descends to a finite coefficient extension. To see this, exhaust $\overline{\mathbf{Q}}_\ell$ by an increasing sequence of finite Galois extensions. Such a sequence exists by Krasner's lemma: approximate the coefficients of a separable minimal polynomial by rational coefficients and lift a nearby root in the same finite extension. Krasner's inequality places the original root in the field of the nearby root. The fields generated by roots of rational polynomials thus suffice. The preimages of the general linear groups over these fields are closed subgroups that cover G_E. One is open by the Baire theorem and has finite index. Adjoin the entries of the matrices of finitely many coset representatives. The O-span of the translates of O^d is then an invariant lattice: it contains O^d and lies in a bounded multiple of O^d by compactness. This works over one sufficiently large finite coefficient DVR for any finite list of packets and additional scalar or root data. An unramified coefficient enlargement also gives a residue field of size greater than $4d + 2$ for each row

in the list. Finite free enlargement preserves normalized length. Above ℓ the inverse avatar ψ used here is trivial on the decomposition groups, so the local relations of $T\psi$ use the same coefficient model as those of T.

At a chosen primitive type prime $q = g_i \neq \ell$ of a full row, the roots are a primitive q-Frobenius orbit of length $r = 3^{a+1}$. Scaling by the label of ω^{-j} preserves primitivity. In a cyclic group of order $q^r - 1$, a factor of order dividing M_ω, with $M_\omega^2 \mid q - 1$, is a power of a generator whose exponent is divisible by each prime dividing $q^r - 1$. Multiplying by that factor preserves generator order. On reduction modulo the coefficient maximal ideal an ℓ-power factor can disappear. Write $s = r/3$. If the residual orbit had size less than r, it would be fixed by the q^s-power map. This would force $q^{2s} + q^s + 1$ to divide a power of ℓ. Since $q \equiv 1 \bmod 9$, this integer is 3 mod 9 and exceeds three. It has the prime 3 and a different prime divisor. Hence the residual orbit still has length r.

Let τ, ϕ be tame inertia and geometric Frobenius at this prime. In the residual lattice, the eigenvalue of τ and its q-power orbit give all r distinct lines, and ϕ cycles these lines. Interpolation in the single matrix $T(\tau)$ gives the projections on the residual eigenlines. Multiplying them by powers of $T(\phi)$ gives all matrix units up to invertible scalars. Nakayama therefore gives $O[T(G_E)] = \mathrm{Mat}_r(O)$ as a linear span, for any stable lattice over a coefficient DVR containing the eigenvalues. The rank-one span holds directly. In the full row the orbit and the tame relation imply that $T(\phi^{-r})$ commutes with $T(\tau)$ and $T(\phi)$, thus is a scalar over the field, and a unit scalar in the lattice. In rank one use ϕ^{-1}. For $d = r'$ the cyclotomic value of $g_* = \phi^{-d}$ is q^d.

In the full row $T(\tau^{(q^r-1)/(q-1)})$ is a scalar of order $q - 1$. In the sharp line the tame scalar comes from the label ω^{-j}, of order $M_\omega / \gcd(M_\omega, j)$. For each ℓ, one of the prime factors $2, 5$ of the fixed modulus is different from ℓ; its $(s_0 + 6)$-th power divides M_ω, where s_0 is the parameter of the weighted cyclotomic datum. Thus the residual scalar has order greater than $4d + 2$, by $m(4r + 2) < 2^{s_0+6}$. Denote this scalar inertia element by δ and its scalar by y. Both $y - 1$ and $y + 1$ are units. The polarization on traces of T is $\mathrm{tr}\, T(cgc) = \chi(g)^{1-t'} \mathrm{tr}\, T(g^{-1})$. The perfect matrix trace pairing with $T(G_E)$ gives $T(cg^{-1}c) = \chi(g)^{t'-1}\alpha I$ whenever $T(g) = \alpha I$. In particular it gives the same scalar y at $(c\delta c)^{-1}$, and the polar term for g_* is

$$d\, v_\ell(\alpha - \chi(cg_*c)\alpha') = d\, v_\ell(q^{dt'} - 1) \leq d\, v_\ell(G).$$

Here $q^{dt'} - 1$ divides the indicated general linear group order because $dt' \leq rt$.

At each old level place above $g_i \neq \ell$, the twisted label lacks root 1 and its tame eigenvalues λ have $\lambda^{q^a-1} = 1$, $a \leq r$. Since $\lambda - 1$ divides $q^a - 1$ in O for $\lambda \neq 1$, the two readings of a chosen pair satisfy separately

$$v_\ell \det(T(\tau) - 1) \leq d\, c_\ell, \qquad v_\ell \det(T((c\tau c)^{-1}) - 1) \leq d\, c_\ell. \tag{16.8}$$

For the divisibility expand $X^{q^a-1} - 1$ at 1 and at λ, and use that the value of the quotient by $X - 1$ at 1 is $q^a - 1$. This applies also to the rank-one label scaled by ω^{-j}. The tame relation bounds the local cocycle length at such a prime by $d \,\mathrm{length}_O(O/a)$ plus the length of $\ker(T(\tau) - 1 \bmod a)$. Evaluation at τ accounts for the first term. For a cocycle zero there, the relation shows that its value at a Frobenius lies in that kernel. These two evaluations determine the cocycle because the wild pro-g_i inertia acts trivially and contributes no continuous maps to the finite ℓ-module. Smith form bounds the kernel length by the respective determinant valuation.

At a place over ℓ, impose the linear relations among evaluations of integral crystalline cocycles of the lattice. For $a \mid \ell^M$, the space obeying these relations has length at most $\mathrm{rank}_O Z^1_f \,\mathrm{length}_O(O/a)$. In fact its coordinates give a homomorphism to O/a on the O-module spanned by the coordinate evaluations in the dual of Z^1_f. This module has rank at most $\mathrm{rank}_O Z^1_f$. For the rational local representation V, Bloch–Kato's dimension formula is

$$\dim H^1_f(E_v, V) = \dim H^0(E_v, V) + \dim(D_{\mathrm{dR}}(V)/\mathrm{Fil}^0),$$

with dimensions over the coefficient field. This holds for the de Rham representations here, with finite condition defined by B_{cris}. Adding the coboundary dimension $d - \dim H^0(E_v, V)$ gives the rank of the

integral cocycles as d plus the number of negative jumps over that place. Summing over places above ℓ counts each relevant embedding once and adds at most fd. There are at most $4f$ old level places off ℓ. The excess over the balanced baseline, the polar term and the two additional determinants (16.8) at the helper pair total at most $(4f+3)d\,c_\ell \le 10fd\,c_\ell$.

At the chosen coefficient place, take the union of the full rows indexed by $(j, \text{good packet on } W_m)$ and the sharp rows indexed by (j,p). These are untwisted lower lattices. The local conditions and the balanced and helper bounds are fixed on that union. The comparison of ambient and crystalline conditions is uniform on the normalized packets of any allowed trial twist, including the genuine direct-sum lattice.

16.2 Local finite loss at the two coefficient precisions

Fix one auxiliary pair and tuple label in the finite normalized catalogue N_x at the base and deeper Taylor–Wiles levels. The local generic fibres at each $v \mid \ell$ have the same labelled Hodge multiset. They are crystalline away from g_0g_1. At $\ell = g_i$ they are potentially crystalline, with the Weil inertia eigenvalues prescribed by the shifted squarefree labels and the t' trivial slots. At an inertia element, there are at most $r'+t'$ such eigenvalues with multiplicity. Each is the Teichmüller lift of an element of $\mathbf{F}_{\ell^a}^\times$ for some $a \le r$. The product of their minimal polynomials over $\mathbf{Q}_\ell$ is a common monic annihilator of degree at most $(r+t)r$. The eigenvalues also satisfy $X^{L_{\mathrm{tame}}} = 1$, where $L_{\mathrm{tame}} = \mathrm{lcm}_{1\le a\le r}(\ell^a-1)$. The auxiliary twist primes and Taylor–Wiles primes lie away from ℓ. The auxiliary twist is locally trivial on $T\psi$. The Tate slots are $O(-a)$, of jump a. The direct sum of these blocks occurs as the lattice of a genuine lift, so the blocks have the same allowed jumps and inertia as the catalogue.

For $K = E_v$, let $\mathcal{Q}$ consist of finite free stable O-lattices in the stated local class, allowing any multiplicities of the allowed jumps and imposing the same tame annihilators when $\ell = g_i$. Write $Z^1_f(K,T)$ for the continuous integral cocycles whose rational classes are in $H^1_f(K,T[1/\ell]) = \ker(H^1(K,T[1/\ell]) \to H^1(K, B_{\mathrm{cris}} \otimes_{\mathbf{Q}_\ell} T[1/\ell]))$. The integral cohomology condition includes torsion. For a de Rham representation V over $\mathrm{Frac}(O)$, the rational finite condition has dimension

$$\dim H^1_f(K,V) = \dim H^0(K,V) + \dim\bigl(D_{\mathrm{dR}}(V)/\operatorname{Fil}^0\bigr).$$

Under local cup duality, $H^1_f(K,V)$ and $H^1_f(K,V^\vee(1))$ are exact annihilators. These are the finite-condition statements of [64, Theorem 3.8 and Corollary 3.8.4, pp. 355–357]. They apply to the rationalizations of L, Y and their Tate duals, which are de Rham also at the tame primes. For a block L put $L^\dagger = L^\vee(1)$. Let $\mathcal{H}(L^\dagger) \subset Z^1_f(K,L^\dagger)$ consist of the cocycles c such that the cup of a and c vanishes for every $Y \in \mathcal{Q}$, every $a \in Z^1_f(K,Y)$, every integer $M \ge 0$, and every equivariant pairing

$$(Y/\ell^M Y) \otimes_{O/\ell^M} (L^\dagger/\ell^M L^\dagger) \longrightarrow (O/\ell^M)(1)$$

followed by a $\mathbf{Z}_\ell$-linear coefficient functional $O \to \mathbf{Z}_\ell$ reduced modulo ℓ^M. These pairings are exactly the reductions of O-linear pairings whose coboundary is divisible by ℓ^M. The divided coboundary is rationally a coboundary, so it is itself a finite cocycle. By allowing scalar multiples of the pairings, $\mathcal{H}(L^\dagger)$ is an O-submodule.

The local estimate proved below is

$$\frac{\mathrm{length}_O\bigl(Z^1_f(K,L^\dagger)/\mathcal{H}(L^\dagger)\bigr)}{\mathrm{length}_O(O/\ell O)} \;\le\; [K:\mathbf{Q}_\ell]\,\mathrm{loss}_{\mathrm{crys}}(\ell)\begin{cases} r' & \text{for } L = T\psi,\\ 1/t & \text{for } L = O(-a),\end{cases} \tag{16.9}$$

where

$$\mathrm{loss}_{\mathrm{crys}}(\ell) = 10^6\frac{m}{\kappa}\left(b_1 v_\ell(g_0g_1) + v_\ell(A_{\mathrm{crys}})\right),$$

and A_{crys} is the integer denoted by A in the coefficient budget. The cost is normalized by $\mathrm{length}_O(O/\ell O)$, even when the coefficient fraction field is ramified.

The tuple relations at finite precision. Let B_{mat} be the finite direct sum of the matrix lattices of the N_x, with g acting by left multiplication. Let C be the O-span of the tuples $(N_x(g)-1)_x$. A cocycle in $L/\ell^M L$ respecting their linear relations defines an O-linear map on C, which is equivariant by

$$N_x(g)(N_x(h)-1) = (N_x(gh)-1) - (N_x(g)-1).$$

For $M > 0$, Artin–Rees gives $C \cap \ell^{kM} B_{\mathrm{mat}} \subseteq \ell^M C$ for large k. Extend the map by zero on $\ell^{kM} B_{\mathrm{mat}}$ to $Y = C + \ell^{kM} B_{\mathrm{mat}}$. This is a stable free lattice with the rationalization of B_{mat}, hence it belongs to $\mathcal{Q}$ at v. The tuple cocycle in Y is the rational coboundary of $(1)_x$, so its restriction belongs to $Z^1_f(K, Y)$. The induced pairing $Y \times L^\dagger \to O/\ell^M(1)$ is among those defining $\mathcal{H}(L^\dagger)$. Use the trace functional from a generator of the inverse different of $O/\mathbf{Z}_\ell$. Modulo ℓ^M this is a perfect trace pairing and identifies the Tate dual of $L/\ell^M L$ as a finite $\mathbf{Z}_\ell$-module with $L^\dagger/\ell^M L^\dagger$.

Write $A_f, A_f^\dagger$ for the reductions of the integral finite cocycles of $L, L^\dagger$. They contain the finite coboundaries and are full annihilators of each other, by the rank and local duality calculation below. The tested tuple cocycle annihilates the reduction H of $\mathcal{H}(L^\dagger)$. The quotient of the tested cocycles by their intersection with A_f therefore injects into the finite dual of $A_f^\dagger/H$. Its order is bounded by that of $Z^1_f(K, L^\dagger)/\mathcal{H}(L^\dagger)$. The reductions A_f obey every local linear relation shared by $Z^1_f(K, L)$, so imposing those relations costs at most (16.9). For $M = 0$ the finite coefficient module is zero.

The global quotient by the intersection of the local relation conditions injects into the sum of the local quotients. The row block uses this estimate at $P_0 = \ell^{M_\ell}$ and the Tate blocks use it at $P = \ell^v$ in the first-order calculation. The normalized total is at most $2f(r' + t'/t)\mathrm{loss}_{\mathrm{crys}}(\ell) \le 4fr'\mathrm{loss}_{\mathrm{crys}}(\ell)$, since $\sum_{v|\ell}[E_v : \mathbf{Q}_\ell] = [E : \mathbf{Q}] = 2f$.

16.2.1 Completions and labelled jump windows

Local fields and the cyclotomic letters. Work at a fixed place $v \mid \ell$ of E in the dual defect calculation. Fix its embedding $e : E \hookrightarrow \overline{\mathbf{Q}}_\ell$, and identify the image of the completion group with $\Gamma_v = \mathrm{Gal}(\overline{\mathbf{Q}}_\ell/E_v)$. Write e_v, f_v for the ramification and residue degrees, so $[G_\ell : \Gamma_v] = e_v f_v = [E_v : \mathbf{Q}_\ell]$. With the chosen coefficient-field identification ι_ℓ, the complex embedding attached to $g \in G_\ell$ is $\sigma_g = \iota_\ell^{-1} g^{-1} e$. Two such embeddings are equal exactly when $g'g^{-1}$ fixes E_v. Let τ_g denote the infinite place of σ_g, indexed by the chosen CM embedding as in the thick set construction. Let $\epsilon_g = 1$ for that chosen embedding and -1 for its conjugate. On the cyclotomic generator ζ_E, the power read by σ_g is $\epsilon_g \tau_g$ modulo N_1, using the identification of the place group with $(\mathbf{Z}/N_1)^\times$.

The Kummer coordinate for E_v comes from a compatible tower of ℓ-power roots of a uniformizer π_v. In the period ring set $z = [\underline{\pi}_v]$, with $\theta(z) = \pi_v$. The coefficient ring of the chart is $\Lambda = W(\mathbf{F}_{\ell^{f_v}})$. Its image under θ is the integer ring of the maximal unramified subfield of E_v. The Eisenstein equation of π_v over this ring is $z^{e_v} + \ell v(z)$ with $v(z)$ a unit. Dividing by $v(z)$ makes it $\mathcal{E}_v(z) = z^{e_v} v(z)^{-1} + \ell$. Fixing $\theta(\Lambda)$ and $\theta(z)$ is equivalent to fixing E_v. If E_v is unramified one can take $\pi_v = -\ell$, giving $\mathcal{E}_v(z) = z + \ell$. For the Kummer shifts, the polynomials $X^{\ell^n} - \pi_v$ are Eisenstein over E_v, so Γ_v acts transitively on their roots. Fix a compatible tower of ℓ-power roots of unity $\underline{\varepsilon}$. For each integer $a \ge 0$ the sets of elements of Γ_v taking each chosen π_v^{1/ℓ^n} to $\varepsilon_n^a \pi_v^{1/\ell^n}$ are nested nonempty closed sets. Their intersection supplies the shift $z \mapsto [\underline{\varepsilon}]^a z$.

The letters are the Teichmüller lifts of the nonzero elements of $\mathbf{F}_{\ell^{f_v}}$. Frobenius on a letter is $q \mapsto q^\ell$. Enlarge the coefficient ring O to contain the integer ring of E_v and a primitive root ζ_O of order $q_v = \ell^{f_v} - 1$. Each embedding $O \to \mathbf{C}_\ell$ is the restriction of some $g \in G_\ell$, and its letter is $F(g) = g(\zeta_O)$. Since $E_v/\mathbf{Q}_\ell$ is Galois, the coefficient embedding determines its reading σ_g. Set $T(q) = 1$ if some g with $F(g) = q$ reads a thick place. Letters without such a reading have $T(q) = 0$. In the unramified case two readings of the same letter restrict to the same embedding of E_v, because $g'g^{-1}$ fixing the primitive root fixes every q_v-th root of unity. Thus at unramified places both the thick indicator and the CM sign are constant on each letter fiber. Choose the sign of that reading as the sign of the letter, with an arbitrary choice on empty fibers.

For later use let $N_1' = N_1/\ell^{v_\ell(N_1)}$. The prime-to-$\ell$ roots of unity in E_v inject into its residue field, so $N_1' \mid q_v$. If $F(g') = F(g)^{\ell^n}$, the quotient $g'g^{-1}$ raises every q_v-th root to the power ℓ^n. Applying this to the N_1'-part of $e(\zeta_E)$ gives

$$\tau_g/\tau_{g'} \equiv \epsilon_g \epsilon_{g'}\, \ell^n \pmod{N_1'}. \tag{16.10}$$

In fact $g^{-1}e(\zeta_E)$ on that part equals $(g'^{-1}e(\zeta_E))^{\ell^n}$, so the ratio is in the indicated order. This is the relation on the actual coefficient embeddings that is used in the local defect.

At a ramified place the only possible rational primes divide $N_1 d$. Each prime dividing N_1 splits in K, so its absolute local Galois group fixes the image of K and all readings at this place have the same CM sign. At $\ell = d$, the quotient of any two readings acts on the N_1-th roots of unity by a power of d, since these roots are unramified at d. If two readings are thick, their ratio of infinite places has class one modulo M. By (16.10), opposite signs would give $-d^n = 1 \bmod M$, excluded by $d \equiv M/2 - 1 \bmod M$. Thus the thick readings over d also have a common sign. For $\ell \mid N_1 d$, take this common sign on every letter, choosing one sign if there is no thick reading. All these primes satisfy $\ell \le \kappa t$ for large x, since d and every prime divisor of N_1 are fixed.

Spacing in the actual Frobenius orbits. Put $H = r^8$, the half-gap of the weights, and $Q = \kappa t$ with $t = 5r$. Let $\ell \le Q$ and $J_\ell = \lfloor \log_\ell(H^3) \rfloor$. For large x, $H^3 \le 2^{12s}$, since $r \le 3x^5$ and $2^s = m^{22} \ge x^{11}$. Suppose that $q, q^{\ell^{a_1}}, q^{\ell^{a_1+a_2}}$ are all thick letters, where $1 \le a_1, a_2 \le J_\ell$. Pick thick lifts with places $\tau_0, \tau_{a_1}, \tau_{a_1+a_2}$. Their signed ratios from the last place, by (16.10), are $\pm\ell^{a_1+a_2}$ and $\pm\ell^{a_2}$ modulo N_1'. Here $a_1 + a_2 \le 24s < 27s$.

For odd ℓ, reduce the two ratios using $\phi : (\mathbf{Z}/N_1)^\times \to (\mathbf{Z}/2^{s+1})^\times$. Their squares eliminate the CM signs. The thick triple then satisfies the three-term condition with exponents $2a_2$ and $2(a_1 + a_2)$, both at most $54s$. It forces $\phi(\tau_{a_1}/\tau_{a_1+a_2}) = 1$, hence $\ell^{2a_2} \equiv 1 \bmod 2^{s+1}$. Factoring the difference of squares and successively squaring gives $v_2(\ell^{2a_2} - 1) = v_2(\ell - 1) + v_2(\ell + 1) + v_2(a_2)$. One of the first two valuations is one, so $2^s \le a_2(\ell + 1)$. This contradicts $a_2(\ell + 1) = O(sx^{10}) = o(x^{11})$.

For $\ell = 2$, use the modulus $3^c \mid N_1'$ with $c = \lfloor \log_3 s \rfloor + 5$. All thick places have the same class modulo 3^c. Already the last two places would give $4^{a_2} = 1 \bmod 3^c$ after squaring their ratio. The order of 4 modulo 3^c is 3^{c-1}, since $v_3(4^{a_2} - 1) = 1 + v_3(a_2)$. This order exceeds $27s$. In either case,

$$T(q) = T(q^{\ell^{a_1}}) = 1 \quad \Longrightarrow \quad T(q^{\ell^{a_1+a_2}}) = 0 \qquad (1 \le a_1, a_2 \le J_\ell). \tag{16.11}$$

The argument works for all letters of the chart, since any violation would have thick lifts at all three positions. It also allows repeated positions in a Frobenius cycle.

The two jump windows and the local degree. Let $B = b_1$ and $G = b_2$ be the rounded thick and thin weights. In the row type take the common stage $j \le m$ for all its high jumps. Use the window parameters

$$\delta = (2j + 1)H - 2t, \qquad W = 4t, \qquad \eta = H^2 + 1, \qquad \beta_g = \begin{cases} B + \delta & \tau_g \text{ thick}, \\ G + \delta & \tau_g \text{ thin}. \end{cases}$$

For the original class, let $U(g)$ be its allowed jumps at the coefficient embedding read by g. Each is attained by a given row or Tate block. The row jumps have either $H/2 \le |u| \le H^2$ or $|\epsilon_g u + w_j(\tau_g)| \le t$. The high case puts $-\epsilon_g u$ in $[\beta_g + t, \beta_g + 3t]$. The Tate jumps $0, \dots, t' - 1$, $t' \le t$, fit in the low window. The actual Tate-dual weights are $u' = -u - 1$. In the high case $\epsilon_g u'$ belongs to $[\beta_g, \beta_g + 4t]$, since the translation changes the signed value by at most one. Its low window is $[-\eta, \eta]$, which also contains all the dual Tate jumps. Thus the original sign is $-\epsilon_g$, and the dual sign is ϵ_g. Use for $U'(g)$ precisely the dual jumps of the block being tested, a subset of the reflected set $\{-u - 1 : u \in U(g)\}$.

The limit on low jumps for a row of rank r' gives at most $\lfloor r'm/\kappa \rfloor$ jumps of its dual in the low window, counted with multiplicity. Indeed a dual jump of absolute value at most $H^2 + 1$ has original absolute value at most $H^2 + 2 \le 2H^2$. The negative members of $U(g)$ and the negative jumps of this dual row are all at

most $-H/2$. The original Tate jumps are nonnegative. For a tested Tate block instead, $U'(g) \subseteq [-\eta, 0]$. These are the two alternatives in the crystalline bound.

The letter jump sets are the unions of the embedding jump sets over each fiber of F, with zero added. At a thick reading they carry the sign and far window of that letter. For small $\ell \le Q$, enlarge the low window to $[-2G, 2G]$. All the thin far jumps lie in this window because $G + \delta + W \le 2G$. This remains true when a thin reading shares its letter with a thick one. Thus the window on letter q has far center $B + \delta$ if $T(q) = 1$ and $G + \delta$ otherwise, with the sign already chosen. For $\ell > Q$ the chart is unramified and $\ell \ge 7$, and the original low window $[-\eta, \eta]$ suffices. In this case every embedding on a letter fiber has the same thick indicator and sign.

The numerical hypotheses for the local estimate follow from the weighted scales. In fact $Q = O(x^{10})$, $W = O(x^5)$, $H \asymp x^{40}$, $m \asymp \sqrt{x}$, $G \asymp b\,x/x_1$, and $B \asymp b$. Consequently

$$99\big(Q^2 + (W+1)Q + Q + W + m\big) \le H, \qquad 99H^3 \le G, \qquad 99(GH^3 + G) \le B,$$
$$2\log(4B) < m, \qquad 0 \le \delta < \eta, \qquad 99\big((\eta + W)Q + \eta + H\big) \le H^3.$$

Also $e_v \le [E_v : \mathbf{Q}_\ell] \le 2f \le H$ for large x, since $f = \varphi(M)3^c 2^s = O_d(x^{11}\log x)$. The degree used at this place is therefore the completion degree $e_v f_v$ and is independent of the enlarged coefficient ring. At $\ell = g_i$, the congruence $g_i \equiv 1 \bmod N_1 d$ makes $E_v = \mathbf{Q}_\ell$. The Hodge windows then use one embedding of E at v, and the tame case has degree one.

16.2.2 Integral graph lifting

Two projections for a finite pairing. Let $f : Y \to L/\ell^M L$ be the map associated to an equivariant pairing in the definition of $\mathcal{H}(L^\dagger)$. Use its graph lattice

$$\mathcal{G}_f = \{(l, y) \in L \oplus Y : l \bmod \ell^M L = f(y)\}.$$

The map $\mathcal{G}_f \to Y$ is surjective and its kernel is the image of the embedding $i(l) = (\ell^M l, 0)$. The rationalization is $(L \oplus Y)[1/\ell]$, so the graph has the allowed jumps and inertia of the original class. The dual projection $i^\dagger : \mathcal{G}_f^\dagger \to L^\dagger$ is also an integral surjection and the graph dual has the reflected jumps $u \mapsto -u - 1$ and inverse Weil types.

Suppose ℓ^k kills the cokernel on integral H^1_f for $\mathcal{G}_f \to Y$. For $a \in Z^1_f(K, Y)$, the lift of $\ell^k[a]$ projects to an integral finite class in L whose reduction is $\ell^k f_*[a]$. Its cup with any $c \in Z^1_f(K, L^\dagger)$ vanishes over the integral coefficients. Indeed Bloch–Kato orthogonality gives zero rationally and the local invariant identifies $H^2(K, O(1))$ with the torsion-free module O. Reducing and applying any of the coefficient functionals proves that the original cup is killed by ℓ^k.

The same conclusion follows if ℓ^k kills the cokernel for $i^\dagger$. Represent a lift of $\ell^k[c]$ by a cocycle with values d in $\mathcal{G}_f^\dagger$, adjusting it by an integral coboundary so that $i^\dagger(d) = \ell^k c$ on cocycles. The map

$$d \longmapsto \big(y \mapsto -d(0, \ell^M y)\big)$$

is integral and equivariant into $Y^\dagger$. It sends finite cocycles to finite cocycles. For a lift $\widetilde{f}(y)$ of $f(y)$, evaluation on the graph gives

$$d(\ell^M \widetilde{f}(y), 0) + d(0, \ell^M y) = \ell^M d(\widetilde{f}(y), y).$$

The new cocycle therefore reduces to the dual pairing with $\ell^k c$. Its integral cup with a vanishes, as in the first projection. The argument applies to every coefficient functional and gives a bound on the original local field, even if the lifting proof uses a crystalline extension of it.

Zero crystalline defect in the Fontaine–Laffaille range. Let $v \mid \ell$ be a crystalline place in the weighted catalogue with $K = E_v$ unramified over $\mathbf{Q}_\ell$. At each embedding $\tau : K \hookrightarrow \overline{\mathbf{Q}}_\ell$, let I_τ be the smallest integral interval containing zero and all the allowed jumps of the original class at τ. Suppose

$$\max_\tau(\max I_\tau - \min I_\tau) \le \ell - 2.$$

Then the cup-orthogonal defect at this place is zero for each of the tested row and Tate blocks, at every coefficient precision.

Indeed, use one crystalline character ξ of G_K of labelled jumps $-\min I_\tau$. Such a character with integral unit values is obtained from the product of the labelled Lubin–Tate characters with these signed exponents. Its values can be taken in the coefficient ring O containing all images of K, by taking the K-valued Lubin–Tate character and its embeddings and inverses. Keep the convention that the cyclotomic character has jump -1. The arbitrary labelled crystalline character statement is [65, Lemma 5.1.6(i), p. 26]. The product used here has values in $O^\times$, since O contains the images of the unramified field K. Twisting by ξ puts all the allowed jumps and those of the formerly trivial quotient in $[0, \ell - 2]$.

In the range $0 \le h \le \ell - 2$, Fontaine–Laffaille theory for $K/\mathbf{Q}_\ell$ unramified and O splitting K has an exact and fully faithful covariant realization on modules over $\mathcal{O}_K \otimes_{\mathbf{Z}_\ell} O$ with a decreasing filtration in $[0, h]$ and divided Frobenius maps whose images generate the module. Its image is closed under subquotients. For free modules the labelled filtration jumps give those of the crystalline realization, and all stable lattices with these crystalline jumps are in the image. This is the coefficient version in [66, §2.3, pp. 15–16] of [67, Theorem 6.1 and Proposition 7.15]. Here $h = \ell - 2$. The graph projection and the crystalline extension of O by Y that represents an integral finite class both have the required jumps after twisting. In Fontaine–Laffaille modules write $D \twoheadrightarrow D'$ for the graph projection and C for the module of the line $O(\xi)$. In each labelled factor choose a generator e_τ of C and write d_τ for its filtration degree. With τ' the factor reached by divided Frobenius, $\Phi_{d_\tau} e_\tau = u_\tau e_{\tau'}$ for a unit u_τ. Strictness of the filtered extension by D' gives a lift $x_\tau \in \mathrm{Fil}^{d_\tau}$ of each e_τ. On this filtered direct sum the extension is specified by $a_{\tau'} = \Phi_{d_\tau} x_\tau - u_\tau x_{\tau'} \in D'_{\tau'}$. Lift $a_{\tau'}$ to $\widetilde{a}_{\tau'} \in D_{\tau'}$. On the filtered direct sum of D and C set $\Phi_{d_\tau} \widetilde{x}_\tau = u_\tau \widetilde{x}_{\tau'} + \widetilde{a}_{\tau'}$. In degrees $i \le d_\tau$ use $\ell^{d_\tau - i}$ times this value. These maps obey the divided Frobenius relations. Their images together with the images from D generate the entire module because the u_τ are units. This extension by D maps to the given extension by D', with the identity on C. Untwisting proves the surjection

$$H^1_f(K, \mathcal{G}_f) \longrightarrow H^1_f(K, Y).$$

The first graph projection therefore gives zero defect at every precision.

In the weighted windows the upper bound for the labelled width is $\beta_\tau + W + \eta$, since only one of the two signs of the far interval occurs at a given embedding. Hence the test holds whenever $\ell \ge B + \delta + W + \eta + 2$, and the smaller bound with G holds when every label of the completion is thin. Here $\delta + W + \eta = o(B)$, so $\ell > 2B$ suffices for all large squares at crystalline unramified places. This includes $\lambda_\ell = v_\ell(A_{\mathrm{crys}}) = 0$, which implies $\ell > 2B$ by the factorial in A_{crys}.

16.2.3 Numerical conditions for the local estimates

In the notation of the jump windows, $H = h_2$ and $B = b_1$. Write the valuation of the budget factor and its derived parameter as

$$\lambda_\ell = v_\ell\left((2B)!\left(\prod_{\max(B-30H^3,0)<u\le B+30H^3} u\right)^{H^3}\right), \qquad K_\ell = \left\lfloor \frac{2\lambda_\ell}{Q} \right\rfloor.$$

First use the zero defect in the Fontaine–Laffaille width range. For the other crystalline places, the following estimates charge the fixed allowance on the same blocks. If the tested block has rank R, put $x_L = 1$ for

the row and $x_L = t$ for a Tate block. Put $\mu = \lfloor Rm/(\kappa x_L) \rfloor$. The signed windows above give the μ bound for a row, and there are no low-jump restrictions to check for a Tate block in the one-sided case.

For $\ell \le Q$ and $\ell \nmid N_1 d$, the small unramified argument gives the required bound. At a crystalline prime with $\ell \le Q$ and $\ell \mid N_1 d$, use the sparse-letter estimate on the actual letter jump sets. Here its low half-width is $2G$, its far width is W, and its combined width can be taken to be $9G$ since $W + 8G \le 9G$. The far centers are $\widetilde{B} = B + \delta$, $\widetilde{G} = G + \delta$, with $\widetilde{G} + W \le 2G$. For $J_\ell = \lfloor \log_\ell H^3 \rfloor$, the spacing condition is (16.11). The sign condition holds as above. In the estimate take its precision parameter to be K_ℓ. Indeed $\lambda_\ell \ge \lfloor 2B/\ell \rfloor$ gives $K_\ell \ge B/Q^2$ for all sufficiently large x. Since $\ell^{J_\ell} \le H^3 < \ell^{J_\ell+1}$ and $e_v \le H$, we have

$$11(\ell^{J_\ell} + 1)\, 9G \le \widetilde{B}, \qquad 2e_v\widetilde{B} \le \ell^{J_\ell+1}K_\ell, \qquad 8e_v \le \ell^{J_\ell+1},$$
$$4\big(J_\ell + 1 + d_v' + \lfloor \log_\ell(2\widetilde{B}) \rfloor\big)\, 9G \le K_\ell.$$

Here $d_v' = 0$ for $e_v = 1$, and $d_v' = v_\ell(e_v) + 1$ otherwise. The first three inequalities follow from $B/G \gg H^3$ and $H^2 \gg Q^2$; the last follows from $B/G \gg mQ^2$ and $\log(2\widetilde{B}) = o(m)$. The sparse estimate yields the cup annihilator exponent $10^4 K_\ell(J_\ell + 1)$ with $J_\ell + 1 \le m$.

At a crystalline $\ell > Q$ with $\lambda_\ell \ge H^3$, use the unramified estimate with two sizes. Keep the same far centers and take the combined width $W' = W + 4\eta$. We have $2B < \ell(\lambda_\ell + 1)$, obtained from $\lambda_\ell \ge \lfloor 2B/\ell \rfloor$. Since $K_\ell \ge \lambda_\ell/Q$ in this range, the weighted scales give

$$5W' \le \widetilde{G} \le \widetilde{B}, \qquad \widetilde{B} < K_\ell\ell^2, \qquad 2\widetilde{G} < K_\ell\ell, \qquad W' \le K_\ell.$$

For the two inequalities with ℓ, use $\lambda_\ell\ell > 3B/2$ and $B/G \gg Q$. For the width use $H^3/Q \gg H^2$. The exponent supplied is $10^4 K_\ell(\lfloor \log_\ell(2\widetilde{B}) \rfloor + 1)$, also at most $10^4 K_\ell m$.

At a crystalline $\ell > Q$ with $\lambda_\ell < H^3$, the factorial factor implies $\ell > 2B/(\lambda_\ell + 1) \ge 2B/H^3 > 9G$. The interval factor implies that no integer in $(B - 30H^3, B + 30H^3]$ is divisible by ℓ. Put $A = 2(\lambda_\ell + 1)$. All nonzero integers within distance $6(A + \eta + W)$ of $0, \pm\widetilde{G}, \pm\widetilde{B}$ are prime to ℓ. Near $\pm\widetilde{B}$ they lie up to sign in the interval charged by λ_ℓ. Near $0, \pm\widetilde{G}$ their absolute values are less than ℓ. The few-low-jumps estimate applies, since also $6(A + \eta + W) < \widetilde{G}$ and $\widetilde{B} + W + \eta \le (\ell - 3)A$. It gives zero defect for the Tate block. For a row, it gives zero when $\lambda_\ell < H/5$, since then $A < H/2$ and neither jump set has negative members between $-A$ and zero. When $\lambda_\ell \ge H/5$ its defect bound is $11 f_v(\mu(A+1) + R(W+1))$ in normalized length. Here $A + 1 \le 3\lambda_\ell$ and $W + 1 \le K_\ell$, by $H/Q \gg W$.

For the exponent estimates, an annihilator ℓ^k of the defect gives normalized length at most $2e_v f_v R k$, by the local crystalline cocycle rank bound. Indeed the rank is at most $R + e_v f_v R$, and the length of $O/\ell O$ is one after normalization. The crystalline cases above therefore give length at most

$$e_v f_v\, 10^5(RK_\ell m + \mu\lambda_\ell) \le e_v f_v\, \frac{R}{x_L}\, \mathrm{loss}_{\mathrm{crys}}(\ell).$$

The inequality uses $x_L \le t$ and $K_\ell \le 2\lambda_\ell/(\kappa t)$. The crystalline deductions use the cup and length estimates in their stated ranges on the letter sets and dual block above. The tame-place argument gives the same bound at the two chosen primes.

16.2.4 Integral lifting at a prime of tame type

Let Σ be a finite Frobenius-stable set of tame inertia characters at p which contains the trivial character and has size at most N, with $N \ge 2$. Let q be the largest Frobenius period of a ratio of characters in Σ, with period one for the trivial ratio. Let $U \subset [-\eta, H] \cap \mathbf{Z}$ contain zero and have at most $N + 1$ elements. Suppose

$$1 \le \eta < H, \qquad 40(HN^2 + N^3) < p.$$

For an integral surjection of potentially crystalline lattices whose tame types and jumps lie in Σ and U, respectively, the cokernel on $H^1_f(\mathbf{Q}_p, -)$ is killed by $p^{7(Nq+\eta(N+q))}$. The bound is independent of the coefficient field and the tame splitting degree.

Take a common tame crystalline extension $K' = \mathbf{Q}_{p^a}(\varpi)$, with $\varpi^e = -p$, $e = p^a - 1$, where a is the least common multiple of the character periods. For a character of period d, its order divides $p^d - 1$. Thus all the types factor through the tame quotient of order e, and the representations become crystalline by descent on their potentially crystalline period modules. The unramified part contains μ_e, so the extension is Galois of degree ae with $p \nmid ae$. Each period is less than p. The choice $a = (p-1)!$ works as well. The two characters of a ratio belong either to the same cycle or to distinct cycles. The ratio period divides the length of the common cycle or the least common multiple of the two lengths. The trivial character occupies a separate cycle. Consequently $q \le \lfloor N^2/4 \rfloor$.

Breuil–Kisin lattices and maps. After a finite free coefficient extension with field F and integers O, use the product over the unramified embeddings, with $A = O[[u]]$ on each factor. Put $v = u^e$, $E = v + p$, and let $\Delta = \mathrm{Gal}(K'/\mathbf{Q}_p)$. Frobenius sends u to u^p and shifts the factors. Annihilation descends to the original coefficient ring. The integral cochains and the rational period maps commute with the finite free extension of coefficients, which is faithfully flat. For the covariant Kisin module M, write $L = A\Phi(M)$. Let S be the Breuil divided-power ring of E and $\mathcal{L} = SL$. Let R be the rational analytic ring on $|u| \le |p|^{1/e}$; M_R and L_R denote base change. We will use only the finite jets at the roots of E for calculations with principal parts. The comparison data give

$$E^H M \subset L \subset E^{-\eta} M, \quad \delta\mathcal{L} \subset \mathcal{L}, \quad E\delta M_R \subset M_R, \quad \delta\Phi = p\Phi\delta.$$

Here δ is the connection in the direction $u\,d/du$. It commutes with descent and has zero constant matrix coefficient in a Frobenius frame at $u = 0$. The signs are fixed by the rank-one condition $L = E^i M$ for jump i.

Use the effective theory after the common Tate twist with Frobenius E^η, of labelled jump η. The shifted interval is $[0, H + \eta]$ with $H + \eta \le p - 2$. The dual of the contravariant realization gives the covariant convention, in which a line of jump i has Frobenius multiplier E^i. Over a discretely valued field with perfect residue field of odd characteristic p, Liu identifies integral semistable lattices of weights in $[0, r]$ in his contravariant convention with finite free strongly divisible modules of height r, for $0 \le r \le p - 2$. The Kisin module on the Kummer restriction has the divided-power base change and analytic connection of this comparison [68, §2.2, Theorem 2.3.5, Corollaries 3.2.3 and 3.3.2, and Theorem 3.5.4]. Here the restriction to K' is crystalline and its residue field is $\mathbf{F}_{p^a}$. This gives stability under δ on the twisted Frobenius span. For untwisting, the Tate factor has connection coefficient in its Frobenius basis

$$-\eta \sum_{n \ge 1} u \frac{d}{du} \log\bigl(\varphi^n(E/p)\bigr).$$

This follows from its horizontal section $\prod_{n \ge 0} \varphi^n(E/p)^\eta$ in the original basis. Here $\varphi(E)/p = 1 + v^p/p$ is a unit in S, since $(v^p/p)^p \in pS$. The terms of the displayed derivative series are in S and tend to zero p-adically. This yields the integral connection on the signed Frobenius span. For $E\delta M_R$ work at a root of E in a filtration-adapted horizontal basis (d_i) of L_R. In the effective convention the recursive Breuil filtration has step r generated there by $E^{(r-i)_+} d_i$. For r at least the highest effective jump this step equals $E^r M_R$ after completion at the root. Thus the original lattice is generated by $E^{-i} d_i$ at that root. This description gives the stated stability by Euler differentiation. Away from the roots the two analytic spans agree. The weight range is below $p - 1$ for every splitting degree ae.

For a short exact sequence of Galois lattices, the Kisin sequence $0 \to M_0 \to M_1 \to M_2$ is exact through the middle and has finite-length final cokernel. Indeed, for a short exact sequence of free lattices on the Kummer tower in the essential image of finite-height Kisin modules, the covariant Kisin functor with any integral Tate shift gives a sequence exact through the middle with final cokernel killed by a power of (p, u) [69, Proposition 3.4, version 3, pp. 24–25]. The restrictions of the crystalline lattices on K' belong to this range by the common effective twist and the comparison just used. Coefficient endomorphisms commute

with the functor, so the assertion holds with O-coefficients, and finite free change of coefficients preserves exactness. The semilinear tame actions commute with the maps. After inverting p the sequence splits as a sequence of modules since $M_2[1/p]$ is free. This gives exactness in the rational jet rings at all the roots of E.

Two Euler degree intervals. Put $\rho = 2(N-1)q$, $\beta = \rho + 1$, and $m_* = H + \eta + \beta$. Then $2m_* + 2 < p$. For $f : M_1 \to M_2$, let I_τ be $\mathrm{Ann}_A(\mathrm{coker}\, f_\tau)$. The corresponding annihilator on the Frobenius spans is $J_\tau = \varphi(I_{\tau-1})A$, by the decomposition of A into the p residue classes of degrees over its Frobenius image. The height inclusions imply $E^{H+\eta} I_\tau \subset J_\tau$. Let s_τ be least with $u^{s_\tau} \in I_\tau + (p^\beta)$. The binomial expansion $v^{m_*-1} \in (E^{H+\eta}, p^\beta)$ implies $p s_{\tau-1} \le s_\tau + e(m_* - 1)$. Taking the maximum gives $s_\tau \le e$ for all factors. One more application of the same expansion gives

$$v^{m_*} L_2 \subset f L_1 + p^\beta L_2.$$

Indeed $u^e \in I_\tau + (p^\beta)$ and the product of u^e with v^{m_*-1} belongs to $J_\tau + (p^\beta)$.

Write $\mathcal{L}_{i,\mathrm{hi}} = \mathcal{L}_i \cap u^{pe} F[[u]] L_i$. The coefficients of degrees below pe in S are integral because their allowed denominators $\lfloor d/e \rfloor!$ are units. Hence $\mathcal{L}_i = L_i + \mathcal{L}_{i,\mathrm{hi}}$ and the intersection with L_i is $v^p L_i$. The operator δ preserves the high module and acts by d on the graded piece of degree $d < pe$. In the inertia isotypic component c these degrees have the form $d(c, c') + je$, with $c' \in \Sigma$ and $d(c, c')$ the residue of $c - c'$ in $[0, e)$. Let P_1, P_2 be the monic polynomials of these degrees for the respective ranges $0 \le j < m_*$ and $m_* \le j < 2m_*$. Their Euler actions raise the corresponding vectors from degree 0 to em_*, and from degree em_* to $2em_*$, modulo the high module.

In a Lagrange denominator at a fixed root, roots of the same character differ by a p-adic unit. Each other character contributes at most one root with a difference divisible by p, since the whole range of j has length less than p. For such a difference D, the reduced denominator of D/e divides $p^h - 1$, where $h \le q$ is the ratio period. Since $|D/e| < 2m_* + 1 < p$, the numerator has magnitude less than p^{h+1} and $v_p(D) \le h$. The whole Lagrange denominator has valuation at most $(N-1)q$. Interpolating p^ρ on the first set of roots and zero on the second gives $p^\rho \in (P_1, P_2)$ over $\mathbf{Z}_p$.

For x in an inertia component of L_2, both $P_i(\delta) v^{m_*} x$ belong to

$$v^{m_*} f L_1 + p^\beta L_2 + \mathcal{L}_{2,\mathrm{hi}}.$$

For P_1 this follows from the preceding lift modulo p^β and the first degree interval on its source. For P_2 the result is in $v^{2m_*} L_2 + \mathcal{L}_{2,\mathrm{hi}}$, and the same lift applies again. The displayed sum is stable under δ. The Bézout identity puts $p^\rho v^{m_*} x$ in it. The high term lies in $v^p L_2$, and the term in $p^\beta L_2$ is divisible by v^{m_*} in that module by coefficient comparison. Cancel v^{m_*} and use the lift at $v^{p-m_*} L_2$. This yields $p^\rho L_2 \subset f L_1 + p^\beta L_2$. Nakayama gives $p^\rho L_2 \subset f L_1$. Faithful flatness of Frobenius pullback proves $p^\rho M_2 \subset f M_1$. Inertia and tame projections are integral since ae is a unit.

Principal parts of the Euler connection. Write $F = \mathrm{Frac}(O)$. At all roots of $E = u^e + p$ and all unramified embeddings, use the rational E-adic completions M_{an} and L_{an}. Put

$$Y = (E^{-\eta} M)^\Delta, \qquad Y' = vY,$$
$$Z = (E^{-\eta} M / M)^\Delta, \qquad J = (E^{-\eta} M_{\mathrm{an}} / M_{\mathrm{an}})^\Delta.$$

Brackets denote principal parts in J. Division by the distinguished monic E^η identifies A/E^η with the O-module of polynomials in u of degree $< e\eta$. Rational interpolation at its roots identifies the corresponding rational jets with $(A/E^\eta) \otimes_O F$. In particular a vector of $M[1/E]$ regular in M_{an} belongs to M. Since $|\Delta| = ae$ is a unit, averaging shows that $[Y] = Z$, Z is finite free, and $J = Z[1/p]$. The module $Y' = vE^{-\eta} M^\Delta$ is p-adically complete.

The endomorphism $D[y] = [E\delta y]$ of J preserves Z. Indeed the connection preserves M_{an} after multiplication by E, and for $j < p$ every member of $\mathcal{L}$ has an integral representative in L modulo $E^j L_{\mathrm{an}}$. For the latter assertion expand each divided power $v^n/n!$ in E. Its coefficient at E^k for $k < j$, $k \le n$ is $(-p)^{n-k}/(k!(n-k)!) \in O$. Apply this calculation to the convergent divided-power expansions in an L-basis. For $y \in E^{-\eta}M$, put $h = H + \eta$. Then $E^h y \in L$ and

$$E^{2H} E\delta y = E^{H-\eta}\big(E\delta(E^h y) - h(\delta E)E^h y\big) \in \mathcal{L}.$$

Choose $\ell \in L$ representing this last vector modulo $E^{2H+\eta}L_{\mathrm{an}}$, which is possible since $2H + \eta < p$. Thus $E\delta y - E^{-2H}\ell \in E^{\eta}L_{\mathrm{an}} \subset M_{\mathrm{an}}$. Moreover $E^{\eta}E^{-2H}\ell \in M[1/E] \cap M_{\mathrm{an}} = M$, since $E\delta y \in E^{-\eta}M_{\mathrm{an}}$. This gives the required integral principal part.

Let $U \subset [-\eta, H]$ be the common set of jumps and put

$$B = \{i + h \,:\, i \in U,\ 1 \le h \le \eta\}, \qquad d = |B \setminus \{0\}| \le (N+1)\eta.$$

At a root use $t = E$, so $\delta t = e(t - p)$. The elementary divisors of M_{an} relative to L_{an} are t^{-i} for the labeled jumps i. The filtration on $t^{-\eta}M_{\mathrm{an}}/M_{\mathrm{an}}$ induced by $t^s L_{\mathrm{an}}$ therefore has a nonzero graded piece only for $s = -i - h$ with $1 \le h \le \eta$. For $\ell \in L_{\mathrm{an}}$,

$$(t\delta + sep)(t^s \ell) = t^{s+1}(\delta\ell + se\ell).$$

Hence the product of $D + sep$ over the distinct occupied degrees s annihilates the filtered space. It has simple roots, so D is semisimple on this space and on J, with spectrum in epB. In particular $D(J) \cap V = D(V)$ for every D-stable F-subspace $V \subset J$.

Let $W = Z \cap D(J)$, $b = (N+2)\eta$, and write

$$P(T) = \prod_{\beta \in B\setminus\{0\}} (T - ep\beta) = P(0) + TQ(T).$$

The factors use integers $0 < |\beta| < p$, hence $v_p(P(0)) = d$. For $w \in W$, the vector $-Q(D)w \in Z$ maps under D to $P(0)w$, because P annihilates $D(J)$. Also $p^{\eta}Z \subset [Y']$ by the identity

$$v\sum_{j=0}^{\eta-1} p^{\eta-1-j}E^j = E^{\eta} - p^{\eta}.$$

Dividing $P(0)$ by its unit factor gives $p^b W \subset D[Y']$.

The Frobenius series. For $y = y' + \Phi m$, $y' \in Y'$, $m \in M^{\Delta}$, write

$$\zeta(y) = -\sum_{n\ge 1}(p\Phi)^n \delta y, \qquad \mathrm{Mon}(y) = [E(\delta y - \zeta(y))], \qquad c_1 = (q+1)\eta.$$

In a horizontal frame δ is the Euler operator $u\partial_u$, so the formal series converges in positive u-degree. We shall use

$$p^{\eta}\zeta(y') \in \mathcal{L}, \qquad [E\zeta(y')] \in p^{p+1-2\eta}Z, \qquad p^{c_1}\zeta(\Phi m) \in \mathcal{L}.$$

Indeed $\varphi(E)/p = 1 + v^p/p$ is a unit of S, since $(v^p/p)^p \in pS$ by $v_p((p^2)!) = p + 1$. The inclusion $E^{\eta}L \subset M$ gives $p^{\eta}\Phi\mathcal{L} \subset \mathcal{L}$. For $y' = vE^{-\eta}x$, $x \in M^{\Delta}$, one has $\Phi y' \in p^{-\eta}v^p\mathcal{L}$. Using $\delta\mathcal{L} \subset \mathcal{L}$ and $\delta\Phi = p\Phi\delta$ gives for $n \ge 1$

$$(p\Phi)^n \delta y' = \delta\Phi^n y' \in p^{n-1-n\eta}v^{p^n}\mathcal{L}.$$

For $n \ge 2$ the monomial contributes at least $v_p((p^n)!) \ge p^{n-1} \ge (n-1)p$, which gives the first membership and p-adic convergence. Modulo E^{η} the expansion of $v^{p^n} = (E-p)^{p^n}$ is divisible by $p^{p^n-\eta+1}$. Represent the remaining $\mathcal{L}$-factor integrally modulo $E^{\eta}L_{\mathrm{an}}$. The discarded terms disappear in the principal part after

multiplication by E, since $E^{\eta}L_{\mathrm{an}} \subset M_{\mathrm{an}}$. Thus $[E\delta\Phi^{n}y']$ is divisible in Z by $p^{p^{n}+n-(n+1)\eta}$. The minimum occurs at $n = 1$ and the exponents tend to infinity.

For the third membership take $w = \delta\Phi m \in \mathcal{L}$. It is invariant, and has zero constant coefficient in the formal horizontal frame. If $w \neq 0$, let $s = \min(e, \operatorname{ord}_{u} w)$ in the rational formal module at zero, taking the minimum over the components. If $s < e$, a nonzero leading coefficient has fiber inertia character $-s \bmod e$ when inertia acts on u by its tautological character. Its period $h \leq q$ gives $e \mid s(p^{h}-1)$. Thus $sp^{q} \geq e$ in either case. In a Frobenius basis of L every coefficient f of w belongs to $S \cap u^{s}F[[u]]$. For $n \geq 1$, $\varphi^{n}(f)/u^{sp^{n}} \in S$. Indeed a term of degree $j \geq s$ moves to degree $(j-s)p^{n}$. For $j < ep$ its original factorial $\lfloor j/e \rfloor!$ is a unit, while for $j \geq ep$ the new degree is at least j. The new factorial therefore absorbs the denominator and preserves convergence of the divided-power coefficients. It follows that

$$(p\Phi)^{n}w \in p^{-n(\eta-1)}u^{sp^{n}}\mathcal{L}.$$

For $1 \leq n \leq q+1$ multiplication by p^{c_1} suffices. For $n = q+1+t$, $t \geq 1$, $sp^{n} \geq ep^{t+1}$, so the monomial contributes at least $p^{p^{t}}$. The net exponent after multiplying by p^{c_1} is at least $q+1+t+p^{t}-t\eta$, which is positive and tends to infinity since $p^{t} \geq tp$.

Correction in an exact triple. For these y, $\omega = \mathrm{Mon}(y)$ belongs to $D(J)$ and $p^{c_1}\omega \in Z$. The integral assertion follows from $y \in Y$, $DZ \subset Z$, and the existence of integral representatives for the coefficients of $p^{c_1}\zeta(y)$ modulo $E^{\eta}L_{\mathrm{an}}$. For the rational assertion solve $\delta s = \zeta(y) \bmod t^{\eta}L_{\mathrm{an}}$ at each root. In an L_{an}-basis, $\delta = e(t-p)d/dt + C(t)$ with $C(t)$ regular. Set $s_0 = 0$. At degree j, obtain s_{j+1} by division by $-ep(j+1)$, for $0 \leq j < \eta$. Rational interpolation realizes the jets in L_{an}, and averaging gives an invariant solution. Now $s \in E^{-\eta}M_{\mathrm{an}}$ and $E(\delta s - \zeta(y)) \in M_{\mathrm{an}}$, so $D[y-s] = \omega$.

Set $k = p - 14H > b$. On Y' the difference $\mathrm{Mon}(y') - D[y']$ lies in $p^{k}W$, by the series estimate and p-saturation of W in Z. Since $p^{b}W \subset D[Y']$, successive approximation yields $p^{b}W \subset \mathrm{Mon}(Y')$. In detail, for a target $p^{b}w$ first choose y' with $D[y'] = p^{b}w$. Its error is in $p^{k}W$. If an error belongs to $p^{b+j(k-b)}W$, subtract a multiple $p^{j(k-b)}z$ with $z \in Y'$ whose leading term corrects it. The new error belongs to $p^{b+(j+1)(k-b)}W$. The corrections converge in Y' and the errors tend to zero in Z.

Finally let $0 \to M_0 \xrightarrow{i} M_1 \xrightarrow{g} M_2$ be the given exact triple with $p^{\rho}M_2 \subset gM_1$. It splits as a sequence of modules after inverting p, since the target is free. Thus $i : J_0 \to J_1$ is injective. For $w \in Z_1 \cap \ker g$ pick $y \in Y_1$ representing w. Then $gy \in M_2[1/E] \cap M_{2,\mathrm{an}} = M_2$. Lift $p^{\rho}gy$ to an invariant $x \in M_1$ by averaging. Write $E^{\eta}(p^{\rho}y - x) = im_0$. It follows that $p^{\rho}w = i[E^{-\eta}m_0] \in iZ_0$.

Given $y_2 \in Y_2'$, $m_2 \in M_2^{\Delta}$ with $\mathrm{Mon}(y_2 + \Phi m_2) = 0$, the two terms after multiplication by p^{ρ} lift to $y_1 \in Y_1'$ and Φm_1, $m_1 \in M_1^{\Delta}$. The error $\omega = \mathrm{Mon}(y_1 + \Phi m_1)$ belongs to $D_1(J_1)$, satisfies $g\omega = 0$, and $p^{c_1}\omega \in Z_1$. Use the preceding paragraph to write $p^{\rho+c_1}\omega = iw_0$ with $w_0 \in Z_0$. Semisimplicity gives $D_1(J_1) \cap iJ_0 = iD_0(J_0)$, hence $w_0 \in D_0(J_0)$. Choose $y_0 \in Y_0'$ with $\mathrm{Mon}(y_0) = p^{b}w_0$. Then $p^{\rho+c_1+b}(y_1 + \Phi m_1) - iy_0$ has zero monodromy and maps to $p^{2\rho+c_1+(N+2)\eta}(y_2 + \Phi m_2)$.

The target finite class. For $c \in H^{1}_{f}(\mathbf{Q}_p, T)$ the corresponding integral extension of the unit by T is crystalline over K', with its jumps and tame types in the same sets. Use the integral Kisin map to the unit to lift p^{ρ} to an invariant vector z. Then $y = \Phi z - z$ belongs to Y. The extension with Frobenius vector y is the pullback by p^{ρ} of the original étale extension, hence represents $p^{\rho}c$. The Frobenius–connection equation gives $\delta z = \zeta(y) - \delta y$ in positive u-degree. The series represents $\delta(\Phi z)$ in the analytic Frobenius span, and this vector lies in the rational kernel span. Since $E\delta z$ belongs to the analytic extension lattice and maps to zero in the unit, $\mathrm{Mon}(y) = 0$ by exactness at the roots. The correction from y to a vector of the required form uses $m = E^{\eta}y$. The cohomologous Frobenius representative of $p^{\rho+\eta}c$ is

$$p^{\eta}y + (\Phi - 1)m = y' + \Phi m, \qquad y' = (p^{\eta} - E^{\eta})y \in Y', \quad m \in M^{\Delta}.$$

Its monodromy still vanishes because modifying the integral section by m changes the error by the class of $E\delta m$.

Realization of the corrected lift. The corrected vector $y = y' + \Phi m$ has zero monodromy. Multiplying first by p^{c_1}, $c_1 = (q+1)\eta$, puts $\zeta(y)$ in the integral Breuil lattice. With this scaled vector form $M_y = M \oplus Az$ with $\Phi_y z = z + y$ and z fixed under Δ. Its Frobenius span is $L_y = L \oplus A(z + y)$, with

$$E^{\eta} L_y \subset M_y, \qquad E^{H+\eta} M_y \subset L_y.$$

For the second inclusion use $E^{H+\eta} y \in E^H M \subset L$. The common effective Tate twist of this extension has Frobenius span $L_{\mathrm{tw}} = E^{\eta} L_y$ in the twisted basis. It lies in its original lattice M_{tw} and has height at most $r_{\mathrm{tw}} = H + 2\eta \le p - 2$.

Define $\delta_y z = \zeta(y) - \delta y$. The series equation for ζ gives $\delta_y \Phi_y = p \Phi_y \delta_y$. Since $\delta_y(\Phi_y z) = \zeta(y)$, the enlarged Frobenius span over S is stable. The equality $\mathrm{Mon}(y) = 0$ gives $E\delta_y z \in M_R$, so $E\delta_y$ also preserves $(M_y)_R$. The Tate factor preserves both conditions.

Apply the usual Kisin base change $S \otimes_{\varphi, A} M_{\mathrm{tw}}$ in height r_{tw}. Linearized Frobenius identifies it with $\mathcal{L}_{\mathrm{tw}} = SL_{\mathrm{tw}}$. Its filtration is the inverse image of $(\mathrm{Fil}^{r_{\mathrm{tw}}} S) M_{\mathrm{tw}}$. This is tested by rational jets through degree $r_{\mathrm{tw}} - 1$ at the roots of E. Thus the filtration is $\mathcal{L}_{\mathrm{tw}} \cap E^{r_{\mathrm{tw}}}(M_{\mathrm{tw}})_R$. The ordinary height-$r_{\mathrm{tw}}$ construction supplies $\varphi_{r_{\mathrm{tw}}} = p^{-r_{\mathrm{tw}}} \Phi_{\mathrm{tw}}$ on this filtration, with generating image and the required reductions modulo all p^n. The quotient by the filtration is p-torsion-free. Let $N_{\mathrm{Br}} = -\delta_{y,\mathrm{tw}}$. The Leibniz rule and preservation of $(M_{\mathrm{tw}})_R$ by $E\delta_{y,\mathrm{tw}}$ give $E N_{\mathrm{Br}}(\mathrm{Fil}^{r_{\mathrm{tw}}}) \subset \mathrm{Fil}^{r_{\mathrm{tw}}}$. Compatibility of Φ_{tw} and the connection gives

$$\varphi_{r_{\mathrm{tw}}}(E N_{\mathrm{Br}} x) = (\varphi(E)/p)\, N_{\mathrm{Br}} \varphi_{r_{\mathrm{tw}}}(x)$$

on that filtration. These are the conditions for the finite free strongly divisible category. Breuil's criterion gives a weakly admissible filtered module and semistable lattice from these integral data [68, Proposition 2.3.4 and Theorem 2.3.5]. The bound $r_{\mathrm{tw}} \le p - 2$ places $\mathcal{L}_{\mathrm{tw}}$ in its range. Apply the contravariant functor over $\mathbf{Z}_p$ and dualize its realization. This yields a semistable lattice over K' with the commuting coefficient action. Its monodromy on the filtered isocrystal is obtained by evaluating the connection at $u = 0$. It vanishes there in the Frobenius frame, including the coefficient $\zeta(y)(0) = 0$ on the added vector. Thus the lattice is crystalline, and we undo the common twist.

The comparison of the Kisin and Breuil realizations identifies the restriction to $G_{K'_\infty}$ of the dualized semistable lattice with the covariant étale realization of M_y [68, Lemma 3.3.4 and §3.5]. The étale equivalence for the Kummer tower identifies free Galois lattices with the corresponding étale Frobenius modules and is exact. Crystalline integral $G_{K'}$-lattices restrict fully faithfully to $G_{K'_\infty}$ [70, Corollaries 3.8 and 7.10]. Apply exactness to the Frobenius extension to recover the original integral kernel T and unit quotient. Both maps are crystalline maps over K' by full faithfulness. For tame descent, use $\varpi_{n+1}^{p} = \varpi_n$ with $\varpi_0 = \varpi$. Each $\gamma \in \Delta$ acts on the Kummer tower by $\varpi_n \mapsto (\gamma(\varpi)/\varpi)^{p^{-n}} \varpi_n$, using the unique p^n-th root in μ_e. The preimage of these actions in $G_{\mathbf{Q}_p}$ normalizes $G_{K'_\infty}$, maps onto Δ, and intersects $G_{K'}$ in $G_{K'_\infty}$. The semilinear action on M_y intertwines the corresponding actions on the étale restriction. By full faithfulness the intertwiners also intertwine the crystalline $G_{K'}$-lattice. Thus the two actions give an integral $G_{\mathbf{Q}_p}$-action on the exact sequence. Its class is finite by the injectivity of restriction on rational B_{cris}-cohomology. Comparisons of finite extension classes can be checked on $G_{K'_\infty}$ by full faithfulness over K', followed by restriction and corestriction of the p-adic unit degree ae.

The target representative, the monodromy correction, and realization cost respectively $\rho + \eta$, $2\rho + c_1 + (N+2)\eta$, and c_1. Thus the total exponent is

$$3\rho + 2c_1 + (N+3)\eta \le 7\big(Nq + \eta(N+q)\big).$$

This proves the local lifting bound on the given stable lattices.

Application at the two type primes. Let $\ell = g_i$, and fix a place $v \mid \ell$. Then $E_v = \mathbf{Q}_\ell$ and the complex embedding and CM sign are fixed at this place. Write L for the actual row block or scalar Tate block, $L^{\dagger} = L^{\vee}(1)$, and $\mathcal{Q}$ for the class of lattices with the allowed jumps and Weil inertia polynomials of the

catalogue. We use the integral lifting argument at a tame type prime on a graph lattice constructed from an arbitrary pairing in the definition of the cup-orthogonal submodule.

The characteristic-zero Weil inertia of every member of $\mathcal{Q}$ has exponent dividing

$$L_{\mathrm{tame}} = \mathrm{lcm}_{1\le a\le r}(\ell^a - 1).$$

This order is prime to ℓ, so the action of wild inertia is trivial and the tame action is through the cyclic quotient of order L_{tame}. Take an inertial element mapping to a generator of this quotient. At this element the common monic annihilator from the catalogue has coefficients in $\mathbf{Z}_\ell$ and degree at most $N = (r+t)r$. Its roots among the L_{tame}-th roots of unity give one set of eigenvalues for all the lattices in $\mathcal{Q}$. They determine their tame inertia characters at the generator. Adjoin the trivial character to obtain a Frobenius-stable set of cardinality at most $N_* = N+1$. Each character has order dividing L_{tame}, so its period divides $\mathrm{lcm}(1,\dots,r)$. For two roots, their ratio has period dividing the least common multiple of their periods. If they are in one orbit the orbit length bounds the ratio period. If they are in distinct orbits, the product of the orbit lengths bounds it and their sum is at most N_*. Thus $q \le \lfloor N_*^2/4 \rfloor$, using also the trivial orbit in the single-orbit case.

The inverse set serves the Tate duals, since the Tate twist is crystalline and hence has trivial Weil inertia. The common tame crystalline extension in the lifting lemma is obtained by taking an unramified degree divisible by the character periods and adjoining a root of the corresponding prime-to-ℓ Eisenstein equation. The unramified degree can be chosen prime to ℓ, since all periods divide $\mathrm{lcm}(1,\dots,r)$ and $r < \ell$.

All members of $\mathcal{Q}$ have jumps among the row jumps and the scalar Tate jumps $0,\dots,t'-1$, $t' \le t$. At the fixed place there are at most $r+t$ allowed integers. The jump comparison puts the large jumps on a common side. Choose the original orientation if they are positive and the Tate-dual orientation if they are negative, using either orientation when there are no large jumps. In the chosen orientation all jumps lie in $[-\eta_*, H_*]$, where $\eta_* = \eta + 1$ and $H_* = \widetilde{B} + W + 1$. Reflection $u \mapsto -u-1$ is included in those endpoints. Adjoining the trivial jump still gives at most $r+t+1 \le N_*+1$ integers. The prime threshold of the tame lifting lemma holds: $1 \le \eta_* < H_*$ and

$$40(H_* N_*^2 + N_*^3) < \ell.$$

Indeed $H_* = O(b)$, $N_* = O(r^2)$ and $\ell \ge b^{10}$. We may therefore use the uniform exponent

$$k = 7\big(N_* q + \eta_*(N_* + q)\big), \qquad q \le N_*^2.$$

In particular $k = O(r^{20})$.

The chosen graph projection and its kernel have the stated common types and jumps, as do the lattices in any finite extension of the target by the unit. Apply the lifting estimate to that projection and then use the graph cup calculation. For every pairing in the definition of $\mathcal{H}(L^\dagger)$ this gives $\ell^k Z^1_f(K, L^\dagger) \subset \mathcal{H}(L^\dagger)$. If $R = \mathrm{rank}_O L$, the normalized defect is at most $2Rk$ at this degree-one place. The allowance pays this since

$$2k \;\le\; \frac{\mathrm{loss}_{\mathrm{crys}}(\ell)}{t}, \qquad \mathrm{loss}_{\mathrm{crys}}(\ell) \;\ge\; 10^6\,\frac{m}{\kappa}\,b_1,$$

using $v_\ell(g_0 g_1) = 1$, $2kt = O(r^{21})$, $m/\kappa \ge 1/r$ and $b_1 \ge b = \exp(x^\varepsilon) \gg r^{22}$. This gives the required tame bound both for the row and for each scalar Tate block.

16.2.5 Integral lifting on an unramified signed cluster

Let $\ell \ge 11$ be prime and $K_\ell/\mathbf{Q}_\ell$ be unramified of degree n_ℓ. Take a finite coefficient field E'_ℓ splitting K_ℓ, with valuation ring O and valuation normalized by $v(\ell) = 1$. Label its embeddings of K_ℓ by $i \in \mathbf{Z}/n_\ell$ with $\tau_{i+1} \circ \sigma_{K_\ell} = \tau_i$, where σ_{K_ℓ} lifts residue Frobenius. Let $F_\ell : T_1 \twoheadrightarrow T_2$ be an equivariant surjection of finite free stable O-lattices in crystalline representations, and let $r_\ell = \mathrm{rk}_O T_2$. A labelled jump a means $\mathrm{gr}^a D_{\mathrm{dR}}(T_1[1/\ell])_{\tau_i} \ne 0$ in the decreasing filtration.

For every C above an absolute positive threshold, the estimate holds with the following data. Take nonnegative integers b_*, w_*, S_*, η_*, and $1 \le j_* \le n_\ell - 1$. At each label choose an integer $s_i \ge 0$ and a sign, independently of the other labels. Its jumps lie in $[-\eta_*, \eta_*]$ together with either $[s_i, s_i + S_*]$ or $[-s_i - S_*, -s_i]$, according to that sign. The multiplicities are arbitrary. Set $g_* = w_* + S_* + \eta_*$ and $I = \{i : b_* \le s_i \le b_* + \eta_*\}$. Require $s_i \le w_*$ outside I and $b_* \ge C(j_* + 1)(g_* + 1)$. For each $t_0 \in I$, either $t_0 - 1, \dots, t_0 - j_*$ are outside I, or there is $1 \le d'_* \le j_*$ with $t_0 - d'_* \in I$ and $t_0 - a \notin I$ for $1 \le a \le d'_* + j_*$, $a \ne d'_*$. All indices in this condition are cyclic. A pair within a cluster can be adjacent. The same jump windows contain those of the target by strictness.

Put $H_{\mathrm{loc}} = b_* + \eta_* + S_*$ when $I \ne \emptyset$ and $H_{\mathrm{loc}} = g_*$ otherwise. These bounds include the endpoints $s_i + S_*$ of all the prescribed outer windows, even when an outer window contains no jumps. Define $H^1_f(K_\ell, T)$ as the full inverse image in $H^1(K_\ell, T)$ of the rational Bloch–Kato finite subgroup defined by B_{cris}. It includes the integral torsion. The cokernel D_{F_ℓ} of the induced map is killed by ℓ^e, where the nonnegative integer e is bounded by an absolute constant times

$$\frac{H_{\mathrm{loc}}}{(\ell - 1)^{j_* + 1}} + (g_* + 1)(j_* + 1) + \log_\ell \max(1, H_{\mathrm{loc}}).$$

After conversion to normalized length, this gives

$$\frac{\mathrm{length}_O D_{F_\ell}}{\mathrm{length}_O(O/\ell O)} \le C(n_\ell + 1) r_\ell \left(\frac{H_{\mathrm{loc}}}{(\ell - 1)^{j_* + 1}} + (g_* + 1)(j_* + 1) + \log_\ell \max(1, H_{\mathrm{loc}}) \right). \tag{16.12}$$

The denominator is the absolute ramification index of E'_ℓ.

The cut exponent. In the proof write $p = \ell$, $G_0 = g_*$, $g = \max(1, G_0)$, $h = \max(1, H_{\mathrm{loc}})$, and $j = j_*$. Choose the minimal bounds $[-M_i, h_i]$ enclosing the jumps and zero. Their widths are at most $2h$ and at most $2G_0$ at thin labels, while a single signed outer window gives width at most $h + g$. Group successive large labels whose cyclic distance is at most j. The predecessor condition makes each group have at most two large labels and leaves at least j thin labels between groups. Cut at the last label of each group. If the next group's first label lies more than $2j + 1$ steps away, insert cuts at successive distances j. Thus all arcs have length from j to $a_* \le 3j + 1$ and start with j thin labels. An all-thin cycle can be divided into lengths from j to $2j$. This construction also works when two large labels are adjacent.

The rational finite-group dimension is at most $n_\ell r_\ell + d_0$, where $d_0 = \dim(T_2[1/\ell])^{G_{K_\ell}}$. The torsion embeds in the quotient of $T_2[1/\ell]/(T_2[1/\ell])^{G_{K_\ell}}$ by the image of T_2, so it requires at most $r_\ell - d_0$ generators. This proves the factor $(n_\ell + 1) r_\ell$ when a uniform annihilating exponent is converted into normalized length. Write $\mu_{\mathrm{cyc}} = [\varepsilon] - 1$ for the full cyclotomic parameter and $A_K^+ = \mathcal{O}_{K_\ell}[[\mu_{\mathrm{cyc}}]]$. For a crystalline lattice over the unramified field K_ℓ, Berger's integral Wach lattice $N(T)$ is free over A_K^+, has trivial cyclotomic action modulo μ_{cyc}, and is the intersection of the rational Wach lattice with the integral étale module. This gives kernel exactness on stable saturated sublattices by intersecting the rational kernel with the integral étale kernel. These assertions in the effective range are [71, Proposition II.1.1 and Lemma II.1.3]. The weights in Berger's convention are the negatives of our decreasing-filtration jumps. The signed extension of the construction follows by a common Tate shift and uniqueness [71, Proposition III.4.2]. Tensoring with a rank-one crystalline lattice commutes with this construction. After shifting the tensor factors to the effective range the product lattice satisfies Berger's uniqueness conditions, including the congruence modulo μ_{cyc} and generation of the étale module.

The O-action passes to $N(T)$. Its reduction modulo μ_{cyc} is $\mathbf{Z}_p$-torsion-free, so each labelled factor there is free over O. Lifting a basis and using μ_{cyc}-adic completeness and injectivity of multiplication by μ_{cyc} makes each factor free over $O[[\mu_{\mathrm{cyc}}]]$. In the full parameter the factor at cut i is the factor of e_{i+1}, with e_i the idempotent of τ_i. Thus Frobenius maps the factor at cut $i - 1$ to that at i, and specialization by φ^{-s-1} at cyclotomic level s on the latter reads τ_{i-s}. Write P_i^{tar} for the target factor and P_i^{im} for the image of the source. Let s'_i be the least exponent of μ_{cyc} killing their quotient. For labelled bounds $[-M_i, h_i]$ put

$H_0 = \max_i(M_i + h_i)$. A common rank-one crystalline twist shifts the i-th jump interval by M_i to $[0, M_i + h_i]$. Take the product of the conjugate Lubin–Tate characters with the prescribed labelled exponents. The rank-one crystalline period calculation for these characters is [72, Example 6.2.6, pp. 1095–1096]. The Wach lattice of the twist is invertible in each factor, so tensoring preserves the cokernel exponent.

Lemma 16.1 (Integral scalar retraction). *Let $S = W(k)[[u]] \to A$ be a unital map, with k a perfect field. Write $I = (p, u)A$ and $I_S = (p, u)S$. Suppose A is complete and separated for I, that p is a nonzero divisor in A, that u is a nonzero divisor in A/p, and that $k \to A/I$ identifies k with a subfield. Then $S \to A$ admits an S-linear retraction $A \to S$.*

Proof. For $n \geq 1$ one has $I^n : p = I^{n-1}$. If $pa \in I^n$, write $pa = a_0 u^n + pb$ with $b \in I^{n-1}$. Reduction modulo p puts a_0 in pA, so division by p gives $a \in I^{n-1}$. The reverse inclusion is immediate. In degree n the classes of $p^j u^{n-j}$ give an (A/I)-basis of I^n/I^{n+1}. They generate by the definition of I^n. For independence, reduce a relation first modulo p. Cancelling u^n there shows that the coefficient of u^n is in I. Remove that term and divide by p, using the colon identity to reduce to a relation in degree $n-1$. This proves independence by induction.

Choose a k-basis of A/I containing 1 and choose lifts $b_\alpha \in A$, taking 1 as its own lift. Each element of A has a unique expansion $\sum_\alpha s_\alpha b_\alpha$ with $s_\alpha \in S$ and only finitely many s_α nonzero modulo I_S^n for every n. To construct it, expand the leading residue successively in I^n/I^{n+1} using the monomials just proved independent and the chosen residue basis. The degree-n corrections are finite sums with coefficients in I_S^n. Completeness of S gives the s_α, and completeness of A evaluates the sum. If two expansions differ, their first nonzero degree gives a relation in I^n/I^{n+1}, which is impossible. Taking the coefficient of 1 gives the required S-linear map. □

For the Wach application take $S = A_K^+ = \mathcal{O}_{K_\ell}[[\mu_{\mathrm{cyc}}]]$ and $A = W(\mathcal{O}_{C_p^\flat})$. Here $\mathcal{O}_{K_\ell} = W(\mathbf{F}_{p^{n_\ell}})$ and the residue field embeds in $A/(p, \mu_{\mathrm{cyc}})$. The Witt ring is p-torsion-free and p-adically complete, and μ_{cyc} mod p is nonzero of positive valuation in the complete rank-one valuation ring $\mathcal{O}_{C_p^\flat}$. Its powers define a complete separated topology on A/p. These facts give completeness and separation for I. Indeed, pass to a subsequence a_n whose differences $a_m - a_n$ for $m \geq n$ belong to I^n. Lift the limit modulo p to $a_0 \in A$ and write $a_n = a_0 + u^n b_n + pc_n$. By the colon identity $c_m - c_n \in I^{n-1}$. Iterating constructs successive p-adic digits of an I-adic limit. The sum of the digits converges by p-adic completeness, and each finite stage agrees with the sequence modulo the next power of p up to a term tending to zero I-adically. For separation, an element in every I^n reduces to zero modulo p. Its quotient by p still lies in every I^n, so iteration and p-adic separation prove the claim. The same verification applies to a Witt ring with a parameter of positive finite tilt valuation in the Kummer chart below.

For lattices with jumps in $[0, H_0]$, Berger's integral comparison places the Wach lattice after Witt base change between the representation lattice and its multiple by $\mu_{\mathrm{cyc}}^{H_0}$ [71, Theorem III.3.1]. Apply this to the effective twists of the source and target. An integral representation-lattice lift of a target vector, scaled by $\mu_{\mathrm{cyc}}^{H_0}$, lies in the source Wach lattice after Witt base change. Apply Lemma 16.1 to its coefficients in a free frame to descend the lift to A_K^+. With O-coefficients use the retraction tensored with O, which is linear over $A_K^+ \otimes_{\mathbf{Z}_p} O$ and commutes with the original lattice map. Thus $s_i' \leq H_0$. For the map from a crystalline extension of T_2 to the trivial quotient, use the target jumps together with zero and the same widths.

If $c\mu_{\mathrm{cyc}}^{k+1} P_i^{\mathrm{tar}} \subset P_i^{\mathrm{im}}$, apply the Wach congruence $\gamma - 1$ on the image and divide by μ_{cyc}. Its leading term gives $c(\chi(\gamma)^{k+1} - 1)\mu_{\mathrm{cyc}}^k P_i^{\mathrm{tar}} \subset P_i^{\mathrm{im}}$, after subtracting the correction in degree $k+1$. Choose $\chi(\gamma)$ generating $\mathbf{Z}_\ell^\times$. Descending from s_i' to power $a \leq s_i'$ costs at most

$$\sum_{\substack{a<k\leq s_i' \\ \ell-1\mid k}} (1 + v_\ell(k)) \leq (s_i' - a)\frac{\ell}{(\ell-1)^2} + c_H, \qquad c_H = \lceil 1 + \log_\ell(H_0 + 1) \rceil.$$

For the inequality count multiples of $(\ell-1)\ell^j$ in that interval, allowing one endpoint error at each power.

Let $q_{\mathrm{cyc}} = \varphi(\mu_{\mathrm{cyc}})/\mu_{\mathrm{cyc}}$ and $u_i' = \lceil s_i'/(\ell-1)\rceil$. Since $q_{\mathrm{cyc}} = \mu_{\mathrm{cyc}}^{\ell-1} + \ell B$ with B integral, the binomial expansion shows $\mu_{\mathrm{cyc}}^{u_i'} q_{\mathrm{cyc}}^{u_i'+c_H+2} P_i^{\mathrm{tar}} \subset P_i^{\mathrm{im}}$. Indeed a term with exponent v on $\mu_{\mathrm{cyc}}^{\ell-1}$ has degree $a = (\ell-1)v + u_i'$. When $a < s_i'$, $(s_i'-a)\ell/(\ell-1)^2 \le u_i' - v$, which is within its available ℓ-exponent after paying c_H. Otherwise the definition of s_i' applies. If the labelled jumps at i are in $[-M_i, h_i]$, the two Frobenius height inclusions with $q_{\mathrm{cyc}}^{M_i}$ and $q_{\mathrm{cyc}}^{h_i}$ now yield

$$s_{i-1}' \le M_i + h_i + c_H + 2 + u_i'.$$

For this descent the height inclusions follow by testing poles at the first cyclotomic roots. Quotients by powers of their monic polynomial are coefficient-free, so the bounds descend integrally. The linearized Frobenius is injective and the parameter ring is free over its Frobenius image with basis $1, \mu_{\mathrm{cyc}}, \dots, \mu_{\mathrm{cyc}}^{\ell-1}$. Thus membership of a Frobenius image in the span of the previous image descends to membership before Frobenius.

In cut notation the recurrence uses the width $M_{c+1} + h_{c+1}$ when descending to c. Iteration around the orbit, retaining H_0 as a second upper bound, gives the minimum of H_0 and the integer part of the geometric sum with ratio $1/(p-1)$ of the successive widths plus $c_H + 3$. At a cut the first j widths are at most $2G_0$. Inserting the resulting bound in the γ-descent at degree zero gives one common exponent r_0 for the two maps under consideration, with

$$r_0 \ll \frac{h}{(p-1)^{j+1}} + \frac{g}{p} + 1 + \log_p h.$$

For the mixed-sign margin, stop the recurrence after the first j thin widths and use $s_{c+j}' \le H_0 \le h + g$. The h-term in the bound on s_c' is at most $h/(p-1)^j$. The γ-descent multiplies it by $p/(p-1)^2$, so the coefficient of h in r_0 is at most $p/(p-1)^{j+2}$. This descends to the tame invariant modules by averaging a preimage over the tame group. The same exponent lifts constant residues modulo its parameter.

Truncation on an arc. Use the tame invariant parameter

$$X = \frac{1}{p-1}\sum_{\delta\in\Delta} \delta(q_{\mathrm{cyc}}) - p, \qquad P_s = \varphi^s(X+p), \qquad R = O[[X]].$$

The tame average of μ_{cyc} against the inverse Teichmüller character is a parameter transforming by that character. Its $(p-1)$-st power generates the invariants, and X has that initial degree modulo p. The same averaging lifts an invariant Wach basis from modulo μ_{cyc}. The first cyclotomic roots are zeros of $X+p$, and monic division makes $X+p$ associated to q_{cyc} in the full parameter ring. Thus $\varphi^a X$ is associated to $X\prod_{s<a} P_s^{p-1}$. Each P_s has constant term p, is congruent to X^{p^s} mod p, and is associated to a monic distinguished polynomial of degree p^s. Use γ now for a generator of the pro-p cyclotomic group and put $u = \chi(\gamma)$, with $v_p(u-1) = 1$. It preserves the divisors and $\gamma(X)/X \equiv u^{p-1} \bmod X$.

Set

$$M = \left\lfloor \frac{h+2g}{p-1} \right\rfloor + 1, \qquad \sigma = \left\lceil \left(3g + \frac{M}{p^{j-1}}\right)\frac{p}{(p-1)^2} + 1 \right\rceil, \qquad J_j = X^M \prod_{s<j} P_s^{3g}.$$

For $b \le pk/(p-1)$, $p^\sigma \varphi^j(X)^k \in (p^b, J_j)$. For $b > \sigma$ put $e = \lceil M/p^{j-1}\rceil$. The choice of σ gives $(p-1)k - 3g \ge b - \sigma + e$. Retain $\prod_{s<j} P_s^{3g}$ in the factorization of $\varphi^j(X)^k$, and expand the remaining power of $P_{j-1} \in (p, X^{p^{j-1}})$. Each term of its power $b-\sigma+e$ either has p-exponent at least $b-\sigma$ or X-exponent at least $p^{j-1}e \ge M$. This proves the congruence. Applying φ^{a-j} gives the same estimate modulo $J_a' = \varphi^{a-j}(J_j)$ on an arc of length a.

The horizontal formal frame for γ at the preceding cut is specified by its constant vector at $X = 0$. The equation in degree k divides by $u^{(p-1)k} - 1$. Thus the k-th coefficient of the frame and of its inverse has denominator at most $p^{k+v_p(k!)}$. The same estimate solves $(\gamma-1)v = q$ for $q \in XN$ with zero prescribed constant. Since $k + v_p(k!) \le pk/(p-1)$, the congruence makes each transported frame and solution integral modulo its transported ideal after scaling by p^σ. For all sufficiently large k, $(p-1)k - 3g \ge 2k - \sigma + \lceil M/p^{j-1}\rceil$

also holds. Using $b = 2k$ in the same congruence proof shows that the valuations of the truncated terms tend to infinity after division by $p^{k+v_p(k!)}$. The quotients by those ideals are coefficient-free by monic division, so these comparisons are compatible with the lattice maps.

Index the levels $s = 0, \dots, a-1$ backwards from the arc endpoint c. The rational completion at P_s uses the transported field at label τ_{c-s}. Its logarithmic parameter t_s is scaled by u under γ, and the coefficient action has order dividing p^s after tame invariants. The endpoint lattice $N_{i,s}$ for source or target i has coefficient of t_s^k in Fil^{-k} of the transported constant space. In the effective range the filtered de Rham comparison gives the containing lattice for the Wach image, since the Wach lattice is contained in the regular representation lattice after Witt base change. The cyclotomic elementary divisors of Berger's rational Wach comparison have orders equal to the de Rham jumps at every positive cyclotomic level and order zero at $\mu_{\mathrm{cyc}} = 0$ [71, Propositions II.2.1, II.3.1 and III.2.1]. After splitting coefficients the containing lattice and the image have equal total colength, so they agree at each factor. Apply this with the label at level s read by φ^{-s-1} and untwist the filtered comparison.

At the preceding cut, Berger's inclusion of D_{cris} in the analytic Wach module and its specialization at the origin give the crystalline frame with prescribed reduction modulo X [71, Proposition II.2.1 and Corollary III.4.5]. For signed jumps apply the comparison after a common Tate twist. Removing the twist changes the crystalline-to-Wach period ratio by a power of t/μ_{cyc}, a unit at the origin. The Taylor expansion of the normalized crystalline frame is the unique formal γ-fixed frame with that constant term. Transport by φ^a transports the coefficients and substitutes $\varphi^a X$ for X. For $s < a$ the substituted parameter has order $p-1$ at P_s. The change of frame and its inverse are therefore regular at each such divisor. Their values at its roots are the transported constant terms, and their higher jets use the whole substituted expansion.

Round the jumps to 0 at thin levels, to $b_s, 0$ at up levels, and to $0, -b_s$ at down levels. The b_s line of the rounded filtration is the subspace Fil^{b_s} for an up level. The zero part at a down level is $\mathrm{Fil}^{-\eta_*}$. Denote their saturated constant lattices by U_s. The rounded local lattice $B_{i,s}$ differs from $N_{i,s}$ by multiplication by P_s^g in either direction. Put $\Lambda = \prod_s P_s^{H_s}$ with $H_s = g + g_*, g + b_s + S_*, g + \eta_*$ at thin, up, and down levels respectively. The arc lattice W_i is $\Lambda^{-1}\,\mathrm{span}_R(\varphi^a N_{i,c-a})$. Its constant partner has the same Λ^{-1} factor on the free constant-lattice series.

The preceding truncation gives partners in both directions, each with cost p^σ and the same local errors of order at least $\mathrm{ord}_{P_s} J_a' - H_s$. At the first $a-j$ levels the first term is $(p-1)M$; at the last j thin levels it is $3g$. These orders are at least the lattices' required orders, namely η_* at an up level, $b_s + S_*$ at a down level, and g_* at a thin level. Thus every partner error belongs to $N_{i,s} \cap B_{i,s}$. An error on the constant side forces divisibility by $P_s^{H_s}$ before dividing the series by Λ. The partners commute with the lattice maps modulo these same errors, since maps between horizontal frames are constant in the initial parameter. Their constant map F_0 is precisely reduction of the Wach map at the preceding cut modulo X.

Write $\Theta = \prod_s P_s^g$. The linearized Frobenius matrices and inverses have denominators only at the first divisor, with pole orders given by the labelled jumps. Iterating gives exactly the height tests at the P_s. Weierstrass saturation shows that $\Theta^{-1} N_{i,c} \subset W_i$ consists of the elements whose localizations lie in $P_s^{-g} N_{i,s}$ for all s. On this submodule the map lifts p^{r_0} times any target by the cut bound. Membership in $XN_{i,c}$ is tested by membership in XW_i and local membership in all $N_{i,s}$, since X is a unit at the rational cyclotomic divisors. Finally $(\gamma - 1)W_i \subset XW_i$ by the Wach action on the previous cut and the constant terms of the divisor multipliers.

A residue fixed modulo $XN_{2,c}$ has localizations in constants plus $N_{2,s}$. Apply γ^{p^s} to each nonzero degree to see that the coefficient in its filtration quotient is zero. In degree zero average the finite coefficient action to take a constant. To pass from $P_s^{-g} B_{2,s}$ back to $B_{2,s}$, kill the nonzero degrees in the rounded boundary windows. For the reverse comparison use the union of the intervals $[-w-g, -w-1]$ for actual jumps w of the source. Each union has one central interval and at most two translated intervals of total width $O(g+1)$. For an arc of length a, the polynomial is the product of $Z^{p^a} - u^{kp^a}$ over the indicated nonzero degrees k, times $\sum_{j_0=0}^{p^{a-1}} Z^{j_0}$ to make degree zero constant. It kills the desired degrees since the coefficient

action has order dividing $p^s < p^a$. Its valuation at one is $a + \sum_k (a + 1 + v_p(k))$. Counting multiples of successive powers of p in each interval gives a common bound $\kappa_{\mathrm{arc}} = O((g+1)(a_*+1) + \log_p(2h))$ for the two polynomials P, P'.

Integral constants on the signed levels. Let L be one of the free constant lattices and write $L_R = L \otimes_O R$. In these coordinates, before division by Λ, the rounded condition $\mathcal{N}$ is membership at each s in

$$\begin{array}{ll} P_s^{H_s} L_R & \text{at a thin level,} \\ P_s^{H_s - b_s}(U_s)_R + P_s^{H_s} L_R & \text{at an up level,} \\ P_s^{H_s}(U_s)_R + P_s^{H_s + b_s} L_R & \text{at a down level.} \end{array}$$

The condition $\mathcal{I}$ allows subtraction of Λd_s with $d_s \in L[1/p]$ at each level. Divisibility over $R[1/p]$ by products of the distinct distinguished powers descends to R by monic division, also after projecting to L/U_s.

Without a down level the conditions coincide. With one down level and no up level, write $y = \Lambda z$. The class of z modulo $P_s^{b_s}$ in L/U_s is a constant, integral because 1 occurs in the integral remainder basis. Lift it to $\lambda \in L$. Subtracting $\Lambda\lambda$ puts y in $\mathcal{N}$. With two down levels $i = 0, 1$ put $g_i = P_i^{b_i}$ for their two powers, relabelling these two indices. The quotient $R/(g_0, g_1)$ is finite torsion over O. If it is killed minimally by ϖ^v for a uniformizer ϖ, the constant 1 has order ϖ^v. A largest Smith-length coordinate of 1 is primitive, so its constant submodule O/ϖ^v splits. The two integral constants in L/U_i therefore agree modulo $U_0 + U_1 + \varpi^v L$. Take lifts $\lambda, \lambda + \varpi^v \lambda'$ in L, respectively for the second and first constants. A scalar series B congruent to zero modulo g_1 and ϖ^v mod g_0 exists by the definition of v. Subtracting $\Lambda(\lambda + B\lambda')$ suffices. The coefficients of this scalar series are integral, so only the two ordinary vectors have to be lifted.

For an up level u and a down level d, put $g_u = P_u^{b_u}$, $g_d = P_d^{b_d}$, and write $y = (\Lambda/g_u)z$. Then $z \in (U_u)_R + g_u L_R$, and $q_d z \equiv g_u d_0 \bmod g_d$ over the fraction field, where $q_d : L \to L/U_d$. Write the remainder of g_u modulo g_d as $\varpi^v \alpha$ with α primitive. Applying an integral functional taking α to one shows that the remainder of $q_d z$ is αv_0 with $v_0 \in q_d(U_u) + \varpi^v(L/U_d)$. Consequently one can remove $\Lambda\lambda + (\Lambda/g_u)\alpha\lambda'$, where $\lambda \in L$ and $\lambda' \in U_u$. If $u < d$, then $b_u < p b_d$ gives $\deg(P_u^{b_u}) < \deg(P_d^{b_d})$. The monic remainder is primitive and one can take $\lambda' = 0$. The order $d < u$ can require that filtered constant.

Put $A_0 = 2\sigma + r_0 + 2g a_*$. For $d < u$ we need $2b_d \le p b_u$ and $2A_0 < b_u$. To lift the filtered constant, let $v_2 \in U_{2,u}$ and take $y_2 = (\Lambda/g_u) g_d v_2$ in $\mathcal{N}_2$. Its arc partner belongs to $\Theta^{-1} N_{2,c}$ by the level tests. Lift p^{r_0} times that partner and return to the constant model. The two partner factors are $p^{2\sigma}$. Their local errors force divisibility by Λ on the constant side. Thus there is y_1 with $\Theta^2 y_1 \in \mathcal{N}_1$ and $F_0 y_1 \equiv p^{2\sigma + r_0} y_2 \bmod \Lambda$. Divide $\Theta^2 y_1$ by Λ/g_u and reduce modulo g_u. The residue belongs to the source up filtration and maps to $D v_2$, with $D = p^{2\sigma + r_0} \Theta^2 g_d \bmod g_u$.

For its scalar content use Gauss valuation assigning X the value θ/p^u for $0 < \theta < 1$. The values of P_u and P_d are θ and θp^{d-u}, by the leading terms of their distinguished polynomials. Every P_s has value at most one. With $A' = 2\sigma + r_0 + 2ga$, choose $A'/(b_u - b_d/p) < \theta < 1$. This gives

$$v_\theta(p^{2\sigma + r_0} \Theta^2 g_d) \le A' + \theta b_d/p < \theta b_u = v_\theta(g_u).$$

If the remainder has content ϖ^v, the Gauss inequality for an expression in the ideal (ϖ^v, g_u) gives $v/v_\varpi(p) \le A' + \theta b_d/p$. Letting θ decrease to the lower endpoint bounds the content by $2A_0$ in normalized units. A linear functional on the integral remainder module taking its primitive part to one now gives $p^{2A_0} U_{2,u} \subset F_0(U_{1,u})$. This also lifts the required filtered constant when the coefficient valuation is ramified.

Together the cases show that for any $y_2 \in \mathcal{I}_2$ there is $y_1 \in \mathcal{I}_1$ with $p^{2A_0} y_2 - F_0 y_1 \in \mathcal{N}_2$, using the same integral scalar series for the lifted constants. Start with an invariant target residue w mod $X N_{2,c}$. Apply the first polynomial $P(\gamma)$ so its localizations belong to constants plus $B_{2,s}$, and take its constant partner. Lift its constants as above and return to an arc partner z_1. The discrepancy $p^{2\sigma + 2A_0} P(\gamma) w - f z_1$ belongs

locally to $B_{2,s}$, by the two partner comparisons. Lift it in $\Theta^{-1}N_{1,c}$ with cost p^{r_0} and add the lift to $p^{r_0}z_1$. This gives an exact arc lift whose localizations are in constants plus $P_s^{-g}N_{1,s}$. The second polynomial $P'(\gamma)$ takes these into constants plus $N_{1,s}$. Its γ-difference thus belongs locally to $N_{1,s}$, and globally to XW_1, so it lies in $XN_{1,c}$. Modulo $XN_{2,c}$ the image of the resulting invariant residue is

$$p^{2\sigma+2A_0+r_0}P(1)P'(1)\,w.$$

The valuation cost is at most $2\sigma + 2A_0 + r_0 + 2\kappa_{\mathrm{arc}}$.

The finite classes represented by the cut residues. For a finite p-primary representation of the local Galois group, Herr's theorem computes the Galois cohomology by the three-term complex of its étale (φ,Γ)-module [73, Theorem 2.1, p. 572]. Apply it to the reductions of the integral étale modules. Passing to inverse limits computes the integral cohomology, since the finite cohomology groups satisfy the Mittag–Leffler condition. The tame group of order $p-1$ is removed by exact averaging, leaving the pro-p generator used above. Use the integral Herr cocycle convention $d_\gamma y = d_\varphi z$, with $(d_\varphi v)_i = \varphi(v_{i-1}) - v_i$ and $d_\gamma = \gamma - 1$. On the direct sum of the $X^M N_i$ the operator d_φ is invertible by its negative geometric series. One Frobenius step multiplies by $P_0^{(p-1)M}$ in the parameter part, which clears the possible pole of order h and leaves a factor in (p, X). Hence if $d_\gamma y$ is in this direct sum there is a unique second cochain there. Changing y by a member of the same direct sum changes the pair by a coboundary.

Such a pair is finite when the first cochain has only finitely many cyclotomic denominators. To verify this, write the associated integral étale extension with quotient vector e fixed by the tame group. If the pole order of the first cochain at P_s is at most m_s, choose nonnegative integers c_s eventually zero with $c_{s-1} - c_s \geq m_s$ for $s \geq 1$. Put $D = \mu_{\mathrm{cyc}}^k \prod_s P_s^{c_s}$ with $k \geq c_0 + m_0$, large enough for $\mu_{\mathrm{cyc}}^k N$ to be Frobenius-stable. In the enlarged full-parameter free lattice with added generator De, the quotient multiplier $\varphi(D)/D$ is integral and the off-diagonal entry $\varphi(D)y/\mu_{\mathrm{cyc}}^k$ is integral. The cyclotomic action has scalar χ^k modulo μ_{cyc} because the second cochain is in $\mu_{\mathrm{cyc}}N$. This lattice generates the étale module, since its cyclotomic factors are units in the étale period ring. To see that it lies in the positive period module, write the frame in $W(C_p^\flat) \otimes_{\mathbb{Z}_p} T_{\mathrm{ext}}$ as the matrix U, where T_{ext} is the extension lattice. Then $\varphi(U) = UA$ for an integral Frobenius matrix A over A_K^+. If the least valuation of the entries of U mod p were $\alpha < 0$, the left side would have least valuation $p\alpha$ and the right side at least α, a contradiction. After the first m Witt digits are integral, subtract their integral lift and divide by p^m. The resulting relation has an extra integral term, so the same minimum argument proves integrality of the next digit. This puts the frame in the positive étale module. A rational finite-height representation is crystalline when its positive module contains a full-rank free cyclotomic-stable lattice on which Γ acts modulo μ_{cyc} by a power of the cyclotomic character [71, Proposition II.1.4]. The lattice constructed here has these properties with exponent k, so its integral extension class is finite.

At a cut c, take an invariant residue w with $q = d_\gamma w \in XN_c$. Solve $(\gamma-1)v = q$ in the horizontal formal frame with zero constant, and take an integral approximation $\tau(q)$ to $p^\sigma \varphi^j(v)$ modulo J_j in the Frobenius pull span. Its class modulo that ideal is unique by saturation and depends only on q mod $X^M N_c$. The first j steps are thin, so J_j times this span lies in $X^M N_{c+j}$. The first cochain supported at $c+j$ with value $p^\sigma \varphi^j(w) - \tau(q)$ has cyclotomic difference in $X^M N_{c+j}$. Solve for its second cochain as above. This defines a finite class $\mathcal{R}_c(w)$ functorially, since changing the integral approximation adds a coboundary. It kills XN_c, by taking $v = w$ in that case.

Represent a target class ξ by a crystalline lattice extension of the trivial quotient. At each cut c choose x_c mapping to p^{r_0} in the quotient Wach lattice. The vector $q_c = (\gamma-1)x_c$ is divisible by X in the extension's Wach lattice, by tame invariance and the cyclotomic congruence. Its quotient by X projects to zero in the étale quotient. Kernel exactness therefore gives $q_c \in XN_{2,c}$. The vector $R_c = p^\sigma \varphi^j(x_c) - \tau(q_c)$ lifts $p^{r_0+\sigma}$ at $c+j$. Propagate each R_c by Frobenius to the next shifted cut. These quotient lifts give a pair representing $p^{r_0+\sigma}\xi$, with second cochain in $X^M N_2$. For the arc from c to $d = c + a$, put $w_d = \varphi^{a-j}(R_c) - p^\sigma x_d$. The

quotient cancels. The height test and exactness of the Frobenius pull spans place w_d in the target arc lattice. Its γ-difference is congruent to $-p^{\sigma} q_d \bmod X^M N_{2,d}$. The first difference of the propagated quotient lifts at $d + j$ is $\varphi^j w_d + \tau(q_d)$. By the functorial truncation of τ the first cochain of $\mathcal{R}_d(w_d)$ is p^{σ} times that difference modulo $X^M N_2$. Second cochains there are uniquely determined, so

$$\sum_d \mathcal{R}_d(w_d) = p^{r_0+2\sigma}\xi.$$

The simultaneous arc lifts and functoriality of $\mathcal{R}_d$ therefore lift this class after the common additional power $p^{2\sigma+2A_0+r_0+2\kappa_{\mathrm{arc}}}$. The resulting annihilating exponent for the cokernel is

$$e = 2r_0 + 4\sigma + 2A_0 + 2\kappa_{\mathrm{arc}}.$$

The common factor is used on every summand of the same class decomposition, so e has no dependence on the number of cuts.

The two mixed sign inequalities hold for the actual clustered windows. Indeed $b_u, b_d \in [b_*, b_* + \eta_*]$ and $b_* \geq C(j+1)(g_*+1)$ imply $b_u < p b_d$ and $2b_d \leq p b_u$. For $2A_0 < b_u$, use $j \geq 1$ and $H_0 \leq h + g$. The h-term in 4σ is bounded by $4hp/(p-1)^3$. The h-term in $2r_0$ is bounded by $2hp/(p-1)^3$ by the stopped recurrence. Their sum is at most $6hp/(p-1)^3 \leq 0.066h$, since $p \geq 11$. The rest is $O((g+1)(j+1)+\log_p(2h))$; the assumed separation and $h \leq b_* + \eta_* + S_*$ make it smaller than the remaining margin for large enough absolute C. Finally $a_* \leq 3j + 1$ gives

$$e \ll h/(p-1)^{j+1} + (g_*+1)(j+1) + \log_p h.$$

When $H_{\mathrm{loc}} = 0$, the change to $h = 1$ is absorbed by $(g_*+1)(j+1)$. The finite-group generator bound at the start now proves (16.12). [64, Lemma 3.8.1 and Corollary 3.8.4]

Application at the small unramified primes. Let $\ell \leq Q = \kappa t$ and $\ell \nmid N_1 d$. Here $E_v/\mathbf{Q}_\ell$ is unramified, $\ell \geq 13$, and the catalogue at the fixed place $v \mid \ell$ is crystalline. Write $n_v = f_v = [E_v : \mathbf{Q}_\ell]$. The coefficient ring contains the integer ring of E_v. The component at q has $1 \otimes \zeta_O = q \otimes 1$, so its left-factor label $\tau_i : E_v \hookrightarrow O[1/\ell]$ satisfies $\tau_i(q) = \zeta_O$. For $q = F(g)$, this identifies $\tau_i = g^{-1}|_{E_v}$, the local embedding in $\sigma_g = \iota_\ell^{-1} g^{-1} e$. The ℓ-power orbit of a starting letter runs through all these labels. As q advances to q^ℓ, left Frobenius σ_{E_v} gives $\tau_{i+1} \circ \sigma_{E_v} = \tau_i$.

Use the signed cluster lifting argument with

$$j_* = J_\ell = \lfloor \log_\ell(H^3) \rfloor, \qquad b_* = \widetilde{B}, \quad w_* = \widetilde{G}, \quad S_* = W, \quad \eta_* = \eta.$$

The chosen centers $\widetilde{B} = b_1 + \delta$ and $\widetilde{G} = b_2 + \delta$ are integral. Each coefficient embedding has large jumps in one allowed signed interval at its own center. At the thick labels this center is b_*, and elsewhere it is $w_* < b_*$. Put $g_* = w_* + S_* + \eta_*$. The scale inequalities give $g_* = O(G)$ and $b_* \geq C(j_*+1)(g_*+1)$ for any fixed constant C from the cluster lifting argument and all sufficiently large x.

We have $j_* \geq 1$, since $\ell \leq Q \leq H^3$, and $j_* < n_v$. For the second assertion, all N_1-th roots are prime to ℓ here and lie in E_v, so $\ell^{n_v} \equiv 1 \bmod 2^{s+1}$. The valuation calculation in the spacing argument gives $2^s \leq n_v(\ell+1)$. For large x this implies $n_v > 12s \geq j_*$, since $2^s \geq x^{11}$ and $Q = O(x^{10})$. If a thick label has a thick predecessor at distance $a \in [1, j_*]$, the spacing relation excludes a second thick predecessor at a smaller distance. It also excludes every thick label in the j_* positions immediately before that predecessor. These are exactly the cyclic predecessor conditions of the cluster lifting lemma. Either sign and either order of two different signs in a cluster are allowed by its proof.

Apply the lifting to the coefficient projection $\mathcal{G}_f \twoheadrightarrow Y$ of the integral graph construction. Both lattices are crystalline, because $\mathcal{G}_f[1/\ell]$ is the direct sum of $L[1/\ell]$ and $Y[1/\ell]$. Its allowable windows are the same windows for every member of the catalogue, including when L is one of the Tate blocks. The theorem uses the decreasing-filtration jump convention in which the cyclotomic twist has jump -1. It can also be

applied to the dual graph projection. For that orientation take $b_*^\dagger = \widetilde{B} - 1$, $w_*^\dagger = \widetilde{G} + 1$, $S_*^\dagger = W$, and $\eta_*^\dagger = \eta + 1$. The actual far magnitudes before reflection lie in $[\beta + t, \beta + 3t]$. After $u \mapsto -u - 1$ they lie in $[\beta + t - 1, \beta + 3t + 1]$, which is contained in both $[\beta - 1, \beta - 1 + W]$ and $[\beta + 1, \beta + 1 + W]$ because $t \geq 2$ and $W = 4t$. These choices preserve the thick set, the predecessor conditions, and the separation for sufficiently large x.

The cluster lifting calculation gives an annihilating exponent for the cokernel on integral H^1_f, uniformly in the ranks of the two lattices. Its bound is

$$k_{\mathrm{cl}} \ll \frac{H_{\mathrm{loc}}}{(\ell - 1)^{j_* + 1}} + (g_* + 1)(j_* + 1) + \log_\ell \max(1, H_{\mathrm{loc}}),$$

where $H_{\mathrm{loc}} = b_* + \eta_* + S_*$ if a thick label occurs, and $H_{\mathrm{loc}} = g_*$ otherwise. The same formula holds with the parameters for the dual orientation. It is the annihilator in the proof that is used here, before conversion to a length bound. Since $\ell - 1 \geq \sqrt{\ell}$ and $\ell^{j_* + 1} > H^3$, the denominator in the first term exceeds $H^{3/2}$. Thus the first term is $o(B/Q^2)$ uniformly, since $H_{\mathrm{loc}} = O(B)$ and $H^{3/2} \gg Q^2$. The second term is $O(G \log x) = o(B/Q^2)$ by the thin-to-thick scale ratio. The same is true of the logarithmic term, which is $O(\log b)$. Hence for all sufficiently large x, the exponent k_{cl} is at most $K_\ell = \lfloor 2\lambda_\ell / Q \rfloor$, using $K_\ell \geq B/Q^2$.

The graph cup calculation gives annihilation by ℓ^{K_ℓ} for every pairing in the one-place defect, including its arbitrary integral coefficient functional. In particular the rank of Y plays no role in the final length. The normalized defect for the tested block of rank R is at most $2n_v R K_\ell$. This is at most $n_v R \, \mathrm{loss}_{\mathrm{crys}}(\ell)/t$, since $2K_\ell \leq 4\lambda_\ell/(\kappa t)$. This proves the required crystalline bound at the small unramified places on the same row and Tate lattices.

16.2.6 The Kummer chart cup calculation

Placement of cuts for the sparse-letter estimate. Write $\varphi(q) = q^\ell$ for letter Frobenius and $J = J_\ell$. A cut is a letter whose next J successors are thin. Fix a letter q. If $\varphi^{-(J+1)}q$ is a cut, it lies in the required previous $2J + 1$ positions. Otherwise some $\varphi^{-a}q$, $1 \leq a \leq J$, is thick. If it is a cut, use it. Otherwise a thick successor occurs within J steps of $\varphi^{-a}q$. By the spacing condition, $\varphi^{-(a+J+1)}q$ is then a cut. Indeed any thick successor of this last letter within J steps would lie before $\varphi^{-a}q$ at distance at most J, giving three thick letters with successive gaps at most J. This proves existence of a cut in the specified distance.

Let $\varphi^{-l}q$ be the nearest strictly preceding cut. If $\varphi^{-a}q$ is thick for some $0 < a < l$, it has a thick successor within J steps, since it is not a cut. The preceding argument gives a cut at $\varphi^{-(a+J+1)}q$ and thus $l \leq a + J + 1$. On the other hand $l - a > J$, since $\varphi^{-l}q$ has only thin successors for J steps. Thus $a = l - J - 1$. The only possible thick positions after the cut, counted backwards from q, are 0 and $l - J - 1$. If at most one thick position occurs or their signs agree, all far jumps between the cut and q outside the enlarged low window have a common sign. Thin jumps are absorbed by the enlarged low window.

At the conductor-dividing places the thick sign is common. Mixed-sign clusters at small unramified primes use the signed Wach bound. In the large-prime calculation with two sizes the mixed-sign case uses the two-letter lift under its own size bounds.

The chart and the integral pairing. Fix a completion $K = E_v$, put $p = \ell$, and write e_K, f_K for its ramification and residue degrees. Denote the completion of $\overline{\mathbf{Q}}_p$ by C_p, its tilt by F, and put $A = W(\mathcal{O}_F)$, $W_F = W(F)$. Choose a uniformizer π of K and compatible p-power roots π_a, and write $z = [\pi^\flat]$, where $\pi^\flat$ is the element of the tilt formed by this tower. If $e_K = 1$ take $\pi = -p$. Put $S = W(\mathbf{F}_{p^{f_K}})[[z]] \subset A$. The homomorphism $\theta : A \to \mathcal{O}_{C_p}$ maps z to π. Scale the monic Eisenstein polynomial by a unit of the coefficient ring so that its constant coefficient is p. Write the result as $E_0(z) = pU_0(z) + a_e z^{e_K}$, with $U_0(0) = 1$ and a_e a unit, and set $E = E_0/U_0$. Then $E = p + z^{e_K} u(z)$ for a unit $u(z)$ in S. The kernel of θ is generated by E. Indeed its reduction mod p has tilt valuation 1, the valuation of the reduction of a generator of $\ker \theta$. Hence the quotient of the two generators is a unit of A. In particular $E = z + p$ in the

unramified chart. Write $E_a = \varphi^a E$ and $\theta_a = \theta \circ \varphi^{-a}$. Then the E_a generate distinct prime ideals, since π_a has valuation $1/(e_K p^a)$. The fixer of $\theta(z)$ and of the Witt coefficient ring is $\mathrm{Gal}(\overline{K}/K)$. For each integer j there is a g_j in this group with $g_j z = [\varepsilon]^j z$, where $\varepsilon = (1, \zeta_p, \zeta_{p^2}, \ldots)$ is a compatible cyclotomic system. At each finite stage the Eisenstein polynomial $X^{p^n} - \pi$ is irreducible over K, so some automorphism sends π_n to $\zeta_{p^n}^j \pi_n$. The resulting closed sets of automorphisms are nested and their intersection is nonempty by compactness.

We use $\tau = \varphi^{-1}([\varepsilon] - 1)$ and $\Delta = z\tau$. Thus $E\tau$ is associated to $[\varepsilon] - 1$. For all $g \in G_K$ the difference $gz - z$ is a multiple of ΔE. In fact $gz = [\varepsilon]^u z$ for some $u \in \mathbf{Z}_p$, and $[\varepsilon]^u - 1$ is divisible by $[\varepsilon] - 1$ by the binomial expansion in the complete period ring. It follows that $gs - s$ is in $\Delta E A$ for $s \in S$. The product Δ is a unit of W_F and its image is a unit at each de Rham divisor θ_a with $a \geq 0$. For $0 \leq a < q$, $E_a \mid \varphi^q \tau$, by factoring successive differences $[\varepsilon]^{p^j} - 1$. These are the scalar chart properties used below. The Frobenius permutation $\sigma : i \mapsto i^p$ on the nonzero residue letters and the change of complex embedding are as in the local chart calculation.

Let Y range over the crystalline catalogue $\mathcal{Q}$ and write $L_2 = L^\vee(1)$ for the tested dual block. At an arbitrary modulus p^M take an O-linear form $\psi : Y \otimes_O L_2 \to O(1)$ whose coboundary divided by p^M is an integral finite cocycle c_3. Every equivariant map $Y \to L/p^M L$ gives such a form by lifting linearly, since its divided coboundary is rationally a coboundary. Let $h : O \to \mathbf{Z}_p$ be any $\mathbf{Z}_p$-linear functional. The pairing used in the cup is $b = h \circ \psi \bmod p^M$, with values in $T_M = \mathbf{Z}/p^M(1)$. It is equivariant. This includes the scaled block maps and the inverse-different trace functional of the congruence argument. The zero modulus exponent gives a zero pairing.

Use the convention $d_g V = (g-1)V$ for an ordinary vector. The finite cocycles on Y, L_2 give extensions of their trivial parameter lattices by Y, L_2. We write $\widetilde{Y}$ and $\widetilde{L_2}$ for these extensions, using a module splitting to write their action as $(v, c) \mapsto (gv + a_c(g), c)$ and $(v, c') \mapsto (gv + a'_{c'}(g), c')$. One can take the parameter lattices to be the entire finite cocycle modules. The extensions are crystalline, since evaluating the finite cocycles in a crystalline period ring splits them there. The third extension $\widetilde{T_3}$ is the extension of O by $(Y \otimes_O L_2)^\vee(1)$ for c_3. Its equivariant pairing with $Y \otimes_O L_2$, with values in $O(1)$, is

$$((d, c), y \otimes v) \longmapsto p^M d(y \otimes v) + c\psi(y \otimes v).$$

Substitution of $p^M c_3(g) = g\psi - \psi$ verifies equivariance. The allowed third jumps at a letter comprise 0 and $-u - v - 1$ for the corresponding allowed jumps u, v on the first two sides. In each case the zero jump includes the trivial parameter. This construction uses a lattice extension even when the tested cup pairing has finite coefficients.

A Witt-vector cochain for cup annihilation. Let $x, w \in W_F \otimes \widetilde{Y}$ project to $1 \otimes p^{k_1} c$, and let $y \in W_F \otimes \widetilde{L_2}$ project to $1 \otimes p^{k_2} c'$. Call a pair of vectors *safe* if the pairing of their first components, via the W_F-bilinear extension B of $p^{k_3} b$, is in $N = W_M(\mathfrak{m}_F)(1)$ inside $D = W_M(F)(1)$. Here M denotes the length of truncated Witt vectors, $\mathfrak{m}_F$ is the tilt maximal ideal, and $W_M(\mathfrak{m}_F)$ consists of vectors with every coordinate in $\mathfrak{m}_F$. Suppose that for all $g, h \in G_K$ the four pairs

$$(gx - x,\ g(hy - y)), \qquad (w - x,\ gy - y),$$
$$(gw - w,\ g(\varphi y - y)), \qquad (\varphi x - w,\ g\varphi y - \varphi y)$$

are safe. This implies that the cup of a_c and $p^{k_1+k_2+k_3} a'_{c'}$ is zero.

Here is the cochain calculation. Write X, W, Y_0 for the first components of x, w, y, and $A_g = p^{k_1} a_c(g)$, $C_g = p^{k_2} a'_{c'}(g)$. The first components of their Galois differences are $\delta_g X = (g-1)X + A_g$, $\delta_g W = (g-1)W + A_g$, and $\delta_g Y_0 = (g-1)Y_0 + C_g$. Frobenius acts on the scalar factor. Set $T = \varphi Y_0 - Y_0$ and choose $q \in D$ with $(\varphi - 1)q = B(W, T)$. Surjectivity of $\varphi - 1$ on D follows from the Artin–Schreier equation in the algebraically closed tilt, then induction on the Witt length. Define

$$d_g = B(A_g, gY_0) - B(X, \delta_g Y_0) + (g-1)q.$$

Since φ commutes with Galois and fixes the constant coefficient vectors, expansion gives

$$(\varphi - 1)d_g = B(\delta_g W, gT) - B(W - X, \delta_g Y_0) - B(\varphi X - W, \delta_g(\varphi Y_0)).$$

All three terms are in N. Also the inhomogeneous cochain differential is

$$g d_h - d_{gh} + d_g = B(A_g, gC_h) - B(\delta_g X, g\delta_h Y_0),$$

by the two cocycle laws. The last term is in N. The fixed elements of D/N under φ are precisely T_M, embedded by constant Witt vectors. Indeed $\varphi - 1$ is surjective on N by the convergent sum $-\sum_{j\geq 0} \varphi^j$ there, the fixed elements in D are the Witt vectors over $\mathbf{F}_p$ with the Tate twist, and these intersect N trivially. Thus d_g mod N is a T_M-valued cochain with differential equal to the claimed cup. The Witt vectors of length M carry the product of the valuation topologies of their coordinates. Their Galois action is continuous and N is open, so the reduced cochain is continuous. If the chart is taken over a subgroup of index prime to p, restriction followed by corestriction multiplies the finite cup class by that index, which is invertible on $H^2(G_K, T_M)$. For the chart above the subgroup is G_K. Finally $T_M \simeq \mu_{p^M}$ maps to $\overline{K}^\times$, giving the required continuous multiplicative coboundary. In the other direction the map on H^2 is injective by the Kummer sequence and Hilbert 90.

The integral period modules used in the cut calculation. Keep the Kummer chart and work on underlying $\mathbf{Z}_p$-lattices. For a crystalline lattice T let $\mathcal{E}_T$ be its prismatic F-crystal and let $M_T = \mathcal{E}_T(S, (E))$, a finite free S-module. Use $P_T = A \otimes_S M_T$ inside $W_F \otimes T$, via the map of prisms $(S, (E)) \to (A, (E))$. Write S_E and A_E for the rings with all E_a, $a \geq 0$, inverted. The localized modules for the points and their exponent tests are $\mathcal{N}_T = S_E \otimes_S M_T$ and $A_E \otimes_A P_T$. For a complete discretely valued K with perfect residue field, the étale realization of a finite locally free prismatic F-crystal on $\mathcal{O}_K$ is a stable lattice in a crystalline G_K-representation, and this is an equivalence of categories [70, Theorem 5.6]. Evaluation at the Kummer prism gives a free module with linearized Frobenius invertible after inverting E, and its étale realization on the Kummer tower is the restriction of T [70, Theorem 7.9 and Remark 7.11]. Here $K = E_v$ has finite residue field, and each of the three integral cocycle extensions has crystalline rationalization. The functors apply on underlying $\mathbf{Z}_p$-lattices, and functoriality keeps the O-action.

Over $\mathcal{O}_{C_p}$, the prismatic F-crystal gives the pair consisting of the étale lattice and the B^+_{dR}-lattice in its de Rham period space obtained by Frobenius pullback [70, Theorem 5.2]. The proof of Proposition 5.3 identifies the crystalline period comparison with evaluation at the residue Witt prism, and Remark 5.4 identifies the pulled-back de Rham lattice with $B^+_{\mathrm{dR}} \otimes_K D_{\mathrm{dR}}(T[1/p])$. For a free Breuil–Kisin–Fargues module over $W(\mathcal{O}_{C_p^\flat})$, the Frobenius-fixed étale module over $W(C_p^\flat)$ gives the integral identification after inverting $\mu = [\varepsilon] - 1$ [74, Lemma 4.26, preprint p. 41]. The Hodge–Tate map in Theorem 5.2 is $\theta\varphi^{-1}$, at which μ has nonzero residue. Pullback by φ^{-1} in the chart replaces μ by τ and the divisor by $\ker\theta = (E)$. It gives

$$P_T[1/\tau] = A[1/\tau] \otimes_{\mathbf{Z}_p} T, \qquad B^+_{\mathrm{dR}} \otimes_A \varphi P_T = B^+_{\mathrm{dR}} \otimes_K D_{\mathrm{dR}}(T[1/p])$$

in the common period representation at θ. Since τ is a unit at θ, the unpulled lattice at θ is $B^+_{\mathrm{dR}} \otimes_{\mathbf{Z}_p} T$. Both comparisons are functorial for the coefficient ring and G_K.

The Frobenius map on M_T is an isomorphism after inverting E. Thus the frames $\varphi^j M_T$ and M_T differ by an invertible matrix over $S[1/E_0, \dots, 1/E_{j-1}]$. At the divisor E_a the φ^{-a}-reading of $\varphi^j P_T$ is the constant de Rham lattice when $a < j$, since all further intervening divisors are units at θ. It is the representation lattice if $a \geq j$, since the image of $\varphi^{j-a}\tau$ is a unit at θ in this range. An invariant basis of forms adapted to the decreasing filtration reads these two lattices by orders 0 and $-u$, respectively, where u is the jump of the corresponding type. A form of type u is a combination over K of basis forms with jumps $\leq u$. This also holds on a letter after applying its Teichmüller idempotent. The tensor product filtration gives the tensor order rule in the cut calculation.

These order tests characterize $\varphi^j P_T$ inside $A_E \otimes_A P_T$, with the profile st_j. Indeed in a basis of $\varphi^j P_T$ an element of the localized module has coefficients with denominators bounded by a finite product of the E_a. The lattice test at E_a says the order of a numerator is at least the corresponding denominator exponent. The rule $\ker \theta_a = (E_a)$, applied repeatedly in the de Rham DVR, gives divisibility by that power in A. Distinct prime factors cancel simultaneously. This proves the standard profile test, including the profile 1 in the Tate lattice of jump -1. For $\mathbf{Z}_p(1)$ the vector $\tau^{-1}(1 \otimes 1)$ generates the prismatic lattice over A. Its Frobenius factor is $\tau/\varphi\tau$. It has the étale comparison and the constant de Rham lattice on its Frobenius pullback, which determine it by Theorem 5.2 in the shifted convention. Consequently multiplication by τ carries the standard lattice into the integral Tate lattice, also after extending scalars to O.

There is a second description of P_T. A vector of $A[1/\tau] \otimes T$ belongs to P_T if and only if all its positive Frobenius translates belong to the constant de Rham lattice at θ. For the reverse implication, take coordinates in a basis of P_T, with denominator a power of τ. The translates force divisibility of the numerators by that power at each prime $\varphi^{-j}E$, $j \geq 1$. Divisibility by all these primes implies divisibility by τ. In fact $\tau = \varphi^{-N}(\tau) \prod_{j=1}^{N} \varphi^{-j}(\xi)$, where $\xi = ([\varepsilon] - 1)/\tau$ is associated to E. Modulo p the valuation of the finite product increases to that of τ. Subtract a multiple of τ with the same reduction as the numerator, then divide by p. All the prime divisibilities survive since none of these primes divides p. Iteration in the complete ring gives divisibility by τ; repeating removes its powers.

For a kernel lattice T in a cocycle extension $\widetilde{T}$, the underlying sequence of $\mathbf{Z}_p$-modules is split, and the rational de Rham sequence is strict for the filtration. Apply this description after clearing the finitely many E_a-denominators to get

$$(W_F \otimes T) \cap (A_E \otimes_A P_{\widetilde{T}}) = A_E \otimes_A P_T.$$

Filtered strictness also passes the order tests to the kernel, since the typed forms on the kernel extend to forms on the extension with the same order bound. The functorial tensor maps take bilinear images into the corresponding localized period modules.

For $x \in M_T$, division of $gx - x$ by $\Delta = z\tau$ leaves a vector of P_T. Use the two maps $(S, (E)) \rightrightarrows (A/(z\tau), (E))$ from the inclusion and from $g \in G_K$. The ideal $(z\tau)$ is stable under δ, since $\delta(z) = 0$ and $\delta(\tau) = ((1+\tau)^p - 1 - \tau^p)/p \in (\tau)$. The divisor E stays a nonzero divisor in the quotient, because it is prime in A and does not divide $z\tau$. Derived (p, E)-completeness follows from the cofiber sequence for multiplication by $z\tau$ on A, and $p \in (E, \varphi(E))$ descends. Its quotient by E is $\mathcal{O}_{C_p}/(\pi(\zeta_p - 1))$, killed by p^2 since $v_p(\pi) + v_p(\zeta_p - 1) \leq 2$, including at 2. This is a bounded prism over $\mathcal{O}_K$. The two maps from S coincide by the chart congruence. Thus the crystal base-change identifications agree there, giving the divisibility for the action through étale comparison. The scalar chart congruence gives $gE/E - 1 \in \Delta A$. Applying φ^a and using $\Delta \mid \varphi^a\Delta$ gives $gE_a/E_a - 1 \in \Delta A$. Hence each denominator transforms by a unit congruent to one modulo Δ, and the divisibility remains true on $\mathcal{N}_T$. Division by Δ preserves all orders at nonnegative divisors. The normalized Galois difference is therefore a kernel vector with the depth profile. Frobenius gives Φ_q^0 on the ordinary profiles and Φ_q^1 on depths by the factor $E_a \mid \varphi^q\Delta$ for $0 \leq a < q$. This gives the integral inputs for the cut walk.

Crystalline coordinates and their finite jets. The crystalline comparison gives a scalar retraction on jets and an explicit index bound for the long gap. Put $\Lambda = W(\mathbf{F}_{p^{fK}})$. Let $\mathscr{R}$ be the p-adic completion of the divided-power polynomial ring over Λ with basis $y^{[n]} = y^n/n!$. An element is $\sum b_n y^{[n]}$ with $b_n \in \Lambda$ tending p-adically to zero. Give it Frobenius extending Witt Frobenius with $\varphi(y) = p^{p-1}y^p$. This preserves the ring and reduces to Frobenius. For $n > 0$, both $\varphi(y^{[n]}) = p^{(p-1)n}(pn)!\,y^{[pn]}/n!$ and $(y^{[n]})^p = (pn)!\,y^{[pn]}/(n!)^p$ vanish modulo p. For the latter coefficient the valuation is the positive base-p digit sum of n. The specified Frobenius thus defines a δ-structure. The ring is p-torsion-free and p-complete, so $(\mathscr{R}, (p))$ is a bounded crystalline prism. The completion of the divided-power envelope defining A_{cris} has the same δ-ring property. Its Frobenius is inherited from A and sends E to a multiple of p. It therefore sends each positive divided power of E to zero modulo p, as does taking its p-th power.

For $j \geq 1$ the two maps from S have the same coefficient map φ^j on Λ and send z to py and 0, respectively. They commute with Frobenius and send E to p times a unit. Here we use the normalized generator $E = E_0/U_0 = p + z^{e_K}u(z)$ from the chart, with $U_0 \in S^\times$ and y topologically nilpotent. The base map $\mathcal{O}_K \to \mathscr{R}/p$ is reduction to the residue field followed by its p^j-power automorphism and inclusion. Both prism maps induce this same base map. The second map factors through the crystalline prism $(\Lambda, (p))$ over $\mathcal{O}_K$ with that residue map. For a basis m of M_T, choose a basis of the evaluation on this second prism. The crystal identifies their pullbacks to $\mathscr{R}$ by a matrix $V(y) \in \mathrm{GL}(\mathscr{R})$.

For sufficiently large j put $y = z^{p^j}/p \in pA_{\mathrm{cris}}$, where A_{cris} is the p-complete divided-power envelope of (E). The expansion of $(E-p)^q/q!$ in the products $(E^i/i!)((-p)^{q-i}/(q-i)!)$ belongs to A_{cris}. Since $z^{e_K} = u(z)^{-1}(E-p)$, this gives divisibility of z^{p^j} by p^3 whenever $v_p(\lfloor p^j/e_K \rfloor!) \geq 3$. Substitution on the divided powers is integral, is continuous for the p-adic topology and commutes with Frobenius. This is a map of the crystalline prisms over $\mathcal{O}_K$, where the base map to A_{cris}/p has the same p^j-power residue twist. The first structural map $S \to A_{\mathrm{cris}}$ is now precisely φ^j. The inclusion with Frobenius $\varphi^j : (A, (E)) \to (A_{\mathrm{cris}}, (p))$ is a prism map over the indicated base. In fact $\varphi(E)/p = \delta(E) + (p-1)!E^{[p]}$ is a unit, since $\delta(E)$ is a unit and $E^{[p]}$ is nilpotent modulo p. The same holds for $\varphi^j(E)/p$, $j \geq 1$. The evaluation of $\mathcal{E}_T$ there is $A_{\mathrm{cris}} \otimes_{\varphi^j, A} P_T$ by transitive base change. After extending to B_{cris}, the element $\varphi^j\tau = \varphi^{j-1}([\varepsilon] - 1)$ is invertible. The étale comparison therefore identifies this pullback with $B_{\mathrm{cris}} \otimes T$. Under that identification a vector in the frame m is sent to its scalar Frobenius translate. The resulting relation is

$$\varphi^j m = C\, V(z^{p^j}/p). \tag{16.13}$$

Here C comes from $(\Lambda, (p))$. Its columns are fixed by G_K because G_K fixes the embedding of Λ and the comparisons commute with the Galois action. An invariant de Rham form therefore takes values in K on C. These are the maps of the crystalline comparison in [70, Proposition 5.3 and Remark 5.4]. The canonical map from the shifted Fontaine prism corresponds to φ from $(A, (E))$. The additional Frobenius iterations account for the residue twist p^j fixed above. The left side is the actual frame of the level-j module in the common étale realization.

For qualitative estimates at the a-th divisor put $K_a = K(\pi_a)$. The completion of $\varphi^{-a}(S)[1/p]$ along E has coefficient field K_a and parameter E. Weierstrass division gives the residue field since $\theta(\varphi^{-a}z) = \pi_a$ and π_a has Eisenstein degree $e_K p^a$. Hensel's lemma lifts this separable residue field to the coefficient field, and the completion maps into B_{dR}^+. Both $V(z^{p^j}/p)$ and its inverse have finite jets in this scalar completion. For a fixed order H, write $z^{p^j n} = z^r u(z)^{-q}(E-p)^q$ with $q = \lfloor p^j n/e_K \rfloor$ and $0 \leq r < e_K$. Modulo E^H, the last factor is divisible by p^{q-H+1} for $q \geq H$. The denominator of the n-th divided power has valuation $n + v_p(n!) \leq 2n$. For large enough j, $q - H + 1 - 2n$ tends to infinity. This gives convergence in each finite jet of the scalar ring with p inverted, including at $p = 2$. If $\varphi^{j+a}m = mH_{j+a}$ is the iterated height matrix, then

$$\varphi^{-a}m = (\varphi^j m)\varphi^{-a}(H_{j+a}^{-1}).$$

The last matrix has meromorphic jets over K_a by the finite height inclusions. Equation (16.13) now shows that the evaluation of $\varphi^{-a}x$ for $x \in \mathcal{N}_T$ by an invariant form has meromorphic jets over K_a.

In the unramified chart, if $p \geq 5$ and $j \geq 2$, the evaluation $y = z^{p^j}/p$ already works at this j, since $z = E - p$ and $v_p((p^j)!) \geq 3$. Suppose $H \leq p^j\delta$ for integers $H, \delta \geq 0$. Both $p^{2\delta}V$ and $p^{2\delta}V^{-1}$ then have order-H jets that lift to matrices over S. For $n < \delta$, the denominator of the corresponding divided power has valuation at most $2n$, paid by 2δ. For $n \geq \delta$, the coefficient of E^h for $0 \leq h < H$ in $(E-p)^{p^j n}$, including this scalar multiplier and denominator, has valuation at least $p^j n - h - 2n + 2\delta \geq 0$. It tends to infinity with n. When $H = 0$ there is no jet restriction. Take matrices P_1, P_2 over S lifting the jets of $p^{2\delta}V^{-1}$ and $p^{2\delta}V$. The columns a of P_1 have Va constant modulo E^H. The columns of $p^{4\delta}I - P_1P_2$ have $Va \equiv 0 \bmod E^H$. The S-span of the two sets contains $p^{4\delta}S^{\mathrm{rank}\,T}$. By (16.13) this proves that $p^{4\delta}$ times the level-j module over S is in the S-span of vectors all of whose form values are G_K-invariant modulo t^H. The same holds at each letter after applying the Teichmüller idempotent. This is the constant needed in the unramified long

gap. The ordered exponent test there is precisely this level- j module over S, by the coefficient retraction below.

The integral coefficient retraction. Apply Lemma 16.1 to the Kummer parameter z. Its reduction in $A/p = \mathcal{O}_F$ has valuation $1/e_K > 0$. Thus the same Witt-ring completeness and regularity verification applies, and the residue subfield is $\mathbf{F}_{p^{f_K}}$. Write $\lambda : A \to S$ for the scalar retraction. Localize and extend the scalar map in an S-basis to obtain a functorial retraction $\mathcal{R} : A_E \otimes_A P_T \to \mathcal{N}_T$.

In particular $S_E \cap A = S$ inside A_E, since the localized retraction fixes S_E and maps A to S. For a point of $\mathcal{N}_T$ with the level-j standard profile, its coordinates in the frame $\varphi^j m$ are in A by the order tests and in S_E by the finite-height change of basis. Thus this point already belongs to the level-j module over S. In particular the radius bound applies to the points coming from the cutoff operators.

The retraction preserves form orders after subtracting any constant in K_a. Evaluate at the a-th divisor after expressing the vector in an S-basis and clearing the finitely many scalar denominators. Approximate the form values on the basis and the constant, to the needed order, by elements of $\varphi^{-a}S[1/p, 1/E]$, using the meromorphic jet calculation. After clearing powers of p and of E_a, the claimed vanishing to a prescribed order is a divisibility by a power of E_a of a scalar $\sum s_j r_j - s_0$, with $s_j, s_0 \in S$ and $r_j \in A$. The equivalence holds in A because $\ker \theta_a = (E_a)$. The divisibility survives λ by S-linearity. Undoing the denominators proves the assertion. In particular the retraction preserves the exponent profiles and commutes with projection to the trivial parameter. If a vector of the S-module maps to zero in that parameter, the strictness at the localized A-level and the natural retraction place it in the S-module of the kernel.

Integral operators for raising depth. Use the Kummer chart and the integral period modules for one of the three extensions, denoted by T in this paragraph. Write $\iota_a : A \to B^+_{\mathrm{dR}}$ for the map through φ^{-a}, $F^n = t^n B^+_{\mathrm{dR}}$ for its de Rham filtration for any integer n, and $Y_a(x)$ for the value of a form of type u at $i[-a]$ on $\varphi^{-a}x$. The form is Galois-equivariant, and its value is ι_a-linear in scalars. Recall two properties of the crystalline coordinates. The finite-jet calculation puts the residue of an integral form value of a localized point in $K_a = K(\pi_a)$. This also holds after subtracting a constant in K and shifting the filtration by a power of $\iota_a(E_a)$. The coordinate projection $\lambda : A \to S$ induces a retraction $\mathcal{R}$ from the ambient localized module to $\mathcal{N}_T$. It is functorial for maps of the lattices and preserves the congruence $Y_a(y) - c \in F^n$ for $c \in K$. This includes preservation of every order bound. In particular the correction below is taken in $\mathcal{N}_T$ at every step.

Let $d = 0$ when $e_K = 1$, and $d = v_p(e_K) + 1$ otherwise. For each pair (a, n) with $a \geq 0$ and $n \in \mathbf{Z}$, there is an additive operator on $\mathcal{N}_T$ of cost

$$c(a, n) = v_p(n) + d + 3a, \qquad v_p(0) = 0.$$

On points over constants it multiplies the projected parameter by $p^{c(a,n)}$. It preserves all orders and all previously valid depths and raises depth n to $n + 1$ at level a. Every form at this level whose value is invariant modulo F^n has invariant value modulo F^{n+1} on the output. The same operator works for every such form.

For $n \neq 0$, write $\gamma = [\varepsilon]$ and choose $g \in G_K$ with $gz = \gamma^{p^a} z$. Set $\omega_0 = gE/E$, which is a unit of A with $\omega_0 - 1 \in \Delta A$, and let $\omega = \varphi^a \omega_0 = gE_a/E_a$. Under θ_a, γ^{p^a} maps to one, so g fixes the residue π_a of z, and hence fixes K_a. The identity $\gamma - 1 = E\tau U$ for a unit $U \in A$ gives

$$\theta\big((gz - z)/E\big) = p^a \theta(U)\theta(\Delta).$$

Taylor expansion of the Eisenstein polynomial at z, followed by division by E and evaluation at θ, gives

$$\theta(\omega_0 - 1) = E'(\pi)\, p^a \theta(U)\theta(\Delta).$$

For a unit normalization $E = E_0/U_0$ of the polynomial E_0, division by E_0 in Taylor's formula replaces $E_0'(\pi)$ by $E_0'(\pi)/U_0(\pi) = E'(\pi)$. The multiplying factor U_0/gU_0 in the Galois quotient has image one

under θ. The derivative of the Eisenstein polynomial is nonzero at π. The valuations of its nonzero terms of degrees $j-1$, $1 \le j \le e_K$, have distinct fractional parts $(j-1)/e_K$. Its valuation is at most that of the leading term, $v_p(e_K) + (e_K-1)/e_K$, and is zero if $e_K = 1$. In particular $v_p(E'(\pi)) \le d$.

Here $v_p(\theta\Delta) = 1/e_K + 1/(p-1) > 1/(p-1)$. If $v_p(q-1) > 1/(p-1)$, the binomial expansion for $q^p - 1$ has its unique lowest-valuation term $p(q-1)$. Raising to a positive exponent prime to p preserves the valuation of $q-1$, and inversion preserves it as well. Applying these facts to $q = \theta(\omega_0)$ proves

$$v_p\big(\theta(\omega_0^n - 1)\big) = v_p(n) + a + v_p(E'(\pi)) + v_p(\theta\Delta).$$

At $p = 2$ the strict bound $v_2(q-1) > 1$ gives $v_2(2(q-1)) = 1 + v_2(q-1) < 2v_2(q-1)$, so the same identity holds. Since $\Delta \mid \varphi^a\Delta$ and $\Delta \mid \omega_0^n - 1$, set $X = (\omega^n - 1)/\Delta \in A$. The valuation of $\theta_a(\Delta)$ is $p^{-a}v_p(\theta\Delta)$, computed from π_a and $\zeta_{p^{a+1}} - 1$. Thus

$$v_p(\theta_a X) \le v_p(n) + d + a + (1-p^{-a})\big(1/e_K + 1/(p-1)\big) \le c(a,n).$$

Surjectivity of θ_a lifts $p^{c(a,n)}/\theta_a X$ to $\nu \in A$. We have obtained $\nu X \equiv p^{c(a,n)}$ modulo E_a, with integral ν.

The congruence improvement uses $(B_{\mathrm{dR}}/F^n)^{G_K} = \mathrm{im}(K)$ for finite n. The successive graded pieces are $C_p(j)$. Tate's fixed-field result gives $C_p^{G_K} = K$. His twisted vanishing gives $C_p(j)^{G_K} = 0$ for every nonzero integer j, by working first over a finite extension where the cyclotomic tower is totally ramified with group $\mathbf{Z}_p$. This applies to the dyadic and odd completions [75, Theorems 1–2, pp. 175–177]. Starting at the first nonzero term, eliminate the order-zero term by subtracting its constant in K, then apply vanishing at the other orders below n. If $Y_a(x) = C + y$ with $C \in K$ and $y \in F^n$, then $\iota_a(E_a)^{-n}y$ has residue in K_a. It is fixed modulo F^1 by the chosen g. Since $gE_a = \omega E_a$, this yields

$$gy - \iota_a(\omega^n)y \in F^{n+1}.$$

The scalar $\iota_a(\Delta)$ is a unit for the de Rham filtration. It follows that the value on $(gx-x)/\Delta$ is congruent to $\iota_a(X)y$ modulo F^{n+1}. Consequently the value of

$$p^{c(a,n)}x - \nu(gx-x)/\Delta$$

is congruent to $p^{c(a,n)}C$ modulo F^{n+1}. The retraction $\mathcal{R}$ preserves this congruence and gives the desired output in $\mathcal{N}_T$.

For $n = 0$ take $N = p^{d+3a}$ and automorphisms g_j with $g_j z = \gamma^j z$, $0 \le j < N$. Before retraction use $\sum_j g_j x = Nx + \Delta\sum_j (g_j x - x)/\Delta$. If $Y_a(x)$ is invariant modulo F^0, it is integral and has residue in K_a. Write this residue in the basis $1, \pi_a, \dots, \pi_a^{p^a-1}$ over K, which exists because $X^{p^a} - \pi$ is Eisenstein over K. The residue of the sum on each term with exponent $1 \le m < p^a$ is zero by the geometric sum in $\zeta_{p^a}^m$, since $p^a \mid N$. Its constant term belongs to K. The retraction preserves congruence to that constant modulo F^1. This also applies when $a = 0$, where every residue already lies in K.

In both cases, a normalized difference $(g_j x - x)/\Delta$ lies in the ambient localized module and has zero projection when x lies over a constant. It retains the original orders, since the order tests are Galois-stable and division by Δ leaves their bounds unchanged. It has orders at least the previous depth bounds as well. Its multiplying coefficient belongs to A, so after retraction the correction lies in both order submodules. The latter is Galois-stable, which proves retention of depth. All operations are integral and commute with the coefficient-lattice action by functoriality. No infinite sum of depth operators is used.

Suppose a point over $p^m c$ has orders o and depth D, and the target is depth $D' \ge D$. Take a finite set I of pairs containing every (a,n) with $D(a,u) \le n < D'(a,u)$ for an allowed jump u at $i[-a]$. Applying the operators at each level in increasing order of n gives depth D' and costs exactly $C_I = \sum_{(a,n)\in I} c(a,n)$ in the multiplier of the parameter. A single operation serves all the forms and jumps using its pair (a,n). Multiply by the remaining power of p to obtain any exponent at least $m + C_I$. The resulting point has orders at least

$$\begin{cases} \max(o(a,u), D'(a,u)), & o(a,u) > 0, \\ \max\big(o(a,u), \min(0, D'(a,u))\big), & o(a,u) \le 0. \end{cases}$$

In fact each value is congruent to a constant modulo the new depth. When both bounds are positive the constant must be zero. When the old order is nonpositive, the constant has order at least zero, giving the other bound. This applies also to a localized point whose initial orders below the Frobenius level are zero. The only additional cost comes from the specified steps. At level a, the nonzero terms have the usual $v_p(n)$ charge and every term, including zero, has the $d + 3a$ charge. In particular the unramified steps at level zero have zero cost if their range contains no nonzero multiple of p. For an interval of real length L in $[-H, H]$ with $H \geq 1$, the sum of the $v_p(n)$ over its nonzero integers is at most $L/(p-1) + \lfloor \log_p H \rfloor$. For each $h \geq 1$ up to $\lfloor \log_p H \rfloor$, count at most $L/p^h + 1$ multiples of p^h. The zero term in the step count incurs only its stated $d + 3a$ charge.

Lifting ideals and propagation along Frobenius. Let C be the trivial parameter lattice of any side, write c_i for the letter component of $1 \otimes c$, and denote the projection to the parameter by pr. For a nonnegative integer modulus r, let $J_l(i; r)$ be the sum of $(r) \subset A$ and the ideal of scalars taking every c_i into the image of the ambient module at letter i with standard orders st_l. The second term is an ideal because the ambient order conditions are A-linear. Equivalently $f \in J_l(i; r)$ if one $b \in A$ works for all c, with a lift of $(f - rb)c_i$ of standard orders for each c. If $f \in S$, applying $\mathcal{R}$ gives lifts in $\mathcal{N}_T$ with common witness $\lambda(b) \in S$, since the retraction commutes with projection and is the identity on the S-localized module. The same argument works for lifting $(f - r\lambda(b))y$ for any standard localized point y in the trivial quotient at i. Such y is integral in $A \otimes C$, by the level-zero standard test for the trivial lattice. Its letter part is generated over A by the c_i, so an ambient lift of $(f - rb)y$ exists before retraction.

The exact ideal $J_l(i; 0)$ records the cokernel multipliers. If $p^c \in J_l(i; p^{c+1})$, then $p^c(1 - pb)$ is in the exact ideal, and $1 - pb$ is invertible in the p-complete ring A by the geometric series. Thus p^c is in the exact ideal and $J_l(i; p^{c+1}) = J_l(i; 0)$. If $p^c E^s \in J_0(i; 0)$, lift $p^c E^s c_i$ and divide by E^s in $\mathcal{N}_T$. The result lies over $p^c c_i$ with orders $-u - s$ at level zero and $-u$ above. The depths satisfy the same bounds by Galois stability of the order tests. This keeps the p-charge at c.

The standard profile test identifies the ambient lattices of orders st_l with the Frobenius translates of the level-zero lattice, so

$$\varphi^l\big(J_0(i; r)\big) = J_l(\sigma^l i; r).$$

In particular $p^c z^{pa} \in J_1(\sigma i; r)$ if and only if $p^c z^a \in J_0(i; r)$. If the jumps at i are all at most $h \geq 0$, multiplication by E^h sends $J_0(i; r)$ into $J_1(i; r)$, since it raises the level-zero order to $h - u \geq 0$ and preserves the others.

The operation $\alpha(g - 1)/(\Delta E)$ on a multiplier $f \in J_1(i; r) \cap S$ stays in $J_1(i; r)$ for $\alpha \in A$ and $g \in G_K$. Indeed start with a localized lift of $(f - rb)c_i$ at standard level one and $b \in S$. Its level-zero values are integral for the de Rham filtration and have residue in $K_0 = K$. The Galois difference increases that order by one. Dividing by Δ preserves the orders and is integral in the ambient localized module by the frame calculation. Division by E uses up the gain at level zero, leaving the other orders unchanged. The result after multiplication by α lies over

$$\left(\alpha \frac{gf - f}{\Delta E} - r\alpha \frac{gb - b}{\Delta E}\right) c_i.$$

Both scalar quotients lie in A by the chart congruence on S, and the second is independent of c. This proves the assertion. Any integral multiple of f can be subtracted from the resulting multiplier within the ideal.

To lower a positive power z^n at level one, take $g_1 z = \gamma z$, let $j = v_p(n)$, and put $\xi = (\gamma - 1)/\tau$. The prime-to-p part of n gives a unit geometric sum in A, since $\gamma \equiv 1$ in the residue field of the local ring A. The factorization

$$\gamma^{p^j} - 1 = \tau \prod_{h=0}^{j} \varphi^h(\xi)$$

and $\xi \sim E$ show that $(\gamma^n - 1)/(\tau E)$ is a unit times $\prod_{h=1}^{j} E_h$. Here $E_h \equiv p$ modulo z. Invert that unit as the coefficient α. On $f = p^c z^n$ the operation gives $p^{c+j} z^{n-1}$ up to an A-multiple of f, which can be subtracted. Iterating and then using Frobenius gives

$$p^c z^N \in J_1(\sigma i; r),\ pM \le N \quad \Longrightarrow \quad p^{c+v_p(N!)-v_p((pM)!)} z^M \in J_0(i; r).$$

Taking $M = \lfloor N/p \rfloor$ adds zero cost. Starting and ending at level zero, one can first apply Frobenius and lower from pN to pM; the added cost for $M \le N$ is $N - M + v_p(N!) - v_p(M!)$, by Legendre's formula.

To exchange one factor of E for z, let $f = p^c z^b E^n$ with $n > 0$ lie in $J_1(i; r)$. Use the scalar depth calculation at $a = 0$ and this n, so $gE = \omega E$, $\omega^n - 1 = \Delta X$, and $\nu X = p^{d+v_p(n)} + EB$ with $B \in A$. The chart identity gives $gz^b = z^b(1 + \tau E Q)$ for $Q \in A$. Expanding $gf - f$ yields

$$z\nu \frac{gf - f}{\Delta E} = p^{c+d+v_p(n)} z^{b+1} E^{n-1} + (zB + \nu Q \omega^n) f.$$

Subtract the last multiple of f. Iteration gives $p^{c+dn+v_p(n!)} z^{b+n} \in J_1(i; r)$. These two divided differences apply at every prime, and the E-exchange uses the chart bound on the derivative through d.

For $r = p^{\kappa+1}$, a common power of z belongs to $J_0(i; r)$ at all letters. The period-lattice comparison after inverting τ supplies one power τ^b taking the integral lattice of the extension into the standard level-zero module. Lifting a constant in the underlying integral splitting and projecting to its letter gives τ^b in the exact ideal at every letter. Modulo p, the ring A/p is a rank-one valuation ring and z, τ have positive finite valuation. Thus some power of z belongs to (p, τ^b); raising it to the power $\kappa + 1$ gives the claimed common multiplier. Suppose all jumps are bounded above by h at every letter, and take the least B with this property for z^B. At level one the ideals contain $z^B E^h$. The identity $z^{e_K} \in (E, p)$ implies $z^{B+e_K(h+\kappa)} \in J_1(i; r)$ for all i, since each binomial term is a multiple of $z^B E^h$ or of $p^{\kappa+1}$. The descent with no added cost gives the common power with exponent $\lfloor (B + e_K(h + \kappa))/p \rfloor$. Minimality therefore yields

$$(p - 1)B \le e_K(h + \kappa).$$

Finally suppose $K/\mathbf{Q}_p$ is unramified with $p \ge 5$ and that every jump at every letter is at most $(p - 3)(s + t)$, where s, t are nonnegative integers. Take $\kappa = s + t$ and $h = (p - 3)\kappa$. The last inequality gives $(p - 1)B \le (p - 2)\kappa$. Since $(p - 1)v_p(B!) \le B$, it follows that

$$(p - 2)\big(B + v_p(B!)\big) \le (p - 1)B, \qquad B + v_p(B!) \le \kappa.$$

For $0 \le m \le t$, if $m \le B$ the level-zero lowering cost is $B - m + v_p(B!) - v_p(m!) \le \kappa - m$. If $m > B$, multiply by z^{m-B} within the ideal. In either case $p^{\kappa-m} z^m \in J_0(i; p^{\kappa+1})$. The case $m = 0$ removes the modulus by the unit argument. With $E = z + p$ the binomial expansion of $p^s E^t$ now uses only these exact multipliers. This proves $p^s E^t \in J_0(i; 0)$ for every letter and each of the three extensions to which the jump bound applies.

Exponents and safety at a cut. Use any of the three extensions above, with its allowed letter sets. A point at i over a parameter $p^k c$ has exponents $o(a, u)$ if each de Rham form of type u at $\sigma^{-a} i$, evaluated on $\varphi^{-a} x$, has order at least $o(a, u)$. Type u means that on the original lattice the form has order at least $-u$. Equivalently, it belongs to Fil^{-u} of the filtered de Rham dual in that coefficient factor. Let its depth $D(a, u)$ mean that $gx - x$ has those exponents for every $g \in G_K$. We use points in the localized period module of the comparison, where the exponents and depth test the indicated vectors before reduction. Write $i[-a] = \sigma^{-a} i$ and

$$(\Phi_q^j o)(a, u) = \begin{cases} j, & a < q, \\ o(a - q, u), & a \ge q, \end{cases} \qquad \mathrm{st}_q(a, u) = \begin{cases} 0, & a < q, \\ -u, & a \ge q. \end{cases}$$

The integral period comparison gives the following properties. Differences of points over the same parameter belong to the kernel, with the minimum of their two exponent profiles. On a point of depth D,

$(g-1)x/\Delta$ belongs to the kernel with exponents D. The Frobenius iterate $\varphi^q x$ is a point at $\sigma^q i$, with exponents $\Phi_q^0 o$ and depth $\Phi_q^1 D$. For the latter assertion at $a < q$, use the vanishing of $\varphi^q\Delta$ at E_a in $(g-1)\varphi^q x = \varphi^q(\Delta)\varphi^q((g-1)x/\Delta)$. In the localized period module of the Tate lattice, the standard level-zero exponents 1 at every E_a imply that multiplication by τ gives an integral A-vector. These are the precise integral inputs to the following walk.

Suppose that a point of the third extension over p^{k_3} has exponents $Z(a,w)$ antitone in the jump w, and that kernel vectors on the two sides have exponents D, D' at the same letter. If for every pair of allowed jumps

$$Z(a,-u-v-1) + D(a,u) + D'(a,v) > 0, \tag{16.14}$$

then multiplying either kernel vector by Δ gives a safe pair. To see this, use filtered de Rham bases on both sides. In a bilinear map, the pullback of a form of type s uses only pairs of basis forms with $u+v \le s$, by the tensor product filtration. Thus the tensor vector has exponents at least $1 - Z(a,-s-1)$ for type s, by antitonicity and (16.14). In the final pairing to the Tate lattice a third type w can contribute together with tensor type s only if $w+s \le -1$. The inequality $Z(a,w) \ge Z(a,-s-1)$ gives order at least 1 on each term. Hence τ sends the paired vector into the integral Tate lattice, and multiplying by $z = [\pi^\flat]$ sends it into Witt vectors with coordinates in $\mathfrak{m}_F$. Reduction of the third pairing and application of h then give the claimed safety for $p^{k_3} b$. Distinct letters pair to zero, since their Teichmüller eigenvalues have unit difference and the pairing is balanced for the action of ζ_O. Thus safety on matching letters gives safety for sums over all letters.

Let C be a set of cut letters meeting each cycle and let q_i be the least positive distance from i to a preceding cut. At each cut i take two points x_i, y_i over $p^{k_1}c, p^{k_2}c'$, with exponents o_i, o_i' and depths D_i, D_i'. At each cut also take a third point over p^{k_3}, with antitone exponents Z_i in the jump. For each of the two points at i, allow its own distance $r \ge 0$ to a cut for use of that third point. For the first point require, at every a, u, v,

$$Z_i(a,-u-v-1) + D_i(a,u) + \min(-v, o_i'(a,v)) > 0. \tag{16.15}$$

For the second point require the interchanged condition. Designate one side at each cut. If the first is designated, require also $o_i \ge \mathrm{st}_{q_i}$ and

$$(\Phi_r^0 Z_{i[-r]})(a,-u-v-1) + D_i(a,u) + \min\big(o_i'(a,v), (\Phi_{q_i}^0 o'_{i[-q_i]})(a,v)\big) > 0. \tag{16.16}$$

If the second is designated, require the interchanged condition with its corresponding r. This includes the case $r=0$ in which the direct third point serves both inequalities.

At all letters put X_i, Y_i equal to the direct points at the latest cut, transported by Frobenius. The distance is zero at each cut and is q_i elsewhere. At a cut i, $(\varphi X)_i$ and $(\varphi Y)_i$ are the transports from the preceding cut at positive distance q_i. Away from the cuts they equal X_i, Y_i. Choose $W_i = X_i$ where the first side is designated and $W_i = (\varphi X)_i$ where the second is designated, with $W_i = X_i = (\varphi X)_i$ off the cuts. Summing gives x, w, y with the required constant projections.

The first safety pair in the cochain calculation follows from (16.15) at the latest cut, using the first Galois difference divided by Δ. At levels before that cut's transport distance the three orders are $1, 0, 0$. The other three safety pairs need only be tested at cuts. If the first side is designated, $W_i - X_i = 0$. For $(gW_i - W_i,\ g((\varphi Y)_i - Y_i))$ use (16.16), since the second difference has the minimum of its two exponent profiles. For $((\varphi X)_i - W_i,\ (g-1)(\varphi Y)_i)$ use the second point's condition (16.15) at the preceding cut and transport it. At $a < q_i$ the orders are $0, 1, 0$, since $o_i \ge \mathrm{st}_{q_i}$. At $a \ge q_i$ the same inequality gives $o_i(a,u) \ge -u$, so the minimum required for the difference follows from the $\min(-u, o)$ in (16.15). If the second side is designated, $(\varphi X)_i - W_i = 0$. The pair with $W_i - X_i$ then uses the interchanged condition (16.16), and the pair with $gW_i - W_i$ uses the transported first point's condition (16.15). At the levels $a < q_i$, safety again follows from the order 1 of the transported depth and order 0 of the difference on the designated side. This proves the four safety tests and hence the scaled cup assertion. A larger exponent works by bilinearity. The same proof with every letter a cut applies to estimates using only the first preceding step.

The length of the resulting local quotient. Suppose the cup assertion with exponent k holds for every $Y \in \mathcal{Q}$, every M, every integral pairing above and every finite cocycle of Y. A tested cocycle $a' \in Z^1_f(L_2)$ then satisfies $p^k a' \in \mathcal{H}(L_2)$ by the definition of the orthogonal module. The same conclusion with exponent zero holds on any fixed $\mathbf{Z}_p$-submodule of finite cocycles for which the uniform cup assertion applies. Such a submodule may be used to bound the length even if it was constructed as a $\mathbf{Z}_p$-module.

For $R = \operatorname{rank}_O L_2$ and coefficient field $F_O = O[1/p]$, put $V_O = L_2[1/p]$. The finite cocycles form a free O-module. They are finitely generated by the finiteness of $H^1(K, L_2)$ over O and that of the integral coboundaries, and are torsion-free as a submodule of the continuous cochains. Bloch–Kato's dimension formula gives

$$\operatorname{rank}_O Z^1_f(L_2) = R + \dim_{F_O}\bigl(D_{\mathrm{dR}}(V_O)/\mathrm{Fil}^0\bigr) \le (1 + [K : \mathbf{Q}_p])R.$$

The invariant term in $H^1_f(K, V_O)$ cancels against the invariant term in the rank of the coboundaries. Thus an exponent k bounds the p-normalized length of the defect by $(1 + [K : \mathbf{Q}_p])Rk \le 2e_K f_K Rk$. Indeed the defect is a quotient of $Z^1_f(L_2)/p^k Z^1_f(L_2)$, whose O-length is the rank times k $\operatorname{len}_O(O/p)$. For a finite O-module D_0 the same normalized length is $\operatorname{len}_{\mathbf{Z}_p}(D_0)/[F_O : \mathbf{Q}_p]$. This follows by taking the cardinality of both modules from the residue fields. If a $\mathbf{Z}_p$-submodule of finite cocycles of index at most p^A lies in the cup-orthogonal part, the normalized length is therefore at most $A/[F_O : \mathbf{Q}_p]$. The ramification degree of the coefficient ring does not give a further factor.

For the finite cocycle application, $Z^1_f(L_2)$ is p^M-saturated in the cochains with coefficients in L_2: division by p^M preserves the cocycle identity and the rational finite condition. Hence its reduction has cardinality $|O/p^M|^{\operatorname{rank}_O Z^1_f(L_2)}$. The corresponding rank for L and that for L_2 sum to $(2 + [K : \mathbf{Q}_p])R$, since the complementary jumps are u and $-u - 1$. Both reductions contain all finite coboundaries. The cup of finite rational classes vanishes after inverting p, and integral cup vanishes as well because $H^2(K, O(1)) \simeq O$ is torsion-free. Reduction gives orthogonality of the two reduced subgroups. Tate's finite local duality pairs $H^i(K, V)$ perfectly with $H^{2-i}(K, V^D)$ when V is a finite p-primary G_K-module and $V^D = \operatorname{Hom}(V, \mathbf{Q}_p/\mathbf{Z}_p(1))$. The Euler characteristic of V is $|V|^{-[K:\mathbf{Q}_p]}$ [76, Theorems 2.1–2.2, pp. 289–290]. Use $V = L/p^M L$ with the trace dual $L^\dagger/p^M L^\dagger$. The product of the two coboundary orders times $|H^1(K, V)|$ is $|O/p^M|^{(2+[K:\mathbf{Q}_p])R}$. The two orthogonal reductions thus have the orders of full annihilators in finite cocycles. It follows that the quotient of any tested cocycle group by its intersection with the finite subgroup has order bounded by the index of the reduction of the integral cup-orthogonal part in the reduction of $Z^1_f(L_2)$. This index is bounded by the integral defect just estimated.

Lifting across thin Frobenius successors. Use the Kummer period coordinates over a local field $K/\mathbf{Q}_p$. Write e_K for its ramification index and d_K for the chart bound $d_K = 0$ if $e_K = 1$, and $d_K = v_p(e_K)+1$ otherwise. Denote Frobenius on letters by σ. Consider one crystalline extension of a trivial lattice by a crystalline lattice, and include zero in its allowed jumps at each letter. Suppose all its jumps are at most ω, and all jumps at $\sigma i, \dots, \sigma^J i$ are at most ω'. If

$$p^{J+1} \ge 8e_K, \qquad k_0 \ge 2J d_K \omega' + \frac{(2J+4)\omega'}{p-1} + \frac{8e_K\omega}{p^{J+1}}, \tag{16.17}$$

then p^{k_0} lifts every constant at i to the period module with the standard exponents of level zero.

Here and below use the lifting ideals at the fixed modulus p^{k_0+1}. The least exponent A such that z^A belongs to the level-zero ideal at every letter satisfies $(p-1)A \le e_K(\omega + k_0)$, by the initial multiplier calculation with global jump bound ω.

Start with z^A at $\sigma^J i$. If a multiplier $p^c z^n$ lies in the level-zero ideal at a successor whose jumps are at most ω', multiply by $E^{\omega'}$ to reach level one. Exchanging the Eisenstein factors there gives $p^{c+d_K\omega'+v_p(\omega'!)} z^{n+\omega'}$. Descending again costs no additional p-power and replaces the exponent by $\lfloor (n+\omega')/p \rfloor$. After the J steps, there is a multiplier $p^c z^n$ at level zero at i with

$$c \le J\bigl(d_K\omega' + v_p(\omega'!)\bigr), \qquad n \le A/p^J + \omega'/(p-1).$$

Frobenius transports this multiplier to $p^{c} z^{p n}$ at level one at σi. Descending to exponent zero at i costs $v_p((p n)!) = n + v_p(n!) \le 2n$. Legendre's formula and (16.17) bound the total by

$$\begin{aligned} c + 2n &\le J d_K \omega' + (J+2)\frac{\omega'}{p-1} + \frac{2 e_K (\omega + k_0)}{(p-1) p^{J}} \\ &\le J d_K \omega' + (J+2)\frac{\omega'}{p-1} + \frac{4 e_K \omega}{p^{J+1}} + \frac{k_0}{2} \le k_0. \end{aligned}$$

Thus p^{k_0} belongs to the ideal taken modulo p^{k_0+1}. The p-adic unit argument for the lifting ideals gives the same multiplier in the exact ideal.

16.2.7 Crystalline bounds from the cup calculation

The small conductor-dividing places. Fix an actual place $v \mid \ell$ with $\ell \mid N_1 d$ and $p = \ell \le Q$. The fixed block L is the row of rank r' or a scalar Tate block. For any crystalline lattice Y in the catalogue class, pair its integral finite cocycle with one on $L^{\dagger} = L^{\vee}(1)$. Use an arbitrary integral lift $\psi : Y \otimes_O L^{\dagger} \to O(1)$ of the pairing modulo p^M, followed after reduction by the given $\mathbf{Z}_p$-linear coefficient functional. The first two extensions of trivial lattices are those of the two cocycles. The third has kernel $(Y \otimes_O L^{\dagger})^{\vee}(1)$ and cocycle $g \mapsto (g\psi - \psi)/p^M$. This is integral by equivariance modulo p^M, and rationally a coboundary, so the third extension is crystalline. This construction applies, in particular, to the congruence maps from the row and Tate graph lattices at their given moduli.

Keep the letters and Frobenius order of the actual completion. Put $b_* = B + \delta$, $g_* = G + \delta$, $V = W = 4t$, $\eta_0 = 2G$, $D = 9G$, and $J = J_\ell$. At letter i write $\beta_i = b_*$ if it is thick, and $\beta_i = g_*$ otherwise. The common thick sign over v lets us designate one side as the negative side at every letter, and one as the positive side. Their allowed jump sets are in

$$[-\eta_0, \eta_0] \cup [-\beta_i - V, -\beta_i], \qquad [-\eta_0, \eta_0] \cup [\beta_i, \beta_i + V],$$

respectively. Thin jumps are included in the low interval even in a letter fibre that also contains a thick reading. With u on the negative side and v' on the positive side, the third extension has jumps in the set comprising zero and $-u - v' - 1$. Since $v' \ge -\eta_0$ and $u \ge -\max(\eta_0, \beta_i + V)$, all three jump sets are bounded above by $\beta_i + D$, where $D \ge V + 4\eta_0$.

Call a letter a cut if its next J successors are thin. Every letter i has a cut at positive backward distance at most $2J + 1$. In fact if $\sigma^{-(J+1)} i$ is a cut we can use it. Otherwise some thick letter $j = \sigma^{-a} i$ occurs at a distance $1 \le a \le J$. If j is a cut we can use it. If it has a thick successor within J, each of its preceding J letters must be thin by (16.11). Then $\sigma^{-(a+J+1)} i$ is a cut. This proof also applies if positions in an orbit repeat. Write l_i for the least positive backward distance to a cut.

At every cut use (16.17) on the three extensions with $\omega = 2b_*$, $\omega' = 2D$ and $k_0 = 16(J+1)K_\ell$. Indeed $\beta_i + D \le 2b_*$ everywhere and $g_* + D \le 2D$ on thin successors. The small-prime inequalities give $p^{J+1} \ge 8 e_K$, $2 e_K b_* \le p^{J+1} K_\ell$, and

$$4\big(J + 1 + d_K + \lfloor \log_p(2b_*) \rfloor\big) D \le K_\ell.$$

Using $p - 1 \ge 1$, the cost on the right side of (16.17) is at most $(2J + 10)K_\ell \le k_0$. Thus all three extensions give points over p^{k_0} times the respective constants, with standard exponents and the same Galois-difference depth $-u$ on jump u at every level.

For the first two points raise the depth at level a to $-u + C_a$, where

$$A_a = \eta_0 \, \mathbf{1}_{a \le 2J}, \qquad C_a = \eta_0 \, \mathbf{1}_{a \le 4J+1}.$$

Here use the integral depth operators. At each affected level the required integers lie in two intervals of radius D, centred at zero and at the signed opposite of the far centre of the jump window at $\sigma^{-a} i$. They

have absolute value at most $2b_*$. Their total cost at that level, with the cost of zero defined using $v_p(0) = 0$, is at most

$$(4D + 2)\big(\lfloor \log_p(2b_*) \rfloor + d_K + 3a\big) \le 18K_\ell.$$

The last inequality uses $D \ge 1$ and $a \le 4J + 1$. The sum over at most $4J + 2$ levels is at most $72(J + 1)K_\ell$. Hence both points can be taken over p^{25k_0} times the corresponding constants. Keep the third point over p^{k_0} with exponents and depth $-w$ on its jump w.

We check the strict inequalities in the Frobenius passage between cuts. The negative and positive points at cut i may be read with exponents

$$o_-(a, u) = \max\{\mathrm{std}_{l_i}(a, u), -u - A_a\}, \qquad o_+(a, v') = -v' - A_a,$$

where $\mathrm{std}_l(a, u) = 0$ for $a < l$ and $-u$ for $a \ge l$. Indeed the improved exponent once the depth has been raised is $-u + C_a$ if $u < 0$, and $\min(0, -u + C_a)$ if $u \ge 0$. On the negative-far side any nonnegative jump is at most η_0. When $a < l_i$ we have $A_a = C_a = \eta_0$, so the improved exponent is at least $\mathrm{std}_{l_i}(a, u) = 0$ in that case as well. For $a \ge l_i$ it is at least $-u$. These facts give both displayed exponents.

In the third point take $w = -u - v' - 1$, so its exponent is $u + v' + 1$, an antitone function of w. For each of the two sides, add its depth to this third exponent and the minimum of the exponent of the other side and that side's standard exponent −jump. The result is at least $1 + C_a - A_a \ge 1$. Let the negative side lead the passage, since it retains std_{l_i}. For its additional inequality replace the other side's standard exponent by its exponent transported from the preceding cut. When $a < l_i$, the transported exponent is zero and $o_+(a, v') = -v' - \eta_0 \le 0$. When $a \ge l_i$, the transported exponent is $-v' - A_{a-l_i}$. Both required lower bounds follow from

$$C_a \ge A_a, \qquad C_a \ge A_{a-l_i} \quad (a \ge l_i),$$

the second since $l_i \le 2J + 1$. In all comparisons the third point is taken at the cut itself. The walk permits either original side to lead. Graded symmetry in cohomology and bilinearity also allow the scaling power to be placed on the tested cocycle of $L^\dagger$.

The cup is therefore killed by p^{k_ℓ} for $k_\ell = 10^4 K_\ell(J + 1)$, since the two point exponents and the third one add to $51k_0 \le k_\ell$. This applies to every pairing, and hence to the actual cup-orthogonal quotient. At $p = 2$ the spacing is the 3^c-congruence, the chart has $v_2(\theta\Delta) = 1 + 1/e_K > 1$, and the zero-index depth step uses the same root sum. The thin-successor and depth costs use $p/(p - 1) \le 2$ and $p - 1 \ge 1$, both valid at 2. For $R = \mathrm{rank}_O L$ and $n_v = [E_v : \mathbf{Q}_\ell]$ the crystalline cocycle rank bounds the normalized length of the defect by

$$2n_v R k_\ell \le 4 \cdot 10^4 n_v R\, m\lambda_\ell/Q \le n_v R\, \mathrm{loss}_{\mathrm{crys}}(\ell)/t.$$

Here $J + 1 \le m$ and $K_\ell \le 2\lambda_\ell/Q$, with $Q = \kappa t$. This is the required allowance for a Tate block and also fits the row allowance. The bound is uniform in the modulus exponent M, including the row modulus p^{M_ℓ} for $M_\ell = \lfloor r b_\ell(X) \rfloor + 1$ and the scalar modulus p^ν in the first-order polynomial. The local degree is the same as in the crystalline duality calculation.

The two sizes at large unramified primes. Fix the actual unramified $E_v/\mathbf{Q}_p$, where $p = \ell > Q$ and $\lambda_\ell \ge H^3$. Write

$$b = B + \delta, \qquad g = G + \delta, \qquad W_1 = W + 4\eta, \qquad k = K_\ell, \qquad \beta_i \in \{g, b\}.$$

Use the opposite windows at the letters of the chart constructed above, with far width W and low half-width η. We use the point notation of the depth calculation. In this paragraph E denotes the chart's polynomial $z + p$, and i_{-a} is the letter a Frobenius steps before i. An admissible pairing has a third side with jumps 0 and $w = -u - v - 1$, where u, v are jumps of the two sides at the same letter. All three sets of jumps are bounded above by $\beta_i + W_1$. Indeed when both jumps are far, their sum is in $[-W, W]$, and when only one is far the low jump changes it by at most η.

Put

$$q = \lfloor b/p \rfloor, \qquad s_i = \begin{cases} \max(0, q - 7k) & \beta_{\sigma i} = b, \\ 0 & \beta_{\sigma i} = g, \end{cases}$$
$$s_* = q + 12k, \qquad C = g + W_1, \qquad D = 2(b + W_1).$$

Take i as a cut when $s_i < \beta_i$. The inequalities in the allowance calculation give

$$b < kp^2, \qquad 2g < kp, \qquad W_1 \le k,$$
$$C \le (p-1)k, \qquad \beta_i + W_1 \le (p-3)s_*, \qquad D \le 4kp^2.$$

For the penultimate inequality use $b < pq + p$, $q < kp$, $W_1 \le k$, and $p \ge 7$ in $(p-3)(q+12k) - pq = 12k(p-3) - 3q$. We also have $s_i + W_1 \le pk$ and $\beta_i + s_i + W_1 \le D$. Every thick letter is a cut, and the predecessor of every thin letter is a cut. Thus the distance $l = l(i)$ to the nearest strict predecessor cut is either 1 or 2, including on Frobenius orbits of length one or two.

Let $L_p = \lfloor \log_p D \rfloor$, $\Delta = 4k$ and $\kappa_* = 120k + 24L_p + 1$. At any cut on each of the three sides, $p^{\kappa_*}E^{s_i}$ kills the cokernel at level zero. If the successor is thick, $s_i + \kappa_* \ge s_*$ because $s_* - s_i \le 19k$. Apply the unramified calculation of the lifting ideals. If the successor is thin, $s_i = 0$ and the thin-successor calculation with one successor uses D and C, and its bound is

$$6C/(p-1) + 8D/p^2 \le 38k < \kappa_*.$$

Choose the third point at each cut over $p^{\kappa_*} \cdot 1$ with orders $z_i(0, w) = -w - s_i$ and $z_i(a, w) = -w$ for $a \ge 1$. These orders are antitone in w.

Lifts for the two-letter gap. We record the part of the quantitative period-frame argument used here. A localized point of standard orders of level $l > 0$ has orders 0 below l and $-u$ on the jump u from level l on. We use $l = 2$, $p \ge 7$, and $D \le p^l\Delta$. The quantitative frame at level l spans $p^{4\Delta}$ times these points by points of the same orders with depth at least D at level zero. This is the quantitative clause of the crystalline frame. Choose a finite family of such points whose span contains the $p^{4\Delta}$-multiple of the standard module. For a standard point y, express $p^{4\Delta}y$ in this family and expand its scalar coefficients in $E = z + p$ using $\Lambda[[z]] = \Lambda[[E]]$, where $\Lambda = W(\mathbf{F}_{p^{f_0}})$. This gives

$$p^{4\Delta}y = \sum_{a=0}^{N} E^a \xi_a + E^{N+1} y_{N+1}$$

with the remainder of standard orders. The coefficients ξ_a are Λ-linear combinations of the same points, so they have standard orders and depth at least D at level zero. The chart group fixes Λ.

If y has extra order $d(u) \le D$ at level zero, then ξ_N has order D there on $d(u) > N$. For the constant term reduce the expansion modulo E on the indicated jump. In the truncated de Rham value of an almost invariant coefficient, the positive-order part must vanish modulo t^D by the fixed period calculation of the improved-order rule. Thus an order of at least one improves to D. Subtract the coefficient and divide by E. This lowers $d(u)$ by one at level zero and leaves the constraints at every higher level intact. Repeat through N. The same argument applies to the projection of the expansion into the trivial quotient.

Suppose now that E^{s_*} kills the cokernel at i, that p^{κ_*} kills it at i_{-l}, and that all jumps on the side being lifted are at most C between these two letters, including i and excluding i_{-l}. Start with a point over c at i with orders $-u - s_*$ at level zero and $-u$ above. Multiply it by

$$F(E) = E^{s_*} \prod_{a<l} \left((E-p)^{p^a} + p\right)^C.$$

This is monic of degree $N = s_* + C\sum_{a<l} p^a$. The new orders are at least $C - u$ below l and $-u$ above, so the point is standard of level l. If $N < D$, the N-th coefficient of its expansion, after multiplication by

$p^{4\Delta}$, has order D at level zero for $u < -N$. In the standard module of the trivial quotient, the projection of each ξ_a is constant modulo E^D by its depth at level zero and monic division. Compare the coefficient of E^N with that of $p^{4\Delta}F(E)c_i$. Since $N < D$ and F is monic, the difference between the projection of ξ_N and $p^{4\Delta}c_i$ is a multiple of E^D in that module. Divide by E^D there and lift p^{κ_*} times the quotient using the lifting ideal transported by φ^l from i_{-l}. Multiplying the lift by E^D gives order and depth at least D at level zero and preserves the standard orders at all other levels. Subtract it from $p^{\kappa_*}\xi_N$. The resulting point over $p^{\kappa_*+4\Delta}c_i$ has depth D at level zero, 0 at $0 < a < l$, and $-u$ above. Its orders are at least D at level zero for $u < -N$ and standard of level l elsewhere. All terms are in the localized Kisin module, so the removal of the error is an exact lift.

Depth choices and the cup. For each cut i, call the side with negative far sign at i the leading side. Set $A_a = \eta$ for $a \le 2$ and $A_a = 0$ for $a > 2$. Use orders $-u - A_a$ on the other side and $\max(\mathrm{std}_l(a, u), -u - A_a)$ on the leading side, where std_l is the standard order of level $l(i)$.

First consider the ordinary point on either side at i. The cut supplies an initial lift with orders and depth $-u - s_i$ at level zero and $-u$ above. Raise its depth to

$$d_i(a, u) = -u + \eta\,\mathbf{1}_{a\le 5} + s_i\mathbf{1}_{a=0}.$$

At level zero the steps are within $s_i + W_1$ of 0 or the negative of the far center of that side. At levels $1, \dots, 5$ the radius is W_1. There are no other steps. The sum of $v_p(n)$ on the nonzero integers of an interval of radius pk in $[-D, D]$ is at most $2pk/(p-1) + L_p$, by counting the multiples of each p^h. For the two intervals at any level use $8k + 2L_p$ as a bound. The extra cost $3a$ per step at positive levels is at most $90k$ at each level, since there are at most $4W_1 + 2$ possible steps. Each of the six levels therefore costs at most $98k + 2L_p$, and $\kappa_* + 4\Delta + 6(98k + 2L_p) \le 25\kappa_*$. This bound pays for the ordinary point and for the additional steps of the gap point below.

The improved orders of the ordinary point supply the prescribed nonleading orders. At level zero, if the initial order is nonpositive, $s_i < \beta_i$ excludes a far negative jump, and then $-u - \eta \le 0$. At higher levels the initial order is $-u$. If all letters from $i_{-(l-1)}$ through i have the same sign, the leading side has $u \le \eta$ at all levels below l. Its improved orders then supply zero below l and $-u$ at levels $a \ge l$.

The sole remaining case for the leading side has $l = 2$ and opposite signs at i and i_{-1}. The noncut i_{-1} is thin and is followed by the thick letter i. It follows that $s_{i_{-2}} = 0$ and $g + 7k \le q$. On the leading side the jumps at i are at most η and those at i_{-1} are at most C. The unramified lifting-ideal estimate gives the exact multiplier E^{s_*} at i, taking exponents $0, s_*$. At each letter j all its upper jumps are at most $\beta_j + W_1 \le (p-3)s_*$. The preceding construction applies with $l = 2$, since

$$s_* + C(1 + p) \le (p + 2)q - (6p - 6)k < pq \le b < D.$$

The strict inequality uses $q < kp$ and $p \ge 7$. Thus the far negative jumps at i are $u < -N$. Raise the depth to D at $a = 0$, 3η at $a = 1$, and $-u + 3\eta\,\mathbf{1}_{a\le 5}$ for $a \ge 2$. The new steps at level one are near zero, and those at levels two through five are within W_1 of the same centers as for the ordinary points. The improved orders give the required leading orders. At $a = 0$ the coefficient argument gives the order on negative far jumps, and the order is at least zero on low jumps. At $a = 1$ all leading jumps are $\ge -\eta$, so order zero suffices there. From level two on the required order is $-u$. Multiply each point by the remaining power to put both sides over $p^{25\kappa_*}$ times their constants.

For the strict inequalities of the Frobenius walk, write $w = -u - v - 1$. For a comparison at the same cut, the other order in the minimum with $-v$ is at least $-v - A_a$. The ordinary depth therefore gives a sum at least 1. For the gap depth, at $a = 0$ the sum is at least $D + u + 1 - s_i - \eta > 0$ by $-u \le b + W$. At $a = 1$ it is at least $u + 2\eta + 1 > 0$ by $u \ge -\eta$. At higher levels it is again at least 1. For the leading comparison with the preceding cut in the ordinary case use the third point at i. Below l, the moving side has $v \ge -\eta$, which gives $\min(-v - A_a, 0) = -v - A_a$. At $a \ge l$ the two moving orders are bounded below by $-v - \eta$ if $a \le l + 2$, and by $-v$ otherwise. The extra η in the ordinary depth covers these levels. In the gap case use the third

point from i_{-2}, whose shift is zero. At level zero its transported order is zero and the moving orders have minimum at least $-v-\eta$, as $v \ge -\eta$. The sum is at least $D-v-\eta > 0$. At level one the moving jumps have $v \le \eta$, so their minimum with zero is at least -2η, less in absolute value than the depth 3η. From level two on the transported third order is $-w$, and the additional 3η in the gap depth covers both moving orders until both become standard. This proves the strict inequalities on every letter, for arbitrary choices of the two crystalline cocycles and the finite pairing.

The walk now gives the cup bound with exponent $51\kappa_*$. Since $W_1 \le b$, $L_p \le \lfloor \log_p(2b) \rfloor + 1$. With $k \ge 1$, the exponent is bounded by

$$51\kappa_* \le 10^4 k\big(\lfloor \log_p(2b) \rfloor + 1\big) \le 10^4 K_\ell m.$$

For the last inequality use $2b \le 4B$ and $\log_p(4B) \le \log_2(4B) < m$. It acts on the integral dual row or Tate block, so every reduction of this bound has the original precision of the finite pairing. Here $e_v = 1$, so the conversion to normalized length is at most $2f_vR$ times this exponent. Since $K_\ell \le 2\lambda_\ell/(\kappa t)$, the normalized length is at most $4 \cdot 10^4 f_v R m \lambda_\ell/(\kappa t)$. This fits $f_v R \,\mathrm{loss}_{\mathrm{crys}}(\ell)/x_L$ for $x_L = 1$ on the row and $x_L = t$ on a Tate block. The case $\lambda_\ell = H^3$ uses these same inequalities.

The loss from low jumps at large unramified primes. Keep $p = \ell > Q$, now with $\lambda = \lambda_\ell < H^3$. Let $A = 2(\lambda + 1)$, and keep η, W, β_i from the actual jump windows. The factorial gives $pA > 4B$ and $p > 2B/H^3 > 9G$. Since $A \le 2H^3$, the scale inequalities give

$$\beta_i + W + \eta \le (p-3)A, \qquad 6(A + \eta + W) < \beta_i.$$

For example $(p-3)A > 4B - 6H^3$, and $\beta_i + W + \eta \le B + \delta + W + \eta < 2B$. No nonzero multiple of p is within $6(A+\eta+W)$ of $0, \beta_i, -\beta_i$. Near 0 or the thin center the absolute value is less than p. Near the thick center the absolute value belongs to $(B - 30H^3, B + 30H^3]$, since $\delta < \eta$ and $6(A+\eta+W)+\delta < 30H^3$. One multiple there would contribute at least H^3 to λ.

Write T for an arbitrary crystalline side in the original class and S for the side with the actual tested dual lattice. The far sign on S is ϵ_i. The third side of a pairing P has jump $w = 0$ or $w = -u - v - 1$ for a jump u of T and v of S. These jumps lie within $2\eta + W + 1$ of $0, \beta_i, -\beta_i$, and are bounded above by $\beta_i + W + \eta$. The terms with both jumps far and those with both jumps low give the neighborhood of zero. With one jump far, the low jump and the translation by one give the neighborhood of its signed far center. In particular E^A kills the level-zero cokernels of all three sides by the unramified lifting calculation.

Write $r(t)$ for the orders $t - u$ at level zero and $-u$ at higher levels, with u the jump under test. An initial point over any constant at a letter has orders and depth $r(-A)$. The levels above zero will be kept unless explicitly changed. The depth can be raised with cost zero to $r(2A+\eta)$ on T or S, and to $r(A)$ on the third side. Every step is at level zero, and is in the three neighborhoods just considered. The unramified step cost is $v_p(n)$ at level zero with the cost at $n = 0$ defined as zero. By the improved-order rule, a point of depth $r(2A + \eta)$ obtained in this way on either side has level-zero order

$$o_0(u) = \begin{cases} 2A + \eta - u & u < -A, \\ \min(2A + \eta - u, 0) & u \ge -A. \end{cases}$$

The third side, when deepened to $r(A)$, has antitone order $z_0(w) = A - w$ for $w < -A$ and $z_0(w) = \min(A - w, 0)$ otherwise. We will use $o_0(u) \ge -u - A$, with $o_0(u) \ge -u$ on the far jumps, and $z_0(w) \ge -w - A$.

Two cases of zero defect. Suppose first that the allowed jumps on both sides miss $[-A, 0)$. Then $o_0(u) \ge -u$ on both sides. All letters are cuts separated by one step. At a cut the side with negative far sign leads and has $u \le \eta$, so its level-zero orders are also nonnegative. Leave the third point at orders $r(-A)$. The same-cut depth inequality at level zero has sum at least $A + \eta + 1$, using $w = -u - v - 1$. For the comparison with the previous cut, use the third point at the present cut. The moving side's level-zero

order has minimum at least $\min(-v, 0)$ with the shifted order from the preceding cut. Since its present jumps satisfy $v \geq -\eta$, the sum is at least $A + 1$. At higher levels all bounds are the standard ones and the sum is at least 1. Thus the walk gives orthogonality at exponent zero.

Suppose instead S has all jumps in $[-\eta, 0]$, as for the tested Tate block. For S, the unramified cokernel bound with upper jump bound zero supplies a point with orders and depth $r(0)$ at every letter. Raise its depth to $r(2A + \eta)$ with cost zero. Transport the point from the previous letter by Frobenius. This gives depth 1 at level zero and $2A + \eta - v$ at level one. Raise the level-zero depth to $2A + \eta - v$ by steps near zero, again with cost zero. Use these points on S, with orders 0 at level zero and $-v$ above, and with depth $2A + \eta - v$ at levels zero and one and $-v$ above. Use the preceding points on T and a third point of orders $r(-A)$ at each cut. Now S leads at every letter. For the same-cut comparison at level zero the T-depth sum is at least $A + \eta + v + 1 \geq A + 1$, and the S-depth sum is at least $\eta + 1$ since the T-order is at least $-u - A$. Above level zero the sums are at least 1.

When comparing the T points across one step, use the third point at the current letter if the T-far sign is positive there. Then $u \geq -\eta$. At level zero the moving order is at least $\min(o_0(u), 0) \geq -u - A - \eta$, so its sum with the third order and S-depth is at least 1. At level one the third order is standard and the moving order is at least $-u - A$. If the T-far sign is negative, use the third point from the preceding letter. At level zero this gives order zero and $u \leq \eta$, so the sum is at least $A - v > 0$. At level one the transported third order is $-w - A$, and the moving order is at least $-u - A$, giving a sum at least $\eta + 1$. Above level one the orders and depths are standard. This proves orthogonality at exponent zero on the one-sided dual block.

The price of increasing a single order. For a fixed letter i, write $m_i(u)$ for the jump multiplicity on that letter. Use the multiplicities summed over the coefficient embeddings on the letter, as in the chart. We first record a relation that applies to a kernel lattice and to an extension lattice. Suppose $c_a(u) < c_a(v)$ for every jump $v < u$ at i_{-a}. Among $f_v m_{i_{-a}}(u) + 1$ localized points of orders c there is a $\mathbf{Z}_p$-linear combination with one coefficient one which splits as the sum of a point whose orders are raised by one at level zero and a point whose order is raised by one only at (a, u).

To see this, use a basis of leading forms for jump u modulo forms of smaller jump at i_{-a}. The number of forms is at most $m_{i_{-a}}(u)$. Evaluating φ^{-a} of the points and removing the power $t^{c_a(u)}$ gives a matrix over B_{dR}^{+}. Upon reduction to C_p, a nonzero minor of largest absolute value among the full-rank minors expresses every row in terms of at most $m_{i_{-a}}(u)$ rows with integral coefficients. Lift the coefficients to A_{inf} and apply φ^{a} so that the evaluation at level a reads the chosen lifts. The residual point has gained the single order at (a, u). Forms of smaller jump have already gained that order by the strict inequality for the exponents. Apply the order-preserving coefficient retraction to make the coefficients lie in $\Lambda[[z]]$. It fixes the localized points and preserves the orders of their residual points. Modulo E, each coefficient is in Λ. The corresponding vectors of remainders lie in at most $m_{i_{-a}}(u)$ copies of Λ, each of $\mathbf{Z}_p$-rank f_v. Take a relation over $\mathbf{Q}_p$ between these vectors and divide by a coefficient of least valuation to get the stated primitive relation. The remaining multiples of E give the summand with all level-zero orders raised by one.

Let $D_i(c)$ be the constants at i that admit points of depth c and of orders at least $r(0)$. Assume $c_b(v) \geq -v$ at every level b and jump v. Fix (a, u) with $c_a(u) < c_a(v)$ for every jump $v < u$ at i_{-a}, and let c' be obtained by raising only $c_a(u)$ by one. If no multiple of p is strictly between $c_0(v)$ and $c_0(v) + p^a$ for each jump v at i, then

$$\operatorname{length}_{\mathbf{Z}_p}\big(D_i(c)/D_i(c')\big) \leq f_v m_{i_{-a}}(u)\big(v_p(c_a(u)) + 3a\big).$$

Here the constants are read in the trivial quotient of the side. To prove the inequality, let X_c be the points defining $D_i(c)$ and M_c the kernel points of orders c. The constant quotient is a quotient of $X_c/(X_{c'} + M_c)$. The single step operator gives $p^{v_p(c_a(u))+3a}$ times any element of this group equal to zero. It raises the depth and differs from the indicated scalar multiplication by a point of M_c, which has no pole.

Take $f_v m_{i_{-a}}(u) + 1$ lifts of elements killed by p in $X_c/(X_{c'} + M_c)$. For each, write its p-multiple as a sum of a point of depth c' and a kernel point of orders c. Apply the primitive relation to the kernel

points. For the resulting combination x of the lifts, $px - w$ is of depth c', where w is over zero and has orders c raised by one at level zero. An operator J of cost zero raises this level-zero depth from $c_0(v) + 1$ to $c_0(v) + p^a$. The steps avoid multiples of p by hypothesis. The improvement rule gives order $c_0(v) + p^a$ for Jw at level zero. If $c_0(v) + 1 > 0$ use the positive initial order, and otherwise the interval avoids zero and $c_0(v) + p^a \le 0$. Thus $Jw = E^{p^a} w'$ with $w' \in M_c$. On this unramified chart $E^{p^a} = E_a + pr$ for $r \in \Lambda[[z]]$. Put $y = rw' - (Jx - x) \in M_c$. Then

$$p(x - y) = J(px - w) + E_a w'.$$

The first summand keeps depth c', since a cost-zero operator differs from the identity by a kernel point of the input depth. The second has orders at least c', so its Galois differences also have depth c'. No E_b divides p. It follows that $x - y$ has depth c' and still has no pole. This is a primitive relation on the p-torsion quotient. Its dimension is therefore at most $f_v m_{i_{-a}}(u)$. Filtering a module killed by p^q by the kernels of successive powers gives length at most q times that dimension. This proves the estimate also when the individual point modules have infinite length.

When raising several orders at one level to an antitone target, raise the smallest jump whose current order is still below its target. Each lower jump is already at least its own target, which is at least the target at the chosen jump. This gives the strict inequality needed for each one-step estimate. Orders at the other levels may be lowered after this calculation without adding to the length.

The colength of the useful constants. Apply the calculation to the universal side on $Z^1_f(S)$ that sends each parameter to its own cocycle. Set

$$n_i = f_v \sum_{|v| \le \eta} m_i(v), \qquad r_i = f_v \sum_{v} m_i(v)$$

for the kernel lattice S. All constants have points of depth $r(2A + \eta)$ and orders at least $r(-\eta)$ by the initial calculation. Indeed negative far jumps are below $-A$. Also p^A times every constant has a point of depth $r(2A + \eta)$ without a pole, by starting with the cokernel killer p^A and then making the free steps from $r(0)$.

Remove the pole in stages $r(t)$ to $r(t + 1)$ for $-\eta \le t < 0$. The quotient of projected constants at each stage is killed by p^A. Lift any $f_v m_i(t) + 1$ constants, in particular representatives of elements of the p-torsion in the stage quotient, to points with orders $r(t)$ and the assigned depth. The primitive relation for the extension lattice gives a combination that gains the order at the jump t at level zero. Its other summand gains all level-zero orders, and hence also gains at t. On jumps $v < t$ the original order $t - v$ is already positive, so the depth improves it to $2A + \eta - v$. For $v > t$, the required $t + 1 - v \le 0$ follows by the same improvement rule and the assigned depth. Thus the combination has orders $r(t + 1)$ and retains depth $r(2A + \eta)$. The extension by the trivial parameter has the same multiplicity at $t < 0$ as its kernel. The resulting primitive relation on the projected quotient bounds the dimension of its p-torsion by $f_v m_i(t)$. If t is absent there is no order to raise and the same argument gives zero loss. The length of this stage is at most $A f_v m_i(t)$. Consequently $D_i(r(2A + \eta))$ costs at most $A n_i$ in colength.

If $\epsilon_i = 1$, retain this depth. If $\epsilon_i = -1$, take the least multiple N_i of p at least $\beta_i + W + 2A$. First impose depth N_i at level zero, using the single-depth estimate at $a = 0$. Lower any initially stronger bound to N_i. All these level-zero depths still exceed or equal $-v$ since $N_i \ge \beta_i + W + 2A$. The hypothesis excluding multiples in an open interval of length $p^0 = 1$ is automatic. On a far negative jump the steps $[2A + \eta - v, N_i)$ cost zero. A divisible integer there would be at most $N_i - p$, hence would be nonzero and within $6(A + \eta + W)$ of β_i. On a low jump the cost is at most $v_p(N_i!) \le A + 1$. Indeed $N_i < \beta_i + W + 2A + p \le (p - 1)A + p$, so integrality gives $N_i \le (p - 1)(A + 1)$ and the bound follows by Legendre's formula. The cost is at most $(A + 1)n_i$.

Next raise the depth at level one to

$$W + 1 - v + 2A\mathbf{1}_{|v| \le \eta}$$

and lower the level-zero depth to $\beta_i + W + 2A$. The hypothesis on multiples for the single-depth estimate at $a = 1$ holds because $(N_i, N_i + p)$ contains no multiple of p. The new depth at level one is antitone on the jump set. A crossing from a far negative jump to the low window changes $-v$ by more than $2A$. All steps are within the prescribed three neighborhoods at i_{-1}, so each has cost 3, including the step at zero if present. They cost at most $3(W + 1)r_{i_{-1}} + 6An_{i_{-1}}$. Thus in either sign the final constants with no pole have colength at most

$$(2A + 1)n_i + 3(W + 1)r_{i_{-1}} + 6An_{i_{-1}}.$$

Intersect the inverse images in $Z_f^1(S)$ of these spaces of constants at all letters. Denote the result by V. The predecessor permutes the letters, so the length of the quotient by V is at most

$$11 \sum_i \big((A + 1)n_i + (W + 1)r_i\big).$$

The construction of V uses only S and its jump windows. In particular it precedes the choice of T and the finite pairing.

The cup bound on V. Use all letters as cuts, and take the T points of depth $r(2A + \eta)$ and the third points with orders $z_0(w)$ at level zero and $-w$ above. For $d \in V$, the chosen S points have the two depths just computed. Their orders may be taken as $-v$ at all levels, except that when $\epsilon_i = -1$ the level-zero order can be taken as $\max(0, -v)$ by the improved-order rule. At the same cut the inequalities at level zero for the T and S depths give respectively lower bounds $A + \eta + 1$ and $\eta + 1$ when the corresponding depth is $r(2A + \eta)$. When the S depth at level zero is $\beta_i + W + 2A$, the same-cut lower bound there is $\beta_i + W + v + 1 \geq 1$. At higher levels the assigned depths are at least the standard orders, giving sums at least 1.

When $\epsilon_i = 1$, the T side leads. Its jumps at level zero are $\leq \eta$ and its orders are nonnegative there. Use the third point at the current cut. The moving order at level zero is at least $\min(-v, 0)$ with $v \geq -\eta$, giving a sum at least $A + 1$. Above level zero the sums are at least 1. When $\epsilon_i = -1$, the S side leads at i. Use the third point transported from i_{-1}. At level zero its order is zero, the T move has order at least $\min(-u - A, 0)$, and $u \leq \beta_i + W$. The sum is at least A. At level one the three orders in the inequality are

$$z_0(-u - v - 1), \qquad \min(-u, o_0(u)), \qquad W + 1 - v + 2A\mathbf{1}_{|v| \leq \eta}.$$

If v is low, use the lower bounds $-w - A$ and $-u - A$ for the first two terms. If v is far and u low, $|w| > A$ so $z_0(w) = A - w$, on either branch of the definition. The sum is again at least $W + 2$. If both are far, $|w| \leq W + 1$ gives $z_0(w) \geq -w - (W + 1)$, and the moving order is at least $-u$. The resulting sum is at least 1. Above level one the orders are standard with sum 1. The Frobenius walk now gives cup products that are coboundaries with exponent zero for every integral T and every finite pairing on the indicated lattices.

The cup-orthogonal quotient on the actual dual row is therefore a quotient of $Z_f^1(S)/V$ as a $\mathbf{Z}_p$-module. Each coefficient embedding contributes its actual multiplicity to a single letter. Thus $\sum_i n_i \leq f_v[O : \mathbf{Z}_p]\mu$ and $\sum_i r_i \leq f_v[O : \mathbf{Z}_p]R$. For an O-module of finite length its $\mathbf{Z}_p$-length divided by $[O : \mathbf{Z}_p]$ is its normalized O-length. Applying this to the cup-orthogonal quotient itself gives the charge

$$11 f_v\big(\mu(A + 1) + R(W + 1)\big).$$

The integral argument uses the given finite pairing throughout, hence also its reductions at the finite precision in the dual defect calculation. For the row, if $\lambda < H/5$ then $A < H/2$ for sufficiently large x. Hence both allowed sets miss $[-A, 0)$, so the defect is zero. The Tate branch always has zero defect here. For the remaining row case $\lambda \geq H/5$, $A + 1 \leq 3\lambda$ and $W + 1 \leq K_\ell$, using $H/Q \gg W + 1$. These inequalities include the boundary $\lambda = H/5$ when it is integral. For a row, $\mu \leq Rm/\kappa$ and $K_\ell \leq 2\lambda_\ell/(\kappa t)$. Thus the charge is at most $11 f_v(3Rm\lambda_\ell/\kappa + 2R\lambda_\ell/(\kappa t))$, within $f_v R \operatorname{loss}_{\mathrm{crys}}(\ell)$.

16.3 Paired ray coordinates and finite character sums

Fix a coefficient prime ℓ, a ray order $q_0 = \ell^{N_{\mathrm{ray}}}$ with $N_{\mathrm{ray}} \geq 1$, and a conjugation-stable finite set T_0 of places whose decomposition groups the next coordinate must kill. Include the places over ℓ. Suppose the proposed Frobenius is in the closed normal subgroup generated by inertia at finitely many places away from T_0. Include these places and the other exclusions in a finite set B. Enlarge T_0 to a conjugation-stable set T whose classes generate the ideal class group. Choose the extra prime-ideal generators outside B and its conjugate, using Chebotarev in the Hilbert class field. Adjoin μ_{q_0} and the q_0-th roots of a finite set of generators of the T-units. At each place supporting the proposed Frobenius these are roots of units of order prime to the residue characteristic. The resulting Galois extension is unramified there. Thus its trivial Frobenius prescription and the finite-quotient prescription for the proposed Frobenius agree on the intersection. The Chebotarev set has positive Dirichlet density over E. The places of absolute degree at least two have density zero, since their sum of q_P^{-s} is bounded at $s = 1$. Thus one can choose a place P of absolute degree one outside the exclusions and the ramified rational primes. It is split over the real field and satisfies $q_P \equiv 1 \bmod q_0$.

The ray exact sequence for the ring with the primes of T inverted is

$$\mathcal{O}_{E,T}^{\times} \longrightarrow k_P^{\times} \longrightarrow \mathrm{Cl}_P(\mathcal{O}_{E,T}) \longrightarrow \mathrm{Cl}(\mathcal{O}_{E,T}) \longrightarrow 1.$$

The last class group is trivial by construction. Under geometric reciprocity the ray group is therefore a quotient of $k_P^{\times}$ by the reductions of the T-units. Their images are q_0-th powers, so the cyclic quotient of order q_0 descends. Subtract its conjugate. Since $P \neq \overline{P}$, this gives a character h_P onto $\mathbb{Z}/q_0\mathbb{Z}$ on inertia at P, the inverse character at $\overline{P}$, and a trivial character on the decomposition groups at T. Normalize the coordinate to give 1 on the chosen tame generator δ_P. Change the new Frobenius lift by a power of δ_P to give $h_P(\phi_P) = 0$. It has the same image in the finite quotient used for its prescription, since P can be chosen outside that quotient's ramification.

The quotient characters are anticyclotomic and kill the ideles extended from $F = E^+$, also for $\ell = 2$. To see the latter claim, choose complex conjugation c of order two in G_F. Transfer to G_E^{ab} sends $g \in G_E$ to $g(cgc^{-1})$ and sends hc, $h \in G_E$, to $(hc)^2 = h(chc^{-1})$. Both have zero image under the coordinate. Global reciprocity identifies this transfer with extension of ideles. Local Hilbert 90 identifies the quotient of the E-ideles by the F-ideles with the norm-one ideles via $x \mapsto x/\bar{x}$. Principal norm-one ideles lift to principal ideles by field Hilbert 90. Consequently a torsion tuple (y_P) gives the norm-one character

$$\vartheta(x/\bar{x}) = j_\iota^{-1}\left(\prod_P y_P^{h_P(\mathrm{Art}_E^{\mathrm{geom}}\, x)}\right)^{-1}.$$

Thus the tuple gives the inverse avatar, and its vanishing coordinate at each decomposition group above ℓ makes the inverse avatar trivial there. At a split pair a diagonal unit dies by the transfer, and the principal units die by tameness. Hence the character can be read on the order classes with the installed finite support. This is Artin reciprocity and its transfer law, with geometric uniformizer symbols.[4, V, §3, especially Theorem 3.6]

Reciprocity for a tied Selmer group. Use the finite local conditions of one balanced row ρ on O^n, where O is a coefficient DVR finite over $\mathbb{Z}_\ell$. Write c for complex conjugation on the absolute Galois group G_E, and χ for the cyclotomic character of G_E. Let $a \mid \ell^{M_\ell}$ be a nonunit and $V = (O/a)^n$. At the conjugation-stable set of fixed places S use the prescribed modules of local cocycles containing coboundaries. At an existing auxiliary pair use the same slope t for the two relations

$$z(\delta_P) = tz(\phi_P), \qquad z((c\delta_P)^{-1}) = tz((c\phi_P)^{-1}).$$

Here cg means conjugation of g by complex conjugation. The ties contain coboundaries. Tame inertia acts trivially at an auxiliary place, and a nonzero slope is allowed only where both Frobenius readings act

trivially on V. Away from these places require unramifiedness. Normalize the global cocycles by $z(\gamma) = 0$, where the distinguished inertia element γ has $\chi(\gamma) = 1$, $\rho(\gamma) = sI$, $\rho((c\gamma)^{-1}) = sI$, and $s^2 - 1$ is a unit. This normalization is available and unique on each cohomology class, since the coboundary at γ is $(s-1)$ times its defining vector. Denote the resulting module by X.

Assume X contains a cocycle z having a unit coordinate in some value. Suppose that at every admissible new pair the two Frobenius values of this z are the two tame values of a single cocycle with the old conditions and ramification allowed at the new pair. We show that $\operatorname{length}_O X \le \beta$, the balanced-row allowance. The contrapositive gives the new pair used in the slope argument.

Take the images of the prescribed finite local cocycle modules in local H^1. If there is a choice of decomposition group in the local length bound, use one attaining the least length at each place. The global Selmer group with these local conditions contains the image of X. At an auxiliary place the local module of tied cocycles injects into V by evaluation at Frobenius. Indeed the tie determines the value at a generator of tame pro-ℓ inertia, wild inertia gives zero, and the representation is unramified there. Write Y for the dual Selmer group with the annihilator local conditions on $V^{\mathrm{D}} = \operatorname{Hom}(V, \mu_{\ell^\infty})$. The Greenberg–Wiles count gives

$$\operatorname{length}_O X \;\le\; \operatorname{length}_O Y + \sum_{v \in S} \operatorname{length}_O Z_v - (|S| + f) \operatorname{length}_O V,$$

where Z_v is the chosen local cocycle module and f is the number of complex places of E. The finite-place contribution with local cocycle module Z_v is $\operatorname{length}_O Z_v - \operatorname{length}_O V$, after subtracting the local coboundaries and invariants. This is nonpositive at the tied places. Global invariants vanish on both modules by the scalar at γ. The complex places give $-f \operatorname{length}_O V$. These are the terms obtained by taking orders in Poitou–Tate and inserting the local and global Euler characteristics. In the length identity the local Euler characteristic contributes $-[E_v : \mathbb{Q}_\ell] \operatorname{length}_O V$ at each $v \mid \ell$, and the global Euler characteristic contributes $f \operatorname{length}_O V$ with the opposite sign. Their sum is the stated complex-place term because $[E : \mathbb{Q}] = 2f$. This applies also at $\ell = 2$ over this totally imaginary field.[20]

Let U be the common kernel of the actions on O^n/ℓ^{M_ℓ}, its conjugate, and $\mu_{\ell^{M_\ell}}$. Let H be the closed normal subgroup generated by the commutators $[g, \gamma]$ and their conjugates under c. The two scalar conditions give $H \subset U$. Let δ_0 be the tame generator at the helper prime. It lies in U as well. For each $h \in H$ take a new pair with Frobenius arbitrarily close to $\delta_0 h$, using the preceding ray construction. The target is in the normal closure of inertia at the helper and the fixed places away from ℓ. The coordinate just constructed kills the groups at the nonhelper auxiliary places. The old coordinate values at the target are 1 at the helper and zero at the other places, since old coordinates kill commutators. Require the new Frobenius to lie in the open subgroup U. The pair may therefore be assigned any slope in this row and zero in the other rows.

Write x for a dual cocycle representing an element of Y. Identify the roots of unity of order ℓ^{M_ℓ} with their additive cyclic group via a fixed generator. On U the restrictions of x, z are homomorphisms, and the restriction of x is independent of its cocycle representative. At the new pair, x is unramified. By the hypothesis on z, let z' be a relaxed cocycle with the two tame values equal to the two Frobenius values of z. The global sum of local invariants of $tz' \cup x$, for $t \in O$, vanishes. All terms outside the new pair vanish by the local conditions and unramifiedness. With a normalized tame generator δ the remaining term at P is $x(\phi_P)(tz'(\delta))$, up to the common sign of the local cup convention. Indeed x is unramified and both actions at the new place are trivial. The local invariant of the cup of an unramified character with the Kummer class of a uniformizer evaluates that character at Frobenius, with the sign determined by the order of the factors. At $\overline{P}$ the corresponding normalized generator is $(c\delta)^{-1}$ since c inverts the roots of unity. The actual chosen generator and its conjugate inverse are the same multiple of these normalized generators

[20] Tate's duality and Euler characteristics in Milne, *Arithmetic Duality Theorems*, 2nd ed., I.2.8, I.4.10, and I.5.1. Compare the Selmer order calculation in Wiles, Proposition 1.6, p. 472.

in tame pro-ℓ inertia, so the relation holds on the two actual values together. Since $z((c\phi_P)^{-1}) = -z(c\phi_P)$, prescribing the values of x, z on the Frobenius and its conjugate yields

$$x(\delta_0 h)(tz(\delta_0 h)) - x(c(\delta_0 h))(tz(c(\delta_0 h))) = 0 \qquad (h \in H).$$

Both prescriptions use finite quotients because the cocycles have finite target. At those quotients the change of Frobenius lift by the new inertia is trivial.

For $h, h' \in H$ put

$$B_t(h, h') = x(h')(tz(h)) - x(ch)(tz(ch')).$$

Taking the mixed difference of the last equality at $hh', h, h', 1$ gives $B_t(h, h') + B_t(h', h) = 0$ by additivity on U. Let $h_\gamma = \gamma h \gamma^{-1}$. Equivariance on U gives

$$B_t(h_\gamma, h') = B_{st}(h, h'), \qquad B_t(h', h_\gamma) = B_{s^{-1}t}(h', h).$$

Indeed $x(ghg^{-1})(v) = x(h)(\chi(g)\rho(g)^{-1}v)$ and $z(ghg^{-1}) = \rho(g)z(h)$. Apply antisymmetry with h_γ and st, then subtract antisymmetry with h and t. Since B_t is additive in t, it follows that $B_{(s^2-1)t} = 0$. The unit condition gives $B_t = 0$ for all t, including at coefficient prime 2.

The values $z(H)$ span V over O. In fact $z([g, \gamma]) = (1-s)z(g)$, so one such value has a unit coordinate, and its orbit spans V because $\rho(G_E)$ spans the full matrix algebra over O. The equality $B_t = 0$ now gives an O-linear map $\eta : V \longrightarrow (V^{\mathrm{D}})^c$ with $\eta(z(h)) = x(ch)$. A linear relation among the $z(h)$ gives a linear combination of the $x(ch)$ vanishing on every $tz(ch')$ by $B_t = 0$. This proves that the combination vanishes everywhere. The same conjugation identities show that η is G_E-equivariant. If $\eta = 0$, then $x(H) = 0$. Comparing $x(u)$ with $x(\gamma u \gamma^{-1})$ for $u \in U$ gives $x(U) = 0$, since $[u, \gamma] \in H$ and $s - 1$ is a unit. The class inflates from $H^1(G_E/U, V^{\mathrm{D}})$. In this quotient γ is central and acts by s^{-1} on V^{D}. For any cocycle on the quotient, comparing its values on γg and $g\gamma$ writes it as a coboundary, using $(s^{-1} - 1)^{-1}$. Thus Y injects into the module of equivariant maps.

Let g_* be the element from the balanced-row bound, with $\rho(g_*) = \alpha I$ and $\rho((cg_*)^{-1}) = \alpha' I$. Equivariance gives $\eta((\alpha - \chi(cg_*)\alpha')V) = 0$. An equivariant map is determined by its value on one vector with unit coordinate, because the orbit of that vector spans V. The length of their module is therefore at most $n\, v_O(\alpha - \chi(cg_*)\alpha')$. The balanced inequality is

$$\sum_{v \in S} \operatorname{length}_O Z_v + n\, v_O(\alpha - \chi(cg_*)\alpha') \;\le\; (|S| + f) \operatorname{length}_O V + \beta.$$

Combining the two bounds proves $\operatorname{length}_O X \le \beta$. This supplies the unseen-class implication in the finite-length induction for the actual crystalline relation modules. The ray order can still be chosen after the bound on the number of auxiliary pairs.

The slope at two successive coefficient lengths. For one row and a new admissible pair let Y be the relaxed module modulo $a = \varpi^s$, and $A, B : Y \longrightarrow (O/a)^{2n}$ its paired tame and Frobenius evaluations. The old tied module is $\ker A$, because vanishing on both tame generators restores unramifiedness at the two new places. Let D be the sum of the kernels of $uA - B$ over $u \in O$, so that $B(D) \subseteq A(D)$. The inverse image P of $A(D)$ in O^{2n} is free of rank $2n$. Choose a lift $\sigma : P \to D$ of reduction through A, and an endomorphism T of P lifting $B\sigma$. Both lifts exist because P is free. If u avoids the roots of the reduction of $\det(U - T)$, then $u - T$ is invertible on P. Applying $uA - B$ to $\sigma(u - T)^{-1}P$ proves that $(uA - B)D = A(D)$. Since its kernel on Y is contained in D, it follows that

$$\operatorname{length}_O \ker(uA - B) \le \operatorname{length}_O(\ker A \cap B^{-1}A(Y)).$$

There are at most $2n$ exceptional residue classes at each length. The balanced residue-field bound allows one u that is a unit and works at both s and $s - 1$. The new slope in this row is u^{-1}.

For a fixed finite list of rows, apply induction on s, starting at $s = 0$ with the helper prime and zero slopes. A pencil with d pairs has length at most the length at the original places plus $2nsd$ in each row, by evaluating at the paired tame generators. Thus a bound for the number of pairs at $s - 1$, independent of ray order, gives such a bound C_s for the number of pairs plus the sum of the lengths at s on a pencil already working at $s - 1$. Among pencils satisfying that sum bound which still work at $s - 1$, minimize the sum of lengths at s. If some row exceeds its β, its module contains a cocycle with a unit coordinate, since otherwise division by ϖ would put the whole module in the one at $s - 1$ by homogeneity of the cocycle equations and local relations. The paired reciprocity implication supplies a new admissible pair with $B(\ker A) \not\subseteq A(Y)$. The new slope makes the length drop strictly at s and prevents an increase at $s-1$. Give the other rows slope zero at this pair, so they retain the previous unramified condition there. The strict integral length drop pays for adding one pair and preserves C_s, contrary to minimality. This proves the simultaneous tied bounds at s with a bound for the number of pairs independent of the ray order. The argument reaches $s = v_O(\ell^{M_\ell})$.

Finite interpolation of the signed values. Fix a critical pair $p = (\nu, z)$ from the ramified signed list. Its tame part is nontrivial at each fixed chosen place and trivial away from those places. For $z \geq 2$, put $n = 2z$, $k_\tau = c_\tau - z$ and $k = \sum_\tau k_\tau$, where $-c_\tau$ is the weight of ν. Fix the modulus I of paired auxiliary places, write $q = N(I)$, and let G be the common multiple of the orders of the class stabilizers. The section at each added place requires the first coordinate to equal 1 in the residue field. Keep this section even if the character on that modulus is imprimitive. For every $\ell \nmid q$ the fixed integer

$$D_z = g^{2fn}(gd)^k$$

clears its point kernel in the integral Weyl gauges ($f = \#\Phi$). In fact the Fourier expansion of the residue indicator on a box of n entries puts g^{2fn} times the normalized vacuum section in the standard-data module of level g. The added part of the modulus has norm prime to ℓ. After moving to integral frames, the integral CM jet at scale gd per applies in degree k. The point kernel uses $\bar{\pi}_G$ per, and $\bar{\pi}_G$ is integral, so multiplying by $(gd)^k$ suffices. Returning to the original frames gives the asserted Weyl gauges.

At the chosen modulus write the lift as $\Xi_\vartheta = \sum_a \vartheta(u_a)^{-1} K_a$, summing over the source classes and the residue lifts. Its squared norm is a sum of $\vartheta(u_a/u_b)\langle K_a, K_b\rangle$, with the inner product linear in the second entry. The common class-level sum over the target of the pairings of the cleared kernels is integral after multiplication by G. Indeed the Weyl pairing is integral in the point gauges and the only class weights are the reciprocal stabilizer orders. Integrality at the conjugate reading covers the conjugated entry of the pairing. Division by q^{n-1} is a unit at the same places. Thus

$$F(I, \vartheta) = GD_z^2 \|\Xi_\vartheta\|^2 / q^{\,n-1}$$

is the evaluation of a single finite sum in the ray group ring with coefficients integral away from q. In the rank-one norm calculation the modulus fixes the source support and the target index. At an added place the support in source units is the kernel of reduction, on which every allowed tame character is trivial. The same local calculation therefore contributes $q_v^{n-1}(1 - q_v^{-n})$ also when the character at v is imprimitive. The spherical Euler factors and the factors at the fixed places are those in the cited calculation. After dividing by q^{n-1} one obtains a fixed nonzero multiple of the incomplete Hecke value at z. For exact additional conductor the quotient by the trivial value is $\rho_{p,+}(\vartheta)$.

The rank-two determinant period. For $z = 1$, the test is the determinant period Λ_ϑ on $U(2)$. Put $\alpha = \nu\, \vartheta\, v_{U(1)}$ and $Q(X) = \prod_\tau \det(X_\tau)^{c_\tau - 1}$, of total degree $h = 2\sum_\tau (c_\tau - 1)$. Each $c_\tau \geq 1$. Here is the clearing factor for this test. Set $\delta = g^2$ and let v_0 be the scalar that normalizes the rank-one vacuum. The scalar

$$N_p = \delta^{8f}\, v_0^4 (\delta d\, \mathrm{per})^h$$

is nonzero and independent of the modulus and the reading. The subscript p indexes the critical pair $(\nu, 1)$. The four entries of the residue matrix give a Fourier denominator dividing δ^{8f} at $\ell \nmid q$, since the modulus

ideal contains $q\delta$. Multiplying the vacuum by v_0^4 and by this integer puts it in the standard-data module of level δ. Integral finite translates preserve that module. Pairing its Taylor series at scale δd per with the integral homogeneous polynomial Q multiplies the kernel by $(\delta d\,\mathrm{per})^h$. Thus N_p times the kernel is integral in every such reading at integral arguments.

For $x = (u, A)$, where u represents a rational double coset at maximal compact level and A is an invertible residue matrix for the fixed modulus, write $g_x = u\widetilde{A}$. Let K_1 be the untwisted kernel for the indicator of the identity matrix at this modulus, and set

$$b_x = \frac{G}{|\operatorname{Aut}(u)|}\, N_p K_1(g_x).$$

These numbers are integral away from q. For each coefficient prime, rational translation can move g_x into the integral subgroup there. For definite $U(2)$ this approximation follows already from the Cayley parametrization: its inverse is defined on a dense open subset of each required local coset, and weak approximation on the space of skew-Hermitian matrices gives a simultaneous rational parameter. The determinant character from a rational translate cancels the determinant weight of Q; the kernel is unchanged. The character of K_1 is $\nu\, v_{U(1)}$, with its original weight and tame part. Expanding the residue character as translates of the identity indicator now gives

$$GN_p\Lambda_\vartheta = \sum_x b_x\, \vartheta(\det g_x).$$

In this identity the modulus is fixed and the character varies. The square $F(I, \vartheta) = |GN_p\Lambda_\vartheta|^2$ has coefficients $b_x\overline{b_y}$ on $\det(g_x)/\det(g_y)$. Each conjugate is integral because the preceding test holds for all coefficient embeddings.

For an exact twist the analytic input here is the normalized complementary exchange on $W = I_2$, with larger space I_2 and smaller space 0. The rank difference is two and the half-difference is one. The finite partners are $\chi(\det X)$, extended by zero on singular matrices, and the function 1 on a single point. On the fixed places χ has the original nontrivial tame component. On the disjoint added places the exact twist has nontrivial tame component. At each real place the exchange weight is $(c_\tau - 1) + 1 = c_\tau \geq 1$, as required in rank two. The partner identity follows from the block partner formula with the rectangular block empty. The real polynomial on the large side is exactly Q; on the small side the empty determinant is one. Consequently that exchange, evaluated at the identity, gives

$$\gamma_\vartheta\Lambda_\vartheta = 1, \qquad |\gamma_\vartheta|^2 = C\big/\Pi_{S_\vartheta}((\nu\,\vartheta)^{-1}, 1).$$

The mass of the zero-dimensional quotient is absorbed in the same fixed constant C by normalizing the first identity. Here S_ϑ consists of the fixed chosen places and the added ones. By definition $\Pi_S(\xi, 1) = (\prod_{v\in S} q_v)\,|L^S(1, \xi)|^2$, where the omitted Euler places include the conjugates. Dividing by the identity at the trivial twist cancels the mass, G, N_p and the contribution of the fixed conductor. This proves

$$F(I, \vartheta)/F(\varnothing, 1) = q\,\rho_{p,+}(\vartheta)^2.$$

Evaluating the finite sums at the inverse character gives the minus reading. The norm relation on enlarging an imprimitive modulus follows from the residue computation below.

Add an unramified pair $P, \overline{P}$. Write $Y = (\nu\,\vartheta)_{\mathrm{bc}}(\varpi_P)$ for the unitary base-change value. At $z \geq 2$ omission of the two Euler factors gives $(1 - Yq_P^{-z})(1 - Y^{-1}q_P^{-z})$. At $z = 1$ the added residue condition is invertibility. On integral matrices X one has

$$\mathbf{1}_{\ker \overline{X}=0} = \sum_{U\subseteq \ker \overline{X}} (-1)^{\dim U} q_P^{\binom{\dim U}{2}}.$$

This is the finite q-binomial theorem evaluated at one. The subspace condition at U is given by translating the split vacuum by the lattice which is the inverse image of the prescribed residue column condition, using

transpose in the matrix coordinates for the right action. If $d = \dim U$, choose bases reducing it to the first d coordinates and use a diagonal matrix with d entries ϖ_P and the rest one. The split Weil action gives the indicator of this column condition multiplied by A^d, where $A = q_P\,\alpha_{\mathrm{bc}}(\varpi_P)v(\varpi_P)^{-2} = q_P Y$. Here q_P is the determinant factor of the vacuum in rank two, and $(v_{U(1)})_{\mathrm{bc}} = v^2$ by restriction to norm-one ideles. Thus the local sum is the split Hecke polynomial at A^{-1} applied to the vacuum. Integrating over the class sum leaves every right translate unchanged. The sum of the Hecke coefficients is $\prod_{d=0}^{1}(1 - q_P^d A^{-1})$, again by the finite q-binomial theorem. Squaring gives $\prod_{d=0}^{1}(1 - Y q_P^{d-1})(1 - Y^{-1} q_P^{d-1})$. This residue-indicator calculation gives the norm relation for the rank-two period.

Pull back each of these finite character sums by the norm-one ray map, using the inverse avatar. Every character argument becomes a monomial in the installed coordinates, with powers read modulo $\ell^{N_{\mathrm{ray}}}$. In the signed value polynomial the coefficient field can be enlarged to contain all the cleared coefficients. At the chosen place $w \nmid \ell$ the fixed tame character is nontrivial on a generator. Its value ξ is a nontrivial root of unity of order dividing $q_w - 1$. Evaluating $(T^{q_w-1} - 1)/(T - 1)$ at 1 shows that $v_\ell(1 - \xi) \le v_\ell(q_w - 1)$, since all other factors are integral. The same applies to the inverse character. The helper Frobenius can be placed sufficiently close to that inertia element for all the finite tests, including the cyclotomic powers in the Euler factors. Consequently the line polynomial for modulus consisting only of the helper is bounded at $u = 1$ with the loss of the helper Euler factors. At a later auxiliary place the value line has trivial tame coordinate. All the nonhelper coordinates are 1, so the inverse avatar χ_ϑ reads u at its Frobenius. For sign $s = \pm 1$ the base-change value to exponent $\epsilon = \pm 1$ reads

$$j_\iota\big(v_{\mathrm{bc}}(\varpi_P)^{\epsilon}\big)\chi_\vartheta(\phi_P)^{-s\epsilon}.$$

The omitted factors are therefore at most four linear guards $1 - \xi' u^{\pm 1}$ with ξ' a unit, up to multiplication by a unit. Include them in the one-variable specialization on the helper line. The helper factors cost at most $4v_\ell(q_w - 1)$, and one can choose the per-guard tolerance so that eight times the coordinate bound times this tolerance fits the additional slack. The coordinate bound precedes the choice of ray order. Then apply the common specialization to the resulting value and Selmer polynomials. The conductor powers are units at ℓ, so the inverse ratio bounds apply to the same fitted character as the local determinant tests.

16.4 Specialization at the first m powers

We choose the order of the paired ray coordinates before choosing their primes. The valuation argument treats all exponents $1 \le j \le m$, including those divisible by ℓ.

Normalize $v_\ell(\ell) = 1$. For a polynomial F over the valuation ring in $\overline{\mathbf{Q}}_\ell$ let $\|F\|_s = \max_i |[T^i]F|_\ell s^i$. This Gauss norm is multiplicative. For instance, in a product the coefficient at the sum of the lowest maximizing indices has a unique contribution of maximal weighted norm. Fix $0 < s < 1$ with $|\ell|_\ell < s^{\ell-1}$ and finitely many positive targets y_i. For all sufficiently large N_{ray}, if $\|F_i\|_s \ge y_i$ for every i, there is one $\zeta \in \mu_{\ell^{N_{\mathrm{ray}}}}$ with $|F_i(\zeta - 1)|_\ell \ge s y_i$ for every i.

Put $v = v_\ell$, $\Delta_s = -\log_\ell s > 0$, and $Q = \prod_i F_i$. For each $n \ge 1$, let X_n consist of $\zeta - 1$ for the primitive ℓ^n-th roots ζ. Its $d_n = \varphi(\ell^n)$ distinct elements have valuation $t_n = 1/d_n$. Consider $n \ge n_s$, where n_s is the least n with $t_n \le \Delta_s$. Let k_n count the roots of Q of valuation t_n, with multiplicity. For a polynomial F with leading coefficient c_F, multiplicativity and linear factorization give

$$-\log_\ell \|F\|_{\ell^{-t}} = v(c_F) + \sum_{F(a)=0} \min(t, v(a)).$$

Each root contributes a nonnegative amount to the difference between this expression at $t = \Delta_s$ and at $t = 0$. A root on the n-th circle contributes t_n. Since all F_i are integral,

$$\sum_{n \ge n_s} k_n t_n \le -\log_\ell \|Q\|_s + \log_\ell \|Q\|_1 \le -\log_\ell \prod_i y_i.$$

Suppose every $x \in X_n$ fails a test. At such an x, a failed factor has $|F_i(x)|_\ell < sy_i \le s\|F_i\|_{\ell^{-t_n}}$, and every other factor is bounded by its Gauss norm at this radius. Thus $v(Q(x)) + \log_\ell \|Q\|_{\ell^{-t_n}} > \Delta_s$. This excess equals $\sum_{v(a)=t_n}(v(x-a)-t_n)$, by ultrametricity for the roots on the other circles. Truncating each nonnegative summand at Δ_s leaves a sum at least Δ_s, including when $Q(x)=0$. For a root a on this circle choose $x_a \in X_n$ maximizing $v(x_a-a)$. At the other points $v(x-a) \le v(x-x_a)$. The differences $x-x_a$ are unit multiples of distinct factors $1-\xi$, $\xi \in \mu_{\ell^n} \setminus \{1\}$. The product of all these factors is ℓ^n, so their valuations are nonnegative and sum to n. The total truncated contribution of a on X_n is therefore at most $\Delta_s + n$. Summing in the two orders gives $k_n/d_n \ge \Delta_s/(n+\Delta_s)$. If every point through level N_{ray} failed, the sum of these lower bounds would grow without bound with N_{ray}, contrary to the fixed bound $-\log_\ell \prod_i y_i$. This gives a threshold independent of the degrees and coefficients of the polynomials.

For several variables fix a coordinate bound q, an integer $a_0 > 0$, a tolerance $\epsilon_0 > 0$, and bounds $B_i \ge 0$. For each i and $1 \le j \le m$, suppose an integral polynomial $F_{i,j}$ has a point $u_{i,j}$ satisfying

$$v_\ell(F_{i,j}(u_{i,j})) \le B_i, \qquad |(u_{i,j})_v - 1|_\ell^{a_0} \le |\ell|_\ell \quad \text{at each coordinate } v.$$

Choose s close enough to one that

$$|\ell|_\ell < \min(s^{\ell-1}, s^{ma_0}), \qquad s^q \ge |\ell|_\ell^{\epsilon_0}.$$

In the next unchosen coordinate replace the variable by $(1+T)^j$. Some j-th root α of $(u_{i,j})_v$ has $|\alpha-1|_\ell \le s$. Otherwise the product of all $|\alpha-1|_\ell$ would exceed s^j, contrary to the assumed bound on $|(u_{i,j})_v - 1|_\ell$. Evaluating at $T=\alpha-1$ shows that the one-variable Gauss norm is at least the lower bound already attained with the unchosen coordinate $u_{i,j}$. The preceding one-variable argument now replaces that coordinate by ζ_v^j for one $\zeta_v \in \mu_{\ell^{N_{\mathrm{ray}}}}$, losing at most a factor s for every i, j. Repeat for all coordinates. The finite list of targets for each coordinate is the original list scaled by the accumulated powers of s, so one threshold for N_{ray} suffices. This gives

$$v_\ell(F_{i,j}((\zeta_v^j)_v)) \le B_i + \epsilon_0 \qquad \text{for all } i, j.$$

Monic integral auxiliary test polynomials D_b can be included by applying the one-variable argument also to $\prod_b D_b(1+T)$. Put $L_* \ge \sum_b \deg D_b$. The product has Gauss norm at least s^{L_*}. Taking $s^{L_*+1} \ge |\ell|_\ell^{\epsilon_0}$ forces $v_\ell(D_b(\zeta_v)) \le \epsilon_0$ for every b, v, since each $D_b(\zeta_v)$ is integral. With $D_b(X) = X^b - 1$ for $1 \le b \le m$ we can take $L_* = m(m+1)/2$. Each required power then has nontrivial inertia on each chosen pair.

Finite polynomial relations at a fixed coefficient depth. Write $p = \ell$, let O be a finite coefficient DVR with $v(p) = 1$, and let $K_0 = M_\ell$ be the positive integer defining the modulus p^{K_0}. Fix an auxiliary pencil with coordinates $h_P : G_E \to \mathbf{Z}/p^{N_{\mathrm{ray}}}$ and a finite allowed set of ramification containing the original places and the auxiliary pairs. For a row $\rho : G_E \to \mathrm{GL}_a(O)$, define the matrix function

$$B_{\mathrm{Y}}(g) = \rho(g) \prod_P Y_P^{[h_P(g)]},$$

where the brackets choose the representative in $[0, p^{N_{\mathrm{ray}}})$. Impose the crossed-homomorphism equations with this matrix function, the original unramified and integral linear relations on the cocycle, and the fixed normalization at the distinguished inertia element. We may also impose a specified subset of the linear ties at the auxiliary pairs. For a finite extension DVR O' and a point $\mathbf{y} \in O'^J$ satisfying $y_P^{p^{N_{\mathrm{ray}}}} = 1 \bmod p^{K_0}$ for all P, these equations define the intended module of continuous cocycles over $A' = O'/p^{K_0}O'$. The defining equations have polynomial coefficients even at other $\mathbf{y}$. Such other points need not define a Galois representation.

One finite set of evaluation coordinates suffices for all points satisfying those congruences. Indeed, first factor out the closed normal subgroup of G_E generated by the inertia groups outside the allowed finite set. The action and every unramified crossed homomorphism factor through the resulting group Π.

Choose an open normal subgroup U of Π in the kernels of all h_P and of the row representation modulo p^{K_0}. Over U a cocycle restricts to a homomorphism into a finite abelian p-group killed by p^{K_0}. The largest abelian quotient of U killed by p^{K_0} is finite. In fact U comes from a finite extension of E, with ramification still confined to a finite set, and its abelian p-quotient of bounded exponent is finite by the ray class description. The finite class group and the local unit groups at the allowed finite places give the bound. At residue characteristic p the deep units are finitely generated over $\mathbf{Z}_p$, and at the other allowed places the p-part of the unit quotient is finite. The kernel of this abelian quotient is characteristic in U and normal in Π. All the cocycles under consideration vanish on the kernel and therefore are determined by their values on one fixed finite set of coset representatives. The original relations are groups of linear equations with coefficients in O evaluated on these same coordinates. The crossed-homomorphism equations are polynomial linear forms in the coordinates by the displayed formula for B_{Y}.

Include the relation rows $p^{K_0} e_i$ as well. If d is the number of coordinates, the O-module generated by the specialized rows at a point $\mathrm{y}_0 \in O^J$ has full rank d. Choose its basis and express each member as a finite O-linear combination of specialized relation rows. Use the same combinations of their polynomial rows to form a square matrix and let $F(\mathrm{Y})$ be its determinant. The matrix at y_0 has Smith exponents at most $v_O(p^{K_0})$, since the row module contains $p^{K_0} O^d$. Hence $v(F(\mathrm{y}_0))$ is the normalized length of the intended coefficient module at that point. At every other admissible point every cocycle lies in the kernel of the specialized square matrix modulo p^{K_0}. Smith form there bounds its normalized length by $v(F(\mathrm{y}))$, with the infinite bound allowed when $F(\mathrm{y}) = 0$. If the center lies in O', first form the same polynomial relations over O' and choose the Smith rows there. All inequalities continue to hold over larger DVRs. This uses $v(p) = 1$ in every extension, and length is divided by the ramification index in whichever coefficient DVR is used.

Cost of removing ties on the helper line. Let P_0 be the helper and J the full auxiliary set. The anticyclotomic ray coordinates give, at $P \neq P_0$ and at the inverse-conjugate reading, character value y_P on tame inertia and y_{P_0} on the selected Frobenius. Write t_P for the slope of an individual row there. A nonzero t_P implies that $\rho(\phi_P)$ and $\rho((c\phi_P c)^{-1})$ are the identity modulo p^{K_0}. For a root u of order dividing p^k set

$$y_{P_0} = u, \qquad y_P = 1 + t_P(u-1) \quad (P \neq P_0).$$

If $N_{\mathrm{ray}} \geq k + 2K_0$, the reductions of these entries define a character modulo p^{K_0}. To see this, raise the formula for y_P to p^k modulo p, obtaining 1, then use the binomial theorem successively modulo higher powers of p. The order of its reduction divides p^{k+K_0}, which suffices for the ray order. In addition

$$|y_P - 1|_p^{p^k} \leq |p|_p.$$

For $u = 1$ this is immediate. Otherwise its order is p^b with $b \leq k$, and $v(u-1) = 1/\varphi(p^b)$.

For a cocycle z at a new pair, use its affine local representation $g \mapsto (B_{\mathrm{y}}(g), z(g))$ over $O'/p^{K_0}O'$. Wild inertia and the prime-to-p part of tame inertia have trivial image. Their linear action is trivial, and their translation action would be a homomorphism into a p-group. The linear action of the tame pro-p part is scalar of exponent dividing p^{k+K_0}. Its affine exponent then divides p^{k+2K_0}. Since $q_P \equiv 1 \bmod p^{N_{\mathrm{ray}}}$, the tame Frobenius relation makes the images of δ_P and ϕ_P commute. The inverse-conjugate readings have the same property. At either reading, set $B = \rho(\phi_P)$ with the matching choice of Frobenius. For $P \neq P_0$ the affine equality is

$$(uB - I)z(\delta_P) = (y_P - 1)z(\phi_P).$$

It follows that $(uB-I)(z(\delta_P) - t_P z(\phi_P)) = 0$, since $t_P(I-B) = 0$ in the coefficient module. The quotient obtained by dropping all these ties injects via the displayed discrepancies into the indicated kernels, one for each of the two readings of each pair. Smith form bounds the increase in normalized length by the sum

of $v\det(uB-I)$. Here B in the determinant can be the fixed integral characteristic-zero Frobenius matrix. At $u=1$, the helper tie is $z(\delta_{P_0})=0$. The same argument uses the kernel of $B-I$ for the unrestricted Frobenius matrix B. This bounds the increase when only that tie is dropped at $u=1$.

Initial points and fixed costs. Fix all row modules, signed families and the bound J_0 on the number of pencil coordinates before choosing a ray order. The bound J_0 comes from the integral length descent in the tied-pencil induction. Work with two separate choices of the finite-relation polynomials. For each row first choose G_i for the relations with the nonhelper ties, taking the sharp center to be the identity tuple. The fully tied untwisted module has normalized length at most its balanced bound β_i^{bal}. Allowing ramification at the helper without its tie costs at most the sum of the two helper determinant valuations. The approximating inertia element gives the upper bound $d_i c_\ell$ for each determinant in rank d_i by the two estimates in (16.8). The helper's Frobenius and inverse-conjugate Frobenius can be chosen in the same determinant-valuation neighborhoods. Thus

$$v(G_i(\mathbf{1}))\le \beta_i^{\mathrm{bal}}+2d_i c_\ell \le 10 f d_i c_\ell.$$

This calculation also requires the original untwisted cocycle relations and the normalization at the same distinguished inertia element used in the tied induction. Test G_i on the row's helper line, with the slopes of that row, using $U-1$ as the polynomial variable. Its value at $U=1$ has the bound just obtained.

For a signed family, at the modulus consisting only of the helper set its single coordinate to U. Its value at $U=1$ is its original nonzero untwisted value times the omitted Euler factors at the helper. The helper can satisfy the finite approximation for all the signed families at once. At the prescribed inertia element each Euler parameter has a nontrivial root of unity with order dividing q_w-1; the Tate factor there is one. The valuation of its difference from 1 is at most $v(q_w-1)$, and continuity permits the same bound at the helper. There are at most four factors in the displayed norm identities. In particular the weaker allowance $8c_\ell$ suffices for the total helper cost. At the extra pairs the signed family uses the helper line with $y_P=1$ for $P\neq P_0$. Adding their modulus when the character on it is unramified multiplies the value by the Euler factors in the signed interpolation calculation. At these pairs the avatar takes value U on the prescribed Frobenius. With the sign s_0 of the signed test and the choice $\epsilon\in\{1,-1\}$ of local reading, each factor is $1-aU^{-s_0\epsilon}$ for an integral unit a. For unit U this is either $U^{-1}(U-a)$ or $-a(U-a^{-1})$.

For the row polynomials the extra factors to control are $\det(U\rho(\phi_P)-I)$ at both readings of each nonhelper pair. Each is a unit multiple of a monic polynomial since $\rho(\phi_P)$ is integral and invertible. The signed factors are the preceding linear factors. If there are a signed families, the total degree of the guard factors is at most $D_*=2J_0\sum_i d_i+4J_0 a$. The one-variable torsion lemma proved above applies to each of the row line polynomials and the signed helper polynomials with the target given by its displayed valuation bound at $U=1$. The Gauss norm on the disc in $U-1$ is at least the absolute value at its center. Include also the product of the monic factors after translating $U=1+T$. Its Gauss norm at radius s is at least s^{D_*}. Take s sufficiently close to 1 that $(D_*+2)(-\log_p s)<c_\ell/2$ and $|p|_p<s^{p-1}$. The common lemma then chooses one $u\in\mu_{p^k}$. The sum of the valuations of the monic factors is at most $(D_*+1)(-\log_p s)$, and the additional loss for each line value is at most $-\log_p s$. Its threshold k depends only on the fixed lower targets and the degree bound for the guards. Thus k can be fixed before N_{ray} and before the primes. The row line at this u gives the untied module of normalized length at most $10 f d_i c_\ell + c_\ell/2$ by the removal-of-ties calculation. The signed line at this u gives the full-modulus value with valuation at most the untwisted valuation plus $8c_\ell+c_\ell/2$ by the imprimitive norm identities. These are separate line points for the separate inputs to the common specialization, even though their helper coordinate u is common.

Now construct the second row polynomials F_i for the untied relations, sharp at the specific points just obtained over the valuation ring of a finite coefficient extension containing O, u and the signed coefficients. Use the full-modulus signed polynomials. The inequalities $|y_P-1|_p^{p^k}\le |p|_p$ hold for every center. Take $N_{\mathrm{ray}}\ge k+2K_0$ also beyond the degree-free threshold in the multivariate lemma above, with $a_0=p^k$, the

preceding valuation bounds as targets, $q = J_0$, and losses at most $c_\ell/2$ in that lemma. This N_{ray} can be chosen before building the pencil, because the targets and the bound on the number of coordinates are then already fixed. The second lemma gives one torsion tuple whose first m powers have every required bound simultaneously. Taking its monic factors $X^b - 1$, $1 \le b \le m$, with loss per coordinate less than c_ℓ/J_0 gives exact support at every power, and bounds the sum of the inertia-generator valuations at each power by c_ℓ. The ray construction identifies the tuple with one inverse avatar, trivial at the prescribed local groups above ℓ. Its powers give one character and the required common conductor on both coefficient readings.

The total row bound is $(10 f d_i + 1)c_\ell$. For each signed family the total increase in valuation is at most $9c_\ell$. The conductor norm is a p-unit. By the exact-conductor value relation this gives $\log |j_\iota(\rho_{p,s}(\theta^j)^{-1})|_\ell \le 9c_\ell \log \ell = v_\ell$ for each critical pair and sign, including the square of the ratio at rank two.

16.5 Split Hecke tests at the selected reading

Fix one of the auxiliary pairs with the corresponding power of the common fitted twist. Let $\xi_w = (A_w q_w^{t'-1})^{-1}$ be the test parameter at an eligible split place. The spherical kernel identity gives $H_w(\xi_w) = 0$ on the genuine lift image. We prove a local determinant bound for this image Z and label idempotent e_c. The left vectors are integral at all readings above ℓ and the right vectors are integral at j_ι. Each right vector is orthogonal to $Z^\perp \cap \operatorname{im} e_c \cap \ker H_w(\xi_w)$ at the applicable split tests.

The local Hecke polynomial. At a split place w write $q = \mathrm{N}\, w$ and let π be the unitary local action of $\mathrm{GL}_n(E_w)$. Let I_w be the upper pro-p_w Iwahori, where $p_w = \operatorname{char} k_w$, and P_I its orthogonal averaging operator. For a uniformizer ϖ, put $t_w = \operatorname{diag}(\varpi, 1, \dots, 1)$. On I_w-fixed vectors the unnormalized Hecke operator U of $I_w t_w I_w$ is $q^{n-1} P_I \pi(t_w)$. The q^{n-1} translations are indexed by the freely varying first-row entries of $I_w/(I_w \cap t_w I_w t_w^{-1})$.

Let x be a nonzero spherical eigenvector at w, and s_j the eigenvalue of the sum of the translations indexed by sublattices of $\mathcal{O}_w^n$ containing $\varpi\mathcal{O}_w^n$ and having residue codimension j. For a lattice L, its translate x_L is well defined because x is spherical. Write T_j for the corresponding partial sum over the lattices whose reduction contains the first line. Thus $T_0 = x$, $T_n = 0$, and $s_n \pi(t_w^{-1})x = T_{n-1}$, since the full top spherical translation is by ϖI_n. It follows that $s_n \ne 0$.

These sums are I_w-fixed. For $0 \le j < n$,

$$U T_j = q^j (s_{j+1} x - T_{j+1}).$$

Indeed t_w takes a lattice containing the first coordinate line modulo ϖ to a lattice whose reduction lies in the tail k_w^{n-1} and has dimension $n - 1 - j$. The upper row unipotents in I_w add all linear maps from that tail subspace to the first line. Each one occurs with relative measure $q^{-(n-1-j)}$ under the average. The tail action permutes the summed subspaces. All subspaces of codimension $j + 1$ avoiding the first line occur as these graphs.

The adjoint on the I_w-fixed subspace is $U^* = q^{n-1} P_I \pi(t_w^{-1})$. For $1 \le j < n$, applying t_w^{-1} to all the lattices of residue codimension j gives T_{j-1} together with the translates at

$$\mathcal{O}_w(\varpi^{-1} e_1 + \widetilde{c}) + L_W, \qquad W \subseteq k_w^{n-1}, \quad \dim W = n - 1 - j, \quad c \in k_w^{n-1}/W,$$

where L_W is the inverse image of W in the tail lattice. For subspaces containing the line, this is $c = 0$. For a graph of a nonzero functional, W is its kernel and c is the coset of the vectors of functional value 1. The graphs of the zero functional give T_{j-1}. The lower column unipotents with entries in $\varpi\mathcal{O}_w$ translate c uniformly through its q^j possibilities, and the tail action permutes W. Averaging the transforms of T_j gives q^{-j} times the sum over the displayed family. Thus

$$s_j U^* x = q^{n-1} T_{j-1} + q^j U^* T_j, \qquad U^* x = q^{n-1} T_{n-1}/s_n, \qquad U U^* x = q^{2n-2} x.$$

Let $P_j(Z) = \sum_{i=0}^{j} (-1)^i q^{i(i-1)/2} s_i Z^{j-i}$, where $s_0 = 1$, and let $P_x = P_n$ be the monic spherical Hecke polynomial. The forward recurrence gives $P_j(U)x = (-1)^j q^{j(j-1)/2} T_j$. If $P_x(c) = 0$, put $z = (P_x(Z)/(Z - c))(U)x$. Then $Uz = cz$. Expand $P_x(Z)/(Z - c) = \sum_{j<n} c^{n-1-j} P_j(Z)$. The adjoint recurrence gives

$$cU^*z = q^{-n} P_x(qc) U^*x + q^{n-2} z.$$

For this calculation the term with $s_n U^*x$ cancels the missing last T_{n-1} term by the displayed endpoint identity. Taking the inner product of U^*x with this formula, with the second argument linear, yields

$$c(q^n - 1)\langle x, z\rangle = P_x(qc)\langle x, x\rangle.$$

Every root c is nonzero since $s_n \neq 0$. Taking the inner product with z and using $Uz = cz$ gives

$$(|c|^2 - q^{n-2})\|z\|^2 = \frac{q^{-n}|P_x(qc)|^2}{q^n - 1}\,\|x\|^2.$$

If $|c|^2 < q^{n-2}$, then $P_x(qc) = 0$. Whenever $P_x(qc) = 0$, the adjoint identity gives $U^*z = q^{n-2}z/c$ and $\langle x, z\rangle = 0$. It follows that z is orthogonal to every $U^i x$, and hence zero. The cyclic space generated by x under U is the span of the T_j, so both U and U^* preserve it.

The roots of P_x are also closed under $c \mapsto q^{n-1}/\overline{c}$. Indeed, the inverse of the i-th minuscule double coset is the $(n - i)$-th minuscule double coset multiplied by $(\varpi I_n)^{-1}$. The adjunction on spherical vectors gives $s_{n-i} = s_n \overline{s_i}$. Substituting the coefficients of P_x then gives

$$Z^n \overline{P_x(q^{n-1}/\overline{Z})} = (-1)^n q^{n(n-1)/2} s_n^{-1} P_x(Z).$$

Cofinality on the fixed denominator packet. Fix the auxiliary pair of ranks r' and n, its base level and label, the coefficient isomorphism ι, and the chosen auxiliary twist ϑ with inverse avatar ψ. Put $t' = n - r'$. Denote the attached normalized representations of the finite set of base packets by N_x, and let x_0 be the genuine packet under test. Its main block is $U_0 = \psi T$ for the same row T as in the Selmer estimate. Its characteristic polynomial at g is the product of that of $U_0(g)$ and $\prod_{0 \le i < t'} (X - \chi(g)^{-i})$. This follows from the trace formula by applying it at the powers of g and using Newton's identities. The denominator test uses the fixed base packet space and these normalized attachments, with the local polynomial annihilation and the generic Frobenius coefficient approximation of those attachments. A packet is weak when each $N_x(g)$ has eigenvalue 1. The normalized representations have integral characteristic polynomials. For this assertion, boundedness of the matrix entries on the compact Galois group bounds all matrix powers for each g. Applying the powers to an eigenvector shows its eigenvalue has ℓ-adic norm at most 1. This also applies to the main block U_0.

Write $q = \mathrm{N}\, w$ and ϕ for the geometric Frobenius of an off-level split datum. Let A_w be the scalar of the spherical test. Thus the multiplier in $N_x(\phi)$ relative to the attached representation is $j_\iota(A_w)^{-1}$, and $|A_w|_{\mathbb{C}}^2 = q^{r'}$. Its ℓ-adic image has norm 1, by continuity of the normalizing character on the compact Galois group. If x is weak, exact local annihilation implies $P_{x,w}(A_w) = 0$ for its spherical Hecke polynomial $P_{x,w}$. Here the weak eigenvalue 1 is an eigenvalue of the normalized representation, so the unnormalized eigenvalue is precisely $j_\iota(A_w)$.

For the spherical eigenvector on this packet, the local Hecke polynomial P_x is $P_{x,w}$, with coefficient $(-1)^i q^{i(i-1)/2} s_i$ in degree $n - i$. Starting with the root A_w, the local norm identity gives the roots $A_w q^i$ through $i = \lfloor (t' - 1)/2 \rfloor$. At each preceding index i one has $2i < t' - 2$, so

$$|A_w q^i|^2 = q^{n-t'+2i} < q^{n-2}.$$

The reciprocal-conjugate root symmetry sends $A_w q^i$ to $A_w q^{t'-1-i}$ because $|A_w|^2 = q^{n-t'}$. This gives the distinct exact roots $A_w q^i$ for $0 \le i < t'$. In particular the reciprocal test at $\xi_w = (A_w q^{t'-1})^{-1}$ vanishes on weak packets.

Use the distinguished row's scalar inertia element γ_{sc} at a chosen level prime off ℓ and write $T(\gamma_{\mathrm{sc}}) = yI$. The label calculation above gives residual order $> 4r' + 2$ for y. The auxiliary twist and its conjugate conductor avoid all the chosen level primes, and χ is unramified at this prime. Thus $\psi(\gamma_{\mathrm{sc}}) = \chi(\gamma_{\mathrm{sc}}) = 1$. The powers $y^{i\ell^{\nu}}$, $0 \le i \le r'$, have distinct residues for every ν, since taking an ℓ^{ν}-power is injective in the residue field. For any $g \in G_E$, the residual polynomial $\det(1 - ZU_0(g))$ has degree at most r' and constant term 1. Choose i for which its value at the residue of $y^{i\ell^{\nu}}$ is nonzero. Then $\delta = \gamma_{\mathrm{sc}}^{i}$ lies in $\ker \chi$, and $\det(1 - U_0(\delta^{\ell^{\nu}} g))$ is an ℓ-adic unit. This unit condition persists throughout an open neighborhood of $\delta^{\ell^{\nu}} g$.

On the same Frobenius datum put

$$B_w(X) = \prod_{0 \le i < t'} (X - q^i), \qquad C_{x,w}(X) = j_\iota(A_w)^{-n} j_\iota\big(P_{x,w}(A_w X)\big).$$

Here the isomorphism acts on coefficients. The chain just proved gives the exact factorization $C_{x,w} = B_w T_{x,w}$ for each weak packet. The generic coefficient approximation of the attached representation, with accuracy < 1, gives $C_{x,w} \equiv \operatorname{charpoly} N_x(\phi)$ modulo the maximal ideal of the ℓ-adic integers. The normalization preserves this accuracy since it scales each coefficient by a unit. In particular $C_{x,w}$ is integral. If its prescribed residual comparator is $B_w S_w$, with $|S_w(1)|_\ell \ge 1$, division by the integral monic B_w makes $T_{x,w}$ integral and $T_{x,w}(1)$ a unit: first $B_w S_w$ is integral by the coefficient congruence, so S_w is integral, and then the reduced monic factor cancels to give $\overline{T}_{x,w} = \overline{S}_w$. For every weak packet residually congruent to x_0, use $S_w = \operatorname{charpoly} U_0(\phi)$, since $\chi(\phi)^{-1} = q$ and the unit test has been imposed. Each polynomial $T_{x,w}/T_{x,w}(1)$ is integral with value 1 at 1. Their finite product e_w has the same value and is divisible by each exact Hecke quotient in the rescaled variable, which differs from $T_{x,w}$ by the nonzero scalar $j_\iota(A_w)^n$.

For any neighborhood of $\delta^{\ell^{\nu}} g$, intersect with the open unit condition. For each of the finitely many weak packets, the generic Frobenius coefficient approximation gives a nonempty open subset of the current open set on which all degree-one Frobenius data off a finite set satisfy that approximation. Refine successively through the finite packet list. Exclude the base-level bad places and their conjugates, the coefficient prime, and the conductor of the fixed twist and its conjugate. Exclude also any additional finite set required by the triangle. There is a degree-one Frobenius datum in the resulting open set by Chebotarev. This is the cofinality after scalar perturbation required at the prime-selection step. In the triangular application $g \in G_\nu$ and the open cyclotomic and evaluation conditions are included in the prescribed neighborhood. The center $\delta^{\ell^{\nu}} g$ satisfies them because $\delta^{\ell^{\nu}}$ contributes zero to the additive evaluations on G_ν.

For each packet whose Galois characteristic polynomial differs from that of x_0 on residue coefficients, choose a coefficient function h and a group element witnessing the difference. Approximate the group element at generic spherical Frobenii, where the normalized monic Hecke coefficients give the limits. The integral Hecke algebra is closed since preservation of the Weyl integral class values is an intersection of closed norm inequalities. Thus h is in this algebra, and $h - h(x)$ vanishes at that packet and has unit value at x_0. Write A for the product over all residue-incompatible base packets of these integral functions. Its value at x_0 is a unit.

The triangular and first-order inputs. Let t' be the rank difference of the auxiliary pair in this denominator test. Choose a cyclotomic separator γ_{cyc} with $\chi(\gamma_{\mathrm{cyc}})^j - 1 \mid \ell^{h_0}$ for $0 < j < t'$. For $\ell \le h_2^2$ use Frobenius above one of the designated primes $g_i \ne \ell$: every $g_i^j - 1$ for $0 < j < t$ divides G, so one may take $h_0 \le v_\ell(G)$. For $\ell > h_2^2$ the base is unramified over $\mathbf{Q}_\ell$ because its cyclotomic conductor is smaller than h_2. Use cyclotomic inertia with residue order $\ell - 1 > t - 1$ and $h_0 = 0$. Set $a = \gamma_{\mathrm{cyc}}^{-1}$ and $h_* = H = h_2 h_0$. The assembled auxiliary primes have $q_w \equiv 1 \bmod \ell^{\nu}$ and avoid the conductor of the fitted twist.

Separating the test from the unwanted packets. For the fixed tuple label c, use the base space containing the image of its idempotent e_c. The finite group character formula shows that Ge_c preserves the integral smooth forms at j_ι: the group order divides G, the character values are algebraic integers, and the

tame factors are units. Work packetwise on this base level and then at the indicated Taylor–Wiles level. Take the Ribet normalizer

$$B_x = N_x(\gamma_0 a) - I - y_0(N_x(a) - I) + \ell^{M_\ell + h_*} \sum_g \lambda_g (N_x(g) - I), \qquad b_x = \det B_x,$$

where γ_0, y_0 are the scalar data on x_0 in its short Selmer condition and $a = \gamma_{\mathrm{cyc}}^{-1}$. For every genuine packet $b_x = 0$ by the trivial summand in (16.7). For every nonweak base packet there is g with $\det(N_x(g) - I) \neq 0$. One can choose the finitely supported nonnegative integers λ_g so $b_x \neq 0$ on all such packets. Add the terms successively: at the next packet the determinant as a polynomial in the new integer coefficient has nonzero leading term, and the determinants on the previous packets are nonzero at coefficient zero. Avoid their finitely many roots. Each b_x is integral by the stable lattice. The polynomial

$$p(T) = \prod_{x \text{ nonweak}} (T - b_x)$$

is integral and $p(0) \neq 0$, with the empty product equal to 1. The base normalizer and $p(0)$ here are fixed before choosing the depth of the triangular primes. The locus $p(b_x) \neq 0$ on the base is contained in the weak packets. Thus the fixed base function $T = Ap(b)$ is supported on weak packets residually congruent to x_0. The right-hand orthogonality controls the part of those packets outside the genuine lift image.

Write $l = e_c \Xi'$ and $w = e_c P_Z \Xi'$. The commuting normal Hecke actions and e_c preserve the orthogonal decomposition by packets, and Z is an invariant subspace. Thus the weak-packet components of $l - w$ lie in $\operatorname{im} e_c \cap Z^\perp \cap \ker H_v(\xi_v)$ for an eligible split v. The right-hand test gives zero pairing with each component. Its projection to the Z-part gives zero as well, so these components vanish. Consequently l, w agree after every diagonal operation supported where $p(b_x) \neq 0$. The form w is supported on the genuine packets. This verification allows a complementary subspace within the same weak Hecke packet as a genuine lift.

Uniform detection of the unramified Tate cocycles. Let $E \subseteq \mathbf{Q}(\zeta_N)$ and let χ be the ℓ-adic cyclotomic character on G_E. There is a finite set $S_{\mathrm{det}} \subset G_E$ that detects every continuous $O/\ell^\nu O$-valued cocycle for χ^{-e} unramified off ℓ, for every $\nu \geq 1$, $e \geq 0$, and ring of integers O of a finite extension of $\mathbf{Q}_\ell$. Unramified means zero on the inertia groups at the places outside the allowed set. Detection means that zero values on S_{det} force the cocycle to vanish. The same set satisfies $\#S_{\mathrm{det}} \leq y^{20}$ for $N \leq y$ and $\ell \leq y^2$.

Put $M' = N\ell^2$, $F = \mathbf{Q}(\zeta_{M'})$, and $d = [F : \mathbf{Q}] \leq M'$. Since $M' \geq 4$, the field F is totally imaginary. Take coset representatives for G_E/G_F and lifts in G_F of a basis of the maximal elementary abelian ℓ-quotient of G_F unramified away from ℓ. Global Artin reciprocity bounds the dimension of that quotient by

$$\dim_{\mathbf{F}_\ell} \mathrm{Cl}(F)/\ell + \sum_{v \mid \ell} \dim_{\mathbf{F}_\ell} \mathcal{O}_{F_v}^\times / (\mathcal{O}_{F_v}^\times)^\ell \;\leq\; \log_\ell h_F + 2d.$$

Indeed, in the inverse limit of the ray class groups for moduli supported above ℓ, the kernel of passage to the ordinary class group is the image of the product of local unit groups. At $v \mid \ell$ the deep principal units are isomorphic by the local logarithm to a full additive lattice of rank $[F_v : \mathbf{Q}_\ell]$. The pro-ℓ unit group has this rank and cyclic torsion (the local ℓ-power roots of unity). Summing its generator bounds uses $\sum_{v \mid \ell} [F_v : \mathbf{Q}_\ell] = d$ and $\#\{v \mid \ell\} \leq d$.

The field discriminant is at most $(M')^d$ in absolute value. For this upper bound one can use the power basis for $\zeta_{M'}$: the identity obtained by differentiating $X^{M'} - 1 = \Phi_{M'}(X)((X^{M'} - 1)/\Phi_{M'}(X))$ at $\zeta_{M'}$ shows that the norm of $\Phi'_{M'}(\zeta_{M'})$ divides $(M')^d$ up to sign. Minkowski gives an integral ideal of norm at most $(M')^d$ in every ideal class, since $(4/\pi)^{r_2} d!/d^d \leq 2^d$ and $M' \geq 4$. The number of sublattices of index n in $\mathbf{Z}^d$ is at most n^{2d-2}. In a Hermite form there are at most n^{d-1} choices of the diagonal, and for each diagonal the reduction bounds on the other entries give at most n^{d-1} choices. Counting ideals among the sublattices up to index $(M')^d$ gives $h_F \leq (M')^{2d^2}$. Consequently $\#S_{\mathrm{det}} \leq 3M' + 2(M')^3 \leq 3(M')^3$. If $N \leq y$ and $\ell \leq y^2$, then $M' \leq y^5$ and $3y^{15} \leq y^{20}$.

For detection, fix the module $M = O/\ell^v O$ and induct on k for cocycles with values in $M[\ell^k]$, starting with the zero module at $k = 0$. If z vanishes on S_{det}, so does ℓz, which takes values in $M[\ell^{k-1}]$ and is zero by induction. The resulting ℓ-torsion-valued restriction to G_F is a homomorphism, since F contains μ_ℓ. It is unramified off ℓ and vanishes on the chosen generators, hence is zero. For $b \in G_F$ and a chosen coset representative s, the identity $z(sb) = z(s) + \chi(s)^{-e} z(b)$ gives $z = 0$. This proves detection for a fixed set, including cocycles with values in the ramified coefficient rings. [21]

Triangular primes from the primitive reciprocity laws. Fix the field E, selected coefficient prime ℓ, denominator test and fitted twist of this application. Here E is the CM field containing μ_{5040}. The quadratic extension E/E^+ is unramified at every finite place and its number of complex places is at least $\varphi(5040)/2 \geq 4$. To see the unit conjugate difference used by that law, take a primitive $\zeta_{28} \in E$, put $i = \zeta_{28}^7$ and $\omega = \zeta_{28}^2 - \zeta_{28}^4 - \zeta_{28}^8$. The fourteenth cyclotomic equation gives $\omega + \overline{\omega} = 1$, so $i\omega - \overline{i\omega} = i$ is a unit in $\mathcal{O}_E$. Use N equal to the cyclotomic conductor used for E, $y = h_2$, and $t_{\mathrm{a}} > 0$ equal to the rank difference of the auxiliary pair in the current denominator test. Use the cyclotomic separator γ_{cyc} and loss h_0 described above. Thus $N + t_{\mathrm{a}} < y$, $\ell \leq y^2$ or $h_0 = 0$, and

$$\chi(\gamma_{\mathrm{cyc}})^j - 1 \mid \ell^{h_0} \quad (0 < j < t_{\mathrm{a}}).$$

Here $y \geq t_{\mathrm{a}} + 2$. For $0 \leq e, k < t_{\mathrm{a}}$ with $k \neq e$, every difference $\chi(\gamma_{\mathrm{cyc}})^e - \chi(\gamma_{\mathrm{cyc}})^k$ is a unit times a signed difference $\chi(\gamma_{\mathrm{cyc}})^j - 1$ with $0 < j < t_{\mathrm{a}}$. Thus the character-separation product

$$\prod_{\substack{0 \leq k < t_{\mathrm{a}} \\ k \neq e}} \left(\chi(\gamma_{\mathrm{cyc}})^e - \chi(\gamma_{\mathrm{cyc}})^k\right) \mid \ell^H \qquad (0 \leq e < t_{\mathrm{a}})$$

has the required divisibility in $\mathbf{Z}_\ell$, including the empty product if $t_{\mathrm{a}} = 1$. For $t_{\mathrm{a}} > 1$ one also has $v_\ell(4) \leq h_0$. At $\ell = 2$ this uses $\mu_4 \subset E$ and the nonzero difference at exponent 1; at odd ℓ it holds because $v_\ell(4) = 0$. It follows for $0 < e < t_{\mathrm{a}}$ that $4(\chi(a)^{-e} - 1)^2 \mid \ell^H$ in $\mathbf{Z}_\ell$, hence also in O.

For each $v \geq 1$, put $A = \mathbf{Z}/\ell^v$ and $G_v = \ker(\chi \bmod \ell^v)$. An even cocycle of type e has values in $O/\ell^v O$ with action χ^{-e}, is invariant under $g \mapsto g^c$ for the fixed complex conjugation c, and vanishes on the intersection of each decomposition group above ℓ with $\ker \chi$. Its other possible ramification off ℓ is at the indicated set of split primes and their conjugates.

The primitive law uses the earlier places of the same type. For a list of s earlier type-e Frobenius data, off ℓ, at degree-one primes over distinct primes of E^+, with Frobenii in G_v, prescribe at most s elements of G_v. There is a locally constant $z : G_E \to A$ that vanishes on these elements and takes a unit value on G_v. Its restriction to G_v is additive and satisfies

$$z(g b g^{-1}) = \chi(g)^{e+1} z(b), \qquad z(b^c) = -z(b) \qquad (g \in G_E,\ b \in G_v).$$

It comes with a finite exceptional set of primes of E^+. For a further family of v-split Frobenius data (w_i, ϕ_i) over distinct primes outside this set, if $z(\phi_i) = 0$ except possibly at j, every even type-e cocycle b ramified only at ℓ, w_i, $\bar{w}_i$ and vanishing at the earlier listed type-e Frobenii satisfies

$$n\, b(\tau_j) = 0 \qquad \text{whenever } 2z(\phi_j) \mid n \text{ in } A.$$

Here τ_j is the tame generator at the chosen place above w_j in its Frobenius datum. We use this assertion with $n = \ell^H$, including its factor 2 in residue characteristic 2.

[21] For Artin reciprocity and ray class fields use Artin and Tate, *Class Field Theory*, Chapter VII §3 and Chapter VIII §§1–2. For the convex body theorem use Minkowski, *Geometrie der Zahlen*.

Primitive reciprocity on the real subfield. Put $K = E^+$ and retain the fixed complex conjugation c. At every finite place of E the unit conjugate difference above excludes the nontrivial element of $\mathrm{Gal}(E/K)$ from inertia. Thus E/K is unramified at all finite places, and every prime of E of prime norm lies over a split prime of K. Write χ_K for the cyclotomic character on G_K. Extend its restriction χ from G_E to the character $\kappa(gc^a) = \chi(g)$ of G_K, for $a = 0, 1$. The cyclotomic character is invariant under conjugation by c, so this is a character, unramified away from ℓ at finite places. Every real complex conjugation has κ-value 1. At odd ℓ, $\mathbf{Z}_\ell^\times$ is procyclic. At $\ell = 2$, the inclusion $\mu_4 \subseteq E$ puts the image of κ in $1 + 4\mathbf{Z}_2$, also procyclic.

Fix $\nu \geq 1$, let $A = \mathbf{Z}/\ell^\nu$, and put

$$M = A(\kappa^{-e}), \qquad B = \mathrm{Hom}(M, \mu_{\ell^\nu}), \qquad U = G_{E(\mu_{\ell^\nu})} = G_\nu.$$

Both coefficient modules are trivial on U, and G_K/U is the product of the group of order two generated by c and the cyclic image of χ mod ℓ^ν on G_E. The character of B is $\chi_K\kappa^e$ mod ℓ^ν. Write P_0 for the s distinct primes of K below the earlier type-e data. Let $k \leq s$ be the number of prescribed elements $u_1, \ldots, u_k \in U$. Let S consist of P_0, the primes over ℓ and the real places.

We use local Tate pairings for M, B and Poitou–Tate orthogonality of the two global localization images. For any finite set V of test primes split in $E(\mu_{\ell^\nu})/K$ and disjoint from S, take the following local conditions for primal cocycles. At $v \mid \ell$, require zero on $D_v \cap \ker\kappa$, with D_v the chosen decomposition group in G_K. At P_0 and at the real places require zero on D_v. Allow all cocycles at V and unramified cocycles at other finite places. The conditions include all local coboundaries. At P_0 the chosen Frobenius is in U, so the entire local coefficient action is trivial. At V a local primal cocycle is determined by its unramified and tame coordinates. Evaluation on an ℓ^ν-tame generator τ_v gives the tame coordinate. The local unramified dual group pairs perfectly with this coordinate. Indeed local class field theory and the Kummer class of a uniformizer give the invariant as the product of a tame value and a dual Frobenius value, divided by ℓ^ν, up to a unit depending on the chosen generators. At all finite places off ℓ where the modules are unramified, the unramified subgroups are mutual annihilators. At a real place the degree one local group in Tate duality is the usual H^1, including at $\ell = 2$.[22]

Write $\bar{M} = M[\ell]$ and $\bar{B} = \mathrm{Hom}(\bar{M}, \mu_\ell)$. Choose V using the two finite groups $H^1(G_{K,S}, \bar{B})$ and $H^1(G_{K,S}, B)$, where $G_{K,S}$ is the Galois group unramified outside S. Finiteness follows from the finite-coefficient global duality theorem. For each element of the first group with nonzero restriction to U, choose a witness in U with nonzero value. Add a distinct prime to V with an arithmetic Frobenius in U having this value. For each u_i, add a distinct prime v_i whose arithmetic Frobenius in U has the same values as u_i under all classes of the second group. Restrictions to U are homomorphisms independent of the choice of cocycle representative, so these conditions specify open neighborhoods in G_K. Chebotarev supplies prime-norm primes with chosen places above them and allows any finite exclusions. Each of these primes splits in the fixed field of U. Extra primes split in that field may also be included in V.

Let Y_ν be the group of primal global cocycles with the listed conditions. Its ℓ-torsion is exactly the group of cocycles with those conditions valued in $\bar{M}$. The dual Selmer conditions for $\bar{B}$ are zero at V, since the primal conditions are full there. In particular these dual classes are in $H^1(G_{K,S}, \bar{B})$. By the choice of V, they restrict to zero on U. Inflation–restriction puts them in $H^1(G_K/U, \bar{B})$. This last group has dimension at most $1 + h^0(K, \bar{B})$. A cocycle on the two-generated group G_K/U is determined by its two values on generators, and its coboundaries have dimension $1 - h^0(K, \bar{B})$.

For the $\bar{M}$ local subspaces L_v just specified, the Selmer dimension formula is

$$h^1_L(\bar{M}) - h^1_{L^\perp}(\bar{B}) = h^0(K, \bar{M}) - h^0(K, \bar{B}) + \sum_{v \in S \cup V} \big(\dim L_v - h^0(K_v, \bar{M})\big).$$

[22] Tate, Stockholm ICM (1962), Theorems 2.1, 2.2, 2.4 and 3.1 for the local and global formulas. See Česnavičius, Theorem 5.1 and §6 for the exact sequence with the real-place conventions.

In this formula $h^0(K_v, \bar{M})$ counts ordinary invariants, also at the real places. To see the real-place term directly, let I_M, I_B be the two localization images on $S \cup V$. Poitou–Tate gives $I_M^{\perp} = I_B$. Taking the dimensions of their intersections with the local conditions gives the Selmer difference as

$$h^1(G_{K,S\cup V}, \bar{M}) - \dim \ker(\mathrm{loc}_{\bar{B}}) + \sum_{v \in S\cup V} (\dim L_v - h^1(K_v, \bar{M})).$$

By the degree-two terms of duality, $\dim \ker(\mathrm{loc}_{\bar{B}}) = h^2(G_{K,S\cup V}, \bar{M}) - \sum_v h^2(K_v, \bar{M}) + h^0(K, \bar{B})$. For the rank-one module $\bar{M}$, each finite-place $h^2 - h^1$ is $-h^0$, with the further subtraction $[K_v : \mathbf{Q}_\ell]$ at $v \mid \ell$. Each real-place $h^2 - h^1$ is zero by periodicity. The global Euler characteristic in dimensions is zero since $\bar{M}$ is trivial at all real places of the totally real field K. The sum of the extra finite-place terms is $-[K : \mathbf{Q}]$. It equals the sum of $-h^0(K_v, \bar{M})$ over the real places, giving the asserted formula.[23]

At $v \mid \ell$ the permitted cocycles factor through the procyclic group $\kappa(D_v)$. Evaluation on a generator is injective, giving $\dim L_v - h^0(K_v, \bar{M}) \le 0$ after subtracting the local coboundaries. At P_0 and at each real place the summand is -1. At V it is 1, since the two coordinates give all of H^1. Adding back the global coboundary dimension $1 - h^0(K, \bar{M})$ yields

$$\dim_{\mathbf{F}_\ell} Y_\nu[\ell] \le |V| + 2 - s - [K : \mathbf{Q}] \le |V| - k - 1.$$

This calculation uses the group killed by ℓ at the original depth. It therefore controls the number of generators of Y_ν over A.

Let $D \subseteq A^V$ be generated by the tame evaluation vectors of Y_ν and the k coordinate vectors of the primes v_i. The bound shows that D has at most $|V| - 1$ generators. Smith form over A yields a vector $r \in A^V$ in its annihilator for the coordinate dot product, with a unit coordinate. Indeed, in elementary-divisor coordinates one index has zero cyclic summand, and the dual coordinate vector is primitive. In particular $r_{v_i} = 0$. Use the local unramified dual coordinates at V to represent r by a tuple β_{loc} whose pairing with a primal local tuple is $\ell^{-\nu} \sum_{v\in V} r_v x_v(\tau_v) \bmod \mathbf{Z}$. Set its coordinates on S equal to zero. It annihilates $I_M \cap \prod_v L_v$. For these finite cohomology groups the annihilator of this intersection is

$$I_B + \left(\prod_v L_v\right)^{\perp}.$$

Consequently a global class β has the same localizations as β_{loc} at V, and meets the dual conditions at S. Since it is unramified at V, it comes from $H^1(G_{K,S}, B)$. Restrict a representative to U, identify the value module with A by a generator of μ_{ℓ^ν}, and call the resulting homomorphism z. Then $z(u_i) = 0$ by the choice of v_i. A unit coordinate at V makes z a unit at its arithmetic Frobenius. For $g \in G_K$ and $b \in U$, the cocycle identity gives $z(gbg^{-1}) = \chi_K(g)\kappa(g)^e z(b)$. Its two specializations give the required twist and oddness. Extend z by zero on the open-and-closed complement of U in G_E. The resulting function is locally constant.

Take the finite exceptional set to contain the finite places of S and V. For any further family in the reciprocity statement, first take an A-valued even cocycle b on G_E. It extends to a κ^{-e}-cocycle b^+ on G_K by $b^+(gc^a) = b(g)$, using its evenness and $\kappa(c) = 1$. At $v \mid \ell$, for $g \in D_v \cap \ker \kappa$ one has $g^2 \in G_E \cap D_v \cap \ker \chi$. Hence $2b^+(g) = b(g^2) = 0$. At a real place the same calculation uses the square of a complex conjugation. At P_0, the unramified restriction is zero on the prescribed geometric Frobenius, and hence on the decomposition group. At every remaining finite place where b is unramified, so is b^+, since the inertia group lies in G_E.

Thus the local invariant of $(2b^+) \cup \beta$ is zero at S, by the dual conditions, and at every other place outside the new ramification set, by unramifiedness. At a new place the given geometric Frobenius and tame generator are those of the same decomposition group over K, because the place splits in E. The dual Frobenius value is $z(\phi_i)$, up to switching arithmetic and geometric Frobenius. The sum of local invariants

[23] Tate, Bourbaki exposé 306 (1965/66), Theorem 2.2 for the global Euler characteristic, and the local Euler characteristic in Tate's Stockholm ICM lecture.

is zero. Since all the other $z(\phi_i)$ vanish, the pairing at j gives $2z(\phi_j)b(\tau_j) = 0$ in A, with a unit change allowed for the given tame generator. This implies the assertion for any multiple n of $2z(\phi_j)$. For $O/\ell^\nu O$-valued cocycles, take coordinates in a $\mathbf{Z}_\ell$-basis of O. All the local conditions and equality at the prescribed points hold coordinatewise, and the same z works on every coordinate. Thus no division by 2 enters this reciprocity law.

Order $J = \{1, \dots, t_{\mathrm{a}} - 1\} \times S_{\mathrm{det}}$ lexicographically, using any fixed order on S_{det}. The first coordinate is the type. For the test point $s \in S_{\mathrm{det}}$ write $g_s = (as)^{-1}sa \in \ker \chi$. At depth ν construct the primes, laws z_j and exceptional sets in this order. Before the step $j = (e, s)$, the simultaneous image of G_ν under the earlier type-e laws is an A-submodule of A^q, where q is the number of those laws. Any submodule of A^q is generated by at most q elements, by taking its preimage in $\mathbf{Z}_\ell^q$. Prescribe lifts of these generators as the zeros of the new primitive law. Take an element where the new law is a unit and multiply by powers of the lifts to cancel its image under the earlier type-e laws. The resulting $b_0 \in G_\nu$ still gives a unit value for z_j. For an integer m_0 representing $\chi(\gamma_{\mathrm{cyc}})$ in A, apply, successively for $0 \le k < t_{\mathrm{a}}$, $k \ne e$, the operations

$$b \longmapsto (\gamma b \gamma^{-1}) b^{-m_0^{k+1}}.$$

Each operation stays in G_ν and multiplies the value of a law of type d_0 by $m_0^{d_0+1} - m_0^{k+1}$. The resulting element u kills all earlier laws. Also

$$z_j(u) = z_j(b_0) m_0^{t_{\mathrm{a}}-1} \prod_{k \ne e} (m_0^e - m_0^k).$$

Lift $z_j(b_0)$ to a unit in $\mathbf{Z}_\ell$ and lift m_0 mod ℓ^ν by $\chi(\gamma_{\mathrm{cyc}})$. The resulting lift of $4z_j(u)$ has valuation at most $(t_{\mathrm{a}} - 1)h_0 + v_\ell(4) \le H$. Reducing the divisibility in $\mathbf{Z}_\ell$ gives $4z_j(u) \mid \ell^H$ in the ring A, including when $H \ge \nu$ and the right-hand side is zero in A.

Set $g = g_s g_s^c u (u^c)^{-1}$. Apply the cofinality on the fixed denominator packet at the already chosen twist. On the off-level set where the distinguished unit test and weak-packet comparisons hold, choose $\delta \in \ker \chi$ such that the data are cofinal at $\delta^{\ell^\nu} g$. At this element each of the laws has value $2z_i(u)$, by oddness and additivity on G_ν. Each even type-e cocycle unramified away from ℓ has value $2b(g_s)$, by evenness and additivity on G_ν. The factor δ^{ℓ^ν} contributes zero to all these values. There are only finitely many such scalar A-valued cocycles by detection. The values of all cocycles with coefficients $O/\ell^\nu O$ are then controlled in the same neighborhood, by taking coordinates in a $\mathbf{Z}_\ell$-basis of O. Impose these finitely many scalar conditions, the values of all current and earlier laws, and membership in G_ν, which are open conditions. Choose a degree-one prime w_j off ℓ and its Frobenius datum there. Exclude the conductor primes of the already fixed twist, the primes above the real primes already chosen or in the current and earlier exceptional sets, and those ramified in E/E^+. The same datum gives τ_j and ϕ_j; in particular $\mathrm{N}\, w_j \equiv 1 \bmod \ell^\nu$, by the cyclotomic value $(\mathrm{N}\, w_j)^{-1}$ of geometric Frobenius. [24]

For an even cocycle of type e vanishing at the earlier τ_i and at the earlier ϕ_i of type e, the earlier conjugate pairs can be removed from its ramification set. Its restriction to old tame inertia is additive, its value at the conjugate generator is zero by evenness, and wild inertia at these primes has residue characteristic different from ℓ. Apply the law z_j to the current and later primes. All the later $z_j(\phi_i)$ vanish by construction, and $z_j(\phi_j) = 2z_j(u)$, so $\ell^H b(\tau_j) = 0$. Here the divisor in reciprocity is $2z_j(\phi_j) = 4z_j(u)$, whose valuation was charged above.

For the remaining coboundary assertion, suppose b is even of type $e > 0$, unramified off ℓ, and zero at every ϕ_j of type e. Then $2b(g_s) = 0$ for each test point s. Put $d_e(g) = \chi(g)^{-e} - 1$. The identity $(as)g_s = sa$ and the cocycle law give

$$d_e(a)(2b)(s) - d_e(s)(2b)(a) = -\chi(as)^{-e}(2b)(g_s) = 0.$$

[24] The existence of degree-one Frobenii in nonempty open sets off any finite set follows from Chebotarev's theorem.

The expression $d_e(a)(2b)(g) - d_e(g)(2b)(a)$ is an unramified-off-ℓ cocycle, so detection makes it zero on G_E. The argument works integrally at 2, and is paired with $4d_e(a)^2 \mid \ell^H$, already proved. Finally

$$(\#J + \#S_{\mathrm{det}})H = t_{\mathrm{a}}\#S_{\mathrm{det}}\, yh_0 \le y^{22}h_0.$$

For $h_0 = 0$ the cost is zero at every prime with a finite detecting set. The ordered types and their count are independent of ν. These are the triangular inputs to the Tate-block argument on the current packet at the actual depth. Each prime and its inverse-conjugate reading use the same variable and the already chosen fitted character.

For a Tate character $\xi = \chi^{-e}$, the paired differences at τ, ϕ read $c + c \circ \mathrm{conj}$, since $\xi(\tau) = \xi(\phi) = 1$ modulo ℓ^ν and the conjugate arguments are inverses. This even cocycle is subject to the triangular relation. On the module left after all the prescribed tame zeros and Frobenius equalities, the odd map $c \mapsto c\circ\mathrm{conj} - c$ has kernel killed by $C = \ell^{h_*}$ when $e > 0$. In its kernel c is even. The triangle applied to $c + c \circ \mathrm{conj}$ says $4(\xi(a) - 1)c$ is a coboundary; the Ribet normalizer's coboundary condition kills it on multiplying by $\xi(a) - 1$. The factor $4(\xi(a) - 1)^2$ divides C. This works at 2 without a parity projector. The strict local condition away from the ramified places follows from the local crystalline relations: at ℓ, crystalline cocycles in $O(-e)$ vanish on the cyclotomic kernel for $e > 0$, since their rational classes are coboundaries. For $e = 0$ the finite condition is unramified and the final everywhere-unramified group is bounded by the class number. Therefore the quotient used in the Tate-block cost is bounded by the local defect and the already established odd or class-number bound. The triangular prime construction uses the primitive reciprocity law on the same field and Frobenius conditions as above.

The trace ideal over the integral Hecke algebra adjoins one variable X_i per pair. Its generators include $\ell^\nu \operatorname{tr} N_x(g)$ and

$$\operatorname{tr}((N_x(\tau_i) - 1)N_x(g)) - X_i \operatorname{tr}(((X - 1)e_i)(N_x(\phi_i))N_x(g)).$$

The inverse-conjugate reading uses the same variable. Indeed write $\bar{h} = chc$. Polarization gives

$$\operatorname{tr} N_x((\bar{h}^{-1})^a g) = \chi(g)^{1-t'}\chi(h)^{a(t'-1)} \operatorname{tr} N_x\left(h^a(cgc)^{-1}\right).$$

The factor $\chi(h)^{a(t'-1)}$ is 1 modulo ℓ^ν for $h = \tau_i, \phi_i$. Its error is thus in the trace ideal.

The first-order Ribet polynomial. Fix the coefficient reading, auxiliary pair, tuple, fitted twist and triangular primes of the preceding test. Write x_0 for its distinguished genuine packet, $\mathcal{X} \subseteq Y$ for the genuine packets, and Y for the finite catalogue at the deeper level. Put $P = \ell^\nu$, $P_0 = \ell^{M_\ell}$, $C = \ell^{h_*}$, and $\Lambda = O[X_j : j \in J]$, where O is a common finite coefficient DVR. On $\mathcal{X}$ take the direct-sum lattice $U_x = \rho_x \oplus \bigoplus_{0\le e<t'} \xi_e$, with $\rho_x = T_x\psi$ and $\xi_e = \chi^{-e}$. On the other packets of Y, write U_x for a stable integral model of the normalized attachment. Enlarge O to contain the coefficients of every e_j and the scalar data. The normalized length of an O-module of finite length is its length divided by $v_O(\ell)$. Thus $v_\ell(d) = v_O(d)/v_O(\ell)$ for $d \in O$.

Let S be the image of the integral Hecke algebra with the variables adjoined, read in $\prod_{x\in Y} \overline{\mathbf{Q}}_\ell[X_j : j \in J]$. Its scalars contain the diagonal O, and the normalized characteristic-polynomial coefficients on every Galois element lie in S. Indeed the normalized spherical coefficients come from the integral Hecke operators. Their value tuples lie in a finite product of copies of O after a finite coefficient enlargement. The O-submodule given by the constant Hecke algebra in that product is finite and hence closed. Thus the coefficient tuples at arbitrary Galois elements, obtained as Frobenius limits, still lie in the integral Hecke algebra. Denote the trace ideal of the preceding paragraph by $\mathfrak{I} \subset S$. It contains the two copies of each trace tie, with one X_j for both copies. Write $\tau_{j,+} = \tau_j$, $\tau_{j,-} = (c\tau_j c^{-1})^{-1}$ and use the same convention for ϕ. The e_j satisfy $e_j(\xi_e(\phi_{j,\pm})) \equiv 1 \bmod P$ for every e, since $\xi_e(\phi_{j,\pm}) \equiv 1 \bmod P$ and $e_j(1) = 1$.

A module of tied rows. Write $A = \Lambda[G_E]$, with $[g]$ the basis element for g, and let U_x, ρ_x and η_e denote the induced representations of A, the last given by $\eta_e([g]) = \xi_e(g)$. Choose a finite list of group elements g_i as follows. The first part consists of the two tame readings δ_v at each twist place and old

level place away from ℓ. The second part spans the tuples $(U_x(g) - I)_{x \in Y}$ over O, which is possible since this span is a submodule of a finite free O-module. The last part is a set s formed from the detecting set S_{det} and all $\tau_{j,\pm}$. An element occurring in more than one part has a separate coordinate in each. Choose constant coordinate vectors $z(g)$ with $\sum_i z(g)_i(U_x(g_i) - I) = U_x(g) - I$ for every x. For a row w, put $V(w) = \sum_i w_i([g_i] - 1)$. Write γ_0, y_0 for the scalar data of the short Selmer condition on x_0, and set

$$\omega = z(\gamma_0 a) - y_0 z(a) + P_0 C \sum_g \lambda_g z(g), \qquad B = V(\omega).$$

Here a and λ_g are the separator and the correction already chosen in the normalizer. Thus $b_x = \det U_x(B)$ uses the base-level normalizer evaluated on the deeper catalogue.

Let $\mathfrak{T}$ consist of the elements $b \in A$ for which $(\operatorname{tr} U_x(b[g]))_x \in \mathfrak{I}$ for every g. Multiplication in A, linearity and cyclicity of trace show that $\mathfrak{T}$ is a two-sided ideal. Define the left row module

$$\mathcal{N} = V^{-1}\big(\mathfrak{T} + (\ker \eta_0)B\big).$$

It contains $PA^{\{i\}}$, the rows

$$z(gh) - z(g) - [g]z(h), \qquad z(g_i) - \mathbf{e}_i, \qquad ([g] - 1)\omega,$$
$$z(\tau_{j,\pm}) - X_j e_j([\phi_{j,\pm}])z(\phi_{j,\pm}),$$

and every constant row representing an O-linear relation among the tuples $U_x(g) - I$. These assertions follow by applying V and then the U_x. For the last displayed row the trace is exactly the corresponding generator of $\mathfrak{I}$.

For each old tame coordinate $\delta_v = g_i$ the module also contains $o_v([\delta_v])\mathbf{e}_i$, where

$$o_v(Z) = \begin{cases} \psi(\delta_v) - Z, & \text{at a fitted-twist place,} \\ 1 + Z + \cdots + Z^{G-1}, & \text{at an old level place.} \end{cases}$$

The quadratic relation at the twist places gives $o_v(U_x(\delta_v))(U_x(\delta_v) - I) = 0$. At an old level place off ℓ, the local attachment has finite inertial image. In characteristic zero its matrices on inertia are semisimple. The labelled inertia annihilator restricts their roots to orders dividing some $g_i^d - 1$ with $1 \le d \le r$. Since these integers divide G, we have $U_x(\delta_v)^G = I$. These assertions apply on all of Y, including the nonlift packets. All $\xi_e(\delta_v)$ equal 1.

The main-block determinant. Evaluate $\mathcal{N}$ at all $X_j = 0$ through ρ_{x_0}. Let $L \subseteq O^{r'|\{i\}|}$ be the span of the ordinary rows of the resulting block rows. A vector (w_i) modulo P_0 annihilating L gives $u(g) = \sum_i z(g)_i w_i$, a cocycle for ρ_{x_0} which obeys the tuple relations and vanishes on the two readings of each τ_j. The row $([\gamma_0]-1)\omega$ forces $(y_0-1)\sum_i \omega_i w_i = 0$. Since y_0-1 is a unit and $\rho_{x_0}(\gamma_0) = y_0 I$, the cocycle rule and $P_0 = 0$ on the coefficient module give $\sum_i \omega_i w_i = u(\gamma_0)$. Therefore $u(\gamma_0) = 0$. The generator relations give $u(g_i) = w_i$, so this map of annihilators into cocycles is injective. These cocycles are continuous because they factor by an O-linear map on the finitely generated span of the tuples, which has its lattice topology by Smith normal form. The tuple relations force $u(g) = 0$ whenever every $U_x(g) = I$, giving unramifiedness outside ℓ, the old level, the fitted-twist places and the triangular pairs. The two tame zeros remove the triangular pairs from this list. The Artin–Rees graph construction in the local reduction applies to these cocycles and to the scalar cocycles below.

Let L_0 denote the normalized crystalline defect on this block, at P_0, and L_e the corresponding defects on the Tate blocks at P. The preceding local reduction gives $L_0 + \sum_e L_e \le 4fr'\mathrm{loss}_{\mathrm{crys}}(\ell)$. The short Selmer condition and this reduction bound the length over O of the above annihilator by

$$b_{\mathrm{main}} = v_O(\ell)\big((10fr' + 1)c_\ell + L_0\big) < v_O(P_0).$$

The strict inequality follows from the definition of b_ℓ and $M_\ell > rb_\ell$. Smith normal form says that the length of this annihilator equals the sum of the exponents of L truncated at $v_O(P_0)$. Thus none is truncated and L has full rank and length of its cokernel at most b_{main}.

Full matrix span means $\rho_{x_0}(\Lambda[G_E])|_{X=0} = M_{r'}(O)$. Left multiplication by preimages of matrix units permits placing any ordinary row of L in any one position of an evaluated block row, with the other positions zero. Group a basis of L into r'-row blocks and lift them in this way to rows of a square matrix M_0 over A, all in $\mathcal{N}$. For $x \in \mathcal{X}$ write $D_x = \det \rho_x(M_0)$, the determinant of the flattened main blocks over Λ. Then $v_O(D_{x_0}(0)) \le b_{\mathrm{main}}$.

The scalar rows and their relations. Apply η_e to the rows, discarding the first, old-tame coordinates. Any O-linear form on the constant coordinate space with values in O/P annihilating an indicated projected relation module gives a scalar cocycle $u(g)$ by evaluation at $z(g)$. Here and below annihilation means that the form kills constant vectors whose images belong to the module. The displayed rows supply the cocycle and tuple relations. The old-tame coordinates have been discarded, so the generator relations give $u(\delta_v) = 0$. The cocycles are therefore continuous and unramified away from ℓ and the new pairs.

If one of these cocycles is a coboundary $u(g) = (\xi_e(g)-1)m$, the row $([g]-1)\omega$ gives $(\xi_e(g)-1)u(\omega) = 0$, where $u(\omega)$ is the form applied to ω. For $e > 0$, the triangular separation gives $\xi_e(a) - 1 \mid C$. Since $\xi_e(\gamma_0) = 1$,

$$u(\omega) = \left(1 - y_0 + P_0 \frac{C}{\xi_e(a) - 1} \sum_g \lambda_g(\xi_e(g) - 1)\right) u(a).$$

The scalar in parentheses is a unit. It follows that $(\xi_e(a) - 1)u(g) = 0$ for every g. For $e = 0$ every coboundary is zero. These are the hypotheses of the Tate-block length argument. On every projected module used below, that argument gives the triangular cost at type e with constant C and length bound $b_e = v_O(\ell)(L_e + v_\ell(X))$. In this expression X is the integer bounding the odd cocycle modules in the budget calculation, including the class-number bound for $e = 0$.

If all the last coordinates are also discarded, the cocycle vanishes on s, hence on both tame readings, so it is unramified outside ℓ. The detecting condition now makes the cocycle zero. Duality for O/P shows that the projection of $\eta_e(\mathcal{N})$ to the middle coordinates contains each standard basis row. Indeed a constant row outside this projection has nonzero image detected by an O/P-valued form, by Smith normal form applied to the preimage of the projection in the constant coordinate lattice, which contains P times that lattice. For a middle generator such a form would have a nonzero value on the corresponding g_i. Choose rows $m_{e,i} \in \mathcal{N}$ projecting to the middle unit rows under η_e. They can be made simultaneous. Starting with a zero row for each i, correct a row w at the next e by adding $\sum_j (\mathbf{e}_i(j) - w_j)m_{e,j}$, with j running over the middle coordinates. At the already treated characters these coefficients evaluate to zero. Call the resulting rows v_i.

For each e, use the rows $\eta_e(v_i)$ to clear the middle coordinates and keep the last coordinates s. Let $N_e \subseteq \Lambda^s$ be the image of $\mathcal{N}$ under this linear operation after η_e, and Φ_e its restriction to constant rows. The image of $z(g)$ for $g \in s$ is the corresponding standard basis vector modulo N_e. Also

$$\Phi_e(z(\tau_{j,\pm})) - X_j \Phi_e(z(\phi_{j,\pm})) \in N_e.$$

Here one can remove e_j because its value on $\xi_e(\phi_{j,\pm})$ is 1 mod P and $P\Lambda^s \subseteq N_e$. Every form on the constant lattice vanishing on $\Phi_e^{-1}(N_e)$ still gives one of the cocycles just examined.

Fitting determinants for the types. Let $\mathfrak{q} = (X_j^2 : j \in J)$, $I_e = \{j : \mathrm{ty}(j) = e\}$, $P_e = \prod_{j \in I_e} X_j$, and $P_{<e} = \prod_{u<e} P_u$. Set $N_e^+ = N_e + (\mathfrak{q} : P_{<e})\Lambda^s$. For $j \in I_e$, any earlier $i \notin I_e$ has $\mathrm{ty}(i) < e$ and $X_i\Lambda^s \subseteq N_e^+$. We claim there is $\Gamma_e \in \Lambda$ such that

$$\Gamma_e P_e \ \in \mathrm{Fitt}_0(\Lambda^s/N_e^+), \qquad v_O(\Gamma_e(0)) \le b_e + (|I_e| + |s|)v_O(C). \tag{*}$$

Here Fitt_0 is generated by the full square determinants of rows of the relation module.

For the proof, fix e, abbreviate Φ_e as Φ, and write $Z(H) = \{u(g) = \alpha(z(g)) : \alpha \text{ vanishes on } \Phi^{-1}(H)\}$ for a relation module $H \supseteq N_e^+$. We use two elementary consequences of Smith normal form and alternating multilinearity. First, for an integral vector v, if all these forms kill v, then $\Phi(v) \in H$, since the constant quotient is killed by P and its O/P-dual separates points. Second, if $aw \in H$ then a times the Fitting ideal

after adding w is contained in the Fitting ideal before adding it. Expand determinants in the enlarged rows. Terms using the new row twice vanish, and one scalar a deals with each remaining term.

The strict-local-zero subset of G_E in the triangular cost condition is denoted by $\mathcal{L}_e$. Enlarge $H = N_e^+$ by the vectors $\Phi(z(g))$, $g \in \mathcal{L}_e$, by adding one uniformizer-division step at a time. More precisely, when a required vector $\Phi(v)$ is outside H, adjoin $\varpi^{a-1}\Phi(v)$, where $a > 0$ is least with $\varpi^a\Phi(v) \in H$ and ϖ is a uniformizer. If $v = z(g)$ for $g \in \mathcal{L}_e$, this costs one in the Fitting valuation and strictly reduces the image of $Z(H)$ modulo the zero conditions on $\mathcal{L}_e$. The image drops because some old form has $\varpi^{a-1}u(g) \neq 0$ and every new form has that value zero. The finite relative length λ_1 forces this procedure to terminate with all the prescribed vectors adjoined, and bounds the total cost.

Next, in order for $j \in I_e$ adjoin $v_j = \Phi(z(\phi_{j,+}) - z(\phi_{j,-}))$. Before adjoining v_j, a form in $Z(H)$ has equal values on the two τ_i for every $i < j$. For earlier $i \in I_e$ the difference of their ties gives this because $v_i \in H$. For earlier $i \notin I_e$ it follows from $X_i\Lambda^s \subset H$. The two Frobenius values also coincide at earlier $i \in I_e$. The triangular condition now makes every form kill $C\,(z(\tau_{j,+}) - z(\tau_{j,-}))$. Smith duality and the difference of the ties show $CX_jv_j \in H$, so adjoining v_j costs the Fitting multiplier CX_j.

Call the resulting module H', and temporarily add the ideal of all the variables times Λ^s, obtaining H_0'. The forms in $Z(H_0')$ vanish on both tame readings of every pair and agree on the Frobenius readings of type e. Write λ_2 for the length of $Z(H_0')$ modulo its C-torsion. The triangular cost is $\lambda_1 + \lambda_2 \leq b_e$. Adjoin $C\Phi(z(g))$ for $g \in s$ to H_0', using the same uniformizer-division argument. This time the nonzero test value is $\varpi^{a-1}Cu(g)$, which vanishes on the C-torsion, so the cost is at most λ_2. The enlarged module contains C times every standard basis row, yielding a determinant $C^{|s|}$. After multiplying by the uniformizer to cover this last cost it lies in $\mathrm{Fitt}_0(\Lambda^s/H_0')$. Modulo the ideal of all the variables, this Fitting ideal is contained in the image of $\mathrm{Fitt}_0(\Lambda^s/H')$, by reducing the entries of its generating determinants. Lift an element with that constant coefficient and multiply by $\prod_{j\in I_e} CX_j$ and the first uniformizer cost. This proves $(*)$, also for $I_e = \emptyset$.

Expansion modulo $(\mathfrak{q} : P_{<e})$ gives

$$\Gamma_e P_e = d_e + \zeta_e, \qquad d_e \in \mathrm{Fitt}_0(\Lambda^s/N_e), \quad \zeta_e \in (\mathfrak{q} : P_{<e}).$$

In the difference between $\prod_e(\Gamma_e P_e)$ and $\prod_e d_e$, telescope with the original factors $\Gamma_u P_u$ to the left of the e-th difference and the factors d_u to its right. Each summand has the factor $P_{<e}\zeta_e$ and thus belongs to $\mathfrak{q}$. Consequently

$$\prod_e d_e \equiv \left(\prod_e \Gamma_e\right)\prod_j X_j \pmod{\mathfrak{q}}. \tag{**}$$

Combining the blocks in the trace ideal. Write each $d_e = \sum_i c_{e,i} \det B_{e,i}$ with $c_{e,i} \in \Lambda$, the rows of $B_{e,i}$ in N_e. Lift the rows to $\mathcal{N}$. Choose an element $q_e \in A$ by multiplying $\prod_{u\neq e}([a] - \xi_u(a))$ by the scalar $C/\prod_{u\neq e}(\xi_e(a) - \xi_u(a))$. Then $\eta_u(q_e) = C$ if $u = e$ and zero otherwise. There is also $k \in A$ acting by 1 on every Tate block and by zero on every ρ_x for $x \in \mathcal{X}$. For each such x, use its own scalar pair γ_x, y_x from the short Selmer condition and take the product of $(y_x - [\gamma_x])/(y_x - 1)$. The order of these factors can be fixed arbitrarily.

For a tuple $\mathbf{i} = (i_e)_e$, form a matrix $M_{\mathbf{i}}$ with rows in $\mathcal{N}$. Its first rows are the old-tame polynomial rows, followed by the middle pivot rows v_i. Its last rows are the sums over e of q_e times the chosen lifts of the rows of B_{e,i_e}. Under η_e, its determinant is

$$A_*\ \det B_{e,i_e}, \qquad A_* = \left(\prod_v o_v(1)\right) C^{|s|}.$$

Indeed the first rows give the displayed old-tame diagonal, the middle pivot is identity, and clearing that pivot gives the last block CB_{e,i_e}. Replace $M_{\mathbf{i}}$ by $M_{\mathbf{i}}' = M_0 + k(M_{\mathbf{i}} - M_0)$, with left multiplication entry by entry. On a genuine packet its main block then has determinant D_x and its Tate blocks have the

determinants just computed. The element $h_{\mathbf{i}} = kq_0 \prod_e c_{e,i_e}$ acts on a genuine packet only on the trivial Tate slot, there by $C \prod_e c_{e,i_e}$. In particular every element of $\ker \eta_0$ annihilates $U_x(h_{\mathbf{i}})$ by left multiplication on that packet.

We give the trace calculation that takes these matrix determinants back to the Hecke algebra. For any matrix over A, including a 1-by-1 matrix, its evaluated flattened characteristic coefficients belong to S. This follows from the determinant identity for a sum of matrices. In formal power series, write $\det(1 - u \sum A_i)$ as the product, over primitive cyclic words w in the indices, of $\det(1 - u^{|w|} A_w)$, where A_w is the ordered word in the matrices. Taking formal logarithms expands the traces of all words, grouped by powers of primitive cyclic words, proving the identity in characteristic zero. There are finitely many factors per coefficient. For a matrix over $\Lambda[G_E]$, take the summands to be scalar multiples of $E_{ij}[g]$. A word is zero, off diagonal with square zero, or diagonal in one block with a scalar times a single group element there. Its characteristic coefficients belong to S by the group-element assertion. The product identity therefore proves the stated integrality.

Now let M have rows in $\mathcal{N}$. Write $V(M_i) = r_i + \lambda_i B$ with $r_i \in \mathfrak{T}$ and $\lambda_i \in \ker \eta_0$. Work also over $S[G_E]$, with the central S-action on the evaluated matrices. Let D be the evaluated flattened determinant of M, and W a matrix over $S[G_E]$ whose evaluated flattening is the adjugate of that matrix. Such a lift is a polynomial in M by Cayley–Hamilton, with the characteristic coefficients in S. Multiplying $V(M_i)$ by W and then by the row ω gives, after evaluation in the packet family,

$$(D - \omega W \lambda) B = \omega W r.$$

For $h \in A$ put $f_{M,h} = (\operatorname{tr} U_x((D - \omega W \lambda)h))_x$. It belongs to S by the characteristic coefficient assertion over $S[G_E]$, by the same word argument. Multiply the last identity by the adjugate of $U_x(B)$ and then by $U_x(h)$; taking traces gives $b\, f_{M,h} \in \mathfrak{I}$. Indeed the adjugate of $U_x(B)$ also lifts by Cayley–Hamilton, and the trace of each term containing an r_i belongs to $\mathfrak{I}$, by cyclicity and the trace ideal property. On a genuine packet on which $U_x(\lambda_i h) = 0$ for each i, this formula gives $f_{M,h}(x) = \det U_x(M) \operatorname{tr} U_x(h)$.

Take $f = \sum_{\mathbf{i}} f_{M'_{\mathbf{i}}, h_{\mathbf{i}}}$. Then $b f \in \mathfrak{I}$, and $b \in S$ with $b_x = 0$ on $\mathcal{X}$ because B vanishes on the trivial Tate slot. On every genuine packet, the sum of products is

$$f(x) = C A_*^{t'} D_x \prod_e d_e.$$

By $(**)$, every proper squarefree coefficient of $f(x)$ below $\prod_j X_j$ vanishes. Its full squarefree coefficient at x_0 is

$$d = C A_*^{t'} D_{x_0}(0) \prod_e \Gamma_e(0), \qquad v_O(d) \le b_{\mathrm{main}} + \sum_e b_e + t' \sum_v v_O(o_v(1)) + (2t'|s| + |J| + 1) v_O(C).$$

The constant coefficient of D_{x_0} is the only contribution from that determinant to the full squarefree monomial.

To compare with the prescribed allowance, at the two readings of the twist places the normalized sum of $v_O(o_v(1))$ is at most $2c_\ell$, since $v_\ell(u^{-1} - 1) = v_\ell(u - 1)$ for a unit and the twist comes with this ramification bound. At the two readings of old level places it is at most $8f\, v_\ell(G)$. There are at most $4f$ such places before taking the two readings. The triangular estimate gives $(|J| + |S_{\det}|)h_* \le h_2^{22} h_0$, and $|s| \le |S_{\det}| + 2|J|$. Using $h_* = h_2 h_0$ and $h_0 \le v_\ell(G) \le c_\ell$ therefore gives

$$\begin{aligned} v_\ell(G^3 d) \; &\le 4 f r' \operatorname{loss}_{\mathrm{crys}}(\ell) + t' v_\ell(X) \\ &\quad + \big(10 f r' + 4 + t'(2 + 8f) + (4t' + 1) h_2^{22} + h_2\big) c_\ell. \end{aligned}$$

Here $r' \le r$, and $f, t', r \le h_2$ with $h_2 \ge 2$. The coefficient of c_ℓ is at most $18 h_2^2 + 3h_2 + 4 + 5h_2^{23} \le h_2^{27}$, and $t' \le h_2^{27}$. This proves $v_\ell(G^3 d) \le r' b_\ell$. The count uses the cost inequality multiplied by h_* and applies also when $h_0 = 0$. Both trace relations use the selected coefficient reading and the same fitted twist. The two Galois places in each conjugate pair share one variable. The determinant estimate uses the multiplier at the selected reading. The conjugate coefficient reading enters through the local test forms and integral row recombinations.

Local operators for the two first-order ties. Work at one selected auxiliary place $w = w_j$, with residue size q, at the same coefficient reading and depth as the Fitting polynomial. Put $m_1 = \ell^\nu$, $n = r' + t'$, and let $\Gamma = \mathrm{GL}_n(k_w)$ act by the residue action ρ on the deeper forms. Write $N_w = \sum_{g\in\Gamma} \rho(g)$ and N for the product of the norms at all the selected auxiliary places. The split local action, written π, agrees with ρ on the integral lifts of residue matrices. Let $H \subseteq \Gamma$ stabilize the first coordinate line. Its inverse image in $\mathrm{GL}_n(\mathcal{O}_w)$ is the parahoric P_H. The pro-p Iwahori group I_w reduces to the upper unitriangular group. The chosen Artin data give a uniformizer ϖ and a unit u with residue generating $k_w^\times$, representing the specified geometric Frobenius ϕ and tame inertia element τ, respectively. Write R_x for the attachment before normalization at this step, so $N_x(h) = a_{\mathrm{nor}}(h)R_x(h)$. The normalizing avatar is unramified at this auxiliary place and has unit values at the coefficient place. Thus $a_{\mathrm{nor}}(\tau) = 1$, and the complex Frobenius factor $\alpha_w = j_\ell^{-1}(a_{\mathrm{nor}}(\phi))$ is A_w^{-1} in the preceding spherical test.

Use t_w, P_I and U from the local Hecke calculation. For every character of P_H through its first diagonal residue, U preserves its eigenspace. Indeed for $g = \left(\begin{smallmatrix} a & b \\ \varpi c & D \end{smallmatrix}\right) \in P_H$, the upper row operation with row $-bD^{-1}$ is in I_w and makes the upper-right block of g zero. Call the result g'. Then $t_w^{-1} g' t_w \in P_H$ has the same first diagonal residue as g. On a vector y of character η this gives $P_I\pi(gt_w)y = \eta(g)P_I\pi(t_w)y$. Average this equality against $\eta(g)^{-1}$ over P_H. The projector so obtained has image fixed by I_w, since η is trivial there, and therefore $P_I\pi(t_w)y$ belongs to the η-space.

Apply the cofactor conclusion of the local Hecke calculation to the string $A_w q^i$, $0 \le i < t'$, on each spherical base packet on which $T = Ap(b)$ is nonzero. By the base separator these packets are weak and residually congruent to x_0. For $0 \le i < t' - 1$ the root $c = A_w q^i$ has $P_x(qc) = 0$, so the cofactor obtained by deleting one occurrence of c kills x under evaluation at U. Their greatest common divisor is

$$S_x(Z) = P_x(Z) \Big/ \prod_{0\le i<t'-1} (Z - A_w q^i).$$

It kills the entire U-cyclic space of x. For $t' = 1$ the empty product gives $S_x = P_x$, which kills this space by the forward recurrence. Adjoint invariance makes x orthogonal to the image of $S_x(U)$ on all I_w-fixed vectors.

Recall that $\alpha_w = A_w^{-1}$. The polynomial $S_x(A_w Z)$ is a nonzero scalar times $(Z - q^{t'-1}) j_\ell^{-1}(T_{x,w}(Z))$, with $T_{x,w}$ the quotient by the full string in the residual Hecke test. Since $T_{x,w}$ divides e_j at the selected reading, put

$$Q_j(Z) = (Z - q^{t'-1}) e_j(Z).$$

Here and below the coefficients act on complex forms through j_ℓ^{-1}. The divisibility just proved gives $Q_j(\alpha_w U)$ as a polynomial multiple of $S_x(U)$. Thus x is orthogonal to $Q_j(\alpha_w U)y$ for every I_w-fixed y.

Integral weighted sums and their norm. Let $\lambda : k_w^\times \to \mathbf{Z}/m_1$ be the logarithm with $\lambda(\bar{u}) = 1$, which exists since $m_1 \mid q - 1$. Define $b_H(h) = \lambda(h_{11})$ on H. Its transfer to Γ is $b^\Gamma = \lambda \circ \det$. Indeed for $g = \mathrm{diag}(z, 1, \dots, 1)$ the transfer is computed on the cycles of the action on lines. The first line contributes z before applying λ. Lines in the complementary hyperplane contribute 1, and all other lines have orbit length $\mathrm{ord}(z)$ and also contribute 1, since $g^{\mathrm{ord}(z)} = I$. The transfer is multiplicative before taking the logarithm, by permuting the lines in the transfer product. Its logarithm kills each elementary transvection because the residue characteristic is prime to m_1. Diagonals of the displayed form, their conjugates and elementary transvections generate Γ, proving the assertion. Choose $\gamma_\Gamma = \mathrm{diag}(\bar{u}, 1, \dots, 1)$, so $b^\Gamma(\gamma_\Gamma) = 1$.

For an integer-valued weight f on $\mathbf{Z}/m_1$ define

$$J_f = \sum_{h\in H} f(b_H(h))\rho(h), \qquad D_f = \sum_{g\in\Gamma} f(b^\Gamma(g))\rho(g), \qquad B_f = Q_j(\alpha_w U) J_f.$$

The range of J_f is I_w-fixed, since the pro-p Iwahori residues belong to $\ker b_H$. The coset sum U preserves integrality on this range, since each translation is at $w, \bar{w}$ away from ℓ and preserves the coefficient gauge at ℓ. Its output is I_w-fixed and hence fixed by the principal subgroup at w. The old-level conditions and

the conditions at the other auxiliary places are preserved by commutation. Thus J_f, D_f, B_f are integral operators on the actual deeper space, supported locally at w. They commute with the diagonal Hecke operators by preservation of the packets. In particular reducing their integral weight functions modulo m_1 gives valid operator congruences. The operator B_1 for weight 1 has H-fixed image by the parahoric invariance of U.

We use the following exact norm identity for any subgroup Δ of the full product of auxiliary residue groups. On the full principal-level class set, with maximal compact at the old places, choose a point in each Δ-orbit. Use the multiplier $G/|\operatorname{Stab}_\Delta|$ on that point and zero on the other points of its orbit. Each stabilizer is a subgroup of the reduction of a class automorphism group at maximal compact level. The principal-level torsion calculation injects that group by reduction at a chosen split prime. Its order divides G, since the catalogue ranks are at most rt. Multiplication by this integral scalar function preserves the coefficient lattice. It commutes with the old compact translations and therefore preserves the specified old-type subspace of the deeper forms. Denote the resulting operator by Ψ_Δ. Each chosen orbit point is counted exactly $|\operatorname{Stab}_\Delta|$ times in the conjugate sum, giving

$$\sum_{g\in\Delta} \rho(g)\Psi_\Delta\rho(g)^{-1} = GI.$$

Use Ψ_H here and take the full-product operator as Ψ in the extraction.

Fix representatives c_η for the lines Γ/H, and write $gc_\eta = c_{g\eta}h(g,\eta)$. Put $A_+ = \sum_\eta \rho(c_\eta)$ and $A_- = \sum_\eta \rho(c_\eta^{-1})$. For an integer lift v of the residue coordinate $x \mapsto x$ on $\mathbf{Z}/m_1$ set

$$F = -A_+B_vA_-.$$

This is integral and place-supported. The exact right shift $B_f\rho(h) = B_{f(\cdot - b_H(h))}$ yields

$$F(\rho(g)-1) \equiv A_+B_1 \sum_\eta v(b_H(h(g,\eta)))\rho(c_\eta^{-1}) \pmod{m_1}.$$

Let V_g be the same last sum preceded by $\Psi_H B_1$. By H-invariance of B_1 and the norm identity, $N_w\Psi_H B_1 = GA_+B_1$. Thus $GF(\rho(g)-1) \equiv N_wV_g \bmod m_1$.

To check the required vanishing, write $L = \widehat{T}N$. The norm has image in the base level. On a spherical base eigenpacket on which $T \neq 0$, the string orthogonality shows that the image of B_1 is orthogonal to the base vectors. The other base packets are killed by $\widehat{T}$. Using the invariant positive inner product and $N^* = N$ gives $LB_1 = 0$. Moreover $L\rho(g) = L$ for every g at this auxiliary place, so the exact equation for $N_w\Psi_H B_1$ gives $|\Gamma|L\Psi_H B_1 = 0$ before reduction. Cancel $|\Gamma|$ on the complex form space to get $LV_g = 0$. This is the first tie with its exact vanishing condition.

Frobenius on the residue character spaces. Here we use the tame local–global compatibility for a sorted pro-Iwahori exponent on the actual attachment. Its exact input off ℓ is finite inertia and the annihilating polynomial $\prod_i (Z - j_\iota(\theta_i(v)))$ at an element of the decomposition group with Artin preimage v and nonnegative geometric Frobenius degree. Apply the local exponent calculation to a nonzero vector with a diagonal unit character and scalar action of the permutations preserving it. This gives an exponent $(\theta_i)_i$ whose functional is nonzero on the vector. Its ordered coordinates have the original unit characters, and a simultaneous permutation of its exponents gives a sorted eigencharacter on a vector in the same cyclic representation. The pro-Iwahori fixed part of each such cyclic representation is finite dimensional here, with the other levels and the weight fixed. The permutation condition holds on both character spaces used below. For a determinant character it holds for all permutations. For a nontrivial first-entry character it holds for the permutations fixing the first index. The exponent functional satisfies $\mathcal{L}(Uy) = \theta_1(\varpi)\mathcal{L}(y)$ on I_w-fixed vectors. The factor q^{n-1} in U cancels the first-coordinate power $|\varpi|^{n-1}$ of the normalized torus character.

On a determinant character space $\xi \circ b^\Gamma$ put $\zeta = j_\iota(\xi(1))$. All tame exponents at u have that value, so $N_x(\tau) = \zeta I$ by finite inertia and the annihilating polynomial at τ. On a first-entry space $\xi \circ b_H$, suppose

$\xi(n!) \neq 1$. Now $R_x(\tau)$ is semisimple with possible eigenvalues $\zeta, 1$, and $R_x(\phi)$ commutes with it. The latter assertion follows by applying the local annihilator with unit preimage 1 to the commutator, whose tame exponents all have value 1. For an exponent θ nonzero on a vector in this space, put $\alpha_i = j_i(\theta_i(\varpi))$ and $H_\theta(Z) = \prod_{i>1}(Z - \alpha_i)$. On the ζ-eigenspace of $R_x(\tau)$, the same local attachment at $\phi\tau^b$, with preimage ϖu^b, gives

$$\big(R_x(\phi) - \alpha_1 I\big) H_\theta(\zeta^b R_x(\phi)) = 0, \qquad 0 \le b < n.$$

These n powers of ζ are distinct. Any positive difference of their exponents is at most $n-1$ and divides $n!$. Interpolation in ζ^b and $H_\theta(0) \neq 0$ give $R_x(\phi) = \alpha_1 I$ on that eigenspace.

Let M be any matrix over the coefficient field. For a vector y in the packet and its first-entry space, this implies the operator identity

$$j_\iota^{-1}\big(\operatorname{tr}((R_x(\tau) - I)M)\big)\, Uy = j_\iota^{-1}\big(\operatorname{tr}((R_x(\tau) - I)R_x(\phi)M)\big)\, y.$$

Indeed a nonzero difference would lie in the same character space by the preservation proved above. An exponent functional nonzero on that difference would give zero by the scalar calculation on the ζ-eigenspace and by $\mathcal{L}(Uy) = \theta_1(\varpi)\mathcal{L}(y)$. Testing the identity with $R_x(\phi)^b M$ and iterating substitutes $\alpha_w U$ for $N_x(\phi)$ in any polynomial in this trace identity. These channel calculations apply to the full catalogue at the deeper level.

The second tie. For the group argument g_0 of a trace generator, write $\widehat{r}$ for the operator of $\operatorname{tr}((N_x(\tau) - I)N_x(g_0))$. Write $\widehat{r}'$ and $\widehat{s}'$ for the operators of $\operatorname{tr}((N_x(\tau) - I)Q_j(N_x(\phi))N_x(g_0))$ and $\operatorname{tr}(Q_j(N_x(\phi))N_x(g_0))$, respectively, and $\widehat{s}$ for the trace operator with $(Z-1)e_j(Z)$ in place of $Q_j(Z)$. These traces are integral Hecke operators by the coefficient argument above. The difference $\widehat{s}' - \widehat{s}$ is divisible by m_1 as an integral operator, since $q \equiv 1 \bmod m_1$ and the difference of the polynomials is $(1 - q^{t'-1})e_j(Z)$.

On the range of D_f, decompose over $\mathbb{C}$ into the characters $\xi \circ b^\Gamma$. On each such space $N_x(\tau) - I = (\zeta - 1)I$, and $\rho(\gamma_\Gamma) - 1$ acts by $\xi(1) - 1$. Hence exactly

$$D_f \widehat{r}' = (\rho(\gamma_\Gamma) - 1) D_f \widehat{s}', \qquad N_w \widehat{r} = (\rho(\gamma_\Gamma) - 1) N_w = 0.$$

For $k_0 = n!$, put $\Delta_{k_0} f(x) = f(x - k_0) - f(x)$. On a first-entry channel of character ξ, $J_{\Delta_{k_0} f}$ is $(\xi(k_0) - 1)J_f$. The channels with $\xi(k_0) = 1$ thus have zero weight. On the others the preceding Frobenius identity gives, exactly,

$$B_{\Delta_{k_0} f} \widehat{r} = J_{\Delta_{k_0} f} \widehat{r}'.$$

The character decompositions verify identities between the integral operators. Their construction used finite integer weighted sums.

Choose $f(x) = v(x)(v(x) + k_0)$, so $\Delta_{k_0} f(x) \equiv -2k_0 v(x) \bmod m_1$. The transfer formula gives $A_+ J_v A_- \equiv D_v \bmod m_1$, as the coefficient of $\rho(g)$ in the left sum reduces to $\sum_\eta b_H(h(g, \eta)) = b^\Gamma(g)$. Also $(\rho(\gamma_\Gamma) - 1)D_v \equiv -N_w \bmod m_1$ by the weight shift. Combining the identities gives

$$2k_0 F \widehat{r} \equiv -2k_0 D_v \widehat{r}' = -2k_0 (\rho(\gamma_\Gamma) - 1) D_v \widehat{s}' \equiv 2k_0 N_w \widehat{s} \pmod{m_1}.$$

Thus the same F as in the first tie gives the second with $\kappa_0 = 2n!$. The operators F_i, N_i and the residue translations commute across distinct auxiliary places by their disjoint support. In the first tie V_g contains the class multiplier and sits on the right of the norm. These identities supply the hypotheses for the squarefree coefficient extraction with the previously fixed depth and the same scalar G.

Extracting the diagonal operator from the first-order input. Let $\mathcal{E}$ be the ring of integral endomorphisms at the deeper level. For each triangular prime i use the operators $N_i, F_i, V_{i,g}$ from the local ties, with $N = \prod_i N_i$ and the full-product norm operator Ψ. For an integral diagonal Hecke function a, write $\widehat{a}$ for its operator. The separator is $T = Ap(b)$.

Use the integral vector $l_0 = Gl$ as input and work in $\mathcal{E}l_0/\ell^v\mathcal{E}l_0$. For a squarefree multi-index d use Θ_d with factor F_i when $d_i = 0$, N_i when $d_i = 1$, and zero for any exponent above 1. Thus $\Theta_1 = N$. Put $P = G^2T$ and $k = |I| + 1$, where I is the set of triangular primes. The identities imply

$$\widehat{P}^m\Theta_d z = 0 \quad \text{if} \quad |d| + m \geq k \text{ and } N\widehat{T}z = 0. \tag{16.18}$$

For $m = 0$ this follows from the definition of Θ_d. For the induction step consider the operator $W = \widehat{P}^m\Theta_d G$. If $d_i = 1$, then $W(\rho_i(g) - 1) = 0$ by the averaging factor. If $d_i = 0$, the first tie makes this expression $\widehat{P}^m\Theta_{d+e_i}V_{i,g}$, which vanishes on every vector by the induction hypothesis, since $\widehat{T}NV_{i,g} = 0$. Hence $W\rho(g) = W$. The identity of Ψ yields $WG = W\Psi N$. Apply it to $\widehat{T}z$ to obtain (16.18) at $m + 1$.

The linear evaluation of a Hecke polynomial $\sum a_d X^d$ by $\sum \Theta_d\widehat{a}_d$, followed by multiplication by κ_0, kills the trace ideal in this quotient. For a monomial multiplied by a generator $r - X_i s$, its i-th factor is exactly the last displayed tie or the zero $\kappa_0 N_i\widehat{r}$. The generators with ℓ^v already vanish there. The cardinality of I is fixed independently of depth, and $p(0) \neq 0$ was fixed at the base level. Choose $v \geq M_\ell$ before the triangular primes, large enough that $\ell^v = w_0\kappa_0(G^2p(0))^k$ with w_0 integral. There is an integral Hecke function v with

$$P^k - (G^2p(0))^kA^k = bv.$$

Evaluate vbf at $z_0 = \Psi l_0$ and multiply by $w_0\kappa_0$. The second term in the displayed formula vanishes in the quotient. Each proper squarefree coefficient f_d satisfies $N\widehat{T}\widehat{f_d}z_0 = 0$: in the original space this is $G\widehat{T}\widehat{f_d}l_0$, which vanishes by projection agreement and the lower coefficient vanishing on w. Thus (16.18) kills its term. At the top one obtains $w_0\kappa_0N\widehat{P^k f_1}z_0 = 0$ in the quotient. Since $N\Psi l_0 = Gl_0$ and P^k is divisible by T, the element in that last congruence before reduction is $\ell^v G\widehat{A^k f_1}(Gw)$. Cancellation in the original vector space gives $G^2\widehat{A^k f_1}w \in \mathcal{E}l_0$, which is integral. Its eigenvalue on x_0 is $A(x_0)^k f_1(x_0)$ before the two factors G, and $A(x_0)$ is a unit. The final pairing with the original left test costs one further G, from finite class stabilizers. Thus the pairing is integral after G^3, with $v_\ell(G^3A(x_0)^k f_1(x_0)) \leq r'b_\ell$.

From diagonal tests to the two determinants. For each required Z, use the good idempotent on the full range and the $X - 1$ tuple idempotent on a line. The diagonal operators just constructed give the eigenvalue lower bound $e^{-r_Z\beta_{j,\ell}}$, including G^3, and integral pairings on the indicated tests. Here $r_Z = r$ for the full range and $r_Z = 1$ for a line. The left test is integral at all embeddings above ℓ. The right test is integral at j_ι and orthogonal to $Z^\perp \cap \mathrm{im}(e) \cap \ker H_w(\xi_w)$ for the indicated split places.

Here is the determinant implication for just the test vectors in use. Use the resulting G^3-multiplied diagonal operators. For a given eigenbasis coordinate, divide the corresponding linear form in the combination coefficients by one of its coefficients of largest ℓ-adic norm. This gives a nonzero linear form over the residue field. After a finite unramified enlargement of the coefficient field, its residue field has more elements than the number of forms. Each zero set occupies a fraction at most the reciprocal of the residue field size, so we can choose integral combination coefficients avoiding their union. Each diagonal entry of the combination then has norm at least $e^{-r_Z\beta_{j,\ell}}$. For $D = \dim Z$, write the D projected right test vectors y_b in the common eigenbasis of Z. The matrix pairing the diagonal operator on these vectors with the D left test vectors x_a is integral. Its determinant is the tested Gram determinant times the product of the diagonal entries. Thus

$$\left| j_\iota\Big(\det\big(\langle x_a, P_Z y_b\rangle\big)_{a,b}\Big)\right|_\ell \leq \exp(Dr_Z\beta_{j,\ell}). \tag{16.19}$$

17 The determinant tests at the selected place

Let θ be the character fitted at the selected reading j_ι above ℓ. At every stage j use $\psi = \theta^j$, with the same conductor prime to ℓ. The basis and scalar comparisons were constructed for trial characters. We apply them here using the first-order estimate at j_ι and the two integral evaluations.

17.1 Two integral tests for one stage

On the small block V° take $D_- = D_{k^-}$, $B_- = B_{k^-}$ and $k_j = h_{k^-}$ in its exact rank. The tables are the direct sum of these Hermitian coordinate spaces, one for each member index i and principal $\mathfrak{M}^2$-class c, with the unweighted sum metric. The target form metric is the class sum of the target Gram form with weight $1/|\operatorname{Aut}|$. In contrast, the adjunction with the good source uses the original Petersson measure on $[U(W_j)]$. These are the metrics defining the two adjoints in the comparison.

For each small-block representative c, choose a rational unitary δ_c with $\delta_c c$ integral and invertible at all places above ℓ. Weak approximation gives this by density of the rational points in the finite product of unitary groups, as one sees on the Cayley open set. The small-block gauge at c of a vector v at a reading a is then at most one exactly when the coordinates of $\rho_-(\delta_c)v$ are a-integral. Here ρ_- is the small-block rational weight action. Indeed rational transport moves both the representative and the vector, and integral invertible matrices act in both directions on the small-block coordinate lattice. The same choice of δ_c works for every coefficient reading above ℓ.

For $a_0 = (i, c)$ and each coordinate vector e_b, let $f_{a_0,b}$ be the table supported at a_0 with value $\overline{\rho_-(\delta_c)^{-1}B_-e_b}$. This conjugation fixes the integral matrix B_- and changes the transport to the one required in the second reading. With the pairing linear in the second entry, unitarity of ρ_- and the identity $D_-B_- = k_jI$ give

$$\langle f_{a_0,b}, \zeta\rangle = k_j\,\overline{[\rho_-(\delta_c)\bar{\zeta}_{a_0}]_b}.$$

The transport of $\bar{f}_{a_0,b}$ is the integral vector B_-e_b. Consequently $G_\psi f_{a_0,b}$ is integral at every reading above ℓ, by the gauge bound for the divided point forms. No inversion of k_j occurs in these tests.

On L_ψ use the linear functions $\ell_{a_0,b}(z) = \langle f_{a_0,b}, G_\psi^* z\rangle$. They separate L_ψ because $G_\psi^*|_{L_\psi}$ is injective and $D_-B_- = k_jI$. Among the determinants of R of these functions on any basis, choose a nonzero determinant of maximal absolute value at j_ι, and let z_i be the dual basis. Each $j_\iota(\ell_{a_0,b}(z_i))$ is integral by Cramer's rule. In the formula for the linear functions this is the condition that the conjugate table of $k_jG_\psi^*z_i$ has integral gauges at $j_\iota\circ\mathrm{conj}$. Thus $k_jG_\psi G_\psi^*z_i$ is integral at j_ι. It is orthogonal to (11.1) by the marked-period calculation and the common-adjoint identity. The first frame is the image by G_ψ of the selected test tables. For a selected f_i and each z_j, its paired entry is

$$\left\langle G_\psi f_i, P_{L_\psi}k_jG_\psi G_\psi^*z_j\right\rangle = \ell_i(K_\psi z_j).$$

The matrix is exactly the matrix of the compressed operator in the dual basis. The fitted denominator estimate applies to these two frames, with the first integral at every reading and the second integral at the selected reading and tested. It yields the bound $\exp(Rr\beta_{j,\ell})$ for $|j_\iota(d_j(\psi))|_\ell$.

The paired rank-one bound on the same twist. Put $Z = \mathbf{C}F_\psi$. The normalized point forms $\Xi = \Xi_{j_\iota}$ and $\Xi' = \Xi_{j_\iota\circ\mathrm{conj}}$ are initially orthogonal to (11.2). To apply the determinant estimate the second form must be orthogonal to $\operatorname{im} e_{\mathrm{triv}} \cap \ker H_w(\xi_w) \cap Z^\perp$. Write $h_1 = h'$ for the chosen-level $U(1)$ class number. Choose class representatives g_c integral above ℓ. Let C^* be the sum of the character-corrected central operators at g_c^{-1}. Its summands are unitary, commute with e_{triv} and $H_w(\xi_w)$, and fix Z. Applying their adjoints to a vector $y \in \operatorname{im} e_{\mathrm{triv}} \cap \ker H_w(\xi_w) \cap Z^\perp$ and descending gives the $U(1)$ class sum of Θy. Its value at any translate equals its value at 1. The class-sum adjunction identifies this value with the pairing with F_ψ, up to its nonzero fixed scale, so it vanishes. Thus the sum of the adjoint translates of y belongs to (11.2). It follows that $C^*\Xi'$ is orthogonal to the required intersection with $Z^\perp$. It remains integral because g_c is integral at ℓ and the character factors there are units. Since $P_ZC^* = h_1P_Z$, the fitted determinant on this line gives

$$|j_\iota(J_{jp}(\psi))|_\ell = \left|j_\iota(h_1\langle\Xi, P_Z\Xi'\rangle)\right|_\ell \le \exp(\beta_{j,\ell}). \tag{17.1}$$

The two Fourier selections may give different test forms. The projection and the central operators are those of the same ψ in both readings.

The joint fitted product. In the reciprocal volume ratio the projected pairings and the compressed determinants occur through the sum

$$\frac{1}{R}\log\left|j_\iota\left(\prod_j d_j(\theta^j)\right)\right|_\ell + 2\sum_{j,p}\log|j_\iota(J_{jp}(\theta^j))|_\ell.$$

The first term uses the dimension-R test at each stage. The second uses the $\kappa/2$ sharp pairs at each stage. Their sum is at most $E_\ell = (r+\kappa)\sum_j \beta_{j,\ell}$, whose finite subsums are bounded by $(r+\kappa)mB_0$. The coefficient-place comparison uses this joint upper bound.

The full lift-volume bound. The individual sharp bound of (12.6), the line estimate (17.1), and the $5m$ squared ramified factors per stage give

$$\log|j_\iota(K_\Gamma Q_\Gamma(\theta))|_\ell \le \kappa\sum_j \beta_{j,\ell} + 2m\kappa P_\ell\log\ell + (2m\kappa+10m^2)v_\ell. \tag{17.2}$$

The single-lift ratio identity gives $D_0K_\Gamma^R Z_1/Z_\theta = |c_1/c_\theta|^R\big(\prod_j d_j(\theta^j)\big)(K_\Gamma Q_\Gamma(\theta))^R$. Use the factor counts of (17.2), the joint fitted product at this place and (13.1) to obtain

$$\begin{aligned}\log|j_\iota(D_0K_\Gamma^R Z_1/Z_\theta)|_\ell \le{}& R(r+\kappa)\sum_j \beta_{j,\ell} + 2Rm\kappa P_\ell\log\ell\\ &+ R(2m\kappa+10m^2+m^2t)v_\ell.\end{aligned} \tag{17.3}$$

18 The product formula and the signed arithmetic estimate

We finish the determinant argument on the fixed frame w of Section 15. Suppose first that $|f_Y(0)|<1$ and take the place (ℓ, j_ι) selected in Lemma 15.1. Use the single character θ obtained at that place. Its added conductor is prime to ℓ. The chosen preimages of y therefore admit the two recombinations in (7.10), at j_ι and $j_\iota^\# = j_\iota c$. Extend each recombination on the u completion rows by the scalar matrix $\zeta_\theta I_u$. The endpoint construction gives integral coordinates for these added rows after the same multiplication.

Write U_+, U_- for the recombinations on the first n rows, in the convention of (7.10). In the Hermitian Gram determinant over $\mathbf{C}$, the first frame is conjugated. After applying j_ι, its entries are therefore read at $j_\iota^\# = j_\iota c$, and those of the second frame at j_ι. The resulting pairing is bilinear over $\overline{\mathbf{Q}}_\ell$. Both frames have integral coordinates, the Weyl pairing is integral, and the terminal class stabilizers are trivial. The local mixed Gram matrix is consequently integral. Its determinant is

$$j_\iota\big(H_\theta|\zeta_\theta|^{2u}\big)\, j_\iota^\#(\det U_-)\, j_\iota(\det U_+).$$

Integrality of this determinant and the bound on the sum of the two negative log determinants give $\log|j_\iota(H_\theta|\zeta_\theta|^{2u})|_\ell \le F_\ell$. The endpoint scalar bound (14.15) now gives

$$\log|j_\iota H_\theta|_\ell \le F_\ell + 4um^2t\,v_\ell. \tag{18.1}$$

The norm comparison on the same basis gives $H_\theta/H_1 = Z_\theta/Z_1$. Multiplying H_θ by the quantity $D_0K_\Gamma^R Z_1/Z_\theta$ in (17.3) thus gives exactly Y. Adding the two log bounds gives $\log|j_\iota Y|_\ell \le A_\ell$ of (15.2), as required for the chosen prime.

For either case of Lemma 15.1, apply the rational product formula to $(-1)^e f_Y(0) = \mathrm{N}_{\mathbf{Q}(Y)/\mathbf{Q}}(Y)$. Its archimedean logarithm divided by e is the mean of $\log|\sigma Y|$ over the e conjugates. Its finite logarithms divided by e are q_ℓ, and their sum plus that archimedean mean is zero. Take a finite set of primes containing the divisors of b_Y and all primes at most $\max(N, \lceil x_1\rceil)$. The sum of the normalized finite logarithms q_ℓ is bounded above by the sum of A_ℓ over this set, also when $a_Y \ge b_Y$, because all A_ℓ are nonnegative.

In this sum the contributions of $F_\ell, \sum_j \beta_{j,\ell}$ and v_ℓ are bounded by $2\delta, mB_0, x^{999}$, respectively, before taking the multipliers in (15.2). The sum contains the entire support of P_ℓ. By (12.1), $\sum_\ell 2P_\ell \log \ell = \sum_\sigma \mathrm{Ask}(w(\sigma), \sigma)$. Its coefficient $Rm\kappa$ combines with $-Rm\kappa S_0$ in (15.1) to give $-Rm\kappa \mathcal{U}^\star = -Rh_{\mathrm{main}}$. Consequently the product formula gives

$$R(h_{\mathrm{main}} - B_*) \le V + 2\delta + u(B_z + B_{\mathrm{c}}), \tag{18.2}$$

where

$$B_{\mathrm{c}} = 4m^2 t\, x^{999},$$
$$B_* = 2rmB_0 + (r+\kappa)mB_0 + 3m\kappa B_0 + (2m\kappa + 10m^2 + m^2 t)x^{999}.$$

Here $r/\kappa \to 1$, $t = 5r$, $m = O(\sqrt{x})$, $r = O(x^5)$ and $m\kappa B_0 = \mathcal{B}$. Thus $B_* = O(\mathcal{B}) + O(x^{1005}) \le 2L$ eventually, since $L = \mathcal{B}\log b$, $a \ge xb$ and $\log b = x^\varepsilon$. Use $u/R \le 4b^{-3/2}$, $B_z = x_1^2 b^2 + x^{99}$ and $\log x_1 / \log b = x^{-5\varepsilon/6}$ to obtain, for large x,

$$u(B_z + B_{\mathrm{c}}) \le Rb. \tag{18.3}$$

In particular $b \le L$ implies

$$R(h_{\mathrm{main}} - 3L) \le V + 2\delta. \tag{18.4}$$

To use the signed remainder, we need $h_{\mathrm{main}} \ge 0$ on these data. Indeed, $a/x_1 \le w(\sigma) \le a$. A positive summand of $\mathrm{Ask}(w(\sigma), \sigma)$ requires $\nu \le x_1$. Since $\Lambda(\nu) \le \log \nu$ and $1 - \chi_d(\nu) \le 2$, uniformly in σ one has

$$\mathrm{Ask}(w(\sigma), \sigma) \le 2w(\sigma) \sum_{\nu \le x_1} \frac{\log \nu}{\nu} \ll w(\sigma)(1 + \log^2 x_1) = o\big(w(\sigma)\log b\big).$$

The final equality follows from $(\log x_1)^2/\log b = x^{-2\varepsilon/3}$. By $g \ge b^{20}$, $\mathrm{rd}_E \ge 1$ and $\Omega = (2\pi/\mathrm{rd}_E)e^{\ell_d}$, we have $S(w_0(\sigma)) \ge w_0(\sigma)(20\log b - O_d(1))$. Since the rounded value satisfies $w_0(\sigma) \ge w(\sigma)$, summing over σ proves $h_{\mathrm{main}} = m\kappa\mathcal{U}^\star \ge 0$ eventually.

Substitute $V = n(h_{\mathrm{main}} - m\,\mathrm{rem} + L)$ in (18.4). With $u = R - n$ this gives

$$nm\,\mathrm{rem} \le nL + 3RL + 2\delta - uh_{\mathrm{main}}.$$

We may now drop the nonpositive term $-uh_{\mathrm{main}}$. The rank estimate gives $n \ge R/2 > 0$ for large x. Divide by n and use $R/n \le 2$ and $\delta \le 5R\log(1+b^2)\mathcal{B}$ to get

$$m\,\mathrm{rem} \le 7L + 20\log(1+b^2)\mathcal{B} \le (47\log b + 20\log 2)\mathcal{B}.$$

Since $\mathcal{B} = fa(\log b)^3$, this proves

$$m\,\mathrm{rem} \le fa(\log b)^5 \tag{18.5}$$

on every sufficiently large selected square. Throughout, d and ε were fixed before the sequence of squares. The bounds used here are uniform in the row choices and the embeddings. The finite allowances were bounded by sums on each datum before the choice of the prime and character. Thus a common threshold at this precision applies to all the selected squares. Both the weights in rem and the auxiliary scale x_1 remain those of (3.2) and (3.1).

19 From the signed bound to the zeros of zeta

Keep the choices k, d_k, ε_k and the squares $x = y_j$ of Section 2. The arithmetic comparison gives the signed inequality (18.5), $m\,\mathrm{rem} \le fa(\log b)^5$, on the datum at each such x once j is sufficiently large. The threshold is fixed once the precision and d_k are fixed.

19.1 Two scales in the weighted remainder

Take the datum at x and write $d = d_k$ and $\varepsilon = \varepsilon_k$. The place group is $(\mathbf{Z}/N_1\mathbf{Z})^\times$, where N_1 has the same prime divisors as M. The T top places form the fiber of 1 mod M, and the J thick places belong to the fiber of -1 mod M. Thus $f = \varphi(M)T$. If n is a unit modulo M, multiplication by $[-n]$ maps the top places bijectively to the fiber of $-n$ mod M. If n is a nonunit, $[-n]$ is the identity by the convention of the datum. In either case the number of top places τ for which $[-n]\tau$ is thick is $J\mathbf{1}_{n\equiv 1 \bmod M}$.

In (3.2) the nominal weight w is a/x at the thick places and a/x_1 elsewhere. Since $a \ge xb$ and $1 \le x_1 \le b$, the cutoff $n \le a$ includes all nonzero terms of the ramps at x and x_1. Dividing each $\mathrm{Ask}(a,\tau)$ by a turns its positive parts into $(1/n-1/x)_+$ or $(1/n-1/x_1)_+$ according to the translated place. The fiber count gives

$$\frac{1}{a}\sum_{\tau\in\mathcal{T}}\mathrm{Ask}(a,\tau) = J\big(S_{d,M}(x) - S_{d,M}(x_1)\big) + TS_{d,1}(x_1).$$

Taking the mean over all f places gives $\bar{w} = \log a - \log x_1 - (J/f)(\log x - \log x_1)$. By (2.15), the main term satisfies

$$\frac{T}{a}\,\mathrm{MainAsk} = \frac{J}{\varphi(M)}(\log x - \log x_1) + T(\log x_1 + c_{d,1}).$$

Subtracting, the two occurrences of $c_{d,M}$ in the progression difference cancel. The exact identity for the signed remainder is

$$\mathrm{rem} = a\big[J(R_{d,M}(x) - R_{d,M}(x_1)) + TR_{d,1}(x_1)\big]. \tag{19.1}$$

For $Q \in \{M, 1\}$ define $H_{d,Q}(t) = \sum_{n\le t}\lambda_{d,Q}(n) - t/\varphi(Q)$. In the reduced classes modulo Qd congruent to 1 mod Q, the weights $1-\chi_d$ sum to $d-1$, by the Chinese remainder theorem and $\sum_{u\in(\mathbf{Z}/d\mathbf{Z})^\times}\chi_d(u) = 0$. Among the other integers $n \equiv 1$ mod Q with nonzero von Mangoldt weight, the nonunits modulo Qd are powers of d. They contribute $O(\log t)$. Since $\varphi(Qd) = (d-1)\varphi(Q)$, the fixed-progression estimate (2.2) gives

$$H_{d,Q}(t) = O_{Q,d,B}\big(t(\log t)^{-B-1}\big) \qquad (B > 0).$$

With $b(n) = \lambda_{d,Q}(n)$ and $a_b = 1/\varphi(Q)$, (2.12) identifies $C_b(s) = G_{d,Q}(s)$ in the count and tail formulas (2.3) and (2.4). Here $C_b(1) = c_{d,Q}$. Integrating the tail bound gives

$$R_{d,Q}(y) = -\int_y^\infty H_{d,Q}(t)t^{-2}\,dt = O_{Q,d,B}\big((\log y)^{-B}\big). \tag{19.2}$$

The original second scale satisfies $\log x_1 = x^{\varepsilon/6}$ by (3.1). Choose $B > 12/\varepsilon$. At each of the two fixed moduli, $(\log x_1)^{-B} = x^{-B\varepsilon/6}$ has exponent strictly less than -2. The implied constants in (19.2) are fixed, so $|R_{d,Q}(x_1)| \le x^{-2}$ for all sufficiently large x.

Divide (18.5) by $maJ > 0$, and use (19.1). The datum has $T/J \le \log b$, $m \ge \sqrt{x}$, and $\log b = x^\varepsilon$. Hence

$$\begin{aligned} R_{d,M}(x) &\le \frac{f(\log b)^5}{mJ} + R_{d,M}(x_1) - \frac{T}{J}R_{d,1}(x_1) \\ &\le \varphi(M)x^{-1/2+6\varepsilon} + (1+x^\varepsilon)x^{-2}. \end{aligned}$$

Since $\varepsilon > 0$, the last expression is at most $x^{-1/2+8\varepsilon}$ for sufficiently large x. Thus

$$R_{d_k,M}(y_j) \le y_j^{-1/2+8\varepsilon_k} \qquad \text{for all sufficiently large } j. \tag{19.3}$$

19.2 The real one-sided estimate

Put $\alpha_k = 1/2 + 8\varepsilon_k = 1/2 + 1/(2k) < \beta_k$. The test squares satisfy $y_j \ge X_j \to \infty$. Hence $y_j^{\alpha_k-1} \le \frac{1}{2}y_j^{\beta_k-1}$ for all sufficiently large j. Combining (19.3) with maximality in (2.16) gives $R_{d_k,M}(y) \le \frac{1}{2}y^{\beta_k-1}$ at every square y of each such block. The blocks cover all sufficiently large squares. For real $x \ge 1$ take $y = \lceil\sqrt{x}\rceil^2$. The

kernels $(1/n - 1/x)_+$ are nondecreasing in x and the weights in (2.11) are nonnegative. Thus $S_{d_k,M}(x) \le S_{d_k,M}(y)$. Since $y \ge x$, $\beta_k < 1$, and $\log(y/x) \le 2x^{-1/2}$, this gives

$$R_{d_k,M}(x) \le R_{d_k,M}(y) + \frac{\log(y/x)}{\varphi(M)} \le \tfrac{1}{2}x^{\beta_k-1} + \frac{2}{\varphi(M)\sqrt{x}}.$$

The second term is at most $\frac{1}{2}x^{\beta_k-1}$ eventually, because $\beta_k > 1/2$. Therefore

$$R_{d_k,M}(x) \le x^{\beta_k-1} \qquad \text{for all sufficiently large real } x. \tag{19.4}$$

19.3 The Mellin transform and the sign at real poles

For the general sum $\mathcal{R}_b$ in (2.3), integration of one term gives

$$\int_n^\infty (y/n - 1)y^{-s-1}\,dy = \frac{n^{-s}}{s(s-1)} \qquad (\Re s > 1).$$

For $\sigma = \Re s > 1$, the sum of the absolute ramp integrals is $\sum_n |b(n)|n^{-\sigma}/(\sigma(\sigma-1)) < \infty$. Integrating the centering terms separately gives

$$\int_1^\infty y\mathcal{R}_b(y)y^{-s-1}\,dy = \frac{D_b(s)}{s(s-1)} - \frac{a_b}{(s-1)^2} - \frac{c}{s-1} = \frac{C_b(s)/s - c}{s-1}. \tag{19.5}$$

In particular the right side is holomorphic at 1 when C_b is holomorphic there and $c = C_b(1)$.

For the weights $\lambda_{d,Q}$ one has $C_b(s) = G_{d,Q}(s)$ and $c = c_{d,Q}$ by (2.12) and (2.14). Thus for $\Re s > 1$,

$$\Phi_{d,Q}(s) = \int_1^\infty yR_{d,Q}(y)y^{-s-1}\,dy = \frac{G_{d,Q}(s)/s - c_{d,Q}}{s-1}. \tag{19.6}$$

This continues meromorphically to $\Re s > 0$. It is holomorphic on $\Re s \ge 1$, by the cancellation at 1 and the boundary nonvanishing of the regularized L-functions. At $0 < \Re\rho < 1$, the logarithmic derivatives in (2.13) have at most simple poles, with residue equal to the zero order of the function being differentiated. Dividing by $\rho(\rho-1)$ gives the residue of $\Phi_{d,Q}$ at ρ in the form

$$r_{d,Q}(\rho) = \frac{\sum_{\xi \bmod Q} \left(\operatorname{ord}_\rho \widetilde{L}_\xi - \operatorname{ord}_\rho \widetilde{L}_{\xi\chi_d}\right)}{\varphi(Q)\rho(1-\rho)}. \tag{19.7}$$

The residue is zero at a regular point. For $d = d_k$, $Q = M$, the twisted orders vanish at every real $\beta_k < \sigma < 1$ by (2.6). The untwisted orders are nonnegative and $\sigma(1-\sigma) > 0$, so every real residue of $\Phi_{d_k,M}$ on $\sigma > \beta_k$ is nonnegative.

Lemma 19.1 (Landau's argument for the upper bound)**.** *For the fixed k and d_k, the transform $\Phi_{d_k,M}$ is holomorphic on $\Re s > \beta_k$.*

Proof. Write $\beta = \beta_k$ and $d = d_k$. By (19.4) and continuity on bounded intervals, take $K \ge 1$ so that

$$g(v) = Ke^{\beta v} - e^v R_{d,M}(e^v) \ge 0 \qquad (v \ge 0).$$

It is continuous because the ramp sum is locally finite and each kernel is continuous. The Laplace transform of g converges absolutely on $\Re s > 1$, where $y = e^v$ in (19.6) gives

$$P(s) = \int_0^\infty g(v)e^{-sv}\,dv = \frac{K}{s-\beta} - \Phi_{d,M}(s). \tag{19.8}$$

Let σ_c be the abscissa of convergence of the nonnegative integral. In $\Re s > \sigma_c$ the integral defines a holomorphic function by absolute convergence locally uniformly on this half-plane. Suppose $\sigma_c > \beta$. The

right side of (19.8) agrees with the integral to the right of σ_c by continuation. At σ_c it has at most a simple pole with residue $-r_{d,M}(\sigma_c) \le 0$, where $r_{d,M}$ is zero on the regular portion of the real ray. For real $s > \sigma_c$, $(s-\sigma_c)P(s) \ge 0$. Taking the right limit shows that the residue is also nonnegative. It is therefore zero and the singularity is removable.

Take a disk of radius $r > 0$ about σ_c on which this continuation is holomorphic. For $a = \sigma_c + r/4$ and $h = r/2$, the Taylor series at a converges at $a - h < \sigma_c$. Its derivatives at a are given by integrating $(-v)^j g(v) e^{-av}$, since a real σ' with $\sigma_c < \sigma' < a$ supplies a convergent majorant after multiplication by a constant depending on j. Every term in the Taylor series at $a-h$ is nonnegative. Monotone convergence gives

$$\sum_{j\ge 0} \frac{P^{(j)}(a)(-h)^j}{j!} = \int_0^\infty g(v) e^{-av} \sum_{j\ge 0} \frac{(hv)^j}{j!}\, dv = \int_0^\infty g(v) e^{-(a-h)v}\, dv.$$

The left side is finite, contradicting $a - h < \sigma_c$. Consequently $\sigma_c \le \beta$, so (19.8) gives holomorphy of $\Phi_{d,M}$ on $\Re s > \beta$. □

19.4 Exclusion of zeros off the critical line

Suppose $\zeta(\rho) = 0$ with $1/2 < \Re\rho < 1$. Choose $k \ge 13$ so that $\beta_k < \Re\rho$ and $|\Im\rho| \le k$. For this k, take the same prime d_k in the transform as in (2.6). The principal character ξ_0 mod M has

$$\widetilde{L}_{\xi_0}(s) = (s-1)\zeta(s) \prod_{p\mid M} (1 - p^{-s}).$$

The multiplier of $\zeta(s)$ is nonzero at ρ, so the principal character contributes a positive order to the numerator of (19.7). The other untwisted orders are nonnegative and the twisted orders all vanish. The residue $r_{d_k,M}(\rho)$ is then nonzero, contrary to Lemma 19.1. This excludes every zero with $1/2 < \Re\rho < 1$.

For modulus one the boundary nonvanishing proved in Section 2 gives $\zeta(s) \ne 0$ on $\Re s \ge 1$ away from the pole at 1. The continuation and functional equation of Riemann [83, Wilkins translation, pp. 2–3] give $\zeta(s)$ meromorphic on $\mathbf{C}$ with only the simple pole at 1, and the identity

$$\pi^{-s/2}\Gamma(s/2)\zeta(s) = \pi^{-(1-s)/2}\Gamma((1-s)/2)\zeta(1-s).$$

For $\Re s < 1/2$ and $s \ne 0$, the reflected zeta value is finite and nonzero by the right-half exclusion and boundary nonvanishing. The factor $\Gamma((1-s)/2)$ is finite and nonzero there. The reciprocal of $\Gamma(z)$ is entire with simple zeros precisely at $z = 0, -1, -2, \ldots$ [46, §5.2(i)]. Solving the equation for $\zeta(s)$ therefore gives exactly the negative even integers as zeros on $\Re s < 1/2$ away from 0. At 0, use $\Gamma(s/2)^{-1} = s/2 + O(s^2)$, $\zeta(1-s) = -1/s + O(1)$, and $\Gamma(1/2) = \sqrt{\pi}$ [46, (5.4.6)] in the same equation to obtain $\zeta(0) = -1/2$. This proves Theorem 1.

References

[1] Dimitris Koukoulopoulos. Pretentious multiplicative functions and the prime number theorem for arithmetic progressions. *Compositio Mathematica* 149 (2013), 1129–1149. Theorem 1.1.

[2] D. R. Heath-Brown. A mean value estimate for real character sums. *Acta Arithmetica* 72 (1995), 235–275. Corollary 2, p. 238.

[3] F. A. Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proceedings of the National Academy of Sciences* 32 (1946), 331–332.

[4] J. S. Milne. *Class Field Theory*, version 4.03, 2020. Chapter V, Theorem 3.6.

[5] Jesse Thorner and Asif Zaman. A unified and improved Chebotarev density theorem. *Algebra & Number Theory* 13 (2019), 1039–1068. Theorem 1.4, p. 1042.

[6] J. A. Green. The characters of the finite general linear groups. *Transactions of the American Mathematical Society* 80 (1955), 402–447. Lemma 7.7 and Theorems 12–14, pp. 439–444.

[7] J. S. Milne. Introduction to Shimura varieties. In *Harmonic Analysis, the Trace Formula, and Shimura Varieties*, Clay Mathematics Proceedings 4, American Mathematical Society, 2005, pp. 265–378. §11 and Theorem 11.2 in the author's 2017 revision.

[8] David A. Vogan, Jr. Unitarizability of certain series of representations. *Annals of Mathematics* 120 (1984), 141–187. Theorems 1.2–1.3.

[9] David A. Vogan, Jr. and Gregg J. Zuckerman. Unitary representations with non-zero cohomology. *Compositio Mathematica* 53 (1984), 51–90.

[10] Michael F. Atiyah and Wilfried Schmid. A geometric construction of the discrete series for semisimple Lie groups. *Inventiones Mathematicae* 42 (1977), 1–62. Sections 3 and 8, especially (3.10), (3.12) and Theorem (8.1).

[11] Binyong Sun. Matrix coefficients of cohomologically induced representations. *Compositio Mathematica* 143 (2007), 201–221. See Theorem 6.3 and (25)–(26).

[12] Stephen S. Kudla. Seesaw dual reductive pairs. In *Automorphic Forms of Several Variables*, edited by I. Satake and Y. Morita, Progress in Mathematics 46, Birkhäuser, 1984, pp. 244–268. Section 1, Definition 1.3 and (1.14), p. 249.

[13] Annegret Paul. First occurrence for the dual pairs $(U(p,q), U(r,s))$. *Canadian Journal of Mathematics* 51 (1999), 636–657. See Theorem 2.7 and Proposition 3.4.

[14] Atsushi Ichino. Theta lifting for tempered representations of real unitary groups. *Advances in Mathematics* 398 (2022), article 108188. arXiv v2, §4.1, p. 10, Atobe's criterion in Theorem 6.1(i), p. 20, Proposition 7.1, Remark 7.2 and §7.2, pp. 25–26, and the action conventions in the proof of Lemma 7.10, p. 30.

[15] Harish-Chandra. Discrete series for semisimple Lie groups. II. Explicit determination of the characters. *Acta Mathematica* 116 (1966), 1–111. §33, Corollary 1 of Lemma 65, pp. 77–78.

[16] Nolan R. Wallach. *Real Reductive Groups I*. Pure and Applied Mathematics 132, Academic Press, 1988. §§7.2.1–7.2.2 and 7.7.1.

[17] Dragan Miličić. Asymptotic behaviour of matrix coefficients of the discrete series. *Duke Mathematical Journal* 44 (1977), 59–88. Part III, Theorem 1 and its corollary. For the vanishing of the constant term of the discrete-series coefficient used in the theta lift, see E. P. van den Ban and J. J. Kuit, *Representation Theory* 21 (2017), pp. 467–468, (I.1)–(I.2).

[18] Stephen D. Miller and Wilfried Schmid. On the rapid decay of cuspidal automorphic forms. *Advances in Mathematics* 231 (2012), 940–964. Theorem A, Theorems 2.23–2.24.

[19] Chenyan Wu. Irreducibility of theta lifting for unitary groups. *Journal of Number Theory* 133 (2013), 3296–3318. Author's version.

[20] Roger Howe. Transcending classical invariant theory. *Journal of the American Mathematical Society* 2 (1989), 535–552. Theorems 1 and 2.1.

[21] Allen Moy and Gopal Prasad. Jacquet functors and unrefined minimal K-types. *Commentarii Mathematici Helvetici* 71 (1996), 98–121. Theorem 6.11, p. 118, and Proposition 6.6, pp. 114–115.

[22] Michael Harris and Jean-Pierre Labesse. Conditional base change for unitary groups. *Asian Journal of Mathematics* 8 (2004), 653–684.

[23] Laurent Clozel and Patrick Delorme. Le théorème de Paley–Wiener invariant pour les groupes de Lie réductifs. II. *Annales scientifiques de l'École Normale Supérieure* (4) 23 (1990), 193–228. Corollary, p. 213.

[24] Susana Salamanca Riba. On the unitary dual of some classical Lie groups. *Compositio Mathematica* 68 (1988), 251–303. Theorem 1.2, p. 253.

[25] Diana Shelstad. Characters and inner forms of a quasi-split group over **R**. *Compositio Mathematica* 39 (1979), 11–45. §4.

[26] Martin Olbrich. L^2-invariants of locally symmetric spaces. *Documenta Mathematica* 7 (2002), 219–237. §2.

[27] André Weil. Sur la formule de Siegel dans la théorie des groupes classiques. *Acta Mathematica* 113 (1965), 1–87. Theorem 5 and (41), pp. 76–77.

[28] Wee Teck Gan, Jonathan P. Hanke and Jiu-Kang Yu. On an exact mass formula of Shimura. *Duke Mathematical Journal* 107 (2001), 101–131. Proposition 2.13 and Section 4, especially (4.2)–(4.3) and Propositions 4.4–4.5.

[29] John T. Tate. Fourier analysis in number fields and Hecke's zeta-functions. In J. W. S. Cassels and A. Fröhlich (eds.), *Algebraic Number Theory*, Thompson Book Company, 1967, pp. 305–347. §2.5 and §§4.4–4.5 in the 1950 thesis.

[30] Shunsuke Yamana. The Siegel–Weil formula for unitary groups. *Pacific Journal of Mathematics* 264 (2013), 235–257. Theorem 2, p. 237, and Lemma 3.1, p. 246.

[31] Yifeng Liu. Arithmetic theta lifting and L-derivatives for unitary groups, I. *Algebra & Number Theory* 5 (2011), 849–921. Theorem 2.1(1), p. 859, and §2B, (2-2)–(2-3), pp. 861–862.

[32] Wee Teck Gan, Yannan Qiu and Shuichiro Takeda. The regularized Siegel–Weil formula (the second term identity) and the Rallis inner product formula. *Inventiones Mathematicae* 198 (2014), 739–831. §11.3, p. 52 in arXiv v3.

[33] Walther Landherr. Äquivalenz Hermitescher Formen über einem beliebigen algebraischen Zahlkörper. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 11 (1935), 245–248. Local–global equivalence, p. 247.

[34] Kaan Akin, David A. Buchsbaum and Jerzy Weyman. Schur functors and Schur complexes. *Advances in Mathematics* 44 (1982), 207–278. Theorem II.2.16.

[35] Hongyu He. On the Gan–Gross–Prasad conjecture for $U(p, q)$. *Inventiones Mathematicae* 209 (2017), 837–884. Definition 1 and Theorem 0.1 (arXiv version 2).

[36] Nicolas Prudhon. Translation of Dolbeault representations. *Journal of Lie Theory* 29 (2019), 1103–1118. Proposition 3.6 and (14).

[37] Anthony W. Knapp and David A. Vogan, Jr. *Cohomological Induction and Unitary Representations*. Princeton Mathematical Series 45, Princeton University Press, 1995. Book page.

[38] Joseph Bernstein and Bernhard Krötz. Smooth Fréchet globalizations of Harish-Chandra modules. *Israel Journal of Mathematics* 199 (2014), 45–111. Theorems 1.1 and 10.6.

[39] Alexander Molev, Maxim Nazarov and Grigori Olshanski. Yangians and classical Lie algebras. *Russian Mathematical Surveys* 51 (1996), 205–282. Proposition 2.7(4) and Remark 2.11.

[40] Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen* 296 (1993), 625–635, Theorem (2.2).

[41] Jean-Benoît Bost and Huayi Chen. Concerning the semistability of tensor products in Arakelov geometry. *Journal de Mathématiques Pures et Appliquées* 99 (2013), 436–488. §§1.2–1.3, the Harder–Narasimhan polygon, and §§2.2–2.4, especially (36), p. 457, Definition 2.3, p. 458, and §2.6, p. 461.

[42] Annette Huber and Guido Kings. Bloch–Kato conjecture and main conjecture of Iwasawa theory for Dirichlet characters. *Duke Mathematical Journal* 119 (2003), 393–464. Preprint, Theorems 1.3.1 and 1.4.1.

[43] John Rognes and Charles Weibel, with an appendix by Manfred Kolster. Two-primary algebraic K-theory of rings of integers in number fields. *Journal of the American Mathematical Society* 13 (2000), 1–54. Appendix A, Theorem A.1 and Corollary A.3.

[44] Manfred Kolster. Higher relative class number formulae. *Mathematische Annalen* 323 (2002), 667–692. For the cohomological form, see John Voight, The Gauss higher relative class number problem, *Annales des Sciences Mathématiques du Québec* 32 (2008), 221–232, Proposition 1.3, (6), and the exact abelian case on pp. 224–225.

[45] Manfred Kolster and Jonathan W. Sands. Annihilation of motivic cohomology groups in cyclic 2-extensions. *Annales des sciences mathématiques du Québec* 32 (2008), 175–187. §§2 and 5 and Theorem 6.1, p. 183.

[46] Tom M. Apostol. Zeta and related functions. In *NIST Digital Library of Mathematical Functions*, Chapter 25. Zeta Laurent expansion, Euler product, Dirichlet L-functions. R. A. Askey and R. Roy, Gamma function, Chapter 5, §5.2(i) and (5.4.6).

[47] Charles A. Weibel. *The K-book: An Introduction to Algebraic K-theory*. Graduate Studies in Mathematics 145, American Mathematical Society, 2013. Theorem VI.8.2 and Theorem VI.8.4.

[48] J. S. Milne. *Arithmetic Duality Theorems*, 2nd ed., BookSurge, 2006. Corollary I.2.3, p. 28, Theorem I.2.8, p. 31, Theorem I.4.10(a), p. 57, Corollary I.4.15, p. 62, and Proposition II.2.9, pp. 170–171.

[49] Brian Conrad. Lifting global representations with local properties. Author's version, 11 December 2011. Appendix B, Propositions B.3–B.4.

[50] Gaëtan Chenevier and Michael Harris. Construction of automorphic Galois representations, II. *Cambridge Journal of Mathematics* 1 (2013), 53–73. The authors' version.

[51] Andrei V. Zelevinsky. Induced representations of reductive p-adic groups. II. On irreducible representations of $\mathrm{GL}(n)$. *Annales Scientifiques de l'École Normale Supérieure* 13 (1980), 165–210. Published article.

[52] Colin J. Bushnell and Guy Henniart. The essentially tame local Langlands correspondence, I. *Journal of the American Mathematical Society* 18 (2005), 685–710. §2.2 and Corollary 3.3.

[53] Colin J. Bushnell and Guy Henniart. The essentially tame local Langlands correspondence, III. The general case. *Proceedings of the London Mathematical Society* 101 (2010), 497–553. Proposition 21.

[54] Moshe Adrian and David Roe. Rectifiers and the local Langlands correspondence: the unramified case. *Mathematical Research Letters* 23 (2016), 593–619. Proposition 3.4.

[55] Thomas Barnet-Lamb, Toby Gee, David Geraghty and Richard Taylor. Local-global compatibility for $l = p$, II. *Annales scientifiques de l'École Normale Supérieure* 47 (2014), 165–179. Theorem 1.1, comparison of the monodromy partitions above ℓ.

[56] Jean-Pierre Labesse. Changement de base CM et séries discrètes. Condition $(*)$ and Corollary 5.3.

[57] Jack A. Thorne. A p-adic approach to the existence of level-raising congruences. *Proceedings of the London Mathematical Society* 128 (2024), e12584. §4, equations (4.1)–(4.2), Theorem 4.2 and Corollary 4.3.

[58] Colette Mœglin and Jean-Loup Waldspurger. Le spectre résiduel de $\mathrm{GL}(n)$. *Annales scientifiques de l'École Normale Supérieure* 22 (1989), 605–674. Theorem of the introduction.

[59] Christos Anastassiades and Jack A. Thorne. Raising the level of automorphic representations of GL_{2n} of unitary type. *Journal of the Institute of Mathematics of Jussieu* 21 (2022), 1421–1444. Proof of Corollary 3.4.

[60] Thomas Barnet-Lamb, Toby Gee, David Geraghty and Richard Taylor. Potential automorphy and change of weight. *Annals of Mathematics* (2) 179 (2014), 501–609. Theorem 2.1.1.

[61] François Rodier. Représentations de $\mathrm{GL}(n, k)$ où k est un corps p-adique. *Séminaire Bourbaki*, 1981–1982, exposé 587, *Astérisque* 92–93 (1982), 201–218, §§4.3–4.4 and Theorem 3.

[62] Peter Schneider and Ernst-Wilhelm Zink. K-types for the tempered components of a p-adic general linear group. *Journal für die reine und angewandte Mathematik* 517 (1999), 161–208. §5, Proposition 6.

[63] Jack Shotton. The Breuil–Mézard conjecture when $\ell \neq p$. *Duke Mathematical Journal* 167 (2018), 603–678. §6.5, Proposition 6.9 in the author's version.

[64] Spencer Bloch and Kazuya Kato. L-functions and Tamagawa numbers of motives. In *The Grothendieck Festschrift*, vol. I, Progress in Mathematics 86, Birkhäuser, 1990, 333–400. Lemma 3.8.1, Corollary 3.8.4 and Example 3.9.

[65] Toby Gee, Florian Herzig and David Savitt. General Serre weight conjectures. *Journal of the European Mathematical Society* 20 (2018), 2859–2949. Lemma 5.1.6(i).

[66] Toby Gee, Florian Herzig, Tong Liu and David Savitt. Potentially crystalline lifts of certain prescribed types. *Documenta Mathematica* 22 (2017), 397–422. §2.3.

[67] Jean-Marc Fontaine and Guy Laffaille. Construction de représentations p-adiques. *Annales scientifiques de l'École Normale Supérieure* (4) 15 (1982), 547–608. Theorem 6.1 and Proposition 7.15.

[68] Tong Liu. On lattices in semi-stable representations: a proof of a conjecture of Breuil. *Compositio Mathematica* 144 (2008), 61–88. Author's version.

[69] Pavel Čoupek, Evangelia Gazaki and Adriano Marmora. Exactness property of Breuil–Kisin functors and Bloch–Kato Selmer groups. Proposition 3.4 and arXiv:2603.26035v3 (21 October 2026), Theorem 3.12(2) and Corollary 3.16.

[70] Bhargav Bhatt and Peter Scholze. Prismatic F-crystals and crystalline Galois representations. *Cambridge Journal of Mathematics* 11 (2023), 507–562. Authors' version.

[71] Laurent Berger. Limites de représentations cristallines. *Compositio Mathematica* 140 (2004), 1473–1498. Sections II.1–II.3, III.2–III.4 and IV.1.

[72] Kiran S. Kedlaya, Jonathan Pottharst and Liang Xiao. Cohomology of arithmetic families of (φ, Γ)-modules. *Journal of the American Mathematical Society* 27 (2014), 1043–1115. Example 6.2.6.

[73] Laurent Herr. Sur la cohomologie galoisienne des corps p-adiques. *Bulletin de la Société Mathématique de France* 126 (1998), 563–600. Integral comparison.

[74] Bhargav Bhatt, Matthew Morrow and Peter Scholze. Integral p-adic Hodge theory. *Publications Mathématiques de l'IHÉS* 128 (2018), 219–397. Lemma 4.26.

[75] John T. Tate. p-divisible groups. In *Proceedings of a Conference on Local Fields* (Driebergen, 1966), edited by T. A. Springer, Springer, 1967, pp. 158–183. Published chapter.

[76] John Tate. Duality theorems in Galois cohomology over number fields. In *Proceedings of the International Congress of Mathematicians* (Stockholm, 1962), Institut Mittag-Leffler, 1963, pp. 288–295. §2, Local theory.

[77] Andrew Wiles. Modular elliptic curves and Fermat's last theorem. *Annals of Mathematics* (2) 141 (1995), 443–551. Proposition 1.6, p. 472.

[78] Emil Artin and John Tate. *Class Field Theory*. AMS Chelsea Publishing 366, American Mathematical Society, 2009.

[79] Hermann Minkowski. *Geometrie der Zahlen*. B. G. Teubner, Leipzig, 1910.

[80] Kęstutis Česnavičius. Poitou–Tate without restrictions on the order. *Mathematical Research Letters* 22 (2015), 1621–1666. Theorem 5.1 and §6.

[81] John Tate. On the conjectures of Birch and Swinnerton-Dyer and a geometric analog. *Séminaire Bourbaki*, 1965–1966, exposé 306, Theorem 2.2.

[82] N. Tschebotareff. Die Bestimmung der Dichtigkeit einer Menge von Primzahlen, welche zu einer gegebenen Substitutionsklasse gehören. *Mathematische Annalen* 95 (1926), 191–228.

[83] Bernhard Riemann. Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse. *Monatsberichte der Berliner Akademie*, November 1859, 671–680. Functional equation, p. 2 of David R. Wilkins's English translation.