VI. Construction of the fitted family and its cohomological inputs

1. The theorem and its construction
2. Choosing the primes, field and scale
3. Cohomological length and the choice of a common interval
  3.1. Exceptional twists and uniformity
  3.2. What the automorphic representations supply
  3.3. Conductor-one tower checks
  3.3a. Vanishing in the trivial and quadratic residual cases
  3.4. The conductor-one odd minus bound
    3.4.1. Complex conjugation, coefficient signs and descent
    3.4.2. A single integer controls the weighted sum
    3.4.3. From reciprocity to the character bound
    3.4.4. Placement in the ordinary window
  3.5. Which odd plus estimates does ordinary zeta need?
  3.6. Why two Frobenius traces can bound a sum over all exterior primes
    The modular form at a split point of Count I
    The two off-block Selmer groups
    The one-character cotangent comparison
    Special-fibre finiteness in Count I
    The trace product
  3.7. Trace products for Counts II and V
    Adjoint rigidity of the irreducible rank-three points in I and V
    How a rank-three seed reaches the other irreducible points
    The rank-four seed in Count II
  3.8. Classical local–global compatibility for the linked ordinary towers
    The determinant of the projected rank-four module
    The linked seed at the two definite unitary levels
    Conditional support propagation on the two links
4. From one host to the points at a cutoff
  The primes and the level ideal on the selected field
  Common weights on the chosen host
    A determinant estimate on the original interval of weights
    Counting the initial segment and pairing its columns
    The determinant on a capped prefix
    Internal and terminal averages on the chosen weights
    Integral harmonic splitting on the flat rows
    Contraction on the bounded middle rows

Restoring the bounded weights outside the tops
Realization with prescribed weights and level
The signed row covolume
The mass ratio and the word-level prime
The two automorphic passes and the point ledger
Saturation of the reference support
How the finite types exclude proper blocks
A common normalization of the row periods
The special values in the common row normalization
The field generated by the critical ratios
Conjugates and the signed finite degree
The valuation at the prime selecting the twist
One character for all the theta evaluations
The provisional and final kits on the same host
5. Fitting the comparison and consuming the family

Abstract. The height argument needs a single family at each cutoff which has both the required prime distribution and small cohomology groups. We trace how the choice of two auxiliary primes produces the field and how the localization kernels enter the error in the height. For ordinary zeta, the equivariant Coates–Sinnott theorem bounds the odd minus contribution by $[E:\mathbb{Q}]t^2\log(tN_E)$, including the ramified coefficient primes. For the odd plus contribution, class field theory eliminates two cases, and an integral estimate for the second algebraic K-group settles the bounded-degree cases at 3 and 7. In the three surviving cases the same two rational primes, 11 and 23, suffice for the numerical trace comparisons. Pan's theorem supplies the modular forms for the split rank-two blocks of Count I. The higher-rank constructions and the deformation-ring estimates enter the remaining comparisons.

Let $\mathbf{1}$ be the primitive Dirichlet character of conductor one. At each sufficiently late cutoff in a specified set of real cutoffs, the construction must supply geometric objects, periods and estimates for one host. Different cutoffs can use different hosts. The quantifier of interest is one family for each sufficiently small positive ε, with a single choice of the family constants before the cutoff varies. The following proof reconstructs the assembly from stated inputs. **Proved here** marks a prose deduction. **Classical** marks a named theorem used as input. **Supplied by Papers II–V** refers to the indicated interface of the other papers in this series, including any hypotheses carried by that interface. **Release input** marks a substantive assertion whose proof in the pinned release still needs a prose translation in this series. All source links are to commit `e7c8972c701d3396947808ab20cc16dc14623d33`. Family class and host requirement Quantifiers of the connected class

1. The theorem and its construction

The entry point `riemannHypothesis` calls `rhOfAll`, which applies the ordinary zeta reduction to `allChi`. The latter calls `everyCharacter_of_row` on the release's integral local–global Galois row. This row has two uses: it goes into the representation and reciprocity inputs for the geometric construction, and it goes into the cohomological-window construction via `vwin_of_row`. The final assembly specializes a

method for all primitive characters. Its full window requirement is $\forall \psi\ \mathrm{VWinAt}(\psi, 401)$. When tracing the ordinary specialization, only $\mathrm{VWinAt}(\mathbf{1}, 401)$ is needed by the host supplier for that character. Export Ordinary reduction from all characters Final conclusion Galois row construction Supply for a fixed character

Assembly theorem (proved here from the stated inputs). Fix the target $\psi = \mathbf{1}$, put $\lambda = 3/2$, and assume the pair and cohomological supply of §§2–3 and the object, representation and comparison inputs of §§4–5. There is an $\varepsilon_* > 0$ such that for every $0 < \varepsilon < \varepsilon_*$ there exist a family F, constants $K, c_{\mathrm{HH}}, e_K, C_L$, and an onset x_* with the following properties.

1. Its cutoff set $S \subset \mathbb{R}$ contains arbitrarily late power intervals, in the sense that for every real u there is $Z \geq u$ with $[Z, Z^{3/2}] \subset S$.
2. Every $x \in S$ with $x \geq x_*$ has a nonempty finite index set I_x. All members at x have the scale x, ε, J and the same arithmetic host data specified in §4. They carry the host analytic and cohomological estimates of §3.
3. For every such x the same F has $\mathrm{ChainTied}(F; 2, 0, K, c_{\mathrm{HH}}, e_K, C_L, x)$. This is the comparison data at x, with its field, ranks, weights and selected objects taken from F there. Its constants are independent of x. Window definition Actual shape call Ties at a cutoff

In the proof below e is a dyadic tower index, j is a positive cohomological twist, and ℓ is its coefficient prime. These are distinct from the two auxiliary rational primes q, d, from the canonical extra level prime q_0, and from the cutoff x. The notation n denotes the order of the arithmetic frame, which also indexes the initial point choices. Window field and tower index Level prime and point indices

2. Choosing the primes, field and scale

Pair selection (proved for the ordinary character in Paper II). Given an integer $M \geq 1$, $0 < \varepsilon < 1/52$, and $A \geq 1$, there is Y_0 so that for every $Y \geq Y_0$ one can take primes q, d with $(qd, M) = 1$, $q \equiv 3 \bmod 8$, $d \equiv 7 \bmod 16$, and

$$\tfrac{1}{2} z < \log q \leq z, \qquad \tfrac{1}{2} z' < \log d \leq z', \qquad z = Y^{\varepsilon/5}, \quad z' = z(\log Y)^2.$$

Here $q \mid d + 1$ and $p \mid d + 1$ for every odd prime $p \leq z$. For each character modulo M, twist by one of $J_q, J_d, J_q J_d$, where $J_r(a) = (a/r)$, and use the inducing primitive character ρ of the resulting function on units. The centered ramp and log derivative at one satisfy $|g_\rho(y)| \leq y^{1/2+\varepsilon/2}$ and $|L'(\rho, 1)/L(\rho, 1)| \leq 4 \log y$ for all $Y \leq y \leq Y^A$. A further assertion applies to real primitive bases of conductor dividing $8M$. For each of the three nonempty twists the primitive L-function has no zero in $\Re s > 1 - 1/400$, $|\Im s| \leq 8M e^z e^{j_d z'}$, where j_d is 1 when J_d occurs and 0 otherwise. At $M = 1$ the ramp tests have one base, while the real box tests still include bases of conductor dividing 8. Definition of the pair and sizes Analytic conditions and quantifiers

The pair theorem selects q first and then d in a reduced class modulo $16q \prod_{3 \leq p \leq z} p$, with the product over primes. The inequality $q > z$ makes the congruences compatible by the Chinese remainder theorem. Its estimates include a progression count, a zero-box discard and a log-derivative and centered-ramp discard using short character sums. For example, **proved here from the counting bounds**, if $U_1 = e^z$, $U_2 = e^{z'}$, a zero-box discard of size $O_M(U_1^{1/4})$ among $\geq U_1^{9/10}$ first choices costs $O_M(U_1^{-13/20})$ of those choices.

The corresponding second discard has proportion $O_M(U_2^{-13/20+1/(4(\log Y)^2)})$ since $U_1 = U_2^{1/(\log Y)^2}$. The release also needs the short-sum discard on each stage and its estimates uniform over the relevant bases. Paper II, §2 derives these estimates at conductor one from the published progression and zero-density theorems. Paper I, §6 supplies the two required contour estimates. Prime selection · Contour estimates Two-stage assembly Combined analytic discards Prior density trace

Proved here. For odd $p \le z$ the congruence $d \equiv -1 \bmod p$ gives $(-d/p) = 1$. At 2 the same splitting follows from $-d \equiv 1 \bmod 8$. The congruence $d \equiv -1 \bmod q$ also gives $(-d/q) = 1$. For any fixed $J \ge 1$ and $Y \ge x^{1/J}$, the lower bounds on $\log q, \log d$ eventually exceed $B \log x$ for each fixed $B > 0$. This gives $q, d > x^B$ uniformly on these scales, including all fixed-prime exclusions and $q \ne 3$ after increasing the onset. Host congruences Size exclusions used by the supply

Let $K_a = \mathbb{Q}(\zeta_{2^a})$ and let K_a^+ be its maximal real subfield. For the ordinary route the field specified by $\mathrm{IsChiField}(\mathbf{1}, e, q, d, E)$ is, up to $\mathbb{Q}$-isomorphism,

$$E = K_{e+2}^+(\sqrt{-q}, \sqrt{-d}).$$

Proved here. In a cyclotomic field containing these three fields, fixing K_{e+2}^+ is equivalent to $a \equiv \pm 1 \bmod 2^{e+2}$ for the cyclotomic automorphism with exponent a. Since $q, d \equiv 3 \bmod 4$, the discriminants of the quadratic fields are $-q, -d$. Quadratic reciprocity, or the quadratic Gauss sum, gives their two fixing conditions $J_q(a) = J_d(a) = 1$. The subgroup fixing a compositum is the intersection of its fixing subgroups. The fourth condition $\psi(a) = 1$ in the general definition imposes no restriction for $\mathbf{1}$. Thus the field is CM, of degree $4 \cdot 2^e$, since the two quadratic radicals are independent over K_{e+2}^+ by their disjoint odd ramification. The tower conductor parameter QS in the host equals 2^{e+2} at conductor one, by the two inclusions of divisors in the frame construction. The frame order is $n = c\,2^e$, with $1 \le c \le |\langle \psi \rangle|$ eventually. Hence $c = 1$ for $\mathbf{1}$. In the general construction the integer factor is the eventual order of a descending sequence of subgroups of a finite character quotient. Fixing subgroup Tower factor and stabilization Host field construction

Proved here (rescaling a cohomological interval). For $x \ge 1$ put $e(x) = \lceil (2 + 8\varepsilon) \log_2 x \rceil$, so $x^{2+8\varepsilon} \le 2^{e(x)} \le 2x^{2+8\varepsilon}$. Given the window input defined in §3, take J from that input, and take the pair theorem with $A = J$ at $Y = X^{1/J}$. It gives one q, d for all cutoffs up to X, and the cohomological input gives $[W, W^2] \subset [Y, X]$ for that pair. For the ordinary case $c = 1$ use $[W, W^{3/2}]$. In general set $k = c^{1/(2+8\varepsilon)}$ and use $x \in [kW, (kW)^{3/2}]$ with tower index $e(x/k)$. Once $W \ge k^3$, this interval lies in $[Y, W^2]$ and $x/k \in [W, W^2]$. Also $c\,2^{e(x/k)}$ is in the required dyadic bracket for $x^{2+8\varepsilon}$. Once $W \ge k^{400}$, $(x/k)^{401} \ge x^{400}$. The bound $\exp((x/k)^\varepsilon/2)$ may be enlarged to $\exp(x^\varepsilon/2)$. The output $\mathrm{WinDatum}$ at x consists of $x^{1/J} \le Y \le x$, the analytic pair up to x, the bracket on $n = c2^e$, and the cohomological estimate with twist limit $\lfloor x^{400} \rfloor$. Taking X arbitrarily large makes the starting point kW arbitrarily large. Dyadic bracket and monotonicity Window rescaling Output at one cutoff

3. Cohomological length and the choice of a common interval

For a number field E, a rational prime ℓ, and an integer $j > 0$, let S_ℓ be the places above ℓ and let G_{E,S_ℓ} be the Galois group unramified outside those finite places. Put

$$Ш_\ell(E,j)=\ker\left(H^1(G_{E,S_\ell},(\mathbb{Q}_\ell/\mathbb{Z}_\ell)(-j))\longrightarrow\prod_{v\in S_\ell}H^1(E_v,(\mathbb{Q}_\ell/\mathbb{Z}_\ell)(-j))\right).$$

Here the twist $(-j)$ uses the inverse j-th power of the cyclotomic character. Write $\lambda_\ell(E,j)$ for its $\mathbb{Z}_\ell$-length, allowing $+\infty$. The release's $\mathrm{VVle}(E,t,B)$ requires, for every choice of the closures and localization data in this definition, finite length for every ℓ and $1\le j<t$, and

$$\sum_{\ell\in P}\log\ell\sum_{1\le j<t}\lambda_\ell(E,j)\le B\quad\text{for every finite set of primes }P.$$

In particular an infinite length cannot be hidden by a convention for converting extended integers to ordinary integers. For a CM E the release also separates the odd-prime lengths into the two complex-conjugation parts. Coefficient and localization definitions Quantifiers of the length budget Combination of the parts

Release input (the window to be consumed). For every $0<\varepsilon<1/52$ and every function $e(x)$ in the dyadic bracket for all $x\ge1$, there are $J\ge1$ and X_0 such that for every $X\ge X_0$ and every pair satisfying the size and congruence conditions of §2 at $Y=X^{1/J}$, there is $W\ge Y$ with $W^2\le X$. For all $x\in[W,W^2]$ and all E satisfying $\mathrm{IsChiField}(\psi,e(x),q,d,E)$,

$$\mathrm{VVle}(E,\lfloor x^{401}\rfloor,\exp(x^{\varepsilon}/2)).$$

The pair is fixed before W is selected. The onset and J are uniform in the pair. The cohomology result does not demand the zero boxes or ramp tests, so the pair chosen with those tests in §2 also qualifies. This is $\mathrm{VWinAt}(\psi,401)$. Window quantifiers Use of pair conditions in the window proof

3.1. Exceptional twists and uniformity

Classical. Class field theory describes the unramified abelian pro-ℓ module of the cyclotomic tower by norm limits of class groups. Ferrero–Washington gives $\mu=0$ for the abelian base fields used here. **Release input.** The comparison in the release uses the quotient with the primes above ℓ split and identifies the length at layer m with a twisted cokernel on that module. With cyclotomic generator γ and its cyclotomic value v, the operator is $(v^j\gamma)^{\ell^m}-1$. The release proves generator bounds and a length bound with at most R ℓ-adic centers for a module with at most R generators. At the moving dyadic base R is uniform in q,d once the target is fixed. For a fixed odd ℓ the base is the cyclotomic field of conductor ℓM with $M=\psi.N$. The degree, cyclotomic embedding, and layer comparison of the dyadic base are part of the input. In particular the comparison for a host of index $e\ge2$ uses some $m\le e+2$ and adds at most one to the dyadic length. Cokernel descent Center and layer estimates Uniform rank and host comparison interfaces Ferrero–Washington in the construction

Proved here (finite exceptional set). For an integer cutoff $T\ge1$, put $b=\lfloor\log_\ell T\rfloor$. Take the exceptional integers $1\le j\le T$ to be those strictly closer than ℓ^{-b} to a center. Distinct integers i,j in this range have $v_\ell(i-j)\le b$, so the ultrametric inequality allows at most one integer per center. The center estimate bounds the other lengths by $R(\max(m,\lfloor\log_\ell R\rfloor+3)+c+b)$ with the cyclotomic constant c of the base. In the release one chooses $T=\lfloor X^{401}\rfloor$ and unions the exceptional sets at 2 and the finitely many odd reference primes. The exceptional sets may vary with X,q,d. Their total cardinality is bounded by a single R_{tot} fixed before X,q,d. Exceptional-set estimate Union and order of choices

Release input (small exceptional twists). The pointwise estimate called LemmaK takes a totally imaginary field L containing a primitive ℓ-th root of unity for odd ℓ or a primitive fourth root of unity for $\ell = 2$, with local data as above, and any $j \geq 1$. For some absolute C_K, $H^2(G_{L,S_\ell}, \mathbb{Z}_\ell(j+1))$ is finite and the log of its order is at most

$$\exp\bigl((2j[L:\mathbb{Q}])^{C_K} \log(2|\operatorname{disc} L|)\bigr).$$

The release applies this at a bounded reference layer and uses its second layer comparison for the remaining growth. The discriminant estimate at that layer, the local and global duality estimates, and LemmaK are inputs here. Together the release obtains, on cutoffs satisfying $j \leq x^\eta$ for each exceptional $j \leq x^{401}$, where $\eta = \varepsilon/(2C_K)$, individual length bounds of size $C(1+\log x) + 2\exp(Cx^{3\varepsilon/4})$ (with the host comparison absorbed into constants). Multiplying by at most x^{401} still gives an eventual bound $\exp(x^{4\varepsilon/5})$ at each tower comparison, since $3\varepsilon/4 < 4\varepsilon/5$. The cutoff $j \leq x^\eta$ permits the use of the discriminant bound in LemmaK. Large exceptional twists must be avoided in selecting the interval. Statement of the H^2 estimate Two layer bounds and pointwise target Dyadic application

Proved here (common interval). Take $C_K \geq 1$, $C_t = 401$, and try the $R_{\mathrm{tot}} + 1$ starts

$$W_a = X^{1/[2(2C_t/\eta)^a]}, \qquad a = 0, \dots, R_{\mathrm{tot}}.$$

A start can fail only if the union of exceptional sets meets $(W_a^\eta, W_a^{2C_t}]$. Indeed for $x \in [W_a, W_a^2]$, a forbidden integer has $x^\eta < j \leq x^{C_t}$. The bands are pairwise disjoint since $W_{a+1}^{2C_t} = W_a^\eta$ and the lower endpoint is open. At least one start therefore works on its entire interval. All starts satisfy $W_a^2 \leq X$ and $W_a \geq X^{1/J}$ for $J = 2^{R_{\mathrm{tot}}+2}(C_t/\eta)^{R_{\mathrm{tot}}}$. With $Y = X^{1/J}$, the size bounds on q, d give $\log(qd) \ll Y^{\varepsilon/5}(1+\log^2 Y) \leq Y^{\varepsilon/4}$ eventually. Also $q, d > 2X^{401} + 1$ eventually for fixed J. Thus the growth and size conditions apply uniformly for the one selected interval. General interval lemma Uniform window exponent Pair size hypotheses in the comparison

3.2. What the automorphic representations supply

The phrase “Vandiver-like” in this route refers to the finite-length and height estimates for positive twists. The window proof uses quantitative estimates and allowances for exceptional positive integers. **Release input.** Its pointwise odd-prime plus bound $\mathrm{PointHV}$ has the following form. Fix ψ and the reference layer $m(\ell)$ at each $\ell \in B_\psi$. There are a polynomial exponent C and fixed primes r_1, r_2 such that for all qualifying hosts, $t \geq 2$, and all $B \geq 0$, the bound $\mathrm{VVat}(\ell, t, B)$ on those reference layers gives

$$\mathrm{VVplusLe}\bigl(E, t, (\mathrm{cond}(E)t)^C + CB\bigr).$$

Here the host has the prescribed character, both q, d exceed $r_1, r_2, 2t+1$ and the fixed data in the definition of $\mathrm{Host.Big}$, and E is the CM field of that host. VVat is finiteness and a sum of lengths at one prime. The minus bound $\mathrm{ThmBcFree}$ applies to all these CM composita for $e, t \geq 2$ with the prime-pair conditions. Section 3.4 proves the conductor-one budget and traces its character estimate to the reciprocity and Gauss-sum inputs. Thus the reference-layer bound at a finite set of primes is just one input to a bound over all coefficient primes of the host. Pointwise plus input Host size and conductor conversion Minus input Odd-residue assembly

The pointwise plus bound uses character summands of H^2 over $\mathbb{Q}$. The characters are those of the host field with the parity prescribed by j. The following are **release inputs**. For fixed target, CountI bounds the

weighted total at odd $\ell \notin B_\psi$ with $\ell - 1 \nmid 2j$ by a fixed power of the host conductor parameter times t, under the two fixed-prime exclusions above. $\mathrm{CountII}$ and $\mathrm{CountIII}$ give individual polynomial bounds in $\ell N j$, with N the host modulus, for the case $\mathrm{EV}(\ell, \rho)$ at $j \geq 2$ under the case-II divisibility conditions, and for nontrivial odd-order ρ outside B_ψ with $\ell - 1 \mid 2j$, respectively. The EV condition requires even order of $\chi_P(\ell, \rho)$, the factor of the prime-to-ℓ part left after removing its ℓ-conductor component. $\mathrm{CountIV}$ treats odd-order prime-to-ℓ part at $\ell \in B_\psi$ using the reference layer. CountV treats $j = 1$ at the reference primes with $R_1(\ell, \rho) \leq 2^{c_2(\rho)}$, and $\mathrm{CountVs}$ treats the complementary bound using duality and LemmaK. The trivial-character outside case, with $\ell - 1 \mid 2j$, uses CountZ from the release's vanishing input PropZ. These conditions concern characters of the host, which continue to vary even for target $\mathbf{1}$. Character cases Bounds and hypotheses Count IV Assembly into the pointwise estimate

In the exported route, the just-mentioned bounds are obtained through the following representation inputs. **Release input (automorphic and Galois construction).** The integral Galois row and isobaric base change feed the package $\mathrm{LiteratureRankNHypothesis}$. The package uses local GL_n segment classification and irreducibility, depth-zero induction, local converse, regular algebraic purity, Galois compatibility and CM coefficient and period comparisons. The patching step produces $\mathrm{HypothesesH}'$, using $\mathrm{descWeak}$ for the weak descent input. From these hypotheses the proof constructs unramified automorphy ends and a reduced-unitary end, and uses the Kato two-place input under the unramified hypotheses for CountI. The $\mathrm{Sp4In0}$ comparison gives $\mathrm{CountII}$ and $\mathrm{CountIII}$, while the reduced unitary lifting bound gives CountV. The same package feeds $\mathrm{lemmaK_litN_alt}$. These assertions are required in the applicable conductors, weights, and local types of the count statements. Classical base change and local representation theorems occur in the package, but their names alone do not supply the patching, descent, Kato or quantitative conclusions. Local, automorphic and CM interfaces Patching, the two ends and the four counts Duality, towers, odd minus estimate and assembly

Proved here (size of the combined budget given the inputs). With $B = \exp(x^{4\varepsilon/5})$, the dyadic contribution and reference-prime allowance have this exponential scale. The fixed-conductor estimate $\mathrm{cond}(E) \mid \mathrm{lcm}(2^{e+2}, M)qd$ and $t \leq x^{401}$ give $\log(\mathrm{cond}(E)t) = O(x^{\varepsilon/4} + \log x)$ on the selected interval. Thus the polynomial term of $\mathrm{PointHV}$ has log $O(x^{\varepsilon/4} + \log x)$. The release's minus-bound comparison gives an allowance $\exp(K' x^{\varepsilon/4})$ under these sizes for fixed target and K' independent of x, q, d. Each of these allowances and their sum are $o(\exp(x^{\varepsilon}/2))$. All terms are nonnegative, so the sum bounds every finite set of primes with the required finiteness. This explains the slack in the output VVle, and the smaller twist limit $\lfloor x^{400} \rfloor$ survives rescaling as in §2. Conductor, height and sum of parts Rescaled limit

3.3. Conductor-one tower checks

For $\mathbf{1}$, $p_K = 7$, $h(\mathbb{Q}(\sqrt{-7})) = 1$, $p_1 = 3$, and $B_{\mathbf{1}} = \{3, 7\}$. The two odd reference bases are $\mathbb{Q}(\zeta_3)$ and $\mathbb{Q}(\zeta_7)$. **Proved here using classical Minkowski and Chevalley.** The first base is the Eisenstein field. For the second the Minkowski bound is less than 5, while the residue degrees of $2, 3$ are $3, 6$. This gives no prime ideal of norm at most the bound. The same bound for $\mathbb{Q}(\sqrt{-7})$ is less than 2. At each degree-ℓ step of either odd tower, exactly one prime ramifies and is totally ramified. Chevalley's ambiguous class number formula gives that the fixed class group has order prime to ℓ if this holds at the preceding layer. A nontrivial finite ℓ-group of classes under a cyclic action of order ℓ has at least ℓ fixed elements, by counting orbits. Induction gives zero ℓ-primary groups throughout these reference towers. The reference bounds can

therefore use zero lengths. The plus bound for the host still has its separate polynomial term. Definition of the reference primes Reference fields and pointwise host estimate Chevalley's formula

Proved here (character cases at conductor one). The host Galois group in §2 is a 2-group. Every character in the odd-prime counts therefore has 2-power order, and case III, which asks for a nontrivial character of odd order, is empty. At $\ell \in \{3, 7\}$ and on sufficiently late windows, ℓ is coprime to the host modulus $QS\,qd$. There $\chi_P(\ell, \rho) = \rho$, so case IV only uses the trivial character. The even-order condition of case II can still hold, for instance for a nontrivial character at an odd prime coprime to the host modulus. Section 3.4 handles the minus eigenspace. Character factors Host modulus and cases Minus bound

Proposition (dyadic bound for the favorable pair). Let q, d be the pair of Paper II, Proposition 2, so $q \equiv 3 \bmod 8$, $d \equiv 7 \bmod 16$, and $v_2 h(-4qd) = v_2 h(-8qd) = 3$. Here $h(D)$ is the quadratic class number of discriminant D. For the host $E = K^{+}_{e+2}(\sqrt{-q}, \sqrt{-d})$ of degree $4 \cdot 2^e$, every $e \geq 0$ and $j \geq 1$ satisfy

$$\lambda_2(E, j) \leq 3e + 4 + v_2(j).$$

Paper II, Proposition 2 · Full dyadic proof

Write $S_a = K_a(\sqrt{qd})$, $L_a = K_a(\sqrt{-q}, \sqrt{-d})$, and X_S, X_L for their ordinary 2-class groups in the inverse limit under norms. The base degrees are $[S_2 : \mathbb{Q}] = 4$ and $[L_2 : \mathbb{Q}] = 8$. At finite layers Chevalley's formula bounds the 2-rank of S_a by two. The base K_a has odd class number, and for $a \geq 4$ the six primes ramified in S_a/K_a come in three conjugate pairs. The unit $1 + \zeta_{2^a} + \zeta_{2^a}^{-1}$, a conjugate interchanging the pairs above d, and ζ_{2^a} give independent residue sign vectors $(1,1,0), (1,0,1), (1,1,1)$ on those pairs. The same rank bound follows at $a = 2, 3$ from the smaller ramification count and, at $a = 3$, the nonsquare residue of ζ_8 above q. The unramified quadratic extension L_a/S_a gives the bound three for L_a by Schreier's inequality on the unramified pro-2 groups. These generator bounds pass to the norm limits. Chems-Eddin and Müller, Theorems 4 and 5, give $\lambda^{-}(S_2) = 2$ and odd class numbers for the real tower. Ferrero–Washington and the Iwasawa structure theorem then give rank two for X_S.

For $k = \mathbb{Q}(\sqrt{-d})$ the two local unit directions at 2 give a $\mathbb{Z}_2^2$-extension unramified outside 2. Each inertia direction maps injectively to the cyclotomic direction, so over the cyclotomic $\mathbb{Z}_2$-extension of k the remaining $\mathbb{Z}_2$-extension is unramified. Base change gives a rank-one quotient of X_L with trivial cyclotomic action. The involution of L_∞/S_∞ exchanges the local unit directions of k and acts by -1 on this quotient. Norm and extension of ideals identify the $+1$ rational subspace of X_L with $X_S \otimes \mathbb{Q}_2$, so this subspace is the rational kernel of the quotient. Thus the integral generator bounds are attained by the free ranks, and $X_S \simeq \mathbb{Z}_2^2$, $X_L \simeq \mathbb{Z}_2^3$. Finite-layer ranks and the norm maps · Chems-Eddin and Müller, Theorems 4 and 5 · Ferrero–Washington

Take the cyclotomic generator γ of value 5, and let M be its action on X_S. The unique prime above 2 in S_a is totally ramified up the tower. Killing its inertia after abelianizing over S_a gives the ordinary 2-class group $X_S/(M^{2^{a-2}} - 1)X_S$. In layers 2, 3, Hasse unit index two would give a unit u with $u/\bar{u} = -\zeta_{2^a}$ after multiplication by a root of unity. Then $u/(1 - \zeta_{2^a})$ would be real with odd dyadic valuation, whereas real elements have even valuation in S_a. The Hasse indices are therefore one. The relative class number formula gives $h^{-}(S_2) = h(-4qd)/2$ and $h^{-}(S_3) = h(-4qd)h(-8qd)/4$. The real class numbers are odd. The two coinvariant orders give the valuation of $\det(M - I)$ and the valuation of $\det(M^2 - I)$, so the favorable valuations give

$$v_2 \det(M - I) = v_2 \det(M + I) = 2.$$

Put $s = \operatorname{tr}(M - I), p = \det(M - I)$. From $\det(M + I) = 4 + 2s + p$ both $s/2$ and $p/4$ are odd. The characteristic polynomial of $(M - I)/2$ is therefore integral and reduces to $Z^2 + Z + 1$. Each eigenvalue of M has the form $1 + 2z$ with residue $\bar{z} \in \mathbb{F}_4 \setminus \mathbb{F}_2$. For $b = 5^j(1 + 2z)$, the two quantities $b - 1, b + 1$ both have valuation one, since $5^j \equiv 1 \bmod 4$. The first squaring gives $v_2(b^2 - 1) = 2$, and subsequent squarings add one each. Thus the two eigenvalues contribute $2(e + 1)$ to the valuation of the determinant of $T_{j,e} = (5^j\gamma)^{2^e} - 1$. The trivial cyclotomic direction contributes $v_2(5^{j2^e} - 1) = e + 2 + v_2(j)$. On the free lattice X_L the sum is the cokernel length $3e + 4 + v_2(j)$. Low-layer class numbers and determinant calculation

For localization, each of E and $L_{e+2} = E(i)$ has two dyadic primes. Locally $\sqrt{-q}$ generates the unramified quadratic extension and $\sqrt{-d}$ is already in $\mathbb{Q}_2$. Each dyadic prime of E has a unique prime above it, totally ramified of degree two, in L_{e+2}. At either prime the local cyclotomic image equals the global image, namely $\{\pm 1\}(1 + 2^{e+2}\mathbb{Z}_2)$ for E and $1 + 2^{e+2}\mathbb{Z}_2$ for L_{e+2}. At every finite step L_b/L_a, $b \geq a \geq 2$, the index 2^{b-a} is also the local degree and the ramification index at each of the two primes. The local and global cyclotomic quotients above L_{e+2} thus identify with the same procyclic group generated topologically by γ^{2^e}. In degree one its cohomology is zero, since multiplication by $5^{-j2^e} - 1$ on the divisible coefficients is onto. Its degree-two cohomology is zero as well. Inflation–restriction therefore identifies the strict group on L_{e+2} with invariant characters of the quotient X_P of X_L that imposes splitting at 2. Odd inertia already lies in the infinite-level group, and at 2 the local degree-one vanishing makes the split condition equivalent to the original local vanishing. Dualizing gives $X_P/T_{j,e}X_P$, an integral quotient of the ordinary cokernel, including any finite torsion of X_P.

For L_{e+2}/E the global and each dyadic local quotient have order two and the same fixed coefficient group $B = (\mathbb{Q}_2/\mathbb{Z}_2)[2^{e+2+v_2(j)}]$, with involution acting by $(-1)^j$. The kernel of restriction before localization is $H^1(C_2, B)$, of order two. Its localization to either prime is the same inflation map and is injective. Thus restriction on the strict group has zero kernel. For a strict invariant class upstairs the global transgression vanishes by comparison with either local transgression. Its lift downstairs has two local obstructions in $H^1(C_2, B)$, defined up to simultaneous translation. Their difference gives the cokernel into the invariants as a subgroup of $H^1(C_2, B)$. This proves the bound with layer $m = e$ and removes the $+1$ allowance in the general restriction estimate of §3.1. Exact localization comparison · Quadratic restriction sequence

The sum for $1 \leq j < t$ is at most $(t - 1)(3e + 4) + v_2((t - 1)!)$, hence $O(t(1 + \log x))$ when $e = O(\log x)$. The favorable choice bypasses the exceptional-twist bound at 2. With $Y = X^{1/J}$, $J \geq 2$ fixed, the same analytically admissible pair supplies the dyadic and odd reference-layer inputs on $[\sqrt{X}, X]$ at all sufficiently large X. The pointwise odd-prime estimates for the host supply the remaining terms in the budget. Prime selection with the prescribed sign

3.3a. Vanishing in the trivial and quadratic residual cases

One count in §3.2 needs an exact zero for the trivial host character. It has an elementary class-field proof. Fix a prime $\ell \geq 5$ and an even integer $j \geq 2$ with $\ell - 1 \mid 2j$. Let O be the integers in a finite extension $K/\mathbb{Q}_\ell$ and let $A = K/O$, with action $\alpha = \kappa_\ell^{-j}$ of $G_\mathbb{Q}$, where κ_ℓ is the cyclotomic character. Let S consist of the continuous cohomology classes in $H^1(\mathbb{Q}, A)$ whose restriction to each inertia group away from ℓ

and to the decomposition group at ℓ is zero. Then $S = 0$. In the release this is PropZ for every allowed coefficient ring. Definition of the strict group · Statement of the vanishing

Proof. The image of each continuous cocycle in the discrete group A is finite because $G_{\mathbb{Q}}$ is compact. Every cohomology class is consequently killed by a power of ℓ. It suffices to show that every $s \in S[\ell]$ is zero. Take a cocycle z for s. Write $\ell z = \partial a$ and choose a' with $\ell a' = a$, using divisibility of A. Replacing z by $z - \partial a'$ puts all its values in $A[\ell]$. A coboundary is $\partial b(g) = (\alpha(g) - 1)b$. For inertia away from ℓ the full cyclotomic action is trivial, so the strict restriction forces z itself to vanish there.

If $\ell - 1 \mid j$, the action on $A[\ell]$ is trivial, so z is a continuous homomorphism. Its image cuts out a finite abelian ℓ-extension of $\mathbb{Q}$ unramified outside ℓ. By Kronecker–Weber it factors through $\mathbb{Q}(\mu_{\ell^\infty})$. Indeed the conductor has no prime divisor away from ℓ; the extension has odd degree and hence has no ramification at infinity. The action on A factors through the same cyclotomic extension, which is totally ramified at ℓ. Its full Galois group is therefore the image of the decomposition group at ℓ. On that decomposition group $z = \partial b$ for some $b \in A$ by the strict local condition. Both sides factor through the same cyclotomic quotient, so the equality holds globally. Trivial residual character in the release · Class field theory

Otherwise $j \equiv (\ell - 1)/2 \pmod{\ell - 1}$, and evenness of j gives $\ell \equiv 1 \pmod 4$. On the additive ℓ-torsion $A[\ell]$, the action is the quadratic character associated to $k = \mathbb{Q}(\sqrt{\ell})$ with discriminant ℓ. In fact the action factors through $(\mathbb{Z}/\ell)^\times$ and is its unique character of order two. This remains valid for ramified O, since an integer divisible by ℓ acts as zero on $A[\ell]$. At every place above ℓ a decomposition group surjects onto $(\mathbb{Z}/\ell)^\times$, so contains h with $\alpha(h) - 1$ an O-unit. Write the local cocycle as $z = \partial b$ in A. Evaluating at h shows $b \in A[\ell]$, since $\ell z(h) = 0$ and $\alpha(h) - 1$ is a unit. It follows that on restriction to G_k the local cocycle vanishes at each such place. The global restriction $z|_{G_k}$ is a homomorphism to $A[\ell]$ unramified at every finite place. It therefore factors through the class group of k on the odd-primary part.

By Minkowski's theorem every ideal class of this real quadratic field contains an integral ideal of norm at most $\sqrt{\ell}$. For each n the number of ideals of norm n is at most the divisor function $d(n)$, by prime decomposition in a quadratic field. With $v = \sqrt{\ell} > 1$, we have

$$h(k) \le \sum_{n \le v} d(n) \le v(1 + \log v) < v^2 = \ell.$$

The narrow and ordinary class numbers differ only by a power of two. Hence $z|_{G_k} = 0$. Now z descends to a cocycle on the order-two group $\mathrm{Gal}(k/\mathbb{Q})$ with values in $A[\ell]$. Since multiplication by two is invertible on $A[\ell]$, averaging on this finite group makes that cocycle a coboundary. This finishes the proof. Quadratic class-number estimate · Quadratic case and its use

3.4. The conductor-one odd minus bound

Proposition (odd minus bound for the ordinary host). Let $e, t \ge 2$ be integers and let q, d be primes with $q \equiv 3 \bmod 8$ and $d \equiv 7 \bmod 16$. Let $E = \mathbb{Q}(\zeta_{2^{e+2}})^+(\sqrt{-q}, \sqrt{-d})$, $N_E = 2^{e+2}qd$ and $D = [E : \mathbb{Q}] = 2^{e+2}$. The conductor calculation is Paper II, Proposition 1. It only uses the coprime moduli and the indicated discriminants, so the bound below covers all of these pairs, including the ones with $q \mid d + 1$ used in the final host construction. For every odd ℓ and $1 \le j < t$, the minus length $\lambda_\ell^-(E, j)$ is finite. For every finite set P of odd primes,

$$\sum_{\ell\in P}\log\ell\sum_{1\le j<t}\lambda_\ell^-(E,j)\le Dt^2\log(tN_E). \tag{VI.1}$$

Both assertions hold for every permitted choice of algebraic closure and local data. For fixed E,t only finitely many odd primes have a nonzero term. In particular this gives the budget in the conductor-one specialization of $\mathrm{ThmBcFree}$. Paper II, Proposition 1 Exact interface Budget and domain General statement

3.4.1. Complex conjugation, coefficient signs and descent

Classical alternative for ordinary zeta. For the conductor-one field above, Poitou–Tate duality, Greither–Popescu's equivariant Coates–Sinnott theorem and Ferrero–Washington give an unconditional proof of (VI.1). This applies for all the indicated e,t,q,d and every odd coefficient prime, including q,d. Section 3.4.3 records the released derivation and the inputs it uses.

Let c be complex conjugation on E, and fix a complex-conjugation lift $\tilde c$ to an algebraic closure $\Omega\subset\mathbb{C}$. On ordinary conductor-one Artin coordinates of Paper II it has bits $(1,1)$ and trivial dyadic coordinate. The inertia involution at d, with bits $(1,0)$ in Paper II's order, acts on just $\sqrt{-d}$. The cohomological projection uses c. On cocycles write $C_{\mathrm{arg}}z(g)=z(\tilde c^{-1}g\tilde c)$ and $C_{\mathrm{tw}}z(g)=\kappa_\ell(\tilde c)^{-j}z(\tilde c^{-1}g\tilde c)$, where κ_ℓ is the cyclotomic character. These are the argument and coefficient-twisted involutions. Conjugation preserves the Galois quotient unramified outside ℓ and permutes the local conditions. It is compatible with the cocycle relation because κ_ℓ is invariant under conjugation. Since $\tilde c$ inverts roots of unity, $C_{\mathrm{tw}}=(-1)^jC_{\mathrm{arg}}$. Thus $Ш_\ell^-(E,j)$ is the -1 eigenspace for C_{arg} and the $(-1)^{j+1}$ eigenspace for C_{tw}. The plus signs are $+1$ and $(-1)^j$, respectively. At odd ℓ the two projections exist over $\mathbb{Z}_\ell$ by division by 2. Write $\lambda_\ell^-(E,j)=\mathrm{length}_{\mathbb{Z}_\ell}Ш_\ell^-(E,j)$. Artin coordinates Conjugation on cocycles Argument sign modulo a coboundary Eigenspaces and addition of lengths

Let $G=\mathrm{Gal}(E/\mathbb{Q})$ and choose the integers O of a finite unramified extension of $\mathbb{Q}_\ell$ containing the values of every character of G. Such an O exists because $|G|$ is prime to ℓ and roots of unity of order prime to ℓ occur in unramified extensions. Let $K_O=\mathrm{Frac}(O)$. Flat base change commutes with the localization kernel and preserves length with this choice of O. Indeed O is free of finite rank over $\mathbb{Z}_\ell$, so continuous cochains commute with base change by using a finite basis, and $O/\ell O$ has O-length one. In the release's general choice of O, lengths multiply by $v_O(\ell)$. Since $|G|$ is a unit, the projectors $|G|^{-1}\sum_{g\in G}\psi(g)^{-1}g$ split the scalar-extended module into character summands. Precisely those with $\psi(c)=(-1)^{j+1}$ occur in the minus part. The resulting full sum has $D/2$ characters, because $c\neq1$ and evaluation at c is onto $\{\pm1\}$ on the character group. The release bounds this count by D. Scalar extension and eigenspace comparison Character family at one prime

Proved here using classical restriction and corestriction. Inflate each ψ to $G_{\mathbb{Q}}$, and write $\chi=\psi\circ\rho$ on the units modulo N_E, where ρ is the cyclotomic quotient onto G. The ψ-summand equals the strict Selmer group $S_\ell(\alpha)$ over $\mathbb{Q}$ of K_O/O with character

$$\alpha=\psi^{-1}\kappa_\ell^{-j}.$$

Here “strict” means zero restriction to inertia at every finite prime and to the full decomposition group at ℓ. The exponent $-j$ is the original coefficient twist, and the factor ψ^{-1} cancels the ψ-action when taking invariants. The positive-degree cohomology of G on ℓ-primary modules vanishes by averaging. Inflation–

restriction identifies the invariant classes after this twist with the rational classes. The local restrictions to the inertia groups and the decomposition group above ℓ detect the required vanishing downstairs: their subgroups over E have indices dividing D, and restriction is injective in degree one because corestriction multiplies by that index. This proves the descent with the indicated local conditions. Its Tate dual character is $\beta = \kappa_\ell \alpha^{-1} = \psi \kappa_\ell^{j+1}$. In particular $\alpha(\tilde{c}) = -1$, while $\beta(\tilde{c}) = 1$. The Bernoulli index in the bound is $k = j + 1$ and the Bernoulli character is χ, with the same parity as k. Strict local family Character twists in the release Eigenspace descent Cyclotomic and Dirichlet character identification

Strict duality. Let $S_{\ell,E}$ consist of the places of E over ℓ and its archimedean places. Put $A = (\mathbb{Q}_\ell/\mathbb{Z}_\ell)(-j)$ and $T = \mathbb{Z}_\ell(k)$ with $k = j + 1$. The finite Tate dual of $A[\ell^m]$ is $T/\ell^m T$. Finite-coefficient Poitou–Tate pairs the degree-one localization kernel for $A[\ell^m]$ with the degree-two kernel for $T/\ell^m T$. The archimedean terms in these degrees vanish. Filtered colimits commute with the finite localization diagram, and the finite-coefficient global and local cohomology groups are finite, so the inverse systems satisfy Mittag–Leffler. Taking limits gives the G-equivariant identification

$$Ш_\ell(E, j)^\vee \simeq Ш^2_{S_{\ell,E}}(E, T) = \ker\left(H^2(G_{E,S_{\ell,E}}, T) \longrightarrow \prod_{w \mid \ell} H^2(E_w, T) \right),$$

where the Pontryagin dual has the contragredient action. The identification is natural under transport of closures and localization data. In particular the character ψ on the strict group corresponds to $\eta = \psi^{-1}$ on its dual, with $\eta(c) = \psi(c) = (-1)^k$ for the minus part. Milne, *Arithmetic Duality Theorems*, I.4.10(a) and I.5.2(b)

Put $U_0 = \operatorname{Spec} \mathcal{O}_E[1/\ell]$. The comparison map $H^2(G_{E,S_{\ell,E}}, T) \to H^2_{\text{ét}}(U_0, T)$ is injective. For $T/\ell^m T$ this is the degree-two edge injection of the Cartan–Leray spectral sequence over finite étale covers of U_0. The direct limit of their degree-one cohomology vanishes, since each class is a torsor for a finite locally constant sheaf and trivializes on a finite étale cover. The étale fundamental group is $G_{E,S_{\ell,E}}$ by the unramified extension correspondence. The finite-coefficient groups in degree one on both sides are finite, so passage to ℓ-adic coefficients gives the injection by inverse limits. We therefore have

$$\operatorname{length}_O S_\ell(\psi^{-1}\kappa_\ell^{-j}) \le \operatorname{length}_O \big(H^2_{\text{ét}}(U_0, T) \otimes_{\mathbb{Z}_\ell} O\big)_{\psi^{-1}}.$$

Here subscripts denote character components. The ψ-part of the strict group and the ψ^{-1}-part of its dual have the same length once finite. Finite locally constant sheaves and the fundamental group · Geisser–Schmidt, *Poitou–Tate duality for arithmetic schemes*, Theorem B

3.4.2. A single integer controls the weighted sum

Put $\Pi_k = \prod_{p \le k+1} p$, with p prime. For a character χ of G pulled back to the units modulo N_E, write χ^* for its primitive character of conductor $f_\chi \mid N_E$, extended by zero at nonunits. For $k = j + 1$ define

$$B_{k,\chi^*} = f_\chi^{k-1} \sum_{u=1}^{f_\chi} \chi^*(u) B_k(u/f_\chi), \qquad A_{N_E,k}(\chi) = N_E \Pi_k B_{k,\chi^*}, \qquad b(N_E, k) = 16 N_E^{k+1} k!.$$

The common character estimate is

$$\operatorname{length}_O S_\ell(\psi^{-1}\kappa_\ell^{-j}) \le v_O(A_{N_E,k}(\chi)). \tag{VI.2}$$

Here v_O counts powers of a uniformizer. In our unramified O, v_O agrees with the valuation v_ℓ on the chosen embedding into $\overline{\mathbb{Q}}_\ell$, normalized by $v_\ell(\ell) = 1$. The release calls the Bernoulli factor alN. The construction of $\mathrm{ThmBcFree}$ takes the finite weighted sum from $\mathrm{PrimeBound}$, which uses $\mathrm{PsiBound}$ at each character. The entry $\mathrm{ML.iv}$ applies vChar with the reciprocity input. Normalizations Character estimate Length to norm normalization Construction of the two wrappers

Coates–Sinnott and the torsion correction. Apply Greither–Popescu, *An Equivariant Main Conjecture in Iwasawa Theory and Applications*, J. Algebraic Geom. **24** (2015), 629–692, Theorem 6.11, to $E/\mathbb{Q}$, $S = \{\infty, 2, q, d, \ell\}$, $p = \ell > 2$ and $n = k \geq 2$. The set S contains the ramified, infinite and ℓ-adic places. The field $E(\mu_\ell)$ is abelian over $\mathbb{Q}$ and is already CM, so Ferrero–Washington gives the required $\mu_{E(\mu_\ell)^{\mathrm{CM}},\ell} = 0$ for every odd ℓ. In particular the theorem applies when $\ell \mid N_E$. Write $U = \operatorname{Spec} \mathcal{O}_E[1/(N_E\ell)]$, $W = H^0(E, (\mathbb{Q}_\ell/\mathbb{Z}_\ell)(k))$ and $e_k = (1 + (-1)^k c)/2$. Lemma 6.9(1) identifies W with $H^1_{\mathrm{et}}(U, T)_{\mathrm{tors}}$, and $H^2_{\mathrm{et}}(U, T)$ is finite by §6.2 of the same paper. The theorem gives

$$e_k \operatorname{Fitt}_{\mathbb{Z}_\ell[G],0} H^2_{\mathrm{et}}(U, T) = \operatorname{Ann}_{\mathbb{Z}_\ell[G]}(W)\, \Theta_{S,E/\mathbb{Q}}(1-k).$$

In their convention $\Theta_S(1-k)$ evaluated at η is $L_S(1-k, \eta^{-1})$. Thus $\eta = \psi^{-1}$ from duality gives $L_S(1-k, \psi)$. Greither–Popescu, definition of Θ_S · Lemma 6.9 and Theorem 6.11 · Published paper · Ferrero–Washington, Ann. of Math. **109** (1979), 377–395

The invariant group W is cyclic of finite order at most $\ell^{v_\ell(kN_E\Pi_k)}$. To see this, set $s = v_\ell(N_E)$. Since $E \subseteq \mathbb{Q}(\mu_{N_E})$, the image of G_E on the ℓ-power roots of unity contains $1 + \ell^{\max(1,s)}\mathbb{Z}_\ell$. Any element of this subgroup lifts to a cyclotomic automorphism with other components 1 that fixes $\mathbb{Q}(\mu_{N_E})$. The element $1 + \ell^{\max(1,s)}$ acts on the twist by its k-th power. For odd ℓ the binomial formula gives $v_\ell((1+\ell^{\max(1,s)})^k - 1) = \max(1, s) + v_\ell(k)$. Thus W is contained in the cyclic torsion layer of that exponent, which proves the bound for $s \geq 1$. If $s = 0$, the same lifting argument gives the full image $\mathbb{Z}_\ell^\times$. If $\ell - 1 \nmid k$, a lift of a generator modulo ℓ has k-th power differing from 1 by a unit, so $W = 0$. Otherwise $\ell \leq k + 1$, and Π_k supplies the extra power of ℓ.

Fitting ideals commute with base change to each factor O of $O[G]$. The annihilator of W does as well, by flat base change to the splitting ring followed by its direct-factor projections. In the η-component it is $(\ell^{a_{\ell,\eta}})$, where $a_{\ell,\eta} = \operatorname{length}_O(W \otimes_{\mathbb{Z}_\ell} O)_\eta \leq v_\ell(kN_E\Pi_k)$, since each component is cyclic. For a finite module over O, its zeroth Fitting ideal is generated by ℓ to the power of its length, as follows by decomposing into cyclic modules. For $\eta = \psi^{-1}$ with $\eta(c) = (-1)^k$, we obtain

$$\operatorname{length}_O\big(H^2_{\mathrm{et}}(U, T) \otimes_{\mathbb{Z}_\ell} O\big)_\eta = a_{\ell,\eta} + v_\ell\big(L_S(1-k, \psi)\big).$$

The removed Euler factors. For $f = f_\chi$, the function $L_S(1-k, \psi)$ is the primitive $L(1-k, \chi^*)$ multiplied by $1 - \chi^*(r) r^{k-1}$ at every finite $r \in S$ with $r \nmid f$. Each such factor is nonzero since $|\chi^*(r) r^{k-1}| > 1$. At $r \mid f$ the primitive Euler factor is already 1. At $r = \ell$ the correcting factor, if present, is a unit since $k \geq 2$. For the other removed primes, étale localization gives an exact sequence of G-modules

$$0 \longrightarrow H^2_{\mathrm{et}}(U_0, T) \longrightarrow H^2_{\mathrm{et}}(U, T) \longrightarrow \bigoplus_{\substack{r \mid N_E \\ r \neq \ell}} \bigoplus_{w \mid r} H^1(\kappa(w), \mathbb{Z}_\ell(k-1)) \longrightarrow 0.$$

Indeed, U_0 is regular, the complement of U is a disjoint union of regular closed points of codimension one, and ℓ is invertible on U_0. Gabber's absolute purity theorem identifies the supported term at w with residue-

field cohomology shifted by two degrees and twisted down by one, first for $T/\ell^m T$. All the finite-coefficient groups in this localization sequence are finite, so inverse limits are exact here and compute continuous ℓ-adic cohomology. At the left, $H^0(\kappa(w), \mathbb{Z}_\ell(k-1)) = 0$ since arithmetic Frobenius acts by $|\kappa(w)|^{k-1} \neq 1$. At the right, $H^3_{\mathrm{et}}(U_0, T) = 0$ by the relative-dimension-zero case of Geisser–Schmidt (0.4) for U_0 over itself, with the set of places $S_{\ell,E}$ and each coefficient $T/\ell^m T$. All infinite places are complex and their degree-three groups vanish. Finiteness of the degree-two groups removes the $\varprojlim^1$ term when passing to T. Fujiwara, *A Proof of the Absolute Purity Conjecture (after Gabber)*, Adv. Stud. Pure Math. **36** (2002), 153–183 · Geisser–Schmidt, Theorem B and (0.4)

For $w \mid r$, write $D_w, I_w \subset G$ for the decomposition and inertia groups, $F_r \in D_w/I_w$ for arithmetic Frobenius and $f_w = [\kappa(w) : \mathbb{F}_r]$. Procyclic cohomology identifies the residue term at w with $\mathbb{Z}_\ell/(r^{f_w(k-1)} - 1)$. Inertia acts trivially and F_r acts on this quotient by r^{k-1}, from its action on the twist and the trivial conjugation action on the absolute Galois group of the residue field. The sum over $w \mid r$ is induced from D_w. The η-component vanishes if $\eta|_{I_w} \neq 1$, which is precisely $r \mid f$. Otherwise projection imposes $r^{k-1} = \eta(F_r)$, giving $O/(r^{k-1} - \eta(F_r))$ since $\eta(F_r)^{f_w} = 1$. Its length is $v_\ell(1 - \chi^*(r) r^{k-1})$, as $\chi^*(r) = \eta(F_r)^{-1}$. This accounts for the factors at $2, q, d$ even when the ambient field ramifies and the character does not. If $\ell = q$ or d, that prime is already inverted in U_0 and its Euler factor was dealt with by the unit calculation at ℓ.

Subtracting the residue lengths in the exact sequence gives

$$\operatorname{length}_O \big(H^2_{\mathrm{et}}(U_0, T) \otimes_{\mathbb{Z}_\ell} O \big)_{\psi^{-1}} = a_{\ell,\psi^{-1}} + v_\ell(L(1-k, \chi^*)).$$

The primitive value is $L(1-k, \chi^*) = -B_{k,\chi^*}/k$ by the Hurwitz-zeta formula, including the zeta case $f = 1$. In the matching parity $\chi^*(-1) = (-1)^k$, the functional equation relates it by a nonzero factor to $L(k, \overline{\chi^*})$, which is nonzero by the Euler product. The torsion bound pays for the division by k, giving $a_{\ell,\psi^{-1}} + v_\ell(L(1-k,\chi^*)) \leq v_\ell(N_E \Pi_k B_{k,\chi^*})$. This proves (VI.2) by §3.4.1. The primitive L-value can have negative valuation, with the torsion correction supplying the integral Fitting generator. Hurwitz expansion and Dirichlet functional equation · Hurwitz special value

Proved here using classical Bernoulli identities. Each $A_{N_E,k}(\chi)$ is integral. Expand $f_\chi^{k-1} B_k(u/f_\chi)$ as $\sum_{i=0}^k \binom{k}{i} B_i f_\chi^{i-1} u^{k-i}$. The factor N_E clears the $i = 0$ denominator, and Π_k clears all denominators of B_i by the von Staudt–Clausen theorem, $B_1 = -1/2$, and vanishing of the odd B_i for $i > 1$. The values of the character are integral as roots of unity or zero. The preceding nonvanishing argument applies under every embedding, since it preserves the parity. Von Staudt–Clausen Release denominator calculation Release nonvanishing

The classical Fourier series of $B_k(u)$ for $k \geq 2$ gives $\sup_{0 \leq u \leq 1} |B_k(u)| \leq 2k!\zeta(k)/(2\pi)^k \leq 4k!/(2\pi)^k$. Therefore $|B_{k,\chi^*}| \leq 4 f_\chi^k k!/(2\pi)^k$. The elementary inequality $\prod_{p \leq m} p \leq 4^m$ follows by induction. For odd $m = 2u+1 \geq 3$, the primes in $(u+1, 2u+1]$ divide $\binom{2u+1}{u} \leq 2^{2u}$, the latter bound using the two equal middle coefficients. At even $m \geq 4$ the product is the previous odd product. Thus $N_E \Pi_k |B_{k,\chi^*}| \leq b(N_E, k)$, using $4/(2\pi) < 1$. The same bound holds under every complex embedding, since it depends only on $f_\chi \leq N_E$ and k. In particular the absolute norm from $\mathbb{Q}(\chi^*)$ is a positive integer bounded by $b(N_E,k)^{[\mathbb{Q}(\chi^*):\mathbb{Q}]}$, uniformly in ℓ. Fourier series of Bernoulli polynomials Archimedean bound in the release

Take the product $M_j = \prod_{\chi(c)=(-1)^{j+1}} A_{N_E,j+1}(\chi)$, where $\chi(c)$ denotes the value of the corresponding character of G. Construct it in a cyclotomic field containing all character values. Galois automorphisms permute these characters because they preserve the rational parity sign, and preserve the primitive conductor and Bernoulli formula under conjugation. Thus M_j is a rational algebraic integer, hence a nonzero integer. The product computed in any $\overline{\mathbb{Q}}_\ell$ is its image there. Add (VI.2) over the character summands to obtain

$$\lambda_\ell^-(E,j) \le v_\ell(M_j), \qquad \sum_{\ell\in P} \log \ell \, \lambda_\ell^-(E,j) \le \log|M_j| \le \tfrac{D}{2} \log b(N_E, j+1). \qquad \text{(VI.3)}$$

Here integrality gives nonnegative valuations for all factors, and the first inequality establishes finiteness before any weighting. Unique factorization gives the middle inequality and shows that only primes dividing M_j can contribute. One may take P to contain all odd prime divisors of $\prod_{j=1}^{t-1} M_j$. Rational integer product Valuation comparison and prime sum Release prime bound

For $2 \le k \le t$, use $16 \le tN_E$ and $k! \le t^{k+1}$ to bound $\log b(N_E,k)$ by $(k+2)\log(tN_E)$. Interchanging the finite sums in (VI.3) gives the unconditional bound

$$\sum_{\ell\in P} \log \ell \sum_{1\le j<t} \lambda_\ell^-(E,j) \le \frac{D}{4}(t-1)(t+6)\log(tN_E) \le Dt^2 \log(tN_E),$$

as asserted in (VI.1). Every odd ℓ is prime to D here. For a general target the release also bounds the odd primes dividing D by its separate $\mathrm{StickBound}$. Exact inequalities in the release Complementary degree-dividing terms

3.4.3. From reciprocity to the character bound

The released derivation of (VI.2) uses the Gauss-sum and Kolyvagin derivative system stated below, together with the reciprocity and length-descent steps. The dependence on this system is specific to that derivation.

For the strict group in (VI.2) let $f = f_\chi$ and use the primitive character χ^*. Inflating it gives the same ψ on $G_{\mathbb{Q}}$. In particular $\alpha = \psi^{-1}\kappa_\ell^{-j}$ and all its local conditions are unchanged by this passage. Since ψ has order prime to ℓ, its reduction is injective on its values. At a prime $p \mid f$ with $p \ne \ell$, some inertia element has $\alpha(g) - 1$ a unit, by primitivity at p. Outside $f\ell$, α is unramified. The release applies the character argument at modulus f and then uses $f \mid N_E$ to enlarge the Bernoulli factor from $f\Pi_k$ to $N_E\Pi_k$. Primitive inertia and the character estimate Passage to the ambient conductor

Proved here using Kummer theory (residually cyclotomic case). Suppose $\bar\alpha = \bar\kappa_\ell$, where bars denote reduction modulo the uniformizer of O. Then ψ is unramified outside ℓ, by injectivity on its values and the equality $\bar\psi^{-1}\bar\kappa_\ell^{-j} = \bar\kappa_\ell$. Put $W = K_O/O$ with action α and let π be a uniformizer. The sequence $0 \to W[\pi] \to W \xrightarrow{\pi} W \to 0$ identifies the strict π-torsion classes with a subgroup of the residual strict group. In detail, globally and on the decomposition group at ℓ, the invariants in W vanish because some inertia element at ℓ has $\kappa_\ell(g) \equiv -1 \bmod \ell$, and then $\alpha(g) - 1$ is a unit. On inertia away from ℓ, W is trivial and multiplication by π on invariants is surjective. The long exact sequence therefore gives the required injectivity on these local groups in degree one. The residual group for $\mathbb{F}_\ell(1)$ is zero: by Kummer theory, a rational class unramified away from ℓ has valuations divisible by ℓ away from ℓ, hence is represented by ℓ^a modulo ℓ-th powers, since ℓ is odd. Local triviality at ℓ forces $\ell \mid a$ as well. Extending

coefficients to the residue field of O preserves this argument by linearity. Each W-valued cocycle has finite image by continuity and compactness, so its class is killed by a power of π. Vanishing of the strict π-torsion forces the entire strict group to vanish. This gives (VI.2) in the branch called exc. Classical Kummer theory Residual case in the main proof Residual vanishing

Proved here (the Bernoulli valuation in the other branch). If $\bar{\alpha} \neq \bar{\kappa}_\ell$, choose g where $\alpha(g) - \kappa_\ell(g)$ is a unit. For an integer $m \geq \max(1, v_\ell(f))$, put $f_0 = f/\ell^{v_\ell(f)}$, $L = f_0\ell^m$, and choose $b \in [1, L-1]$ representing the action of g on μ_L. Then $\chi^*(b)b^k - 1$ is a unit. Indeed modulo π, $\alpha(g) - \kappa_\ell(g)$ is $(\chi^*(b)b^j)^{-1} - b$. For each unit $u \bmod L$ use its representative in $[1, L-1]$, write $y \equiv bu \bmod L$ with $1 \leq y < L$, and put $h_u = (bu - y)/L$. The integer

$$\theta_b(u) = \left\lfloor \frac{bu}{L} \right\rfloor - \left\lfloor \frac{b(L-u)}{L} \right\rfloor = 2h_u - (b-1)$$

is the difference of the two Stickelberger floor coefficients. Choose cyclotomic lifts σ_u acting by u on μ_L, and define $z_b = \sum_u \theta_b(u)\alpha(\sigma_u)^{-1}$. Thus $z_b \equiv S_b = \sum_u \theta_b(u)u^j\chi^*(u) \pmod{\ell^m O}$. On the subgroup fixing μ_L both α, β are $1 \bmod \ell^m O$. Use $\delta = \tilde{c}$ at all depths. It acts by -1 on μ_L, and $\alpha(\delta) - 1 = -2$ is a unit. Floor coefficients Congruences on lifts Choice of the unit and sum

Write $H(u) = L^{k-1}B_k(u/L)/k$ and $d_k = 2k\Pi_k$. Expansion in Bernoulli numbers gives

$$d_k\big(H(u) - u^k/(kL) + u^{k-1}/2\big) \in L\mathbb{Z}$$

because all terms with Bernoulli index $i \geq 2$ contain L^{i-1} and d_k clears the remaining denominators. Substituting $y = bu - Lh_u$ and reducing the binomial expansion of y^k modulo L^2 now gives

$$d_k\big((bu)^j\theta_b(u) + 2H(y) - 2b^kH(u)\big) \in L\mathbb{Z}.$$

Multiply by $\chi^*(bu)$ and sum over units. Since $u \mapsto y$ permutes them, if χ_L denotes the character modulo L induced by χ^*, the result is the congruence

$$d_k\left(b^j\chi^*(b)S_b - 2(\chi^*(b)b^k - 1)\frac{B_{k,\chi_L}}{k}\right) \in \ell^m O. \tag{VI.4}$$

In this formula B_{k,χ_L} is evaluated at the imprimitive modulus L. Removing the Euler factors at primes dividing L but absent from f multiplies B_{k,χ^*} by $\prod_{p|L,\, p\nmid f}(1 - \chi^*(p)p^{k-1})$. This follows from the classical Dirichlet Euler product and the formula at $1-k$. Here the only possible extra prime is ℓ, and its factor is an ℓ-adic unit because $k \geq 2$. Hence B_{k,χ_L} and B_{k,χ^*} have the same valuation. Coefficient expansion in the release Summed congruence Extra Euler factor

Choose m large enough that $|\ell^m|_\ell < |a_0 d_k B_{k,\chi^*}/k|_\ell$, for a specified $0 \neq a_0 \in O$. The right side is positive by §3.4.2, and $|a_0|_\ell, |d_k|_\ell \leq 1$. Equation (VI.4), the unit $\chi^*(b)b^k - 1$, and the oddness of ℓ give $|S_b|_\ell = |B_{k,\chi^*}/k|_\ell$ by the strong triangle inequality after dividing the error by d_k. The same equality holds for z_b by its congruence with S_b. In particular $a_0 z_b \mid \ell^m$ and $z_b \mid f\Pi_k B_{k,\chi^*}$ in O, since $f\Pi_k k$ is integral. This is the numerical input to the cocycle bound. There is no factor polynomial in the auxiliary modulus L in the final valuation. Choice of depth and divisibility outputs Norm comparison and final divisibilities

Classical reciprocity input and release assertions. The remaining step is a bound on lengths of normalized cocycles, using an auxiliary family of Gauss sums and Kolyvagin derivatives. The release obtains its reciprocity statement Rec from Poitou–Tate reciprocity and local unramified orthogonality in the

Tate pairing at primes away from ℓ. Its needed consequence pairs an α-cocycle z with a β-cocycle ν killed by a common nonzero a. Both are unramified off a finite allowed set. At the allowed primes z is locally a coboundary in $W[a]$ except at one auxiliary prime $r' \neq \ell$, and both characters are trivial modulo a on the decomposition group at r'. There z vanishes on inertia. If $\nu(\tau) = yw_0$ for an inertia element τ at r', with w_0 of annihilator (a), then $yz(\phi) = 0$ for Frobenius ϕ. The cup product uses $\alpha\beta = \kappa_\ell$. Its local invariants at the unramified places vanish by orthogonality, and the places where z is a coboundary contribute zero. Global reciprocity leaves the unramified–inertia pairing at r'. Poitou–Tate and local duality, Milne, *Arithmetic Duality Theorems*, Chapter I Unramified orthogonality Reciprocity interface Reduction to finite coefficients and duality Input wiring

Proved here using classical duality (the finite-coefficient step). Identify $W[a]$ with $B = O/(a)$ by choosing w_0. The case of a unit a is trivial. Otherwise, because O is unramified over $\mathbb{Z}_\ell$, write $a = u\ell^h$ with $u \in O^\times$ and $h = v_O(a) \geq 1$, and identify B with $O/(\ell^h)$. Thus $\ell^h B = 0$, and this exponent is exact. Write z, ν in B-coordinates as z_B, ν_B, so that their product cup class has coefficients in $B(1)$. For each $\mathbb{Z}/\ell^h$-linear functional $\lambda : B \to \mathbb{Z}/\ell^h$, push the cup class to μ_{ℓ^h} by applying λ to the underlying coefficient and retaining its Tate twist. The global sum of the local invariants is zero. Outside the allowed finite set both cocycles are unramified, so the cup class factors through the procyclic unramified quotient, whose ℓ-cohomological dimension is one. At the other allowed primes the class vanishes because z_B is a coboundary there in B. At the infinite place the positive-degree cohomology is zero for coefficients of odd order. Thus the pushed local cup class at r' has zero invariant and is zero by local Tate duality.

At r' both cocycles are additive because α and β are trivial on the decomposition group modulo a, so $\kappa_\ell = \alpha\beta$ is also trivial on $B(1)$. For an unramified Frobenius ϕ, write $z_B(s) = \eta(s)z_B(\phi)$, where η is the unramified additive class with $\eta(\phi) = 1$ in $\mathbb{Z}/\ell^h$. Consequently the vanishing just proved is the vanishing of the pairing of η with $s \mapsto \lambda(\nu_B(s)z_B(\phi))$, interpreted with the cyclotomic action on μ_{ℓ^h}. This map is a cocycle there because the cyclotomic action on $B(1)$ is trivial at r'. The unramified classes with trivial $\mathbb{Z}/\ell^h$ coefficients are generated by η. Unramified orthogonality for the dual coefficient forces $\lambda(\nu_B(\tau)z_B(\phi)) = 0$ on inertia. The functionals separate points of B, by the decomposition of finite abelian ℓ-groups into cyclic groups. Thus $\nu_B(\tau)z_B(\phi) = 0$, which is the assertion used in the length descent. Global cup class and removal of the other places Local argument with the unramified generator Separation by finite-coefficient functionals

Release input (Gauss-sum and Kolyvagin derivative system). Fix $\ell, m, L, \alpha, \beta, b$ as above. Take a prime $r \equiv 1 \bmod L$ by classical Dirichlet, choose a place above r, and take $w \in K_O/O$ with annihilator (ℓ^m). Put $U = \mathrm{Gal}(\Omega/\mathbb{Q}(\mu_L))$. For every list I of distinct admissible primes r', where $r' \equiv 1 \bmod L$, $r' \neq r$ and $\ell \nmid \mathrm{ord}_{r'}(r)$, the system gives a continuous additive map $\nu_I : U \to K_O/O$ killed by ℓ^m. These maps are β-equivariant under conjugation and vanish on inertia outside $Lr\prod I$. The initial map takes some inertia element at the chosen place over r to a unit times $z_b w$. For each fresh admissible r', at every place above r' and every Frobenius lift ϕ there, some inertia element has new value equal to a unit times $\nu_I(\phi)$. Here adjoining a prime means prepending it to the list. The decomposition groups at these primes lie in U by splitting. Admissibility The derived system and its local identities Existence statement

The release chooses a degree-one prime in an inverse ideal class, uses a Jacobi sum at that prime, and divides out its contribution from the Stickelberger element for a generator of the product ideal. The resulting elements in $\mathbb{Q}(\mu_{L\prod I})$ satisfy norm and reduction relations up to roots of unity. The release corrects the relations by applying $X \mapsto g_0(X)/X^{a_*}$ and taking a $2f_0\prod I$-th power, where

$a_* \equiv \kappa_\ell(g_0) \bmod \ell^m$ and $\alpha(g_0) - \kappa_\ell(g_0)$ is a unit. It then takes the iterated Kolyvagin derivatives, extracts the associated Kummer maps, and projects by a sum weighted by β. The initial value is $z_b w$ times the unit $2f_0(\alpha(g_0) - a_*)^2$. This is the valuation passed to length descent. The Jacobi-sum identities, the correction of the norm and reduction relations, and the Kummer derivative construction with its local identities are the untranslated inputs. Jacobi-sum and ideal-class construction Norm and reduction up to roots of unity Correction and derivatives Successor value Initial unit and projected floor sum

Proved here using classical Chebotarev (simultaneous auxiliary prime choice). Let $z_U : U \to K_O/O$ be a continuous nonzero additive map killed by ℓ^m and equivariant for α. Let V_I be the O-span of $\nu_I(U)$, a cyclic submodule of $W[\ell^m]$. Choose $\rho^{\ell^m} = r$ and let $H \subset U$ be the kernel of the homomorphism $\vartheta(u) = u(\rho)/\rho$ into μ_{ℓ^m}. For $g \in G_{\mathbb{Q}}$, $\vartheta(gug^{-1}) = g(\vartheta(u))$, since $g^{-1}(\rho)/\rho$ is fixed by u. If an additive map F on U is equivariant for γ, killed by ℓ^m, and zero on H, it factors through the cyclic image of ϑ. Conjugation by g on that image is the $\kappa_\ell(g)$-power map, so $(\gamma(g) - \kappa_\ell(g))F = 0$. Equivariance on the radical quotient Power action computation

Let $B_z = \ker z_U$. The unit difference $\alpha(g_0) - \kappa_\ell(g_0)$ implies $H \not\subset B_z$. If $V_I \neq 0$, let K be the preimage of ℓV_I under ν_I. If $H \subseteq K$, apply the same quotient argument to $\nu_I \bmod \ell V_I$, with $\gamma = \beta$ and $g = \tilde{c}$. The difference $\beta(\tilde{c}) - \kappa_\ell(\tilde{c}) = 2$ is a unit, so all values would lie in ℓV_I, contradicting that they span V_I. Each value outside ℓV_I generates V_I. If $V_I = 0$, use $K = B_z$. Finally let $C = \mathrm{Gal}(\Omega/\mathbb{Q}(\mu_{L\ell}))$. Take $u_1 \in U$ acting by $1 + L$ on $\mu_{L\ell}$. The element $\tilde{c}u_1\tilde{c}^{-1}u_1$ lies in H by inversion of $\vartheta(u_1)$. It acts by $(1+L)^2 \equiv 1 + 2L \not\equiv 1 \bmod L\ell$, since $\ell \mid L$ is odd. Hence $H \not\subset C$. Next cyclotomic layer Nonzero and generating values on the radical stabilizer

Suppose B_z, K, C cover H. A group cannot be the union of two proper subgroups, since multiplying elements exclusive to each gives an element outside both. Choose $x \in H$ outside B_z, C, and $y \in H$ outside K, C. Then $x \in K,\ y \in B_z,\ xy \in C$. By cancellation x^2y belongs to neither K nor C, so $x^2y \in B_z$. It follows that $2z_U(x) = 0$, contradicting odd-order coefficients and $x \notin B_z$. Thus the open set $H \setminus (B_z \cup K \cup C)$ is nonempty. Take a coset of an open normal subgroup contained in it and apply Chebotarev to the class of its image in the finite quotient. Choose the place and Frobenius lift ϕ realizing that coset, avoiding the primes dividing $Lr\prod I$ and any other given finite set of primes. At the resulting prime r', fixing μ_L and moving $\mu_{L\ell}$ gives $r' = 1 + Ls$ with $\ell \nmid s$. The reduction of ρ is a unit in $\mathbb{F}_{r'}$ because ϕ fixes ρ. Hence $r^{f_0 s} \equiv 1 \bmod r'$ by $\rho^{\ell^m} = r$, and $\ell \nmid \mathrm{ord}_{r'}(r)$. This gives a fresh admissible prime with $z_U(\phi) \neq 0$ and $\nu_I(\phi)$ generating V_I. Subgroup avoidance Admissibility Open Frobenius conditions Simultaneous choice

Proved here (extension of the derived maps). The derived maps are additive on U. To pair them with z globally, choose δ' with $\beta(\delta') - 1$ a unit, which exists since $\bar{\alpha} \neq \bar{\kappa}_\ell$. Define

$$n_I(g) = (\beta(\delta') - 1)^{-1}\nu_I(\delta' g \delta'^{-1} g^{-1}).$$

The commutator is in U because $G_{\mathbb{Q}}/U$ is abelian. The commutator of δ' with gh is the product of its commutator with g and the conjugate by g of its commutator with h. Additivity and equivariance give the β-cocycle identity for n_I. For $g \in U$, equivariance and additivity give $n_I(g) = \nu_I(g)$. Continuity follows from that of ν_I. This extension takes values in the cyclic O-span of $\nu_I(U)$, and it vanishes on inertia outside $Lr\prod I$ because that inertia lies in U. Thus the annihilator used in the reciprocity step kills the full cocycle. Global extension and annihilator

Proved here from these inputs (length descent). Normalize a strict cocycle by subtracting the coboundary of $(\alpha(\delta)-1)^{-1}z(\delta)$, so that its new value at δ is zero. Each normalized class is represented at some finite level $a_0 \neq 0$, since it has finite image. Include a local primitive at ℓ among the elements killed by a_0. At each $p \mid f$ away from ℓ, zero restriction to inertia and a unit value of $\alpha - 1$ there force the same local primitive on the whole decomposition group by the cocycle identity and normality of inertia. That primitive is killed by a_0, by evaluating at the unit difference. This gives precisely the normalized a_0-torsion cocycle conditions of $\mathrm{Sh}\mathrm{T}$ at the primes dividing L. These conditions form a directed union as $v_O(a_0)$ increases. A common finite bound on all these modules bounds the Selmer group, since every finite strictly increasing chain in it is witnessed in some finite level. Normalized cocycles and local conditions Directed limit comparison

Fix a_0 and choose m as above. Put $J = v_O(\ell^m)$ and $u = \ell^m/z_b \in O$. Let A_I generate the annihilator of $\nu_I(U)$, and let S_I be the normalized module with extra vanishing on decomposition groups at r and the primes of I. Thus $v_O(A_I) \leq J$. A normalized cocycle vanishing on U descends to the abelian quotient $G_{\mathbb{Q}}/U$ on $W[a_0]$. Commuting any element with δ in the cocycle identity gives $(\alpha(\delta)-1)z(g) = 0$. Hence a nonzero normalized cocycle is nonzero on U. On U it is additive and α-equivariant because α acts trivially on $W[a_0]$ there. At a prime split modulo L and away from L, vanishing on inertia and at Frobenius gives vanishing on the decomposition group. This holds at all primes above it by the cocycle conjugation formula, since α acts trivially on $W[a_0]$ at the split decomposition groups. For a fresh prime from the simultaneous choice, the derived system makes the new inertia value have the same annihilator as $\nu_I(\phi)$. This annihilator is (A_I). Write $A_{I\cup\{r'\}} = A_I B$. The initial inertia value has annihilator (ℓ^m/z_b), so $a_0 \mid A_\varnothing$. The successive divisibilities then ensure that the new annihilator kills both cocycles. In a cyclic coefficient submodule with generator w_0 of annihilator $(A_I B)$, this inertia value equals a unit times Bw_0, by the DVR classification. Reciprocity therefore kills the evaluation $z(\phi)$ by B. Evaluation has kernel contained in $S_{I\cup\{r'\}}$ and image in $W[B] \simeq O/(B)$. Hence

$$\mathrm{length}_O\, S_I + v_O(A_I) \leq \mathrm{length}_O\, S_{I\cup\{r'\}} + v_O(A_{I\cup\{r'\}}).$$

Whenever $S_I \neq 0$, the chosen evaluation is nonzero, so $v_O(B) \geq 1$. The initial inertia value $z_b w$ has annihilator (u), with $a_0 \mid u$. The same reciprocity argument at r bounds $\mathrm{length}_O \mathrm{Sh}\mathrm{T}$ by $\mathrm{length}_O\, S_\varnothing + v_O(A_\varnothing) - v_O(u)$. Along the iteration the annihilator exponents grow strictly and are bounded by J, so a module $S_I = 0$ is reached. The inequalities telescope to $\mathrm{length}_O \mathrm{Sh}\mathrm{T} \leq J - v_O(u) = v_O(z_b)$. Extended lengths pose no problem because the terminal bound is finite. As $v_O(z_b) = v_O(B_{k,\chi^*}/k) \leq v_O(A_{N_E,k}(\chi))$, this gives (VI.2) for the second branch. Cocycle restriction and vanishing criterion Reciprocity step One length increment Initial term and length descent

3.4.4. Placement in the ordinary window

In Paper II's notation the cutoff before rescaling is v, the small positive parameter is ε, and $t = \lfloor v^{401} \rfloor$. The field above has $D = 2^{e(v)+2} \leq 8v^{2+8\varepsilon}$. On the scale in Paper II, $\log(N_E t) \leq 2v^{\varepsilon/4}$ eventually, uniformly in the pair. The unconditional minus estimate $Dt^2 \log(tN_E)$ is therefore at most $16v^{804+33\varepsilon/4}$, within the allowance $C' \exp(8v^{\varepsilon/4})$ for sufficiently large v at each fixed $\varepsilon > 0$. The bound holds pointwise on every such host and contributes no exceptional twists to the interval choice. The total weight is $B_2 \log 2 + B_- + B_+$, with B_2 the unweighted dyadic budget and B_+ the odd plus budget on the same host. The host plus eigenspace has the cohomological sign $(-1)^j$, whereas the idempotent in the Fitting theorem selects $(-1)^{j+1}$. The dyadic and host odd plus bounds are still required under their separate

hypotheses. With those bounds supplied, the total is absorbed into $B_* = \exp(v^\varepsilon/2)$, which becomes $\sqrt{b} = \exp(x^\varepsilon/2)$ after rescaling. Paper II, cohomological budgets Assembly of weighted and unweighted parts Final exponential absorption

3.5. Which odd plus estimates does ordinary zeta need?

Write $a = e + 2 \geq 4$, $N = 2^a qd$, and $D = [E : \mathbb{Q}] = 2^a$. A character of the host is uniquely $\rho = \eta J_q^u J_d^v$, where η is a character of the real dyadic factor and $u, v \in \{0, 1\}$. The three factors are independent by their ramification. Complex conjugation acts on the character by $(-1)^{u+v}$. Thus the plus part in twist j uses $u + v \equiv j \bmod 2$. There are $D/2$ such characters over a splitting coefficient ring. Throughout this section $1 \leq j < t$ and $q, d > \max(2t + 1, 7)$. Host character factors and cases

Count	Coefficient prime, twist and character	Contribution for ordinary zeta
I	$\ell \notin \{2, 3, 7\}$, $\ell - 1 \nmid 2j$	Weighted sum over all such ℓ and all $D/2$ characters
II	$j \geq 2$, ℓ odd, $\ell \in \{3, 7\}$ or $\ell - 1 \mid 2j$, and $\rho \neq 1$	Individual length estimate
III	Nontrivial character of odd order	Empty
IV	$\ell = 3, 7$, $j \geq 2$ even, $\rho = 1$	Zero
V	$\ell = 3, 7$, $j = 1$, $v_2(\operatorname{cond} \eta) \geq b_\ell$	Individual length estimate in growing degree
Vs	$\ell = 3, 7$, $j = 1$, $v_2(\operatorname{cond} \eta) < b_\ell$	Fixed-degree estimate proved below
Z	$\ell \notin \{2, 3, 7\}$, $\ell - 1 \mid 2j$, j even, $\rho = 1$	Zero by §3.3a

Here $b_3 = 14$ and $b_7 = 18$. Indeed the prime-to-ℓ part of ρ is all of ρ. At every prime in II or IV the ℓ-component of its conductor is trivial, because $q, d > \max(7, 2t + 1)$ and an exterior prime with $\ell - 1 \mid 2j$ satisfies $\ell \leq 2j + 1$. Hence the release's even-order test in II selects precisely the nontrivial characters. The reference threshold at $j = 1$ is $R_1 = 16\ell^2 4^{v_2(\ell^2 - 1)}$. The odd-conductor part of ρ has order at most two and the order of ℓ modulo four is two. Substitution gives 9216 and 200704, respectively. The number of Vs characters is $2^{\min(a, b_\ell - 1) - 1}$, since there are 2^{m-2} real dyadic characters of conductor dividing 2^m for $2 \leq m \leq a$, and two choices of odd quadratic factor at $j = 1$. The outside sum in I includes $\ell = q, d$. Thresholds · Weighted and individual bounds

The rational strict group. Let O be a splitting discrete valuation ring of residue characteristic ℓ, with fraction field K, uniformizer π, residue field k, and $e_O = v_O(\ell)$. Set $W = (K/O)(\rho^{-1}\kappa_\ell^{-j})$. Write $S(W)$ for the classes in $H^1(\mathbb{Q}, W)$ trivial on finite inertia away from ℓ and on the full decomposition group at ℓ. The extra decomposition conditions at $p \mid N$, $p \neq \ell$, are automatic. If inertia acts nontrivially, its image has order prime to ℓ and $W^{I_p} = 0$. Otherwise Frobenius minus one acts on W by the nonzero scalar $\rho(\mathrm{Frob}_p)^{-1}p^{-j} - 1$. Multiplication by this scalar is surjective. In both cases the unramified local

cohomology is zero. Taking the ρ-part over E gives $S(W)$ by restriction and corestriction, globally and on the local conditions, since $[E:\mathbb{Q}]$ is invertible in O. Its contribution to the sum is $(\log \ell)\,\mathrm{length}_O\, S(W)/e_O$. Rational strict groups and case bounds

For IV it suffices to treat $\rho = 1$. If $\ell - 1 \mid j$, the trivial residual-character proof of §3.3a applies also at $\ell = 3$. Otherwise the residual character $\bar{\alpha} = \bar{\kappa}_\ell^{-j}$ is nontrivial on inertia at ℓ. A class killed by π lifts to a cocycle with values in $W[\pi]$. On inertia away from ℓ this lift vanishes, since inertia acts trivially on the entire module. At ℓ its restriction is a coboundary ∂b in W. Evaluating at an element where $\alpha - 1$ is a unit gives $b \in W[\pi]$. Over the field cut out by $\bar{\alpha}$, the lifted cocycle is therefore an everywhere unramified additive character, trivial above ℓ. Its field of definition lies in $\mathbb{Q}(\zeta_\ell)$, which has class number one for $\ell = 3, 7$ by §3.3. Extension of ideals followed by norm is multiplication by a divisor of $\ell - 1$, so the residual field also has trivial ℓ-class group. Inflation–restriction gives zero over $\mathbb{Q}$, and the π-primary torsion property gives $S(W) = 0$. This covers every even twist used by IV with the same two reference primes.

The bounded-conductor case at 3 and 7. Let F be the field fixed by $\ker \rho$. Its degree is the order of ρ. In Vs that order is at most $2^{b_\ell - 3}$, since the real dyadic group at conductor $2^{b_\ell - 1}$ is cyclic of that order and the quadratic factor has order two. The degree caps are 2048 and 32768, even as q, d grow. Write Δ_F for the absolute discriminant. Every character of $\mathrm{Gal}(F/\mathbb{Q})$ has conductor dividing N, and conductor–discriminant gives $\Delta_F \le N^{[F:\mathbb{Q}]}$. Restriction to F embeds the strict group into the one with ρ trivialized. For $j = 1$ Poitou–Tate duality identifies the dual of the latter strict group, before extending scalars to O, with the kernel of the localization map on $H^2(G_{F,S_\ell}, \mathbb{Z}_\ell(2))$. Here S_ℓ contains the infinite places and the primes above ℓ, and modified real cohomology vanishes at odd ℓ. Apply duality first to finite coefficients and then pass to the direct limit on strict degree-one classes and inverse limit on degree-two classes. The finite-level cohomology over the ring of S_ℓ-integers is finite, so the inverse systems satisfy the Mittag–Leffler condition. The odd-primary étale K-theory comparison for that ring identifies the ambient H^2 with the ℓ-completion of $K_2(\mathcal{O}_F[1/\ell])$. In total degree two the only other possible descent term is $H^0(F, \mathbb{Z}_\ell(1)) = 0$, since a number field contains only finitely many roots of unity. Localization preserves the ℓ-part of $K_2(\mathcal{O}_F)$ because residue-field K_2 vanishes and the residue-field units above ℓ have order prime to ℓ. A summand of order ℓ^s acquires O-length $s e_O$. We have proved

$$\frac{\log \ell}{e_O}\,\mathrm{length}_O\, S(W) \le \log \bigl|K_2(\mathcal{O}_F)[\ell^\infty]\bigr|. \tag{VI.5}$$

This argument uses prime-to-ℓ restriction and works for ramified coefficient rings. At $\ell = 7$, F is totally imaginary by parity and unramified at 7 by its conductor. Its group of roots of unity has order w_F prime to 7. Emery's bound on K_2 away from $6w_F$ gives a bound $O_{[F:\mathbb{Q}]}(\Delta_F^2 (\log \Delta_F)^{[F:\mathbb{Q}]-1})$ on the right of (VI.5). For 3 the following integral estimate suffices. Poitou–Tate duality, Ch. I §4 · Tate on K_2 of global fields · Weibel, Chapter VI §8, étale comparison for number rings · Emery, Theorem 1.3

Lemma. For each fixed d_0 there are $A, C > 0$ depending only on d_0 such that $\log |K_2(\mathcal{O}_F)| \le A(2\Delta_F)^C$ for $[F:\mathbb{Q}] \le d_0$.

Proof. Put $m = [F:\mathbb{Q}]$, $r = 6$, $L = \mathcal{O}_F^r$ and $\Gamma = \mathrm{GL}_r(\mathcal{O}_F)$. Use Ash's contractible well-rounded retract X of positive forms, with their common scale fixed by making the shortest nonzero lattice vectors have length one. Each open cell is specified by the set of shortest vectors in L, which spans F^r. Stabilizers are finite, and passing to a boundary cell enlarges that set. We give the count that allows small-prime torsion to

be retained. All constants until the end of the proof depend only on m, r. Well-rounded cells and stable homology, Soulé §§3–4

Choose an F-basis $x_1, \dots, x_r$ of shortest vectors and put $L' = \sum_i \mathcal{O}_F x_i$. Packing balls of radius $1/2$ in the real normed lattice shows that its covolume $\mathrm{vol}(L)$ is bounded below. In the chosen basis $\mathrm{vol}(L')$ is a signature constant times $\Delta_F^{r/2} \prod_\sigma \det(H_\sigma)^{1/2}$, with conjugate embeddings counted twice. Here H_σ is the local Gram matrix. Its diagonal entries are at most one, so each determinant is at most one. Hence the index $u = [L : L']$ is $O(\Delta_F^{r/2})$. The lower bound on $\mathrm{vol}(L')$ gives $\det(H_\sigma)^{-1} = O(\Delta_F^r)$ at every embedding. For any other shortest vector $\sum b_i x_i$, we have $u b_i \in \mathcal{O}_F$ and $|\sigma b_i| = O(\Delta_F^{r/2})$ by the bound on H_σ^{-1}. There are polynomially many possible coordinate vectors at each u. Indeed an integral element has an embedding of absolute value at least one unless it is zero, by its norm. Packing in the box prescribed for $u b_i$ proves the count.

The coordinate lattice lies between $\mathcal{O}_F^r$ and $u^{-1}\mathcal{O}_F^r$. There are at most $u^{(mr)^2}$ possibilities, since a subgroup of $(\mathbb{Z}/u)^{mr}$ has at most mr generators. In a fixed coordinate lattice a shortest-vector set has at most 3^{mr} elements, by disjoint balls of radius $1/2$ around them inside the ball of radius $3/2$. Thus the number of distinct sets in the polynomially bounded pool of coordinates is polynomial in $2\Delta_F$. The lattice and the set determine the original cell up to Γ. For each face of this cell the original chosen basis stays shortest, so its enlarged set satisfies the same coordinate and packing bounds. The length of a flag of faces is bounded by the dimension of the space of forms plus one. Barycentric subdivision therefore gives polynomially many simplex orbits in every fixed dimension.

Let Q be the image of Γ modulo $3\mathcal{O}_F$, so $|Q| \le 3^{mr^2}$. The kernel is torsion-free. Indeed restriction of scalars embeds it in $1 + 3M_{mr}(\mathbb{Z}_3)$, where cubing an element of exact depth $s \ge 1$ raises its depth by exactly one, and a power prime to three preserves depth. Let EQ be the homogeneous bar construction, the contractible free Q-complex with ordered p-simplices in Q^{p+1}. The diagonal action of Γ on $X \times EQ$ is free. A stabilizer must lie both in the torsion-free kernel and in a finite stabilizer on X. After subdivision the product cells give a free cellular resolution. In total degree three the quotient has $O((2\Delta_F)^C)$ cells, since at most $|Q|^{p+1}$ choices from EQ accompany any orbit from X. Each product boundary column has Euclidean norm at most five, by the two simplex boundary formulas, also when faces coincide in the quotient. The torsion of $H_2(\Gamma, \mathbb{Z})$ injects into the torsion of the cokernel of this boundary. Hadamard's inequality on its nonzero maximal minors bounds the logarithm of its torsion order by the number of columns times $\log 5$.

Finally the Bass stable rank of a ring of integers is at most two. Rank six lies in the integral isomorphism range for homological stability in degree two, so this H_2 identifies with stable H_2. The elementary stable group is perfect and equals the stable special linear group over a ring of integers. Its universal central extension has kernel $K_2(\mathcal{O}_F)$. In the split determinant sequence for the stable general linear group, units act trivially on elementary-group homology. Indeed on a finite bar cycle conjugation by g agrees with conjugation by $\mathrm{diag}(g, \det(g)^{-1})$ in one larger rank. In the homology spectral sequence, perfectness kills the incoming second differential to bidegree $(0,2)$ and the determinant section kills the incoming third differential. Consequently $K_2(\mathcal{O}_F)$ embeds integrally into stable H_2. This K_2 is finite by finite generation and its rational rank. Combining the estimates and taking the maximum over $m \le d_0$ proves the lemma. Stability range, Soulé §4.2 · Stable elementary groups and K_2, Weibel, Chapters III–IV

Taking $d_0 = 2048$ at 3 in (VI.5) gives $\frac{\log 3}{e_O}\,\mathrm{length}_O\, S(W) \le (3N)^{C'}$ for a fixed C'. There are at most 4096 such characters at 3 and 65536 at 7. These factors can be absorbed into one fixed power of $7N$. This proves the full contribution of Vs at conductor one.

Remaining odd plus estimates. Count I still requires a bound by $(Nt)^C$ on the *sum* over each finite set of exterior primes at fixed j. Counts II and V require individual bounds by $(\ell N j)^C$ with an absolute exponent for this host family. In II the exterior primes are at most $2j+1$, so summing those bounds over $j < t$, the eligible primes and at most $D/2$ characters costs only a further fixed power of Nt. The two reference primes in II and V have the same property. These three estimates would complete the odd plus term of §3.2. Degree grows in V, and both degree and twist may grow in I and II. The exponents of the fixed-degree lemma above depend on degree. Assembly of the exterior sum

3.6. Why two Frobenius traces can bound a sum over all exterior primes

Count I asks for an estimate on a sum over an arbitrary finite set of coefficient primes ℓ, at fixed twist j. On the ordinary family the comparison field in this count is the fixed field $K_0 = \mathbb{Q}(\sqrt{-7})$. Its class number is one. An exterior prime in Case I is prime to $7h(K_0)$ and to every host character order. Its local modulus component in a host character is trivial except possibly at q, d, where it is quadratic at most. Inertia at 7 acts nontrivially on K_0 and trivially on the two cyclotomic fields of conductors N and ℓ. These facts give the field, order and ramification hypotheses in the release's one-character estimate. Parity and the condition $\ell - 1 \nmid 2j$ are exactly those of §3.5. Hypotheses on the comparison field and character

The modular form at a split point of Count I

Proposition. Let U be the rational rank-two block of a type-II point in Count I with the Galois conditions of the release. Put $M = 7N$, $k = 2j+1$ and $\theta = \theta_{\mathbb{Q}(\sqrt{-7})}$. For the chosen coefficient embedding there is a primitive cuspidal newform f of weight k, nebentype θ and level $7\,\mathrm{cond}(\rho)^2$ dividing $7N^2$, attached to $U(-j)$ in the inverse arithmetic Frobenius convention. Its polarized rank-two adjoint block has zero rational finite Selmer group. Rank-two block conditions · Local and global level conditions · Pan, Theorem 1.0.2

Proof. The representation U is continuous, absolutely irreducible already on the cyclotomic kernel over K_0, odd by its determinant θ, and unramified outside ℓM. Its residual semisimplification has characters $\theta\overline{\Psi}, \overline{\Psi}^{-1}$, where $\Psi = \rho\kappa_\ell^j$. These follow from the block conditions, the identity of residual traces and Brauer–Nesbitt (the determinant is fixed and $\ell > 2$). One may work over a finite coefficient extension. For instance the image of the compact global group is covered by its closed subgroups of matrices in the countably many finite extensions of $\mathbb{Q}_\ell$. One such subgroup has interior by Baire, and adjoining the entries of finitely many coset representatives suffices. The same argument applies to the finite list of local characters. Rank-two block conditions

The ordinary flag at ℓ has a line and a quotient whose inertial characters, in that order, are $\xi\kappa_\ell^j, \xi\kappa_\ell^{-j}$, with ξ factoring through the action on μ_ℓ. Untwist by ξ. The two characters are now unramified twists of cyclotomic powers and hence crystalline. Their ratio D, line divided by quotient, has Hodge–Tate weight $-2j$ in the convention $\mathrm{HT}(\kappa_\ell) = -1$. Both $H^0(\mathbb{Q}_\ell, D)$ and $H^0(\mathbb{Q}_\ell, D^\vee(1))$ vanish. The local Euler characteristic gives $\dim H^1(\mathbb{Q}_\ell, D) = 1$. The crystalline dimension formula gives $\dim H^1_f(\mathbb{Q}_\ell, D) = 1$, since $\mathrm{Fil}^0 D_{\mathrm{dR}} = 0$. Thus every extension with this flag is crystalline after the finite twist. This proves the

potentially semistable hypothesis in Pan's theorem and gives the two distinct weights. Bloch–Kato, §3 · Pan, Theorem 1.0.2

Apply Pan's theorem to $U(-j)^\vee$. Count I has $\ell \geq 5$, and its residual representation remains a sum of two characters. The Hodge–Tate weights of this dual are $\{1-k, 0\}$, so a modular attachment in Deligne's arithmetic Frobenius convention has weight k. More explicitly, a one-dimensional geometric twist over $\mathbb{Q}$ is a Tate power times a finite Dirichlet character by class field theory. Absorb the latter into the primitive newform. The gap of the two weights fixes k, and their endpoint fixes the Tate normalization. The determinant of the dual is $\theta\kappa_\ell^{2j}$, fixing the nebentype. Dualizing gives exactly the polynomial of f at inverse arithmetic Frobenius on $U(-j)$. Changing the coefficient embedding replaces f by its conjugate. Pan, Theorem 1.0.2 · Newform convention

At $p \neq \ell$ inertia acts diagonally on U through two characters of $\mathrm{Gal}(\mathbb{Q}(\mu_M)/\mathbb{Q})$. Each has local conductor exponent at most $v_p(M)$, also at 2, by local reciprocity. The monodromy operator is zero and the Artin conductor exponents add. At $p = \ell$ the finite inertial part of the potentially crystalline Weil–Deligne representation is the scalar ξ, with monodromy zero. Its conductor exponent is zero if $\ell \nmid M$. Indeed the level condition gives two finite inertial characters of modulus dividing M, and their cyclotomic exponents $j, -j$ distinguish them on every open inertia subgroup. Thus ξ agrees on inertia with the corresponding finite character. At $\ell \mid M$ its tame conductor exponent is at most one. The exponent of the two-dimensional representation at every p is consequently at most $2v_p(M)$. Tate twist and duality preserve these conductor exponents. Local–global compatibility for classical newforms identifies them with the newform conductor exponents, giving the asserted divisibility. Local and global level conditions · Carayol on local compatibility away from the coefficient prime · Saito on compatibility at the coefficient prime

For the prescribed inertia of the marked point one can compute the exact conductor. At $p \neq 7, \ell$, the restriction to inertia is conjugate to characters ρ_p, ρ_p^{-1} on the outer block, giving exponent $2a(\rho_p)$, where $a(\rho_p)$ is the conductor exponent. At 7, inertia over K_0 acts trivially and the quotient has order two. Its determinant is the nontrivial quadratic character, so the two eigenvalues are $1, -1$ and the exponent is one. At ℓ the scalar ξ on Weil–Deligne inertia equals $\rho|_{I_\ell}$, by the prescribed ordinary characters and their distinct cyclotomic exponents. This again gives exponent $2a(\rho_\ell)$, including $\ell = q, d$. Since $7 \nmid N$, the conductor is $7\,\mathrm{cond}(\rho)^2$. Inertia comparison in the marked frame · Automorphic attachments and the level calculation

The form is non-CM. Otherwise U is induced from a character of a quadratic field L, and its trace vanishes on the nonidentity coset of G_L. The ratio of the two residual characters is then -1 on that entire coset and 1 on G_L, by taking quotients of two elements in the coset. This ratio restricts on I_ℓ to $\overline{\kappa}_\ell^{2j}$, since θ is unramified there and the square of the finite inertia character of ρ is trivial. This restriction is nontrivial by the defining condition of Count I. On the other hand the decomposition group at ℓ preserves the ordinary line. In its intersection with G_L, an inertia element has two distinct eigenvalues by the infinite cyclotomic image and the two different powers. Its eigenlines must be the two inducing lines. The ordinary line is one of them, so the entire decomposition group fixes it and belongs to G_L. This contradicts the nontrivial residual ratio on I_ℓ. Rank-two block conditions

On the scalar matrices the polarized block has character θ. On the trace-zero matrices its action agrees with $\mathrm{ad}^0\, U$. Over G_{K_0} it is conjugation. At complex conjugation the action is $-JX^tJ$ for the antidiagonal identity J, which on trace-zero matrices equals conjugation by $\mathrm{diag}(-1, 1) = U(c)$ in the marked basis. Newton–Thorne's Theorem B(1) applies to the non-CM form over $\mathbb{Q}$ and kills the trace-zero adjoint Selmer

group. Passing to the dual or a Tate twist preserves this adjoint. A finite Selmer class for the scalar character θ restricts to an everywhere unramified additive character on G_{K_0}. At ℓ the crystalline degree-one classes for the trivial coefficient are the unramified ones, by the crystalline dimension formula. The class group of K_0 is finite and restriction detects the rational class by corestriction. This kills the scalar contribution. Action on the rank-two block · Newton–Thorne, Theorem B(1)

The forms in this proposition supply the rank-two portion of the finite list in the trace calculation. There are $O(kM^4)$ of them. All embed as normalized forms in $S_k(\Gamma_1(M^2))$, and distinct newforms give linearly independent simultaneous eigensystems away from M by multiplicity one. The index of $\Gamma_1(M^2)$ is at most M^4. The valence bound in weight k bounds the dimension by $1 + k[\mathrm{SL}_2(\mathbb{Z}) : \Gamma_1(M^2)]/12$, by bounding the vanishing order at infinity of a nonzero form. The list is stable under conjugation. At either split probe $p = 11, 23$, for $\ell \neq p$, its rank-three trace values are $1 + p^{-j}a_p(f)$. Deligne's bound gives absolute value at most three in every embedding, and multiplication by p^j clears the denominator. If two points give the same form, their restrictions $1 \oplus U|_{G_{K_0}}$ are isomorphic. One may therefore use all these forms, keeping repeated trace values with their system multiplicities, uniformly across coefficient primes. The irreducible rank-three points require their own list. Newform convention · Deligne, Weil I

The two off-block Selmer groups

The two vanishings from Kato's Theorem 14.2(1). At a split type-II point, let U be the rational rank-two block of determinant θ, the quadratic character of K_0. The preceding proposition attaches a primitive newform f of weight $k = 2j + 1$ to $U(-j)$. At good primes the attachment compares the Hecke polynomial with **inverse arithmetic Frobenius**. Kato uses this convention in (14.10.3), and (14.10.2) gives determinant $\kappa_\ell^{1-k}\omega_f^{-1}$ for nebentype ω_f. Hence $U \simeq V_\lambda(f)(j)$ after extending the coefficient field at the place $\lambda \mid \ell$ selected by the attachment, by comparison of the unramified polynomials and semisimplicity. Apply Theorem 14.2(1) over the finite abelian extension $K_0/\mathbb{Q}$ with $r = j$. We have $1 \leq j \leq k - 1$ and $2j \neq k$. The finiteness holds at every coefficient place and for every stable lattice in Kato's representation with this twist. Type-II attachment · Attachment convention · Kato, Astérisque 295 (2004), Theorem 14.2(1) · Kato (14.10.2)–(14.10.3)

At every finite $v \nmid \ell$ the local condition is the kernel of restriction to inertia, including at the bad primes. At $v \mid \ell$ it is the kernel of the map to cohomology with coefficients in $B_{\mathrm{cris}} \otimes V_\lambda(f)(j)$, and the degree-one condition at infinity is zero. For a stable lattice T in $V = V_\lambda(f)(j)$, Kato's torsion Selmer group in §14.1 uses the images of these local subspaces in $H^1(K_{0,v}, V/T)$. Reduction therefore maps $H^1_f(K_0, V)$ into that finite group. Every class in the rational Selmer group is unramified outside a fixed finite set S containing the primes above ℓ and the ramification of V. In the exact sequence on $G_{K_0,S}$ for $0 \to T \to V \to V/T \to 0$, the kernel of reduction is the image of $H^1(G_{K_0,S}, T)$. Inflation into full global cohomology is injective in degree one for V/T. The integral group is finitely generated over the coefficient DVR. Its image cannot have finite index in a positive-dimensional vector space over the fraction field, so finiteness of the reduction image gives $H^1_f(K_0, V) = 0$. Finite coefficient extension preserves these kernels. Kato §14.1, pp. 234–235 · Kato §§14.8–14.9

Restriction and averaging over the quadratic group identify this Selmer group, after coefficient extension, with

$$H^1_f(\mathbb{Q}, U) \oplus H^1_f(\mathbb{Q}, U \otimes \theta).$$

Indeed the invariant and anti-invariant projections correspond to the two summands of $\mathrm{Ind}_{G_{K_0}}^{G_{\mathbb{Q}}} \mathbf{1}$. The local identifications preserve the finite conditions. Away from ℓ restriction to finite-index inertia detects the kernel by restriction and corestriction. At ℓ the same restriction test applies after tensoring with B_{cris}. Thus both off-block summands vanish. Equivalently, apply 14.2(1) over $\mathbb{Q}$ to f and the primitive quadratic twist of f at the same twist j. Kato §14.6 · The two off-block projections

The ordinary argument in the release. For either $V = U$ or $U \otimes \theta$, the Galois hypotheses give odd determinant $\delta = \theta$, absolute irreducibility on $H_0 = G_{K_0(\mu_{\ell^\infty})}$, and a local line A and quotient B with inertia characters $\xi\kappa_\ell^{j}$ and $\xi\kappa_\ell^{-j}$, where ξ factors through the action on μ_ℓ. The automorphic witness is an additional input to the local comparison. It gives a spherical source after twisting by $\phi = \delta^{-1}\xi^{-1}$ and twisting down by j. Here δ is unramified at ℓ. The normalized character on A evaluates on Frobenius to the unit root, up to inversion, of the good Hecke polynomial of this source. Both roots have complex absolute value ℓ^j, by purity at good primes. In particular the character of A cannot be κ_ℓ. Equality would force $j = 1$ and trivial finite inertia twist, and would make the normalized Frobenius value a root of unity. The quadratic twist preserves the inertia flag and multiplies the unit root by an unramified sign. Galois hypotheses and witness · Ordinary source comparison · Kato (14.10.5) and Proposition 17.1

For the global class use the newform h attached to $V(-j)$, so $h = f$ or its primitive quadratic twist. Kato's Theorem 12.5(1), at cyclotomic level zero with $r = k - j = j + 1$ and trivial Dirichlet character, identifies the signed period of the local dual exponential with

$$(2\pi i)^{j-1} L^{\{\ell\}}(h^*, j+1)\, \gamma^{(-1)^{j-1}}.$$

Here h^* is the dual newform, the Euler factor at ℓ is omitted, and the period vector has nonzero indicated sign component. Kato's Theorem 13.5(1) (Jacquet–Shalika) gives nonvanishing at $j + 1 = (k+1)/2$, the boundary of the absolute-convergence half-plane. The omitted Euler factor causes no zero there, by the root bound (14.10.5). The geometric reciprocity statement for the smoothed constituents after newform projection is Theorem 9.7. By Theorem 12.6 the integral smoothed families span the rational zeta submodule, so at least one fixed family has nonzero dual exponential after specialization. Proposition 8.12 and Example 13.3 give its integral norm relation (13.1.1). The resulting class y is unramified away from ℓ, and its local projection $y_B \in H^1(\mathbb{Q}_\ell, B)$ is nonzero since the dual exponential on $H^1(A)$ lands in $\mathrm{Fil}^0 D_{\mathrm{dR}}(A) = 0$. For every $n \geq 1$ and every $q \equiv 1 \bmod \ell^n$ outside a fixed finite bad set, corestriction from conductor $q\ell^a$ to the degree-ℓ^n subfield of $\mathbb{Q}(\mu_q)$ gives the tame class with further corestriction $P_q y$. Here $P_q = \det(1 - q^{-1}V(\mathrm{Frob}_q)) \equiv \det(V(\mathrm{Frob}_q) - 1) \bmod \ell^n$, with arithmetic Frobenius. The same family supplies all depths and the classes have the integral unramified condition away from ℓ, q. Kato Theorem 9.7 · Theorems 12.5(1) and 12.6 · (13.1.1) and Example 13.3 · Theorem 13.5(1) · (14.10.5)

Write $V^D = V^*(1) \simeq V \otimes \kappa_\ell \delta^{-1}$ and identify its pairing with V with the wedge pairing in this model. A strict dual class z is unramified away from ℓ and zero locally at ℓ. Scale y, z to make their cocycles and the required local primitives integral. The pinned kernel tests give $y|_{N_0}, z|_{N_0} \neq 0$ for a nonzero strict z, where $N_0 = \ker V \cap \ker \kappa_\ell$. The first test uses the nonzero quotient class and the unequal ℓ-inertia weights on the quotient and the unipotent upper entry. For the second, a class vanishing on N_0 descends to the quotient where an element of $\ker V$ with nontrivial cyclotomic value is central and acts on V^D by a nonidentity scalar. The cocycle identity then makes it a coboundary. The H_0-stable span of the two additive restrictions projects onto both copies of V. If the alternating pairing of the two restrictions vanished, that span would be the graph of a scalar by nondegeneracy and absolute irreducibility. On that graph the pairing

is twice a nonzero scalar times the wedge. Thus there are $g, h \in N_0$ with $b = \langle y(h), z(g)\rangle - \langle y(g), z(h)\rangle \neq 0$. There is $\gamma \in [H_0, H_0]$ with $D_\gamma = \det(V(\gamma) - 1) \neq 0$. A subgroup of SL_2 consisting entirely of unipotents is either trivial or fixes a unique line. Applied to the normal commutator image, either case would contradict absolute irreducibility. Pinned kernel and commutator tests

At each depth n, the cup $y \cup z$ has a primitive u modulo ℓ^n on the global group unramified outside a finite set S containing ℓ and the representation's ramification. Its local classes vanish by strictness at ℓ and the two unramified conditions elsewhere. Poitou–Tate gives the global primitive since the dual obstruction is an everywhere unramified character of $\mathbb{Q}$ with constant coefficients. Put $D_s = \det(V(s) - 1)$ and $H(s) = D_s u(s) - \langle \mathrm{adj}(V(s) - 1) y(s), z(s)\rangle$. The tame norm relation gives $D_s H(s) = 0 \bmod \ell^n$ for $s \in [H_0, H_0]$. In this step Chebotarev approximates s by a Frobenius τ at a good q whose degree-ℓ^n cyclotomic quotient ψ_q at conductor q is trivial at S. The splitting requirement comes from adjoining ℓ^n-th roots of the primes in S and μ_{ℓ^n}, an abelian extension after base change to the cyclotomic tower. The first moment m of the tame Shapiro cocycle, corrected by the coboundary in its augmentation, satisfies $dm = -\psi_q \cup P_q y$. The cup cocycle $m \cup z - \psi_q \cup P_q u$ is locally zero away from q, hence at q by global reciprocity. Evaluating at a tame inertia element t with $\psi_q(t) = 1$ and its conjugate by τ gives $(V(\tau) - 1) m(t) = P_q y(\tau)$ and $\langle m(t), z(\tau)\rangle = P_q u(\tau)$. The adjugate and the congruence for P_q give $D_\tau H(\tau) = 0$. Applying the resulting congruence at γ and $\gamma[g, h]$, whose matrices agree, gives $D_\gamma^2 b = 0 \bmod \ell^n$ because the cup identity gives $H(\gamma[g, h]) - H(\gamma) = D_\gamma b$. As n is arbitrary, this contradicts $D_\gamma b \neq 0$. Kato §14.9 on duality · Pinned tame and commutator calculation

For an ordinary dual class, its localization at ℓ comes from $H^1(\mathbb{Q}_\ell, B^*(1))$. The reciprocity sum for its cup with y gives zero at ℓ since the other local terms vanish. At ℓ this pairs the lift with y_B in the perfect pairing of the two one-dimensional local groups for B and $B^*(1)$. The lift is zero and the dual class is strict, hence zero. The primal condition at ℓ is the image of $H^1(A)$, exactly orthogonal to the dual condition. It has dimension one by local Euler characteristic since $H^0(A) = H^0(B) = H^0(A^*(1)) = 0$. In the Poitou–Tate difference of the ordinary Selmer dimensions, global invariants vanish, unramified local dimensions equal invariant dimensions away from ℓ, and the contribution $+1$ at ℓ cancels -1 from the single real invariant. The primal ordinary Selmer group therefore vanishes too. The finite condition lands in the primal ordinary one since $H^1_f(\mathbb{Q}_\ell, B) = 0$ by the Bloch–Kato dimension formula and $\mathrm{Fil}^0 D_{\mathrm{dR}}(B) = D_{\mathrm{dR}}(B)$. Kato §14.9 and Proposition 14.12 · Pinned ordinary dual argument

The one-character cotangent comparison

Fix a good prime $p \neq \ell$ and let g be geometric Frobenius at $w \mid p$ in K_0. Set $\Psi = \rho \kappa_\ell^j$ and $T_\rho = \Psi(g) + 1 + \Psi(g)^{-1}$.

Write $E_\ell = \mathrm{Frac}\, O$, let ϖ be a uniformizer, and put $e_O = v_\varpi(\ell)$. Alpöge's one-character comparison uses the normalized polarized lift on $G_\mathbb{Q}$ with matrix $\mathrm{diag}(\Psi, 1, \Psi^{-1})$ on G_{K_0}. Complex conjugation has matrix part 1, and its action on the Lie algebra is $X \mapsto -JX^tJ$ with J the reversing matrix. The two marked entries at the separating element and its conjugate inverse are fixed by the normalization, leaving conjugations by $\mathrm{diag}(t, 1, t^{-1})$. The graded presentation A/I of the normalized problem is taken at the split augmentation x with kernel P in A. It comes with lifts of matrix entries, descent of points and comparison of tangents with the complete local ring. The defining ideal $I = (F, G, H, H', Z)$ has

equations of weights $-1, -2, 1, 2, 0$, respectively. The model and its integral lifting property are inputs to this calculation. Normalized split lift · Graded model · Integral lifting clause

Let $C = (P/P^{2,\mathrm{sat}})_{-1}/\langle [F_i] \rangle_O$, where $P^{2,\mathrm{sat}} = \{z \in P : bz \in P^2 \text{ for some } 0 \neq b \in O\}$. The strict coefficient $(E_\ell/O)(\Psi^{-1})$ uses the character of the adjoint lattice line $O(E_{10} - E_{21})$ in indices $0, 1, 2$. Conjugation by the diagonal acts on both entries through Ψ^{-1} and the sign action $-J(\cdot)^t J$ fixes their difference. The source imposes zero decomposition restrictions above ℓN and zero inertia restrictions at every finite prime. At $p \mid N$, $p \neq \ell$, the unramified local group vanishes by the local calculation for $S(W)$, so the two strict kernels agree. Strict lengths

For $0 \neq b \in O$, normalize a class in $S(W)[b]$ by subtracting the unique coboundary that makes its cocycle z zero at the marked g_1, where $\Psi(g_1)^2 - 1$ is a unit. Then $bz = 0$ and $z(c) = 0$ by plus parity and $\ell \neq 2$. The cocycle also vanishes at $cg_1^{-1}c$. Its value on the killed inertia away from ℓ is zero. At ℓ its local primitive already lies in $W[b]$, by evaluation where $\Psi^2 - 1$ is a unit on inertia. Identifying $W[b]$ with $O/(b)$, the cocycle defines a weight -1 deformation in $O \oplus \epsilon O/(b)$ through multiplication of the reference lift by $1 + \epsilon z(E_{10} - E_{21})$. Locally it is an infinitesimal conjugate of the reference lift. Lift the local primitive integrally and use conjugation stability of the ordinary ideal at ℓ. The integral lifting clause of the model then lifts the global derivation on the ambient A to an O-valued derivation with that local compatibility. It kills $P^{2,\mathrm{sat}}$ by torsion-freeness, and its reduction kills the F_i. This gives $S(W)[b] \hookrightarrow \mathrm{Hom}_O(C, O/(b))$, since the functional recovers the derivative of each lifted matrix entry. Every class is killed by some ϖ^n by continuity with discrete coefficients. The length of each $\mathrm{Hom}_O(C, O/\varpi^n)$ is at most $\mathrm{length}_O C$ once C has finite length, so the directed union gives the same bound for $S(W)$. Cocycle comparison · Integral lift · Length bound

The graded calculation requires degree-zero finiteness of $A_0/(A_0 \cap I)$. The special-fibre argument below supplies this input from the linked ordinary support hypotheses.

In characteristic zero the prescribed ordinary ideal at ℓ must hold on dual-number lifts, so their adjoint cocycles satisfy the finite local condition by potential crystallinity. The away-from-ℓ inertia equations give the other finite local conditions. At the split augmentation, rational vanishing for the nonpositive-weight character summands gives cotangent spanning in weights $-2, -1, 0$. This uses class field theory in weight zero and the local finite condition, Poitou–Tate and positive-twist H^2 finiteness in negative weights. The source's point tests also exclude opposite one-sided branches over the same invariant value away from x. At two-sided points, where both positive and negative weights are nonzero, adjoint rigidity reduces the normalized tangent to the Euler direction of the torus. For the split type-II points this uses the modular block and its scalar, trace-zero and off-block Selmer vanishings. For the type-I points it uses a polarized cuspidal attachment and the adjoint vanishing at the chosen coefficient place. The trace comparison uses attachments over K_0 at the Count I levels N in type I and $7N$ in type II, with the local conditions of the respective admissible families. Finite tangent conditions · Split cotangent spanning · Ordinary coefficient wrapper · Admissible families

The graded calculation uses a three-graded Rees algebra $\mathcal{R}$ generated over A_0 in the three unit degrees by $A_{-2}u_0$, A_2u_1, and $(A_0F_*^2 + A_{-2}^2A_2)u_2$. Here Smith form on the weight -1 linear relations gives r classes $\varpi^{c_i}y_i$ and an extra equation F_* with zero linear part, so $\mathrm{length}_O C = \sum_i c_i$. Write r' for the number of weight $+1$ equations and s, s' for the numbers in weights $-2, +2$. Cut the weight-zero equations in the split formal chart, let $J_{\mathcal{R}}$ be the coefficientwise kernel of that chart map, and use the sections $\mathbf{t} = (F_i^2u_0, G_iu_0, H_i^2u_1, H_i'u_1, F_*^2u_2)$. For $a = \mathrm{tr}\, r(g) - T_\rho$ in degree zero, choose a representative

regular in the cut chart, adding a weight-zero term hF_i with $h \in A_1$ if needed. This preserves all point values. In large multidegrees d, the leading terms in $\mathcal{R}/(J_{\mathcal{R}}, a)$ occur in bidegree $(2d_0 + 4d_2, 2d_1 + 2d_2)$ in the negative and positive variables. The Hilbert series for their ranks is $(1-X)^{-r}(1-X^2)^{-s}(1-Y)^{-r'-1}(1-Y^2)^{-s'-1}$. Removing one squared section for the ϖ-layer calculation takes the finite differences with factors $(1-X^2)^{r-1+s}(1-Y^2)^{r'+s'}(1-X^4Y^2)$. The resulting series is $(1+X)^{r-1}(1+Y)^{r'}(1+X^2)/((1-X)(1-Y^2))$, with eventual coefficient $2^{r-1+r'}$ at the specified even positive degree. Each F_i^2 contributes $2c_i$ layers. Additivity and eventual nonnegativity of the Koszul Euler characteristics thus give the lower bound $2^{r+r'} \operatorname{length}_O C \le \chi(a, \mathbf{t})_d$. Leading-term quotient · Euler calculation

For the upper bound, let $M_d = (\mathcal{R}/(J_{\mathcal{R}}, \mathbf{t}))_d$. Replacing each section by its cokernel drops a shifted kernel with nonnegative Euler characteristic. Hence $\chi(a, \mathbf{t})_d \le \theta_a(M_d)$, where $\theta_a(M_d) = \operatorname{length}_O(M_d/aM_d) - \operatorname{length}_O M_d[a]$. In large degrees the generic support lies on the two-sided locus. Indeed a homogeneous prime on the relevant Rees locus with all $A_{-2}A_2$ products zero would give opposite one-sided points by restriction to the two Rees axes. Away from x this is excluded by the point tests, and over x a negative-only point is excluded by the rational spanning in weights $-1, -2$. Over the two-sided locus one can invert fh with $f \in A_{-2}$ and $h \in A_2$. The slice $f = 1$ kills the Euler tangent, so it is finite and reduced over the algebraic closure. Its residual μ_2 action gives at most one dimension per weight at each invariant value. Squaring the $r + 1 + r'$ equations increases this length by a factor at most 2 per equation, by filtering with its first power and using the shifted weight piece. The local multiplicity of $M_d[1/\varpi]$ is therefore at most $2^{r+1+r'}$. Generic support · Torus slices · Squared equations · Euler positivity

With a nonzero at the counted points, the kernel and cokernel lengths on finite torsion cancel in $\theta_a(M_d)$. On the free part θ_a is the valuation of the determinant of multiplication by a. Each geometric invariant point contributes the valuation of its trace difference with the generic multiplicity just bounded. The marked-frame comparison injects the two-sided orbit representatives into distinct admissible isomorphism classes on G_{K_0}. Dividing the upper bound by $2^{r+r'}$ leaves the factor 2. Dividing by e_O converts v_ϖ to v_ℓ, with $v_\ell(\ell) = 1$. Finite extension of coefficient DVRs multiplies the Selmer length and e_O by the same relative ramification index. The source invokes the one-character theorem first over an unramified coefficient DVR and then uses this invariance. At geometric Frobenius, with $Q = \mathrm{N}w$ and $z = \rho(Q)^{-1}$, the probe is $\operatorname{tr}\sigma(g) - (Q^{-j}z + 1 + Q^j z^{-1})$. Multiplying by the ℓ-adic unit Q^j clears denominators without changing valuations. Distinct systems remain distinct indices even if they have identical trace values. This one-character comparison is separate from the support implication on the two Count II links. One-character theorem · Marked-frame comparison · Coefficient independence · Frobenius application

The details of the Rees support and Euler comparisons are given in the one-character source note. Ordinary cotangent length

Special-fibre finiteness in Count I

The degree-zero finiteness deduction uses one abelian CM field L containing K_0, fixed for the marked Count I problem. Put $M = 7\ell N$, write η for the sign of K_0, and $\omega = \bar{\kappa}_\ell$. The link takes $G_{L^+} = U = H_0 \cap H_{\mathrm{th}}(w) \cap H_{\mathrm{lay}}(m_{\mathrm{lay}})$ and $G_L = U \cap \ker \eta$. Here the index-two H_0 contains complex conjugation and agrees with $\ker \eta$ on decomposition groups above M and ramifies somewhere outside M, $H_{\mathrm{th}}(w)$ cuts out the degree-ℓ cyclotomic subfield of conductor w, and $H_{\mathrm{lay}}(m_{\mathrm{lay}})$ cuts out the degree-$\ell^{m_{\mathrm{lay}}}$ cyclotomic layer. The auxiliary $w \nmid M$ splits in the fields of $\eta, \bar{\Psi}, \omega$. The layer

$m_{\mathrm{lay}} = 3C_0 + 7$ uses a uniform class-group rank bound C_0 and gives $[L^+ : \mathbb{Q}] > 3C_0 + 7$ and local degrees above ℓ greater than five. The frame also uses the strict residual classes of characters $\bar{\Psi}^{-1}\eta$ and $\bar{\Psi}^{-2}\eta$. Its quotient $U \twoheadrightarrow \Phi.\Gamma$ kills the closed normal subgroup generated by away-from-ℓ inertia in $\ker \eta \cap \ker \bar{\Psi}$, and restricts to a surjection $G_L \twoheadrightarrow \Phi.G_0$ onto the sign kernel. Linked field · Layer and frame · Common quotient

Work over an unramified coefficient DVR containing the host character values. A local map from the normalized ring to a complete characteristic-ℓ DVR A' with finite residue field kills the coefficient maximal ideal. Its lift has the marked normalization and is ordinary by the local ideal itself. At ℓ the spectral projectors of an inertia element with $\bar{\Psi}^2 \neq 1$ transport the prescribed flag to A'. The Count I conditions $\ell - 1 \nmid 2j$ and quadratic-or-trivial finite inertia give that separation. The prescribed kernel away from ℓ gives descent through $\Phi.\Gamma$, since reduction preserves prime-to-ℓ inertia kernels. For every linked $v \mid \ell$ the three diagonal characters of the flag on inertia are the powers of $\bar{\Psi}$ with weights $(1, 0, -1)$. They equal 1 on $I_v \cap \ker \bar{\Psi}$. Integral flag · Descent and local comparison

For each rank $n = 2, 3$, write B_n for the degree-zero algebra of the linked matrix space, $\phi_n : R_n^{\det} \twoheadrightarrow B_n$ for the twisted determinant comparison and $U_n = \phi_n(J_n)B_n$ for the image of the ordinary Hecke ideal. The residual tuples on G_L before scalar twisting are $(\bar{\Psi}, \bar{\Psi}^{-1})$ and $(\bar{\Psi}, 1, \bar{\Psi}^{-1})$. The rank-two twist restricts the algebraic character κ with $\kappa\kappa^c = \kappa_\ell^{-1}$, and the rank-three twist is κ_ℓ^{-1}. Let $\mathfrak{l}_n$ be the coefficient maximal ideal together with the equations setting the ordered local characters to 1 on $I_v \cap \ker \bar{\Psi}$. On a prime containing $U_n + \mathfrak{l}_n$ the pro-ℓ weight characters are trivial, since the inertial characters then have order dividing the order of the multiplicative group of the coefficient residue field. Finiteness of $R_n^{\det}/J_n$ over weight therefore gives $\mathfrak{m}_{B_n} \subseteq \sqrt{U_n + \mathfrak{l}_n}$. Torsion-freeness gives the expected component dimension $1 + n[L^+ : \mathbb{Q}]$. The separate rank-two support problem uses a rational seed of determinant $\eta\theta$ with residual characters $(\bar{\Psi}\eta, \bar{\Psi}^{-1})$, where θ has order ℓ and conductor w and is trivial on G_L. Scalar twists · Weight-fibre finiteness · Seed and dimension inputs · Seed characters

The rank-three seed. Write $\rho|_{I_\ell} = \widetilde{\omega}^{a_\ell}|_{I_\ell}$, with $2a_\ell \equiv 0 \bmod \ell - 1$, and put $b = j + a_\ell$ and $\chi' = [\bar{\Psi}\omega^{-b}]$, using the Teichmüller lift. Thus χ' is unramified at ℓ and has the host's inertia away from ℓ. Use the fixed ℓ-adic comparison place throughout. The auxiliary character θ of conductor w has order ℓ and is trivial on G_L. Set $\chi = \chi'\eta\theta$ and choose $m \geq 2$ sufficiently large with $m \equiv b \bmod \ell - 1$. The primitive character χ has conductor $7w\,\mathrm{cond}(\chi')$ and parity $(-1)^{m+1}$. The Eisenstein series $E_{m+1}(1, \chi)$ has constant term zero modulo the coefficient maximal ideal. Reducing the order-ℓ factor deletes the multiples of w and introduces the factor $1 - (\chi'\eta)(w)w^m \equiv 0$. Subtract the constant term times the integral correction form of the same weight and character, then apply Deligne–Serre lifting. The resulting newform f has this exact prime-to-ℓ level by primitivity, with $a_p(f) \equiv 1 + \chi(p)p^m$ for every prime p using zero on nonunits. In particular $a_\ell(f) \equiv 1$. Its arithmetic Frobenius attachment V_f has determinant $\chi\kappa_\ell^m$ and $\bar{V}_f^{\mathrm{ss}} = 1 \oplus \bar{\Psi}\bar{\eta}$. Eisenstein congruence and rational attachment · Integral correction and lift · Deligne–Serre, Lemme 6.11 · BLGGT, Theorem 2.1.1

The seed cusp on $\mathrm{GL}_3(\mathbb{A}_L)$ is the Gelbart–Jacquet adjoint of the solvable base change of f to L. Indeed, in a dihedral projective image the trace-zero coset would force the residual ratio $\bar{\Psi}\bar{\eta}$ to have square 1, contrary to the marked element of G_L. The ordinary weight gap makes the projective image infinite, hence Zariski dense in PGL_2 by irreducibility and the exclusion of a torus normalizer. This persists on open subgroups and supplies cuspidality and the irreducibility of the adjoint on the selected L. In geometric normalization the attachment is

$$R = \kappa_\ell^{-1}\,\mathrm{Ad}^0(V_f|_{G_L}), \qquad \bar{R}^{\mathrm{ss}} = \omega^{-1}(\bar{\Psi}\oplus 1\oplus \bar{\Psi}^{-1}).$$

The common Tate twist supplies the factor $|\mathbf{k}_z|$ on the adjoint Satake roots at geometric Frobenius at $z \nmid \ell$. BLGGT's rank-three comparison at ℓ applies by Shin-regularity. Pinned seed and comparison · Arthur–Clozel, Ch. 3 · Gelbart–Jacquet, Theorem 9.3 · BLGGT II, Theorem 1.1

At ℓ let e be the unramified character with arithmetic Frobenius value the unit root of $X^2 - a_\ell(f)X + \chi(\ell)\ell^m$. It reduces to 1 on the decomposition group. The crystalline ordinary line and quotient of V_f have characters $\chi\kappa_\ell^m e^{-1}$ and e. On the marked local coordinate of L, put $\lambda = \chi'\kappa_\ell^m e^{-2}$. The first adjoint line consists of maps from the quotient to the line, so the order for R is

$$(\xi_0, \xi_1, \xi_2) = \kappa_\ell^{-1}(\lambda, 1, \lambda^{-1}).$$

This reduces in that order to $\omega^{-1}(\bar{\Psi}, 1, \bar{\Psi}^{-1})$ on the full local group. Its inertia exponents are $(m-1, -1, -m-1)$, with labeled Hodge–Tate weights $(1-m, 1, 1+m)$. The gaps exclude equality of two ordered characters up to one cyclotomic power. The smooth inducing characters after removing the labeled algebraic factors are unramified. The arithmetic unit-torus character in coordinate i is $\kappa_\ell^{h_i - i}$ on inertia, with h_i the corresponding labeled weight. At a lift g_x of the uniformizer in geometric inverse Artin coordinates the normalized ordinary steps are $\prod_{i<s}\xi_i(g_x)\kappa_\ell(g_x)^i$ for $s = 1, 2, 3$. Their reductions are $(\omega^{-1}\bar{\Psi}, \omega^{-1}\bar{\Psi}, 1)(g_x)$. Seed point checks · Ordinary principal-series operators

Away from ℓ the conductor of V_f equals that of its determinant. At a ramified place the conductor sum forces the codimension of inertia invariants to be one, giving inertia $1 \oplus \chi$. Monodromy commutes with the distinct inertia characters and is therefore zero. On G_L both η and θ are trivial, giving the inertia multiset $(\chi', 1, \chi'^{-1})$ for R and conductor exponent $2a_z(\chi')$ at $z \nmid \ell$, with a_z computed after restriction to L_z. At $z \mid \ell$ the crystalline parameter has conductor exponent zero. The multiset is the Teichmüller lift of the prescribed residual inertia, and injectivity of reduction on prime-to-ℓ roots kills the linked inertia kernel. These checks apply at both base levels of the definite $U(I_3)$ over L^+. They have empty Taylor–Wiles sets, the split tame type at every marked $v \nmid \ell$ over M and the ordinary Iwahori condition at ℓ, spherical level away from the marked places, and $K_c(1)$ at the respective split scalar place v_0 or v_0' above distinct rational primes outside M. The seed is spherical at both of those places. The compact highest weight is $(m-1, 0, 1-m)$. The simple generic packet has trivial component group modulo the fixed dual center, and the pure-inner-twist product formula gives its compact member with the required split local components. Selected realizations · Two runs · Local newform conductor · KMSW, Theorems 1.6.1 and 5.0.5

Proposition (the split coefficient type). At a marked split place $z \nmid \ell$, let $F = L_z$, $K_z = \mathrm{GL}_3(\mathcal{O}_F)$ and $t_z = (\chi', 1, \chi'^{-1})|_{I_F}$, viewed over $\mathbf{C}$ via ι. There is a smooth finite free coefficient representation T_z over the given coefficient DVR O such that, for every smooth irreducible π of $\mathrm{GL}_3(F)$,

$$\left((T_z \otimes_{O,\iota} \mathbf{C}) \otimes_{\mathbf{C}} \pi\right)^{K_z} \neq 0 \quad \Longleftrightarrow \quad r_\pi|_{I_F} \simeq t_z \quad \text{and} \quad N_\pi = 0,$$

where (r_π, N_π) is the Frobenius-semisimplified local Langlands parameter in the algebraic normalization. The external tensor product over the marked places satisfies $\mathrm{TypeForAlt}$. Pinned type test · Integral existence and normalization

Proof. Use local reciprocity to transport t_z to unit characters $(\alpha, 1, \alpha^{-1})$. Group each distinct character μ_a in a single block of size its multiplicity. Thus the trivial case has one block of size three and the nontrivial quadratic case has blocks of sizes two and one. In the distinct case order the characters $(\alpha, \alpha^{-1}, 1)$, since

$a_F(\alpha^2) \le a_F(\alpha)$ for the conductor exponent a_F. In the compact subgroup $J_z \subset K_z$ the lower block (b, a) is divisible by $\mathfrak{p}_F^{a_F(\mu_a/\mu_b)}$. The character λ_z of J_z is the product of $(\mu_a/\mu_{a+1})(\Delta_a)$ over the leading block minors Δ_a, with $\mu_{s+1} = 1$ for s blocks. On the integral Levi it is $\prod_a \mu_a \circ \det$. In particular the quadratic block uses the full integral GL_2 with its determinant character. The exponents refer to the actual local unit filtration, including wild characters at dyadic places. Compact model · Multiplicativity on the conductor subgroup

For the upper block parabolic P, the function $f_z = \mathbf{1}_{J_z}\lambda_z$ belongs to the compact model of induction from this determinant character of $P \cap K_z$, with trivial action on its unipotent radical. When the parameter on the right of the equivalence holds, Frobenius preserves each inertia block because its character extends to the Weil group. Frobenius semisimplification therefore splits it into Weil characters even in a repeated block. Local Langlands and the segment classification express π as the parabolic product of twisted spherical representations on the blocks. Each block supplies a determinant eigenvector, and Iwasawa decomposition embeds the compact function model in $\pi|_{K_z}$ by multiplying its functions by the tensor of these vectors. Conversely, any nonzero map from the cyclic K_z-span of f_z sends f_z to a nonzero λ_z-vector. The conductor calculation gives that vector nonzero image in the upper Jacquet module. The geometric lemma and the separation of the distinct unit-character lines then give the same product of twisted spherical blocks. This argument applies to all smooth irreducibles, including spherical Langlands quotients at reducibility points. The seed itself is a unitary principal series with unlinked inducing characters, so the compact-model embedding applies directly. Induced-model embedding · Jacquet argument · Parameter criterion for the block product · Inertial local Langlands, Theorem 3.7

The Teichmüller values of λ_z lie in $O^\times$. The O-span of the translates of f_z is finite free, since it is a submodule of the functions on a finite congruence quotient of K_z. Take its O-linear dual for T_z. The tensor invariant is equivalently a K_z-map from the complexified cyclic span into π. This is the dual convention in the pinned test. The unramified twist $|\det|^{-1}$ in algebraic normalization leaves the test unchanged. The pinned construction obtains its own integral detector and combines the local tests over the product of integral maximals. Integral detector in the release · Product criterion

The seed on the projected ordinary module. In each run fix the enlarged coefficient DVR $O = c_1.O$, the chosen integral type τ_0, and the base level used to define J_3. Write A for the weight algebra, Θ for its torus character, M^{ord} for the ordinary module before residual projection, and $\mathbb{T}^0$ for the commutative A-algebra generated by the Hecke operators on it. The fixed type and the split local vectors above give a nonzero classical eigenvector at the seed weight x_A. Arithmetic control identifies the characteristic-zero specialization of M^{ord} at x_A with the classical ordinary space at that level, intertwining the spherical operators and all three normalized ordinary steps. The eigenline therefore defines a character of $\mathbb{T}^0$, since every operator relation on M^{ord} acts on its specialization. The character is integral over the arithmetic coefficient DVR by finiteness over A. Its generator reductions agree with the prescribed residual profile, by the good split Frobenius polynomials and the ordered step congruences above. Thus it sends the idempotent e_a of that residual block to 1. It descends to a character y of the image algebra $\mathbb{T}$ on $e_a M^{\mathrm{ord}}$, because the kernel of restriction from $\mathbb{T}^0$ is $(1 - e_a)\mathbb{T}^0$. Coefficient DVR · Arithmetic control · Residual projector · Eigensystem on the projected module

The given map $f : R_3^{\det} \twoheadrightarrow \mathbb{T}$ defining J_3 has the spherical polynomial and ordered unit and uniformizer identities. After applying y, its determinant agrees with that of R at every good split geometric Frobenius. These classes are dense up to conjugacy on the level quotient, so continuity and characteristic-zero

determinant uniqueness give equality of the laws. At $z \mid \ell$ write $u_{z,s} = y(U_{z,s})$ for the normalized step values and $u_{z,0} = 1$. The torus identity gives the i-th character on inertia as $\kappa_\ell(g)^{-i} x_A(\Theta(e_{z,i}(a_z(g))))^{-1}$, where a_z is inverse geometric Artin and $e_{z,i}$ inserts a unit in coordinate i. At the chosen uniformizer lift g_x the character is $\kappa_\ell(g_x)^{-i} u_{z,i+1}/u_{z,i}$ for $i = 0, 1, 2$. The inertia formula and the three ratios give exactly ξ_i on the Weil group and hence, by continuity, on the full decomposition group. The determinant and ordered characters determine the point, so $yf = x$ and $J_3 \subseteq \ker x$ for each of the two base levels. The seed determinant is absolutely irreducible and $\ker x$ is a stable characteristic-zero prime. Hecke identities · Frobenius density · Chenevier, Proposition 1.27 and Theorem 2.12 · Ordered comparison and seed inclusion

This deduction uses one arithmetic weight. The existence of f on the whole image in each run uses the following fixed-level construction.

Proposition (the rank-three determinant on the projected image). Let $X = e_a M^{\mathrm{ord}}$ with the fixed integral type, weight algebra and level of the run, and let $\mathbb{T} \subset \mathrm{End}_A(X)$ be its Hecke image. Use a set of arithmetic weights $x : A \to \overline{\mathbb{Q}}_\ell$ whose kernels intersect in zero. Assume equivariant ordinary control on these fibers and a spanning family of classical parts on each projected fiber on which all the Hecke generators act by scalars. For each contributing system $y : \mathbb{T} \to \overline{\mathbb{Q}}_\ell$, assume a continuous rank-three attachment r_y on the same level quotient $\Gamma = V.\mathrm{frame}.G_0$ of G_L. Its good split polynomials specialize from the formula below. It has conjugate duality with multiplier κ_ℓ^{-2}, and trivial action on $I_u^{\mathrm{kill}} = I_u \cap \bigcap_i \ker \bar{\chi}_i$ at each marked tame place u, where $\bar{\chi}_i$ are the linked residual characters. At each marked $z \mid \ell$ it has a full stable flag whose ordered characters are the specialized torus and step characters below. Then there is a unique continuous degree-three determinant law D on Γ with the spherical polynomials, specializing to $\det r_y$ at these evaluations. Together with the three ordered characters at each z, it gives the weight-compatible surjection $f : R_3^{\det} \twoheadrightarrow \mathbb{T}$ with $J_3 = \ker f$. Hecke kernel construction · Universal pair

Proof. Ordinary freeness and the residual idempotent make X finite free over A. For a nonzero operator in $\mathbb{T}$, some matrix entry in an A-basis survives at a separating weight. Equivariant control transports its action to the projected fiber. Since scalar parts span that fiber, the operator is nonzero at some character y. Thus the product of the y is injective on the actual image and the image is reduced. This argument only uses the induced action on the fiber, which can have a kernel as an action of the abstract scalar extension of $\mathbb{T}$. The image is finite over the complete Noetherian local ring A, being a submodule of $\mathrm{End}_A(X)$. The prescribed generator reductions select one local factor with residue field k_O. Finiteness over A makes the maximal-ideal topology equivalent to the $\mathfrak{m}_A$-adic topology and gives completeness and finite adic quotients. Hence $\mathbb{T}$ is compact Hausdorff. Each arithmetic weight takes values in a finite coefficient DVR, and integrality over the weight scalars puts $y(\mathbb{T})$ in the integers of a finite extension of that DVR's fraction field. The preimage of the coefficient maximal ideal is $\mathfrak{m}_{\mathbb{T}}$, so y is continuous. Consequently the diagonal map into $\prod_y \overline{\mathbb{Q}}_\ell$ with the product of the ℓ-adic topologies is a closed embedding, since its source is compact and its target is Hausdorff. Projected freeness · Arithmetic control · Local image and units · Separation on the image · Compactness and continuity

Write $p : G_L \to \Gamma$ for the quotient, $[g]$ for a group-ring basis element and $c_i(g)$ for the tuple of i-th elementary characteristic-polynomial coefficients of the $r_y(g)$. At a good split w, these tuples are the images of the elementary coefficients, with the alternating signs as in the spherical polynomial

$$D(Z-[p(\mathrm{Frob}_w^{\mathrm{geom}})])=\sum_{s=0}^{3}(-1)^s q_w^{s(s-1)/2}T_{w,s}Z^{3-s},\qquad T_{w,0}=1,\quad q_w=|\mathbf{k}_w|.$$

They are invariant under all conjugations of the Frobenius lifts. The primes above nonsplit places of L^+, apart from finitely many ramified places, have relative degree two and density zero among the primes of L. Chebotarev on finite quotients gives density of the union of the good split classes in Γ. By continuity and closedness each $c_i(g)$ lies in $\mathbb{T}$ for every g and varies continuously there. This extends the coefficients on single group elements, including any products of group elements viewed as single elements. Good split density · Determinant construction

Put $t=c_1$. To obtain the law on a sum $\sum_g Z_g[g]$ of finite support, use the following coefficients for the cubic, with distinct group elements where indicated.

Z_g^3	$c_3(g)$
$Z_g^2Z_h,\ g\neq h$	$c_2(g)t(h)-t(g)t(gh)+t(g^2h)$
$Z_gZ_hZ_k,\ g,h,k$ distinct	$t(g)t(h)t(k)-t(gh)t(k)-t(gk)t(h)-t(hk)t(g)$ $+t(ghk)+t(gkh).$

Use each ordered pair for the second row and each unordered triple once for the third. At each y, the second row is $\mathrm{tr}(\mathrm{adj}(r_y(g))r_y(h))$ by the cubic adjugate identity. For the third row, sum over ordered triples of distinct row indices (a,b,c), taking row a from $r_y(g)$, row b from $r_y(h)$ and row c from $r_y(k)$. Permute rows and columns simultaneously to that order. The resulting determinants have columns indexed by (a,b,c), so repeating an index contributes zero and the sum can run over all triples. Expanding these determinants and summing gives the product of three traces from the identity permutation, the three transposition terms, and the two oriented three-cycle terms. Thus at every y the cubic is the matrix determinant of the sum. Comparing coefficients in polynomial rings via the injective product of the y gives the unit and multiplicativity identities and compatibility when supports are enlarged or terms are combined. Evaluating these polynomials over every commutative $\mathbb{T}$-algebra gives the determinant law. Continuity is the adic continuity of the group-element coefficients. Uniqueness follows from density and the integral determinant-law identity expressing all mixed coefficients in the group-element coefficients (Amitsur's formula). The displayed cubic keeps all multiplicities with integer coefficients, including in residue characteristics two and three. Integral determinant descent · Uniqueness from dense classes

At $z\mid\ell$ put $U_{z,0}=1$. The character $\psi_{z,i}\kappa_\ell^i$ on the decomposition group is defined via inverse geometric Artin by the unit character $\Theta(e_{z,i}(-))^{-1}$ mapped to $\mathbb{T}$ and the uniformizer value $U_{z,i+1}/U_{z,i}$ for $i=0,1,2$. Each step is a unit. The powers of each ratio extend to $\widehat{\mathbb{Z}}$ through the finite adic quotients, so reciprocity gives the continuous characters, uniquely determined on the dense Weil group. The torus identity identifies their inertia restrictions with the universal weight characters. It also gives their prescribed values on the preimage of ℓ-primary torsion in the local topological abelianization, since that torsion has zero image in the unramified $\widehat{\mathbb{Z}}$. Reduction on inertia and the three prescribed residual step ratios give the ordered residual characters on the full decomposition group. For D, reduction of the good spherical polynomials gives the prescribed residual determinant of $\omega^{-1}(\bar\Psi\oplus 1\oplus\bar\Psi^{-1})$ on dense classes. Continuity and the same integral cubic give equality as laws. The global congruence is a determinant identity, so it allows nonsplit residual representations with that semisimplification. The local congruences retain the prescribed order. The tame kernel condition is on the marked typed inertia subgroups. At the separate depth-one place the base construction uses the specified level and scalar residual polynomial.

Weight and residual identities · Local torsion subgroup · Ordered inertia, reduction and torsion · Residual determinant law

Conjugate duality identifies $\kappa_\ell(g)^{-2} r_y(cg^{-1}c^{-1})$ with $r_y(g)^{\mathsf{T}}$ by one change of basis for all g, where c is the frame's complex conjugation. Applied to group-ring sums this proves determinant polarization at each y. On its ordered flag $0 = F_0 \subset F_1 \subset F_2 \subset F_3$, the factor $r_y(g_i) - y(\psi_{z,i}(g_i))\,1$ sends F_{i+1} into F_i. For independently chosen g_i in the local decomposition group, written in r_y via the designated map to Γ, the product in order $0, 1, 2$ is zero, acting from right to left. Its trace after left multiplication by any $r_y(h)$ is zero. The same trace statement holds for $r_y(g) - 1$ with $g \in I_u^{\mathrm{kill}}$. Separation gives the polarized, ordered and tame relations in $\mathbb{T}$. Universality gives f, and equality on the universal inertia characters makes it compatible with A. Since q_w is a unit at each good $w \nmid \ell$, the spherical coefficients recover every $T_{w,s}$ in the image of f. The ordered products of $\psi_{z,i}(g_x)\kappa_\ell(g_x)^i$ recover all three $U_{z,s}$ there. These operators and A generate $\mathbb{T}$. Descent of the deformation relations · Surjectivity on generators

The release uses good weights at fixed admitted algebraic weight $\mathrm{lam} = 0$. Local reciprocity gives power-series coordinates from the free inertia generators with the stipulated torsion values fixed. One can vary the smooth characters through independent infinite sets of torsion twists in the open unit discs, choosing disjoint ranges of ℓ-power orders to distinguish the three unit characters at each coefficient place. Weierstrass preparation in each coordinate gives weight separation. The pinned control and decomposition give scalar action on the good projected fibers at the same tame type, residual profile and away-from-ℓ level. The spectral and local comparison input covers every contributing constituent, including endoscopic sums and any Speh strings before applying a local exclusion. Once full split transfer is available at a coefficient place, a longer Speh string repeats supercuspidal support under unramified shifts and cannot match the good torus support with distinct unit characters. The reverse implication of the split type proposition and local–global compatibility give the exact Teichmüller inertia multiset with zero monodromy at the typed places. At ℓ the toral zero-monodromy comparison, labeled Hodge–Tate weights and normalized ordinary steps give the ordered stable flag. Almost-everywhere spherical transfer by itself leaves these local comparisons open. The release takes $\mathrm{HypothesesH}'$ as input. The parameter-refined trace route below uses stabilization and the quasi-split stable multiplicity formula to obtain the specified local transfer. Good-weight separation · Projected scalar action · Classical covering · Tame comparison · Ordered flag comparison · Hecke kernel construction

Propagation additionally requires separating classical evaluations of the actual projected modules at all required auxiliary levels and $\mathrm{LevPtAlt}$ at the actual split patch completions. Its comparison is with the Frobenius-semisimplified Weil–Deligne parameter, with equality required for toral parameters of zero monodromy. It uses the arithmetic spherical annihilator comparison and classicality at irreducible arithmetic points. The patch modules retain the type, weights and residual data, with ordered residual Frobenius refinements at the patch places. They require freeness over the prescribed diamond group rings, determinant-compatible comparison after augmentation with the same base module, and Frobenius and ordered tame relations after clearing the Frobenius difference factors. The source uses the exponent $3! = 6$ in the integral annihilation relations at rank three. These are the same-module inputs to the passage from $\mathrm{LevelAlt}$ to HasAut. Classical auxiliary-level contract · Arithmetic annihilator comparison · Same-module levels · Determinant-compatible comparison · Split patch identities · Augmentation and annihilation · Passage to patching · Integral rank-three interpolation

Count I rank-three seed calculation · Split type calculation · Base-level eigencharacter and determinant

The arithmetic support condition $\mathrm{P1}$ concerns every generization of a supported nice type-I point. More precisely, if $f : R_n^{\det} \to A_x$ is the determinant map of a characteristic-ℓ curve point with the required irreducibility, distinct ordered local characters and no cyclotomic intertwiner, then $J_n \subseteq \ker f$ implies $J_n \subseteq \mathfrak{q}$ for every prime $\mathfrak{q} \subseteq \ker f$. For the intertwiner test, at the realization's auxiliary place choose g with $\omega(g) \neq 1$ and residual characteristic polynomial $(X - c)^n$, $c \neq 0$. The characteristic polynomial of $r(g)$ evaluated at $\kappa_\ell(g) r(g)$ is invertible by reduction, so any cyclotomic intertwiner is zero by Cayley–Hamilton. The support theorem uses ordinary auxiliary-level modules at all Taylor–Wiles depths, with group-ring freeness, maps to the prescribed base level and the uniform tame inertia and Frobenius relations, together with the cohomological generator and obstruction estimates. In the patched completion the support and ordinary dimension comparisons give propagation. For $\mathrm{P1}$ the chain has length zero and the stable discriminant stays nonzero at every prime below $\ker f$. The chart statement follows by normalizing the one-dimensional quotient to a DVR and contracting the primes through the determinant comparison. The companion chain condition $\mathrm{P3}$ makes the corresponding propagation across nice endoscopic points joined by thick primes, where a nonsplit extension of the middle constituent by the outer block witnesses thickness. Determinant support tests · Absence of the intertwiner · All auxiliary levels · Patching and zero-length chain · Transfer to charts

Rank-two propagation puts every off-triangular prime over U_2. The outer two-by-two corner of a rank-three endoscopic prime is a rank-two point with the matching weight equations. The middle character reduces to the reference on the fixed-weight tests. This transfers the rank-two radical bound to an ideal $I_3 \subset B_3$ controlling all endoscopic off-triangular primes. The rank-three seed, the connected chart comparison with intersections of dimension at least $3[L^+ : \mathbb{Q}] + 1$, and $\mathrm{P1}$ and $\mathrm{P3}$ give support over U_3 on the remaining off-triangular primes. Thus OffFin holds with $I_* = U_3 \cap I_3$. Every off-triangular prime contracts to one containing I_*, and $\mathfrak{m}_{B_3} \subseteq \sqrt{I_* + \mathfrak{l}_3}$ follows by primality from the two component bounds. Rank-two propagation · Corner comparison · Rank-three propagation

For a linked DVR point the fixed-weight ideal $\mathfrak{l}_3$ vanishes. In the off-triangular case OffFin makes its invariant map factor through the residue field. The diagonal entries become the reference entries and the three two-cycles and two three-cycles of off-diagonal entries in the determinant all vanish. In the triangular case divide the three diagonal characters by their reference characters. The ratios take values in principal units. Each finite reduction kills inertia above ℓ by the frame's local generator test, and away from ℓ by prime-to-ℓ inertia in the link. Class field theory bounds all their orders by one exponent, for example $2h_L$. A finite-order principal unit in a characteristic-ℓ domain is 1, by Frobenius for its ℓ-part and reduction for its prime-to-ℓ part. These two cases give the exact polynomial $\prod_{i=-1}^{1}(X - \bar{\Psi}(g)^i)$ on G_L. An injective finite DVR extension handles the coefficient enlargement in the linked realization, and the identity descends coefficient by coefficient. One arithmetic realization in each rank suffices for this finite-fibre test. This is FinEnd with the **same** subgroup G_L for every test point. Off-triangular case · Triangular case · Coefficient descent · Common subgroup

For any $g \in G_{K_0}$ a power lies in G_L, so the eigenvalues are algebraic over the coefficient residue field. Reduction is injective on algebraic constants in A', which recovers the reference polynomial at g. The marked diagonal difference has three distinct eigenvalues, so a stable subspace is spanned by coordinate lines. The polarized transpose pairs the adjacent entries. Opposite signs with an adjacent entry make the directed graph of entries strongly connected, contradicting the split semisimplification given by the full polynomial. If only the two opposite corner entries occur, the middle line has character 1 and removing it gives the same contradiction on the outer block. The fixed complex conjugation extends this one-sign

triangularity to the polarized matrix parts. Each augmented invariant monomial therefore vanishes on every test. A monomial surviving at a prime over ϖ would survive on a one-dimensional complete domain quotient and its normalized DVR. Projection to degree zero then gives $\sqrt{\varpi C^0} = \mathfrak{m}_{C^0}$ for the completed invariant ring C^0. Its special fibre is Artinian and topological Nakayama makes C^0 finite over O, hence also its entry-monomial subalgebra. Algebraic constants and marked triangularity · DVR detection · Weight-zero finiteness

This derives FinA from the linked arithmetic support hypotheses. At either base level the rank-three seed inclusion follows by the single-specialization argument once the projected determinant pair is constructed. The remaining finiteness inputs include the rank-two seed inclusion and the same-family spectral and Galois comparisons needed to apply ordinary patching in both ranks for $\mathrm{P1}$ and $\mathrm{P3}$. The rank-four argument below derives auxiliary-level freeness and the augmentation comparison on the selected Count II links from their fixed-level hypotheses. The characteristic-zero attachments and the controlled-level trace lists enter the one-character inequality separately. Ordinary Hecke family · Linked arithmetic input

Count I special-fibre finiteness

The trace product

Here are the two arithmetic inputs to the deduction. Write $v_\ell(\rho) = \mathrm{length}_O\, S(W)/e_O$ for the strict length in §3.5 with ρ and j understood. For a good rational prime $p \geq 7$, choose $w \mid p$ in K_0, put $Q = \mathrm{N}w$, and let g be inverse arithmetic Frobenius at w.

1. For $\ell \neq p$, there is a finite set X_ρ of representatives of distinct admissible Galois representations over K_0. Write $\Psi = \rho\kappa_\ell^j$ and $T_\rho = \Psi(g) + 1 + \Psi(g)^{-1}$. If the traces of all admissible representations at g differ from T_ρ, the strict length is finite and satisfies

$$v_\ell(\rho) \leq 2 \sum_{\sigma \in X_\rho} \mathrm{ord}_\ell(\mathrm{tr}\,\sigma(g) - T_\rho), \qquad \mathrm{ord}_\ell(\ell) = 1. \tag{VI.6}$$

   The normalized model and arithmetic hypotheses for this comparison are given above. The inequality applies to the two-sided orbit representatives once the admissible systems separate their traces from the diagonal. Application of the one-character bound

2. There is a Galois-stable list of complex traces u_s of full Hecke eigensystems at g. Its two parts are the cuspidal rank-three systems and the systems with trace $1 + \mathrm{tr}\,U$. For the chosen K_0, assume $|S_p| \leq C[N(j+1)]^A$, $|u_s| \leq \sqrt{Q} + 2$ at every embedding, and $Q^j u_s$ algebraically integral. The constants precede the host, twist and test prime. For each coefficient embedding and each ρ, the members of X_ρ have their traces at distinct system indices s in this list. Repeated trace **values** are retained with their system multiplicities. The release obtains this input by local–global compatibility and counting systems at controlled levels. Combination of the two sets of systems

Proposition. These two inputs imply Count I on the ordinary host, with constants independent of the finite set of coefficient primes.

Proof. Take two fixed primes $p = 11, 23$ and enlarge the common threshold so $q, d > 23$. Both split in K_0 since their classes $4, 2 \bmod 7$ are squares. In particular $Q = p$. Work first at one of the two primes. Let $\mathcal{X}_j$ be the complex characters of the host of parity $(-1)^j$. Its cardinality is $D/2$ and it is stable under all

automorphisms of the algebraic numbers. Set $R = p^{j}$ and $z_\rho = \rho(p)^{-1}$. For each ρ and each system index define

$$F_{\rho,s} = R(u_s - 1) - z_\rho - R^2 z_\rho^{-1}.$$

This factor is the integral multiple R of the difference between the trace u_s and the diagonal trace $R^{-1}z_\rho + 1 + Rz_\rho^{-1}$. The last summand in $F_{\rho,s}$ has absolute value R^2, whereas the first two together have absolute value at most $R(\sqrt{p}+3)+1 < R^2$ for $p = 11, 23$ and $R \geq p$. Hence every factor is a nonzero algebraic integer of absolute value at most $2R^2$. Taking their product over both index sets gives a nonzero rational integer n_p, since every field automorphism permutes the factors. Thus

$$\log |n_p| \leq (D/2)|S_p|(2j\log p + 1). \tag{VI.7}$$

Fix an exterior $\ell \neq p$ and an isomorphism from the chosen algebraic closure of $\mathbb{Q}_\ell$ to $\mathbb{C}$. It identifies the ℓ-adic host characters with the complex characters of the same parity after adjoining their values. For a nonsplitting coefficient field, extend it to include the host character values. The normalized length is invariant under this extension. Now each difference in (VI.6) occurs amongst the factors indexed by that character. The factor R is an ℓ-adic unit. Unused factors have nonnegative valuation, and the indexing keeps every occurrence even if two different characters have identical traces. We obtain

$$\sum_{\rho \in \mathcal{X}_j} v_\ell(\rho) \leq 2\,\mathrm{ord}_\ell(n_p). \tag{VI.8}$$

For $\ell = 11$ use n_{23}, and for all other ℓ use n_{11}. The weighted sum over any finite set in Count I is at most $2\log|n_{11}n_{23}|$ by unique factorization in $\mathbb{Z}$. Equation (VI.7), $D \leq N$ and $j+1 \leq t$ give the asserted fixed power of Nt. In the release the construction is made for the characters of each exact kernel. The complex fibers have sizes $\varphi([G_N : H])$ and are disjoint, so their total on the host is at most D. The product proof above uses the union of the parity fibers. Product and subproduct comparison · Two-prime estimate

An explicit form of the two integers. Put $m = 2^{a-2}$, and for $p = 11, 23$ set $h_{11} = m$ and $h_{23} = m/2$. The signs $J_q(p), J_d(p)$ disappear from the character products in (VI.7). To see this, take the representative $-p \equiv 1 \bmod 4$ of the class of p in the cyclic real dyadic group. For $b \equiv 1 \bmod 4$, repeated squaring gives $v_2(b^{2^k}-1) = v_2(b-1) + k$ if $b \neq 1$ and $b \neq -1$. Here $v_2(-p-1)$ is two at 11 and three at 23. Thus the dyadic evaluations at p run over the h_p-th roots of unity, each m/h_p times. In either parity two pairs of quadratic exponents occur. Each multiplies the dyadic values by a sign, which permutes those roots since $a \geq 4$ makes both orders even. The multiplicities for the full set are therefore two at 11 and four at 23.

For $R = p^{j}$ define integral polynomials $C_0 = 2$, $C_1 = X$, $C_{d+1} = XC_d - R^2 C_{d-1}$. If $X = U + R^2/U$, induction gives $C_d(X) = U^d + (R^2/U)^d$. The monic polynomial $C_h(X) - (1 + R^{2h})$ has the h roots $z + R^2/z$, $z^h = 1$. They are distinct since equality for two unequal roots would imply $zz' = R^2$, whereas $|zz'| = 1$. We obtain the formula

$$n_p = \prod_{s \in S_p} \left[C_{h_p}\!\big(R(u_s - 1)\big) - 1 - R^{2h_p} \right]^{2m/h_p}. \tag{VI.9}$$

The eigensystems in this expression still depend on the host level and twist. For any Case I prime other than 11 and 23, (VI.8) can in fact be applied twice. Its length is bounded by twice the valuation of $\gcd(|n_{11}|, |n_{23}|)$. At 11 use the valuation of n_{23} alone, and at 23 that of n_{11} alone. This also identifies a

finite set containing every Case I prime of positive length for the specified host and twist. Classification by system indices · Trace estimate for a Galois-stable family of roots

Count I therefore uses the conditional one-character comparison above and the two controlled-level lists. The split coefficient type gives local occurrence for the linked rank-three seed, and the one-weight comparison gives its support at each base level with the projected determinant pair. The auxiliary-level comparisons for special-fibre finiteness and the propagation and descent giving the controlled attachments for all two-sided type-I points remain part of the Count I input. The normalized model with its ordinary ideal and cohomological hypotheses also enters the comparison. The rank-two proposition and Kato's theorem supply the pointwise modular-block ingredients under their stated Galois conditions. For the rank-three adjoint vanishing given attachment, see §3.7. Automorphy and finiteness hypotheses

3.7. Trace products for Counts II and V

Use the notation of §3.5 and write $L_{\ell,j,\rho} = (\log \ell)\,\mathrm{length}_{O}\, S(W)/e_O$. The bounds below give the two individual estimates once the source's deformation and automorphy inputs hold on this host. Count statements

Counting the systems. For a fixed imaginary quadratic field F, the number of cohomological cuspidal systems in rank r at principal level M with coefficient system V is $O_{F,r}(M^{2r^2}\dim V)$. At fixed level 3 the finitely many arithmetic quotients have finite cell models by Borel–Serre compactification. Passing to level $3M$ lifts each cell at most $(3M)^{2r^2}$ times, by the index bound from reduction into $\mathrm{GL}_r(\mathcal{O}_F/3M)$. The cochain dimension is bounded by the cell count times $\dim V$. Cuspidal cohomology injects into arithmetic cohomology, with each distinct cohomological constituent contributing at least one dimension. Level M invariants are also invariant at level $3M$. In rank four the Weyl formula at each embedding, for $m = (j+1, j, 1, 0)$, gives

$$\prod_{i<k} \frac{m_i - m_k}{k-i} = \frac{j^2(j^2-1)}{12}.$$

The two embeddings give $O((j+1)^8)$ for the coefficient dimension. The count is therefore $O(M^{32}(j+1)^8)$. Rank three over $\mathbb{Q}(i)$ has count $O(M^{18})$ at $j=1$ because the weight $(1,0,-1)$ is fixed. Level index and cohomology Cuspidal injection and weight criterion Weyl construction Borel–Serre, §§8–11 Franke, Ann. Sci. ENS 31 (1998), 181–279

To count rank-four points over $\mathbb{Q}$, the source fixes 18 imaginary quadratic fields with pairwise coprime discriminants. At most one is ramified at the test prime. At most 16 sign characters can be self-twists of a rank-four irreducible representation, since their intertwining lines in its endomorphisms are independent. Thus some remaining restriction is irreducible. A restriction has at most two lifts up to isomorphism, differing by a sign twist. At the test prime one recovers candidate rational Frobenius eigenvalues by taking roots of order at most two, at bounded cost. The resulting system bound has the same powers of M, j. This step uses base change at the stated level, local–global compatibility and Galois conjugation of algebraically normalized Satake data. Equality of full system indices implies isomorphism by Chebotarev and Brauer–Nesbitt. Equal traces at one prime may belong to different systems, and the product retains their multiplicities. CM system classification Rational classification Arthur–Clozel, Ch. 3

Count II. Put $c = v_2(\mathrm{cond}\,\eta)$ and $B = 2^{n_{\mathrm{II}}}$. The source chooses an odd primitive character of conductor B and uses the trivial odd-modulus auxiliary character. Its level parameter is $M = \ell N B$, and the three

finite characters of the split lift factor through modulus M and have orders dividing $k = 2(\ell - 1)B$. The odd-conductor part of ρ has order at most two, so $n_0 = v_2(\ell^f - 1)$ for $f = \operatorname{ord}_{(\mathbb{Z}/4)^\times}(\ell)$. If $\ell \equiv 1 \bmod 4$, then $2^{n_0} \le \ell - 1$. Otherwise $2^{n_0} \le 2(\ell + 1)$. Since $2^c \le 2^a$ and $\ell^2 < qd$ on II, the ceiling in n_{II} gives

$$B < 2\max\{\ell N 4^{n_0}, 64 \cdot 2^c \ell^3\} \le 8\ell N(\ell + 1)^2, \qquad M \ll \ell^4 N^2, \qquad k \ll \ell^4 N.$$

Thus the dyadic cyclotomic field has degree at most B. The dyadic roots required in the coefficient ring lie in an unramified extension of $\mathbb{Q}_\ell$. Even the source's combined exponent $n_{\mathrm{II}} + n_{\mathrm{III}}$ costs degree $O(\ell^6 N)$, since the unused III parameter satisfies $2^{n_{\mathrm{III}}} = O(\ell^3)$ at conductor one. In any coefficient extension the Selmer length and e_O scale by the same relative ramification index. Dyadic orders and levels Ceilings Construction and level Finite-character exponent Coefficient extension

At $p \ge 5$ prime to M, the split trace is $T_4 = p^{j+1} z_1 + p^j z_2 + p z_3 + 1$, with $z_i^k = 1$. The admissible rank-four systems have integral trace s with $|s| \le (3\sqrt{p} + 1)p^{(j+1)/2} \le (7/2)p^j$. Since $p \ge 5$ and $j \ge 2$, the leading term of T_4 exceeds $|s| + p^j + p + 1$ in modulus. Hence $0 < |s - T_4| \le p^{2(j+1)}$. Form A_4 by multiplying $s - T_4$ over the systems and the conjugates of the phase triple. Galois permutes the integral factors, so A_4 is a nonzero integer. Consecutive use of the cotangent estimate, integrality of the unused factors and $-\log|A_4|_\ell \le \log|A_4|$ gives

$$L_{\ell,j,\rho} \le 2\log|A_4| \ll |\operatorname{Orb}(z_1, z_2, z_3)|\, M^{32}(j + 1)^9 \log p. \tag{VI.10}$$

Indeed the cotangent estimate charges twice the sum of the negative ℓ-adic log norms over distinct irreducible points. Each point uses a distinct system factor and all other factors have ℓ-adic norm at most one. Split trace Separation and integral product Product and classification Normalized comparison

The source's least good prime $p \ge 5$ is 5 or 7 on this host. The Galois orbit contains at most $\varphi(k) \le k$ phase triples. Equation (VI.10) gives $O(\ell^{132} N^{65}(j + 1)^9)$. Alternatively take the least $p \equiv 1 \bmod M$. The phases are all one and Linnik gives $\log p \ll \log(\ell N)$. This gives

$$L_{\ell,j,\rho} \ll \ell^{128} N^{64}(j + 1)^9 \log(\ell N) \ll (\ell N j)^{67}. \tag{VI.11}$$

For the last bound at an exterior prime, use $\ell \le 2j + 1$, $\ell^5 < N^{5/2}$ and $\log(\ell N) \ll \sqrt{N}$ to absorb $\ell^{61}\log(\ell N)$ in $O(j^{56} N^3)$. At the reference primes the first bound is $O(N^{64}(j + 1)^9 \log N)$. Good-prime conditions Least good prime Linnik's bound, Xylouris

Count V. Here $j = 1$ and the dyadic conductor threshold in §3.5 gives the source's rank-three local conditions with level N. The sign character cuts out $K = \mathbb{Q}(i)$, of discriminant -4. For $p \ge 7$ inert in K and prime to ℓN, put $Q = p^2$ and let g be inverse arithmetic Frobenius at the prime of K over p. The diagonal characters are $(\rho\kappa_\ell, 1, (\rho\kappa_\ell)^{-1})$, so the trace at g is $T_3 = Q^{-1} z + 1 + Q z^{-1}$, where $z = \rho(p^{-2})$. In the system count $|s| \le \sqrt{Q} + 2$ and Qs is integral. Because $Q > \sqrt{Q} + 3 + Q^{-1}$, every factor $Q(s - T_3)$ is nonzero and has modulus at most $2Q^2$. The product A_3 over the Galois-stable rank-three systems and the conjugates of z is a nonzero integer. The broader source norm includes a rank-two branch. For V, the point classifier sends every tested representation into the rank-three branch. The cotangent bound uses twice the valuation sum and the source allows two embeddings of K. Within each embedding the classification is injective on the points being counted. Since Q is an ℓ-adic unit, this gives

$$L_{\ell,1,\rho} \le 4\log|A_3| \ll |\operatorname{Orb}(z)| N^{18}(4\log p + 1). \tag{VI.12}$$

Host data and dyadic condition Trace and classification Rank-three branch Two embeddings Trace bound Integer product

At $\ell = 3$ take $p = 7$. At $\ell = 7$ take 11 or 19 different from q. Neither can be $d \equiv 7 \bmod 16$. The odd quadratic factors disappear at p^{-2}. In the real dyadic group of order 2^{a-2} the square has order dividing 2^{a-3}, so the phase orbit has size at most 2^{a-4}. Alternatively choose the least prime in the reduced class $p \equiv -1 \bmod 2^a,\ p \equiv 1 \bmod qd$. It is inert, avoids 3 and 7, and $\log p \ll \log N$ by Linnik. Its phase is one. Thus (VI.12) gives $L_{\ell,1,\rho} \ll N^{18} \min(2^{a-4}, \log N) \ll (\ell N)^{19}$. The coefficient field can split the host characters by an unramified extension of degree dividing $\varphi(2^{a-2})$. Source prime and orbit choices Source phase orbit Linnik's bound

Inputs to the cotangent comparisons. In II, the cotangent comparison uses the weight -1 cotangent length at the split lift, a finite degree-zero algebra, exclusion of the specified block end, automorphy of all irreducible points, and rigidity modulo symplectic conjugation. The source proves rigidity by passing to an auxiliary CM field and returning an adjoint coboundary to $\mathfrak{sp}_4$. In V the reduced lifting input gives the marked element, finiteness in degree zero and automorphy of the type-I points over K after a Tate twist. The frame comparison uses local tangent conditions, Poitou–Tate, Euler characteristics and vanishing in the relevant weights. Once the type-I adjoint vanishing below applies, the remaining tangents come from the torus orbit. The frame comparison gives the common bound for the finite-coefficient lengths, which then bounds the divisible strict group. The construction of the reduced lift and the classification at the stated local level remain inputs to (VI.12). II cotangent bound · II rigidity · V reduced lift · V cotangent and finite-coefficient bounds · V assembly

A common choice of rational primes for all three counts. For $q, d > 23$ take $p = 11$ unless $\ell = 11$, in which case take $p = 23$. In I these primes split in $\mathbb{Q}(\sqrt{-7})$. In II both avoid every prime in the ramification set dividing $2\ell N$ when used as indicated. In V, where $\ell = 3, 7, 11$ is inert in $\mathbb{Q}(i)$ and its prime has norm 11^2. The hypotheses on the Galois element used to bound a deformation length in II and V follow from the respective archimedean trace gaps. The uniform field and representation hypotheses in §§3.6–3.7 remain in force.

With this fixed choice, the phase triples of II have at most $k \ll \ell^4 N$ conjugates and $M \ll \ell^4 N^2$. Equation (VI.10) and the bound on the logarithm of each factor give $L_{\ell,j,\rho} \ll \ell^{132} N^{65} (j+1)^9$. In II, $\ell \le \max(7, 2j+1)$. Summing over its primes, over $j < t$ and the host characters gives $O(DN^{65} t^{143})$. For V the phase at 11 has at most 2^{a-4} conjugates, so the bound on a single length is $O(N^{19})$ and both reference layers together cost $O(DN^{19})$. Equation (VI.7) summed over $j < t$ gives $O(DN^A t^{A+2})$ for I, with A the exponent in the count of the systems. Together with §3.5 this proves a uniform bound $(Nt)^{C_*}$ for the entire odd plus contribution from the stated deformation and automorphic estimates. The fixed choices display how all three comparisons apply on the same host and to the same positive twist range.

Adjoint rigidity of the irreducible rank-three points in I and V

Let σ be the rank-three representation on G_K at a type-I point, with $K = \mathbb{Q}(\sqrt{-7})$ in I and $K = \mathbb{Q}(i)$ in V. Write θ_K for the quadratic sign and A for the adjoint of the polarized deformation over $\mathbb{Q}$. The normal form at complex conjugation c has matrix part 1 and sign -1, with $\Psi(c) = \rho(-1)(-1)^j = 1$ for the host character ρ. If J is the antidiagonal identity matrix, the pairing is $\sigma(g)^t J \sigma(cgc^{-1}) = J$. Thus A is the full endomorphism space, with the conjugation action on G_K and the action $X \mapsto -JX^tJ^{-1}$ at c. The attachment required in both counts is $r_{\pi,\iota} \simeq \sigma(-1)$ for a regular algebraic conjugate self-dual cuspidal π

of $\mathrm{GL}_3(\mathbb{A}_K)$. The Tate twist leaves A unchanged and makes the polarized multiplier $\theta_K \kappa_\ell^{-2}$, which has value -1 at c. After killing the finite inertia component of ρ, the ordinary diagonal has cyclotomic exponents $(j, 0, -j)$. With $\mathrm{HT}(\kappa_\ell) = -1$, the weights of $\sigma(-1)$ are $(1-j, 1, 1+j)$ at each embedding, hence $(0, 1, 2)$ in V. These weights and the symmetric pairing give the normalization of Newton–Thorne's adjoint. I attachment · V attachment · Polarized action · Newton–Thorne §1

Image at the specified coefficient prime. In the marked deformation frame σ is absolutely irreducible on G_K by type I. The pinned open-subgroup lemmas use the local flag and the diagonal residual trace to upgrade this to absolute irreducibility on every open subgroup of G_K. To see the test, let U be the normal core of an open subgroup. A power of a local inertia element with cyclotomic value $1 + \ell$ lies in U, kills the finite character and has three distinct eigenvalues. Any U-stable proper nonzero subspace is a sum of its eigenlines. If it is a plane, its intersection with a distinct G_K-translate gives a U-stable line, since a G_K-stable plane would contradict type I. The G_K-orbit of the line then consists of all three eigenlines, so the global matrices permute them transitively. The local decomposition group preserves the ordinary flag, which is spanned successively by these eigenlines in their flag order. It therefore acts diagonally in this basis. Take a global element g inducing a three-cycle and a local element x with $u = \overline{\Psi}(x)$ and $u^2 \neq 1$. All three matrices of gx^n, for $n = 0, 1, 2$, have trace zero. Reduction in the original integral model with $v = \overline{\Psi}(g)$ would make the three distinct values vu^n roots of $T^2 + T + 1$, a contradiction. I open-subgroup lemma · V open-subgroup lemma

In I the square of the finite inertia component is trivial, including at $\ell = q, d$, and $\ell - 1 \nmid 2j$ supplies x on inertia. In V, both 3 and 7 are inert in K and prime to N. Take a local arithmetic Frobenius lift of norm ℓ^2 and adjust it by inertia to have $\kappa_\ell(x) = 1$. Cyclotomic inertia is surjective over the unramified local quadratic extension. The odd quadratic factors of ρ are trivial on ℓ^2. For dyadic conductor 2^b the exact order of the even dyadic factor at ℓ^2 is $2^{b - v_2(\ell^2 - 1)}$, since the character is faithful on the cyclic group $1 + 4\mathbb{Z}/2^b\mathbb{Z}$. Here $v_2(\ell^2 - 1)$ is 3 at 3 and 4 at 7, and $b \geq 14, 18$, respectively. Reduction preserves these orders. The root-counting argument thus applies even in characteristic 3. The same x also witnesses that the upper triangular ratios of the characteristic-zero diagonal characters differ from κ_ℓ in the V tangent test. Count I conductor and inertia conditions · V local element and inertia subgroup

The identity component of the Zariski closure of $\sigma(G_K)$ acts absolutely irreducibly, since its inverse image is open. Its unipotent radical fixes a nonzero subspace, so it is reductive, and its center acts by scalars. Its connected derived group is therefore irreducible in dimension three. The possibilities for its semisimple Lie algebra inside $\mathfrak{sl}_3$ are $\mathfrak{sl}_3$ and $\mathfrak{sl}_2$ acting by the symmetric square. Two simple factors would require dimension at least four for a faithful irreducible action. Both possibilities have torus elements with three distinct eigenvalues. The Zariski closure of the image of $G_{K(\mu_{\ell^\infty})}$ contains this derived group, since the cyclotomic kernel contains commutators and the commutator map respects Zariski density. The Tate twist is trivial on that kernel. After a finite coefficient enlargement to split the characteristic polynomials, Newton–Thorne **Lemma 2.28(2)** gives an enormous image in the sense of **Definition 2.23**. The argument applies over any finite extension of K by the same open-subgroup test. Newton–Thorne Definition 2.23 and Lemmas 2.22, 2.28

The finite tangent conditions. A dual-number deformation $(1 + \epsilon y)R$ of the polarized point R gives a cocycle with values in A. Away from ℓ, the fixed inertia conjugacy class gives an inertia coboundary by differentiation. Where the frame instead kills a finite-index subgroup of inertia, restriction and corestriction in characteristic zero give the same conclusion. At ℓ in I the map of the local deformation ring to the dual

numbers kills the specified local ideal. The crystalline comparison for that ideal, applied to the dual-number lift, gives crystallinity over the open subgroup killing the sign and the finite local character. In V, the frame gives a triangular dual-number lift over that subgroup with the fixed cyclotomic powers on inertia. The diagonal blocks are crystalline, since their remaining deformation is unramified. The graded Hom modules for extending the flag have positive cyclotomic exponents j or $2j$, with the ratio-κ_ℓ case excluded by x above. The Hom module in each extension step has $\mathrm{Fil}^0 = 0$ and $H^0 = H^2 = 0$, by the graded characters and local Tate duality. The local Euler characteristic and Bloch–Kato dimension formula identify H^1_f with H^1 for these extensions, so the dual-number lift is crystalline over that subgroup. In both cases its self-extension class becomes a coboundary over B_{cris}. Restriction and corestriction for the finite local extension give the Bloch–Kato condition over the original field. These are the unramified and crystalline kernels of Newton–Thorne §1. I local ideal and tangent comparison · I inertia and crystalline kernels · V inertia comparison · V crystalline tangent comparison

Vanishing given attachment. Newton–Thorne, *Adjoint Selmer groups of automorphic Galois representations of unitary type*, **Theorem 5.2** (also Theorem A for unitary type), now gives $H^1_f(\mathbb{Q}, A) = 0$. Finite extension of the coefficient field commutes with the local kernels and detects vanishing. If the polarized attachment is supplied over a finite CM extension L/K, the theorem gives vanishing over L^+ and restriction injects the original rational Selmer group there by corestriction. Neither step requires a degree prime to ℓ. The finite tangent cocycle is therefore a global coboundary, corresponding to an infinitesimal change of basis. This gives the type-I rigidity used in (VI.6) and the rigidity test of the V frame used in (VI.12). Newton–Thorne Theorem 5.2 · I cotangent test · V rigidity test

The first construction still needed for this deduction is a polarized cuspidal attachment of $\sigma(-1)$ at the prescribed coefficient place for each type-I point, including $\ell = 3, 7$ in V. The residual semisimplification is a sum of three characters, so restriction or coefficient enlargement leaves it reducible. In particular BLGGT Theorem 4.5.1, which assumes residual irreducibility on the mod-ℓ cyclotomic kernel, cannot provide the attachment directly. At 3 its numerical lower bound $\ell \geq 2(d+1)$ also fails even for residual constituent dimension $d = 1$. The source routes the attachment through an abelian CM field L containing K and then uses solvable automorphic descent. For descent of the attachment, BLGGT Lemma 2.2.2 requires the original polarized representation over K and irreducibility over L, supplied here by the open-subgroup test. The trace count further requires the local level witness over K. An attachment over an arbitrary finite CM extension suffices for Selmer vanishing, while the pinned descent and level comparison use the linked field and its specified local conditions. BLGGT Theorem 4.5.1 · BLGGT Lemma 2.2.2 · Source construction in I and V · Source descent and local conditions

For the trace count the auxiliary field must also carry a uniform collection of systems. To see the distinction, suppose a single finite Galois extension L/K works for all points under consideration, with the probes split completely and the upper levels controlled. At such a probe the upper and lower traces agree. A representation which stays irreducible on G_L has at most $|\mathrm{Gal}(L/K)|$ possible extensions up to isomorphism once its restriction is fixed. Indeed Schur's lemma makes the ratio of two extensions a character of the quotient. Thus a single Galois-stable list upstairs, counted with this multiplicity, can bound the required product downstairs whenever its size and the field degree meet the budget. Potential automorphy over fields depending on the point gives no such numerical conclusion by itself. For the linked field in the release the residual outer characters are paired by duality, with a marked element on which they differ. The source uses support on the prescribed ordinary Hecke module for the rank-three lifts and a

separate symplectic lifting theorem for Count II. Their precise settings and the solvable descent step are set out in the supplement. Higher-rank attachments and descent

How a rank-three seed reaches the other irreducible points

The Count I seed and its base-level support conditions are given in §3.6. For I and V the rank-three argument begins with a supported cuspidal adjoint point over the chosen linked field L. The projected Hecke quotient is finite and torsion-free over its weight algebra. The same support problem then occurs at characteristic-ℓ curves with the prescribed ordered flags. Seed and weight computation

One required characteristic-ℓ vanishing admits a short proof, also for $\ell = 3$. Let A be the matrix of the marked decomposition-group element on a residual curve over $k'[[t]]$, and u its cyclotomic value. The characteristic polynomial P of A reduces to $(T-c)^n$, with $c \neq 0$ and $\bar{u} \neq 1$. If $uAY = YA$, then $P(uA)Y = YP(A) = 0$. Modulo t, $\bar{u}A - cI$ is the sum of the nonzero scalar $(\bar{u}-1)c$ and a nilpotent matrix. Thus $P(uA)$ is invertible modulo t, hence over $k'[[t]]$. This kills Y, also over the fraction field. Marked curves and weight ideals

The support argument has a delicate step at a reducible curve. At adjoining marked curves the procedure completes the same finite patched module. At each completion the regular sequence of patch and weight parameters and freeness of the augmentation fibre give dimension $3[L^+ : \mathbb{Q}] + 1 + d_{\mathrm{patch}}$ for each minimal support component. The Galois deformation estimate at a suitable adjoining prime places its lift above a minimal support prime with irreducible generic determinant. Contract this support prime to the common ring and lift it to the next completion. Flatness after localization supplies a lift. Its residue-space fibre of the common module is nonzero because it is an extension of a nonzero vector space. Passing to a minimal support prime there recovers the dimension needed at the next junction. At the last curve the augmentation-fibre identity puts the desired irreducible prime on the original Hecke support. This uses a common patched module throughout the chain. Dimension, contraction and lifting on support

The local charts used in this argument connect adjacent components in dimension at least $3[L^+ : \mathbb{Q}] + 1$. The source proves the corresponding global propagation using the connected core of the framed space and the rank-two support on the reducible locus. Given these chart and tower estimates, an integral irreducible lift then has nonzero fibre in the finite Hecke module by Nakayama. Arithmetic classicality identifies its determinant and ordered inertia characters with a cuspidal specialization, giving the required weights $(1-j, 1, 1+j)$. Section 3.8 gives the integral comparison on the definite-unitary tower and states the spectral hypotheses for its classical evaluations. From local charts to each integral lift · Precise arithmetic inputs

The rank-four seed in Count II

Let $F = k_0$ be the selected real quadratic field, split at 2 and at the odd coefficient prime ℓ. Label the real embeddings by ι and the two places $\mathfrak{p}_\pm \mid \ell$. The restricted odd dyadic character ξ has primitive conductor $\mathfrak{n} = 2^{n_{\mathrm{II}}}\mathcal{O}_F$, since the dyadic completions are $\mathbb{Q}_2$. The tower gives $\bar{\xi}^2 \neq 1$ and $\bar{\xi}(\mathrm{Frob}_{\mathfrak{p}_\pm}) \neq 1$ at **both** places. The refined vector must have level $K_0(\mathfrak{n}\mathfrak{p}_+\mathfrak{p}_-)$ with the finite nebentype of ξ. The cohomological parameter is fixed at $w = k+1$ while the classical weights become $(k+1+2t_0, k+1)$, with $k, t_0 \geq 2$ divisible by $\ell - 1$. The desired representation r has determinant $\xi\kappa_\ell^k$ and residual semisimplification $1 \oplus \bar{\xi}$. Its ordered local requirements are the following. Detailed Hilbert seed calculation Seed predicate

Place	Inertial cyclotomic exponents, line then quotient	Full residual characters, line then quotient
$\mathfrak{p}_+$	$(k+t_0,-t_0)$	$(\bar{\xi},1)$
$\mathfrak{p}_-$	$(k,0)$	$(1,\bar{\xi})$

Parallel cuspidal congruence. For each unramified $\mathfrak{p}$, define the finite ideal character by $\eta(\mathfrak{p})=\iota(\xi(g_{\mathfrak{p}}^{-1}))$, with $g_{\mathfrak{p}}$ arithmetic Frobenius. The weight-one Hilbert series $E_1(1,\eta)$ has coefficients of positive index $\sum_{\mathfrak{d}\mid\mathfrak{m}}\eta(\mathfrak{d})$, using zero at ideals not coprime to the conductor. Stabilize at both places using $(1-V_{\mathfrak{p}_+})(1-\eta(\mathfrak{p}_-)V_{\mathfrak{p}_-})$, with V shifting ideal coefficients. This gives the ordered U-roots $\eta(\mathfrak{p}_+),1$. The first subtraction kills the constant terms at the standard narrow-class expansions and leaves the coefficient of the unit ideal equal to one. To raise the parallel weight, divide the level-one series $E_k(1,1)$ by its constant term $2^{-2}\zeta_F(1-k)$. The other coefficients are integral. Colmez's residue theorem and Deligne–Ribet interpolation give $v_\ell(\zeta_F(1-k))<0$ for sufficiently divisible positive k on the trivial branch. The regulator is nonzero because the split local logarithm of a generator of the real quadratic unit group is nonzero. The Euler factors at both ℓ-places are units. Thus the raising form is congruent to one at the standard expansions. The Hecke recurrences on the reduced coefficients of positive index agree at the two weights, since $(\mathrm{N}\mathfrak{q})^k\equiv 1 \bmod \ell$ at good primes and the level operators act by shifts. Deligne–Serre lifting on the saturated coefficient lattice retains both $U_{\mathfrak{p}}$ eigenvalues, the dyadic operators and the finite character. Dasgupta–Darmon–Pollack, Proposition 2.1 Colmez Deligne–Ribet Deligne–Serre, Lemma 6.11 Finite ideal character

A characteristic-zero Eisenstein system at this weight and level would have residual inducing characters $\{1,\bar{\xi}\}$ by good traces. At either split ℓ-place the local conductor exponents sum to at most one and the central character is unramified, so both inducing characters are unramified. The unit root at both places would be furnished by the same global inducing character, the one whose root lacks the factor ℓ^k. The two prescribed residual roots and the nontriviality tests at both places exclude this. The lifted packet is cuspidal. Its dyadic U-eigenvalues are units, by reduction to one. This argument and the construction of the parallel packet apply at $\ell=3$. Dasgupta–Darmon–Pollack, Proposition 2.1 Colmez Deligne–Ribet Deligne–Serre, Lemma 6.11

Mixed weight at fixed level. Write $b=k+1$, $m=b-2$, and $t_{v_+}=t_0,\ t_{v_-}=0$ for the two places over ℓ. Every positive multiple t_0 of $\ell-1$ gives a mixed cuspidal eigenpacket at the original $K_1(\mathfrak{n}\mathfrak{p}_+\mathfrak{p}_-)$ with fixed finite-order central character ω and no extra finite character at either ℓ-place. Here K_1 means bottom row $(0,1)$ modulo the level. Equivalently, with fixed ω one may use K_0 with the prescribed lower-right-entry character at the dyadic places. At $v\mid\ell$, $I_v=K_0(v)=Z(\mathcal{O}_v)K_1(v)$ and ω_v is unramified, so these vectors are Iwahori invariant there. Prescribed level and scalar character Detailed ordinary calculation

Take the quaternion algebra D/F ramified at the two real places and split at every finite place. Classical Jacquet–Langlands preserves the finite local representations and sends holomorphic weight $b+2t_v$ to the definite type of symmetric degree $m+2t_v$ on the norm-one group. For an integer ring O of a sufficiently large finite coefficient field, let $\rho_{\mathbf{t}}$ act on

$$M_{\mathbf{t}}=\bigotimes_{v\mid\ell}\det^{-t_v}\operatorname{Sym}^{m+2t_v}O^2.$$

Each local scalar acts with power m. This is the Johansson–Newton algebraic coefficient with common central weight m. At infinity the algebraic type is normalized by $\mathrm{nrd}^{-m/2}$, so its real scalar character is sgn^b. If $f_\infty(g)$ is the vector in the real-type model transported by ι^{-1}, passing to $f(g) = |\operatorname{nrd} g|_f^{-m/2}\rho_{\mathfrak{t}}(g_\ell)^{-1}f_\infty(g)$ makes it left $D^\times(F)$-invariant. Use the product K of the original tame K_1 and the two Iwahori groups. In this model impose

$$f(guz) = \Theta(z)\rho_{\mathfrak{t}}(u_\ell)^{-1}f(g), \qquad \Theta(z) = \omega_f(z)|z|_f^{-m}\prod_{v\mid\ell} z_v^{-m},$$

for $u \in K$ and finite central ideles z. Here $|\cdot|_f$ is the product of the finite normalized absolute values and character values are read in the chosen coefficient field. The norm factor cancels the valuation of the ℓ-adic power, so Θ is unit-valued. The product formula and $\omega_\infty(-1) = (-1)^b$ at both real places give $\Theta|_{F^\times} = 1$. On $K \cap Z_f$ it agrees with the prescribed right action, using the scalar coefficient power and triviality of ω on scalar units in K. Thus the full central character is fixed integrally, including its unramified values. Jacquet–Langlands, Chapter XVI Johansson–Newton, §§2.1–2.4

Evaluation gives fibers over the finite set $X = D^\times(F)\backslash D^\times(\mathbb{A}_{F,f})/(KZ_f)$. For a representative g, the stabilizer is $(D^\times(F) \cap gKZ_fg^{-1})/F^\times$. It is finite because the rational projective group is discrete in the adelic adjoint representation, with compact local bounds here, including at both real places. At a split dyadic place it injects, after conjugation, into the projective image of $K_1(2^a)$, with $a = n_{\mathrm{II}} \geq 1$. The upper Borel over $\mathbb{F}_2$ has order two and the principal congruence kernel is pro-2, so that projective image is pro-2. Each finite stabilizer is therefore a 2-group. Averaging its action on the fiber, including the central scalar, is integral at every odd ℓ. The resulting lattice of functions is finite free and its reduction is the full function space for the reduced coefficients and characters. Norm-two stabilizer condition

In the pinned convention the raw operator h_b is the double-coset sum with prefactor $(\mathrm{N}v)^{b/2-1} = (\mathrm{N}v)^{m/2}$, with b fixed. On passing to the algebraic function model, each right translate contributes $|\operatorname{nrd} h|_f^{m/2}$, which cancels that prefactor on a determinant-uniformizer coset. At $v \mid \ell$ the remaining ordinary normalization is $\ell^{t_v}h_b$. For the upper Iwahori use the pinned representatives

$$h_v(x) = \begin{pmatrix} \ell & x \\ 0 & 1 \end{pmatrix}, \qquad x \bmod \ell.$$

The K_1 and Iwahori sums use the same representatives on our invariant vectors. The term acting on the coefficient is $\ell^{t_v}\det(h_v(x))^{-t_v}\operatorname{Sym}^{m+2t_v}(h_v(x))$. The determinant factor cancels exactly. More generally for $h \in I_v\operatorname{diag}(\ell^r, 1)I_v$ with $r > 0$, the integral factor is $(\ell^{-r}\det h)^{-t_v}$. Modulo ℓ each such h has zero first column and lower-right entry $\bar{d} \neq 0$. Its symmetric power factors through the quotient line generated by $e_2^{m+2t_v}$, and on that quotient the normalized action reduces to $\bar{d}^m$. The Iwahori action on the quotient also reduces to $\bar{d}^m$, since its extra unit powers have exponents divisible by $\ell - 1$. This uses $F_v = \mathbb{Q}_\ell$. In the pinned coordinates the image of an individual representative is spanned by $(\bar{x}e_1 + e_2)^{m+2t_v}$. Pinned double coset and raw operator Weight normalization Johansson–Newton, Proposition 4.1.1

Let C_{t_0} be the reduction of the definite function lattice. Project in each symmetric power to the coefficient of $e_2^{m+2t_v}$, which defines a map $p_{t_0} : C_{t_0} \to C$. For diagonal entries $(\bar{a}, \bar{d})$ of an Iwahori element modulo ℓ, this coefficient quotient has character $\bar{d}^m(\bar{d}/\bar{a})^{t_v} = \bar{d}^m$. Thus C is the same scalar-valued function space at every weight under consideration, with central law $\bar{\Theta}$ and right K-action through $\prod_{v\mid\ell}(\overline{u_{22,v}})^{-m}$. In

particular $\dim C \le |X|$. The coefficient projections are onto, and averaging over the finite 2-group on each fiber commutes with them and preserves surjectivity, including at $\ell = 3$.

The map p_{t_0} intertwines all chosen Hecke operators. At either ℓ-place the normalized representative has bottom-right entry 1, so projection after its coefficient action equals projection before it. Away from ℓ, the translates have identity coefficient component at both ℓ-places and the same central law. Write $U = \bar{U}_{v_+}\bar{U}_{v_-}$ for the product of the two normalized residual operators. Each summand of U acts through the tensor of the two coefficient quotients on the function value at its input, so U kills $\ker p_{t_0}$. This holds on equivariant functions before choosing representatives of X.

For sufficiently large N, $\mathrm{im}(U^N)$ has stabilized on each space, and U acts invertibly on this image. The projection maps the stabilized image in C_{t_0} onto that in C. It is injective there because U kills its kernel. Since the two operators commute, their product is invertible precisely on the summand where both are invertible. The resulting isomorphism intertwines the full Hecke action, including on generalized simultaneous eigenspaces over an algebraic closure of the residue field. A primitive lattice vector for the parallel packet has nonzero reduction in this ordinary summand. Thus the generalized ordinary multiplicity d_{t_0} of its full residual system satisfies

$$1 \le d_{t_0} = d_0 \le |X|.$$

Equivalently, d_{t_0} is the sum of the generic-fiber multiplicities of the systems reducing to this residual system. After extending coefficients to split the generic action, intersect a common invariant flag with the lattice. The intersections are saturated, so the one-dimensional graded pieces reduce with the corresponding residual characters. Their number in the chosen residual system is the dimension of its generalized eigenspace. The isomorphism is at the residual level, and the sum can include several distinct characteristic-zero systems. Independent trace proof with higher congruence precision

Apply Deligne–Serre lifting to the commutative algebra of all these operators on the mixed finite free lattice. The algebra is finite and torsion-free over O as a submodule of its endomorphism algebra, so the simultaneous residual character lifts through a characteristic-zero component after extending coefficients. This gives one eigensystem with $\ell^{t_v}\lambda'_v \equiv \lambda_v$ at both $v \mid \ell$ and $\lambda'_v \equiv \lambda_v$ away from ℓ, where λ'_v is the raw h_b eigenvalue. Both dyadic unit roots survive. All symmetric degrees are positive, which excludes constituents factoring through reduced norm. Reverse Jacquet–Langlands and select the holomorphic real lines to obtain the mixed cuspidal form with the prescribed finite types and refined level. This proves the fixed-character weight change for the Count II field, including at $\ell = 3$. Deligne–Serre, Lemma 6.11

Local readout given the mixed eigenpacket. The source first attaches r_1 with determinant $\xi\kappa_\ell^{-k}$ and then sets $r = r_1 \otimes \kappa_\ell^k$. Put $t_{\mathfrak{p}_+} = t_0$ and $t_{\mathfrak{p}_-} = 0$. Before twisting, the stable line has exponent $t_{\mathfrak{p}}$ and the quotient has exponent $-k - t_{\mathfrak{p}}$. After twisting these exponents become $k + t_{\mathfrak{p}}$ and $-t_{\mathfrak{p}}$. Use $\mathrm{HT}(\kappa_\ell) = -1$. The pinned ordinary comparison reads the **stable line**: after removing its cyclotomic power from r_1, the geometric Frobenius value is the normalized operator $\iota^{-1}(\ell^{t_{\mathfrak{p}}}\lambda'_{\mathfrak{p}})$, where $\lambda'_{\mathfrak{p}}$ is the raw eigenvalue. The normalized operator reduces to $\xi(\mathrm{Frob}_{\mathfrak{p}_+}^{-1}), 1$ modulo the coefficient maximal ideal in the geometric convention. Inertia, with trivial extra finite character and $\ell - 1 \mid k, t_0$, and arithmetic Frobenius then determine the line reductions $\bar{\xi}, 1$ on the full decomposition groups. The determinant gives the opposite quotient reductions in the table. Only the Galois comparison for the classical point enters this readout. Ordinary comparison Tate twist Stable-line identity

At a dyadic place the primitive central conductor and the level give the same lower and upper bound n_{Π} on the local conductor exponent. The unit U on that newvector gives a nontrivial local standard L-factor, forcing a principal series with one unramified character and one of conductor exponent n_{Π}. Local–global compatibility gives inertia $1 \oplus \xi$ and zero monodromy. Away from 2ℓ the representation is unramified. At $\mathfrak{p}_{\pm}$, Iwahori level allows spherical principal series or an unramified twist of Steinberg. The ordinary comparison gives Frobenius slope gap $k + 2t_{\mathfrak{p}} \geq 2$, excluding Steinberg, which has gap one. Hence the exponent-one level at ℓ is the level of the **refined vector** on a spherical representation. These deductions use the classical local attachment and give the prescribed principal-series local conditions used in the induction. Casselman BLGGT, Theorem 2.1.1

Absolute irreducibility of the cuspidal Hilbert attachment upgrades to irreducibility on open subgroups. If an open normal core preserves a line, an inertia element with distinct eigenvalues forces it to preserve at most two lines. Global irreducibility then gives two lines interchanged by the group, and hence a quadratic self-twist δ. Reduction of the twist gives $\{1, \bar{\xi}\} = \{\bar{\delta}, \bar{\delta}\bar{\xi}\}$. Since ℓ is odd, this would imply $\bar{\xi}^2 = 1$. Thus the open-subgroup test also works at 3. Saito, §2, citing Taylor, Proposition 3.1

For $s \in G_{\mathbb{Q}} \setminus G_{k_0}$ representing the marked coset, the induced representation restricts to $r \oplus r^s$ over k_0. In the induced basis s acts by $\left(\begin{smallmatrix} 0 & r(s^2) \\ I & 0 \end{smallmatrix}\right)$. Since ℓ is odd and the two residual eigenvalues of $r(s^2)$ are distinct, the eigenspace projectors are integral in the selected lattice. The change to the four marked lines exhibits the required alternative of residual entries. Before its common twist the four exponents in decreasing order are $k + t_0, k, 0, -t_0$. The two inducing planes remain inequivalent on the relevant CM field. An isomorphism up to a character on the cyclotomic quotient would, by equality of the determinants, have a quadratic twist. Squaring a suitable inertia element with infinite cyclotomic image would then equate the two distinct pairs of powers above. Given the irreducibility of the two restrictions, Clifford theory gives irreducibility of the induction, and quadratic automorphic induction after base change gives a cuspidal rank-four seed. The target point at twist j has exponents $(j+1, j, 1, 0)$. The passage from the seed to all of those target points uses a projected rank-four Hecke module, the marked residual support estimate, and classicality at the target weights, using the ordered Hecke comparison of §3.8. Residual entries and symplectic lattice · Conditions on the projected module and descent

3.8. Classical local–global compatibility for the linked ordinary towers

Integral arithmetic specialization. Fix one Count II link on the definite unitary group over L^+, with CM field L. Write A for the complete local ordered inertia algebra over the coefficient DVR O. At each ℓ-place use the marked split GL_4 coordinate. Fix the away-from-ℓ level U, either at the base or at a prescribed finite patch level, and keep the integral tame coefficient representation τ_0 at the split tame places, including the full dyadic type. Its characteristic-zero type test is the prescribed Teichmüller inertia tuple. On forms it is imposed by equivariance. One can work at the kernel of τ_0 with that action of the finite quotient, without dividing by its order. The auxiliary split place v_0 has principal level one and lies over an odd rational prime $p_0 \neq \ell$ unramified in L, with residue degree one. These are the hypotheses for the stabilizer and finite-level ordinary module used below. Integral tame type · Finite-level setting and character · Auxiliary prime

The stabilizers in the double-coset description are trivial. For a rational stabilizer, its powers are bounded at every finite place by the compact levels and at infinity by definiteness. Their entries lie in a single fractional lattice with bounded archimedean coordinates, hence the matrix has finite order. At v_0 it is conjugate into $1 + \mathfrak{p}_{w_0} M_4(O_{L_{w_0}})$, a torsion-free group by the matrix logarithm over the unramified completion at the odd

prime p_0. This gives arbitrary integral coefficient values at the finite set of representatives. At $\ell = 3$ the proof uses the same principal level and integral coset sums. No averaging over a potentially noninvertible stabilizer or type quotient is required. Trivial stabilizers

Let Θ be the weight character on the diagonal units. With inverse-Artin coordinate y for an inertia element g at w, the i-th diagonal factor is $\Theta(e_{w,i}(y)) = (\psi_{w,i}(g)\kappa_\ell(g)^i)^{-1}$, for $0 \le i \le 3$. At an arithmetic map $x : A \to O'$ into an enlarged coefficient DVR, write $x\Theta = \phi\chi_\lambda$, with χ_λ the lowest character of the algebraic coefficient ξ_λ and ϕ finite. The smooth right translation character is ϕ^{-1}. In particular the scalar-unit law includes $\kappa_\ell(g)^6 \prod_i x\psi_{w,i}(g)$ in the denominator of $x\Theta(yI_4)$. The tame scalar action is included in τ_0, and the full determinant step above ℓ is the normalized central uniformizer translation. For the target exponents $(j+1, j, 1, 0)$ and common linked cyclotomic shift b_v, the lowest tuple $h_{v,i} - i$ is $(-j-1-b_v, -j-1-b_v, -3-b_v, -3-b_v)$, nondecreasing for $j \ge 2$. The seed exponents $(k+t_0, k, 0, -t_0)$ also give a dominant algebraic coefficient in the dual convention after the common shifts. Ordered weight and residual characters · Arithmetic compatibility · Coefficient weights

At $w \mid \ell$, let N be the integral upper unipotents and N_c^- the lower unipotents congruent to 1 modulo $\mathfrak{p}_w^{c+1}$. Put

$$z_w = \operatorname{diag}(\pi_w^3, \pi_w^2, \pi_w, 1), \qquad U_a = \sum_{u \in N/(aNa^{-1})} R(ua)$$

for a diagonal a with nonincreasing uniformizer exponents. The coset sums compose as $U_aU_b = U_{ab}$. In factoring yu for $y \in N_c^-$ in upper, diagonal and lower order, the resulting diagonal is congruent to 1 at depth $c+1$. The permutation of the upper cosets leaves the lower factor conjugated by a^{-1}. For $a = z_w$, conjugation gains π_w^{i-j} in each entry $i > j$. Thus U_{z_w} lowers depth $c+1$ to c once the smooth character is trivial on the depth-$c+1$ diagonal units. The three proper step diagonals have product z_w, so all their normalized operators are invertible on the ordinary part. The fourth step is central and invertible as well. Contracting and step diagonals · Coset multiplication · Level lowering

Let C_c be the classical lattice at this U, τ_0, ϕ and at depth c at least the coefficient-lattice bound, with ϕ trivial on diagonal units congruent to 1 modulo $\mathfrak{p}_w^{c+1}$ at every $w \mid \ell$. It is integral in the untwisted coordinates $\widetilde{F}(g) = \xi_\lambda(g_\ell)^{-1}F(g)$. The control map takes the lowest-weight coordinate in the integral weight basis, keeping the τ_0 index. Upper unipotents fix this coordinate functional. In a normalized translate the coefficient matrix is $\xi_\lambda(u)\chi_\lambda(a)^{-1}\xi_\lambda(a)$. It induces the identity on the lowest coordinate and, for $a = z_w$, its columns away from the lowest line are divisible by the uniformizer of O' in the local coefficient factor. Indeed the difference from the lowest weight is a nonzero sum of positive roots, each pairing positively with $(3, 2, 1, 0)$. Use the product over the marked ℓ-places and all labeled embeddings. Each nonlowest tensor column is divisible in at least one factor, and a scalar factor has only its lowest coordinate. Thus the product sends an integral form whose control value is divisible to a form divisible in every coordinate. The map intertwines all the step operators and good spherical operators, and its torus law is $(\phi\chi_\lambda)^{-1}$. Integral weight basis · Operator intertwining · Filtration contraction

On the ordinary summand, where the contracting product is invertible, this divisibility gives injectivity modulo the uniformizer and then integrally. For surjectivity modulo the uniformizer, go to a depth $c' \ge c$ at which both reduced actions of the lower unipotents are trivial. Prescribe the scalar ordinary values in the lowest coordinate at the finite representatives, with zero in the other coordinates. Equivariance extends these to an integral classical form with the prescribed reduced control values. A sufficiently large power of

the contracting product moves it back to depth c. Choose the power to be a multiple of the period of the scalar ordinary vector modulo the uniformizer, so the control values still agree. Eventual factorial powers give the ordinary projection and preserve that agreement. Nakayama gives integral surjectivity. With finite-level ordinary base change ordBC and freeness $\mathrm{ordStruct}$, this gives the Hecke-equivariant isomorphism

$$e_z^{\mathrm{alg}} C_c \;\simeq\; O' \otimes_{A,x} S_A^{\mathrm{ord}}(U, \tau_0, \Theta).$$

The ordinary lattice is the same under increase of admissible depth by the lowering identity. The comparison applies even when smooth characters coincide. Integral approximation · Ordinary freeness · Integral base change · Arithmetic specialization

Residual projector and characteristic-zero fiber. In the finite commutative image of the chosen operators over A, take the Henselian idempotent e_a for the residual profile a, with zero if that profile is absent. The profile includes the good split spherical coefficients and the ordered cumulative unit values at ℓ, including the cyclotomic corrections to $\psi_{w,i}$. Over a finite field of size Q splitting the residual characters, a finite list of operators separates the desired profile from the others. The product $\prod_i (1 - (\bar{H}_i - \bar{a}_i)^{Q-1})$ has residual values one and zero on the selected and unselected characters. A sufficiently large factorial power gives the idempotent in the residual image algebra, even for nonsemisimple actions. Every residual character of the specialized image pulls back to a residual character of the original algebra with the same values. By uniqueness of idempotent lifting, the integral control map therefore restricts to e_a on both sides. The projected module $M_U = e_a S_A^{\mathrm{ord}}$ is finite free over A. Every characteristic-zero Hecke character of its image algebra above x has a nonzero generalized eigenspace on the classical weight fiber by faithfulness and Nakayama, and hence admits a simultaneous eigenvector. Here the conclusion is a vector on the definite unitary group with the prescribed level, type and normalized ordinary eigenvalues. Integral idempotent · Projected freeness

At a fixed algebraic weight, finite-order variations in the free local unit directions give separating arithmetic evaluations on A. Distinct sufficiently large ℓ-power orders in one direction for each diagonal index at each place give distinct smooth characters and preserve the residual profile, also at 3. Weierstrass preparation applied in each variable gives separation on this grid. Freeness then detects a nonzero vector by one nonzero coordinate, and detects a nonzero operator by its action on a vector. The tame and patch levels are fixed throughout, while the admissible ℓ-depth may depend on the smooth character. Density at a fixed algebraic weight Detailed reader proof and pinned source calls

The cuspidal Galois row. Fix $\iota : \overline{\mathbb{Q}}_\ell \simeq \mathbb{C}$. For a polarized regular algebraic cuspidal π of rank a over a CM field E, use $\pi^c \simeq \pi^\vee \otimes (\chi \circ N_{E/E^+})$ with $\chi_v(-1) = (-1)^a$ at every real place. The required local normalization is $\mathrm{rec}_w(\pi_w |\det|_w^{(1-a)/2})$, with a uniformizer corresponding to geometric Frobenius. Its labeled Hodge–Tate numbers are $(a-1)/2 - z$ for infinitesimal-character entries z. Barnet-Lamb–Gee–Geraghty–Taylor, Theorem 2.1.1, gives those lists, spherical crystallinity above ℓ and Frobenius-semisimplified Weil–Deligne compatibility away from ℓ, retaining monodromy. This proves the two away-from-ℓ tests of the pinned cuspidal row. The toral test has zero monodromy. In the large tame test $E_w = \mathbb{Q}_p$, with $a = 1$ or $p > 16a^2$, a repeated eigenvalue of a tame generator must be a $(q_w - 1)$-st root of unity. The full comparison gives its inertia polynomials and zero monodromy. At $w \mid \ell$, Theorem 1.1 of BLGGT II gives equality of semisimplified Weil representations and a bound on the monodromy partitions in general. If the automorphic monodromy is zero the bound forces Galois monodromy zero, even in even rank without Shin regularity. Rank one follows from reciprocity. Caraiani's full result applies in the conjugate-self-dual case, or after the algebraic character twist of Clozel–Harris–Taylor, Lemma 4.1.4, followed by

untwisting. These results include $\ell = 3$ and concern the characteristic-zero cuspidal attachment. Cuspidal row BLGGT, Theorem 2.1.1 BLGGT II, Theorem 1.1 Caraiani

Constituents of the projected classical space. The linked unitary group is compact at infinity, has hyperspecial level at nonsplit places of the unramified CM extension, and uses the marked general linear coordinate at each split place. Every contributing constituent in its classical evaluations has a standard transfer $\boxplus_b P_b \boxtimes [d_b]$. Here P_b is unitary conjugate-self-dual cuspidal of rank a_b and $\sum_b a_b d_b = n$. Write $t_{b,k} = (d_b - 1)/2 - k$ and $s_{b,k} = t_{b,k} + (a_b - n)/2$ for $0 \le k < d_b$. The regular infinitesimal-character list is the union of the strings of P_b shifted by $t_{b,k}$. Shifting by $s_{b,k}$ instead makes $P_b |\det|^{s_{b,k}}$ regular algebraic in rank a_b. Its polarizing character on E^+ is $|\cdot|^{2s_{b,k}} \eta_{E/E^+}^{a_b}$, of sign $(-1)^{a_b}$. Apply the cuspidal row to each of these regular algebraic blocks. The direct sum has at each split place the block exponent $s_{b,k} + (1 - a_b)/2 = t_{b,k} + (1 - n)/2$. Its Weil–Deligne monodromy is the sum of the cuspidal block monodromies. In particular a string of shifted trivial characters has zero monodromy. Reversing the strings at conjugate embeddings gives the polarization of the assembled rank-n representation with multiplier κ_ℓ^{1-n} in its conjugate-duality relation. The linked comparison character is a separate common algebraic twist of the ordered pair and multiplier. It shifts the labeled weights and twists the entire local parameter. Its central and finite-character data are the ones fixed by the link. No independent change of twist at a patch place is used. Assembly of the classical row Block calculation and common twist

The release supplies this transfer for every irreducible constituent σ at the base level or a fixed auxiliary level. In RunFin, run the comparison with the empty split test or with a single prescribed split place w. In either run the finite candidate list has the compact real infinitesimal characters of σ and the split match almost everywhere. For every finite set of spherical nonsplit places, some candidate matches throughout that set. If each candidate failed somewhere, one failing place per candidate would give a finite set no candidate could match. Thus one can take a datum D from the empty run and a datum D_w from the run at w, both matching at every spherical nonsplit place. They have the same Satake data almost everywhere by $\mathrm{FullSat}$ and are regular by the weight comparison. The pinned $\mathrm{LocRigid}$ identifies their blocks at all finite places up to permutation, so $\mathrm{PermSub}$ transports the matching split subquotient of D_w back to D. This proves the split comparison for one family at every w, together with the inert spherical comparison and the real weight of D. The pinned $\mathrm{baseChangeN}$ supplies $\mathrm{BaseChangeIsobaricN}$ to the ordinary construction. In GuWrap, irreducibility of normalized products of unitary local factors upgrades the subquotient to the full product, and the Speh description of the discrete general linear factors gives BcSpeh. Finite candidates and local rigidity Pinned base change Speh assembly Proof for the base and auxiliary levels

Good arithmetic weights and Speh strings. At one marked split $w \mid \ell$, write $k = L_w$ and let $\eta_1, \dots, \eta_4$ be the inverses of the smooth weight characters on $\mathcal{O}_k^\times$. At a good weight they are distinct on the full unit group, with arbitrarily high conductor allowed. The ordinary Jacquet comparison places each local component of the selected eigenspace in the torus Bernstein class determined by their unordered multiset. Let $z_{\mathfrak{s}}$ be that class's central idempotent over $\mathbb{C}$. For a compact open J fixing a selected vector, $z_{\mathfrak{s}} * e_J$ is a compactly supported Hecke test and fixes the vector. It preserves the selected refinement since $z_{\mathfrak{s}}$ acts as the identity on its whole smooth block. This projector is defined at the individual weight. The integral residual projector e_a continues to specify the projected module. Pinned ordinary Jacquet comparison Roche, §7 and Theorem 7.2 Bezrukavnikov–Kazhdan–Varshavsky, §1.3

At matching split transfer, $P_b[d_b]$ has the supercuspidal support of $P_{b,w}$ repeated at each of the unramified shifts $|\det|^{(d_b-1)/2-j}$, $0 \le j < d_b$. Induction takes the union with multiplicity, and subquotients retain this support. A supercuspidal factor of rank greater than one gives a nontoral support Levi. If all factors are characters, $d_b > 1$ repeats the same unit restrictions. The modulus and rank normalizations are unramified. An auxiliary determinant twist in an endoscopic embedding is common to the copies within one string, so it preserves the repetition. The regular torus projector therefore acts by zero on the whole matched induced representation, including when the trace inserts an intertwiner. The argument also excludes a repeated copy of the same length-one global block. KMSW, §§1.2.2 and 1.3.2 Zelevinsky, §1 and Theorem 4.2

The parameter-refined trace at the marked place. Fix a weak standard parameter Ψ, indexed by its almost-everywhere Satake data and archimedean norm as in KMSW Corollary 3.3.2. Put $v = w|_{L^+}$ for the marked place and take $f_v = z_{\mathfrak{s}} * e_J$. Hold the other factors f^v fixed. Take compact coefficient projectors at infinity, the normalized hyperspecial units at nonsplit places and at good unmarked places, and compact tests fixing the chosen vectors at the remaining split places. These include the kernels or isotypic projectors for the tame types, principal level one at v_0, any other coefficient places with level, and the actual auxiliary patch levels. The set removed from the unramified Hecke refinement contains all these exceptional places and infinity. Transfer uses the chosen pure inner twist of the definite Hermitian group and hyperspecial normalization at the nonsplit places. KMSW, (3.3.1)–(3.3.7) and Corollary 3.3.2 Mok, Corollary 10.0.2 and Proposition 10.0.3

At a split place the proper elliptic datum $H = U(a) \times U(4-a)$ is the block Levi of $\mathrm{GL}_4(L_w)$, and the identity datum uses the same function in the common coordinate. For H take normalized parabolic descent, with compact conjugation average and half-modulus. If the embedding twists the Levi representation by the determinant character ξ_H, multiply the descended test by ξ_H in the convention $\pi(f) = \int f(g)\pi(g)\,dg$. Its trace on a Levi representation is then the trace of the original test on normalized induction of the twisted factors. The matching orbital-integral identity uses the relative discriminant. At the other finite places the same rule applies to split tests, and hyperspecial transfer applies at the nonsplit places. At infinity use the smooth transfer of the compact coefficient projector. KMSW, Proposition 1.5.1, §1.6 and Lemma 1.6.5 Mok, Theorem 3.2.1(a)

Definiteness leaves only the full group in the rational Levi expansion (3.3.1). In (3.3.2), the stable multiplicity formula (3.3.6)–(3.3.7) gives terms only for endoscopic parameters mapping to the same Ψ. KMSW Proposition 1.5.1 and (1.2.4) identify their stable traces at the split place with general linear characters. After the twists and normalized induction, every surviving term uses the character of the same local representation $\pi_{\Psi,v}$ of the composed parameter. For arbitrary h_v in the local Hecke algebra this gives

$$\operatorname{tr} R_{\mathrm{disc},\Psi}(h_v f^v) = C_\Psi(f^v)\operatorname{tr}\pi_{\Psi,v}(h_v).$$

With the other factors projectors, a selected constituent contributes its local character with positive coefficient, and the trace is finite at each compact level. Independence of characters identifies this local constituent with $\pi_{\Psi,v}$. This is Mok Proposition 10.0.3 for elliptic weak indices. Anisotropy removes the proper rational Levi terms for other indices as well. Varying the split place with Ψ fixed gives the same comparison at the tame and patch places. At the marked place the support calculation makes every term on the right zero for $d_b > 1$, whereas a selected vector gives positive trace on the left for $h_v = z_{\mathfrak{s}} * e_J$. KMSW, (3.3.1)–(3.3.7) and Corollary 3.3.2 Mok, Corollary 10.0.2 and Proposition 10.0.3

The trace derivation uses connected-group stabilization and matching transfer in KMSW §3.1, the split character identity, and the quasi-split stable multiplicity formula of Mok Theorem 5.1.2, imported in KMSW Proposition 3.3.1. Mok assumes twisted stabilization in Hypothesis 4.2.1. Atobe et al. §§0.1 and 0.4 record the general twisted weighted fundamental lemma as a premise of that framework. The argument uses the refined distributions for non-generic indices in the length test. After the test, KMSW Theorem* 1.7.1 and Theorem 5.0.5 apply in their generic pure-inner scope. Their local classification hypothesis (3.6.3) and the intertwining relation of Theorem 2.6.2 are established there in that scope. This gives the real infinitesimal-character comparison and the inert hyperspecial comparison, via restriction of the unramified parameter to W_{L_w}. On the pinned route those comparisons are part of BcSpeh. KMSW, §3.1 and Chapter 5 Mok, Hypothesis 4.2.1 and Theorem 5.1.2 Atobe et al., §§0.1 and 0.4

Length-one constituents and their attachments. The surviving generic discrete parameters have standard transfer $\boxplus_b P_b[1]$ with distinct unitary conjugate-self-dual cuspidal blocks of the standard rank-four parity, conjugate-symplectic for the embedding with $\chi_+ = 1$. Proper endoscopic sums can contribute when their combined split support has four distinct unit restrictions after transfer and their packet members satisfy the multiplicity condition and the prescribed levels. Each P_b is individually cuspidal. With several blocks the standard transfer is an isobaric sum. The real, split and inert comparisons above are the inputs to the Galois assembly on this same fixed level. The rank shift $P_b|\det|^{(a_b-4)/2}$ in the classical row above then gives regular algebraic polarized cusps, with geometric normalization $P_{b,w}|\det|^{-3/2}$ for each block. Apply rank-one reciprocity and BLGGT Theorem 2.1.1 for the higher ranks, and take the direct sum. The split Hecke polynomials come from these same local components, and the unramified factors give the representation of the Galois quotient for that level. At a good coefficient place the matching is toral with zero monodromy, so the monodromy bound in BLGGT II Theorem 1.1 gives equality there, including at $\ell = 3$. The ordered flag requires the separate ordinary comparison below. If the attached semisimple rank-four Galois representation is irreducible, equality of the determinant forces a single block, giving a cuspidal standard transfer. KMSW, Chapter 5 BLGGT II, Theorem 1.1 Good-weight Speh test and source anchors

The refined local identity. At a split $w \nmid \ell$ let the marked geometric Frobenius and pro-ℓ tame inertia generator have inverse-Artin coordinates x_w, δ_w, respectively a uniformizer and a unit. Write q_w for the residue size. For a generalized torus eigenspace of character φ on the prescribed Iwahori-level vectors, the geometric lemma gives the normalized inducing characters $\varphi(e_i(z))|z|_w^{(n-1)/2-i}$, with coordinates numbered from zero. The common rank normalization therefore yields

$$\alpha_i = q_w^i \varphi(e_i(x_w)), \qquad \beta_i = \varphi(e_i(\delta_w)).$$

If the local representation is spherical, its Frobenius polynomial is $\prod_i (Z - \alpha_i)$, including repeated roots. When $\alpha_i \neq \alpha_j$ and $\alpha_i \neq q_w \alpha_j$ for all $i \neq j$, the principal series is unlinked. Its parameter is toral with zero monodromy and the distinct Frobenius eigenlines identify the inertia scalars β_i. Thus for the actual Frobenius and inertia matrices P, T the relation needed before inverting any eigenvalue difference is

$$\prod_{j \neq i} (P - \alpha_j I)(T - \beta_i I) = 0.$$

This is the split patch-prime calculation in $\mathrm{LocRelT}$ and $\mathrm{LevPtAlt.lg}$, over the actual split completion. In particular the large tame clause of the cuspidal row is a separate test. For a split spherical double coset the usual Satake factor $q_w^{j(n-j)/2}$ together with the polynomial coefficient factor $q_w^{j(j-1)/2}$ gives the same rank normalization. Pinned local relation Bernstein–Zelevinsky, §2.12 Zelevinsky, Theorem 4.2

The ordered ℓ-adic readout. With distinct smooth unit characters, the toral parameter at a split $w \mid \ell$ and the ordinary operators give the ordered flag. The coefficient convention is the dual of the prescribed highest weight, with labeled i-th Hodge–Tate number $h_{\tau,i} = \mathrm{lam}_{\tau,n-1-i} + i$. If the smooth characters are ω_i, their restrictions to units are inverse to the prescribed torus eigenspace characters. The unscaled j-step operator has eigenvalue $\prod_{i<j} \omega_i(x_w) q_w^{-i}$. In inverse-Artin coordinates the Galois flag character has value $\iota^{-1}(\omega_i(z)) \prod_\tau \tau(z)^{-h_{\tau,i}}$. The unit-normalized ordinary operators at the split places determine the uniformizer factors on the same eigensystem. To see the first line in the potentially semistable module, twist by its prescribed character. Its Weil line is separated by the distinct smooth characters, has Newton number zero, and all labeled Hodge numbers are nonnegative. Weak admissibility forces the induced Hodge numbers on the line to be zero. This gives a weakly admissible subobject, and induction on the quotient gives the ordered flag. This is a local argument at a classical specialization. The small coefficient prime 3 satisfies the same calculation once the prescribed distinct characters and the ordinary operators are supplied. Ordinary operators Ordered flag calculation

The determinant of the projected rank-four module

The Speh test reduces the spectral requirements for constructing the fixed-level determinant to length-one transfers on the separating good weights. Each contributing good eigensystem must still supply a continuous semisimple rank-four representation of the same Γ_U, with the prescribed inertia kernels at this level and all four good split Frobenius coefficients. Good-weight separation also requires scalar Hecke action on those classical fibers. Identifying the determinant with a classical representation at an arbitrary arithmetic weight uses its own pointwise attachment and Frobenius comparison. The ordered local and patch relations use their respective pointwise comparisons before separation. Pinned determinant hypothesis Good-weight interpolation

Assume the projected spectral transfer and the characteristic-zero Galois rows above at this level. Write $\mathbb{T}_U$ for the faithful commutative Hecke image on M_U and Γ_U for the CM subgroup of the corresponding Galois quotient unramified outside the designated base and patch places. For $M_U \neq 0$ the construction gives a continuous degree-four determinant D_U on the group algebra of Γ_U, with values in $\mathbb{T}_U$, and four ordered local characters $\psi_{w,i}$. The level's polarized ordinary determinant ring B_U maps onto $\mathbb{T}_U$, compatibly with the weight map from A. Its kernel J_U is $\mathrm{Ann}_{B_U} M_U$. When $M_U = 0$ take $J_U = B_U$. Determinant construction · Hecke quotient · Auxiliary levels

For $g \in \Gamma_U$ write $P_g(X) = D_U(X - g) = \sum_{j=0}^{4} (-1)^j c_j(g) X^{4-j}$ with $c_0 = 1$. At a good split place w outside the base and patch sets write F_w for geometric Frobenius, the inverse of the source's lift acting on the residue field by the q_w-power map. Let $T_{w,j}$ be the unscaled spherical coset of $\mathrm{diag}(\varpi_w I_j, I_{4-j})$. Then

$$P_{F_w}(X) = X^4 - T_{w,1} X^3 + q_w T_{w,2} X^2 - q_w^3 T_{w,3} X + q_w^6 T_{w,4}.$$

Indeed if the normalized Satake parameters are b_i, the coset has eigenvalue $q_w^{j(4-j)/2} e_j(b_1, \ldots, b_4)$. The local shift $|\det|^{-3/2}$ gives geometric eigenvalues $q_w^{3/2} b_i$, so their j-th symmetric function is $q_w^{j(j-1)/2} T_{w,j}$. These factors are units over O. All four coefficients are used integrally and the fourth recovers the central generator. Trace alone can lose the comparison at $\ell = 3$. For the generator of C_2, the representations $1^{\oplus 4}$ and $1 \oplus \mathrm{sgn}^{\oplus 3}$ have the same trace in characteristic three and different polynomials. At arithmetic Frobenius the roots are reciprocal and $c_j(F_w^{-1}) = c_{4-j}(F_w)/c_4(F_w)$. Spherical normalization · Geometric Frobenius · Integral generators

The linked target is $r|_{G_L} \otimes \lambda_{\mathrm{link}}$ for the representation r of the totally real frame. At the seed the character multiplying the roots also includes the parity factor $\beta^{n_{\mathrm{par}}}$ of the link, with n_{par} specified in the seed descent below. In general a twist by η multiplies $c_j(g)$ by $\eta(g)^j$ and $\psi_{w,i}(g)$ by $\eta(g)$. Returning from the linked comparison to the frame divides by the same factors. The formula above uses the linked system, with conjugate duality $r_y^c \simeq r_y^\vee \otimes \kappa_\ell^{-3}$ at a classical point y. Thus the roots at the conjugate good place are q_w^3 times the inverse roots at w. Root scaling · Polarization

At good arithmetic weights the projected classical decomposition is scalar for the Hecke image. If an element of $\mathbb{T}_U$ vanishes on all these characters, its action on each vector has all basis coordinates vanishing at every good weight. The integral control and weight separation above force those coordinates to vanish in A. This proves $\bigcap_{y\ \mathrm{good}} \ker y = 0$ in $\mathbb{T}_U$ and hence reducedness of the fixed-level image. To construct D_U, embed this compact image algebra as a closed subring of the product of the good evaluation fields. The polynomials of the classical representations in that product lie in the image at every good geometric Frobenius by the displayed formula. Those lifts, with all conjugates, are dense by Chebotarev for degree-one primes outside the excluded set. Continuity and the determinant descent identities then put the product determinant over $\mathbb{T}_U$. In particular this uses the polynomial-law identities on group algebras. At any arithmetic y the same good Frobenius identities give $yD_U = D(r_y)$ for the semisimple representation from its classical constituent. The determinant and flag-character relations impose equations in the modular quotient B_U/J_U. Separation · Frobenius density · Descent of determinants · Arithmetic comparisons

For $w \mid \ell$, let g_ϖ have inverse-Artin coordinate ϖ_w and put $U_{w,0} = 1$. For the normalized step operators, indexed by $1 \le j \le 4$, the character of the i-th successive quotient has the values

$$\psi_{w,i}(g_\varpi) = \kappa_\ell(g_\varpi)^{-i} U_{w,i+1}/U_{w,i}, \qquad \psi_{w,i}(g) = \kappa_\ell(g)^{-i}\Theta(e_{w,i}(y))^{-1} \quad \text{when } g \text{ has unit coordinate } y.$$

Here $\kappa_\ell(g_\varpi) = q_w^{-1} N_{L_w/\mathbb{Q}_\ell}(\varpi_w)$ is a unit. The fourth step includes the correction $\kappa_\ell(g_\varpi)^6$ on the product of flag factors. At good points the ordered flag is the one in the preceding local calculation. By separation its polynomial identity $P_g(X) = \prod_i (X - \psi_{w,i}(g))$ holds over $\mathbb{T}_U$ for all $g \in G_{L_w}$. At arbitrary arithmetic points it gives the local semisimplification. Realizing the prescribed flag order and making a zero-monodromy comparison there requires the applicable local hypotheses. At each typed split dyadic place the transferred inertia is the diagonal multiset of Teichmüller lifts of the four linked residual inertia characters, with zero monodromy. Its polynomial holds in the family by the same separation. The type fixes the inertia at its actual conductor depth. A dyadic Frobenius polynomial requires the determinant evaluated on a choice of lift. Ordered character formulas · Cyclotomic reciprocity · Dyadic type · Inertia action

The annihilator in an arithmetic fiber. Let $x : A \to O$ be an arithmetic O-algebra map and $I = \ker x$. On M_U/IM_U, the image of the determinant coefficient algebra acts through the image of the good spherical algebra. The spherical image is a finite O-submodule of the endomorphisms of the integral arithmetic fiber, so it is closed in the adic topology. Continuity and Frobenius density approximate each determinant coefficient by good spherical coefficients. Over the fraction field of O the spherical action is scalar on each classical constituent, including at weights with coincident smooth characters. Freeness makes the arithmetic lattice inject into this scalar extension. Thus if $\mu(t)^k M_U \subset IM_U$ for t in the determinant coefficient algebra and $k \ge 1$, then $\mu(t)M_U \subset IM_U$. This reducedness concerns the image of that action. An abstract fiber can have extra nilpotents. For example, if a is a nonzero parameter in a domain A, the algebra $A[Z]/(Z(Z-a))$ acts faithfully by $Z \mapsto (0, a)$ on A^2. At $a = 0$ its abstract fiber contains the nonzero class Z of square zero, while its specialized action kills Z. Reducedness modulo ℓ and

the full refined action are separate questions. Spherical action · Frobenius approximation · Specialized determinant coefficient action

Detailed determinant comparison

The linked seed at the two definite unitary levels

Let R denote the marked rank-four seed on the totally real tower field $F = L^+$, with symplectic multiplier θ. The starting input is the attached regular algebraic cusp Π_{seed} on $\mathrm{GL}_4(\mathbb{A}_L)$ for $R|_{G_L} \otimes \kappa_{\mathrm{seed}}$, together with the seed's ordinary flags and irreducibility on the cyclotomic kernel. In the mixed construction this uses cuspidal solvable base change of the Hilbert form, noninvariance of the inducing cusp under the quadratic involution, and normalized quadratic automorphic induction. At the marked element s outside G_{k_0} the induced matrix sends (u, v) to $(r(s^2)v, u)$. If $r(s^2) = \mathrm{diag}(\alpha^2, \beta^2)$, the vectors $(\beta e_2, e_2), (\alpha e_1, -e_1), (-\beta e_2, e_2), (\alpha e_1, e_1)$ have the ordered eigenvalues $(\beta, -\alpha, -\beta, \alpha)$. The paired products equal the pre-correction multiplier $\alpha\beta$. The eigenbasis change is integral since $2, \alpha, \beta$ are units. At a marked element in the tower field, the scalar multiplier correction scales the four eigenvalues together. Attached seed · Marked eigenbasis

Write w_i for the doubled inertia exponents of R and $d_i = (w_i - w_{i+1})/2 \geq 2$. The square of κ_{seed} has exponent a on open inertia, with $a \equiv w_0 \bmod 2$. For each link $K \in \{K_1, K_2\}$ take $n_{\mathrm{par}} = |w_0|$ and $\zeta_K = \lambda_K \beta_K^{n_{\mathrm{par}}}$. The construction uses common λ_K and β_K on both links and distinct auxiliary primes. At each coefficient place write $e = 2a_K + n_{\mathrm{par}} b_K$, where λ_K and β_K^2 have respective exponents a_K and b_K on open inertia and b_K is odd. The quotient $\zeta_K/\kappa_{\mathrm{seed}}$ has integral exponent $(e - a)/2$. Reciprocity therefore gives an algebraic Hecke character whose determinant twist of Π_{seed} is the cusp Π_K attached to $R_K = R|_{G_L} \otimes \zeta_K$. Its ordered Hodge–Tate weights are $h_{\tau,i} = -m_{\tau,i}$ with $m_{\tau,i} = (w_i + e)/2$ in the selected embedding and consecutive exponent gaps d_i. Parity calculation · Algebraic twist · Common characters and distinct auxiliaries

The link gives $\lambda_K \lambda_K^c \theta = \kappa_\ell^{-3}$ and $\beta_K \beta_K^c = 1$. Since complex conjugation acts by an inner conjugation on the restricted frame, the symplectic pairing gives $R_K^c \simeq R_K^\vee \kappa_\ell^{-3}$. At a good place geometric Frobenius has cyclotomic value q_w^{-1}. Inverting the roots and multiplying by q_w^3 keeps the factor $q_w^{3/2}$ in the root normalization and inverts the unitary Satake parameters. Strong multiplicity one gives $\Pi_K^c \simeq \Pi_K^\vee$. The compact real highest weight is $\nu_i = m_{\tau,i} + i$, with infinitesimal entries $z_i = m_{\tau,i} + 3/2$. Each archimedean parameter character $(z/|z|)^{2z_i}$ has conjugate-duality sign -1, its value on $j^2 = -1$ in the real Weil group. The entries are distinct, so the sign of this sum is uniquely conjugate-symplectic. KMSW Lemma 1.2.5 and Propositions 1.3.1 and 1.3.3 select $\kappa = +1$ and $\chi_+ = 1$ for the rank-four embedding. The algebraic highest weight gaps $d_i - 1 \geq 1$ verify Shin's slightly regular hypothesis. Shin gives finite-place temperedness, and BLGGT II Theorem 1.1 or Caraiani gives the full Frobenius-semisimplified comparison at ℓ, with $\mathrm{rec}_{\mathrm{geom}}(\Pi_{K,w}|\det|^{-3/2})$. This comparison permits ramified diagonal characters at the coefficient places. Polarized pairing · KMSW, §§1.2–1.3 · Shin, Theorem 1.2 and Corollary 1.3 · BLGGT II, Theorem 1.1 · Caraiani, Theorem 1.1

The actual unitary group is $U(I_4)$ over F. The isometry torsor from the quasi-split Hermitian space gives a global pure inner twist. At an inert finite place L_w/F_v is unramified quadratic. The determinant of two hyperbolic planes agrees with that of I_4 modulo local norms, so the local group is quasi-split and the integral self-dual lattice gives the specified hyperspecial subgroup. At a split place the marked coordinate

identifies the group with $\mathrm{GL}_4(F_v)$ and the conjugate coordinate uses inverse transpose. At infinity the group is compact and the regular packet supplies the representation of weight ν. The simple parameter $\Pi_K[1]$ has trivial component group after quotient by the fixed dual center $\{\pm 1\}$. The local characters on that center obey the pure-inner-twist product formula. The global generic character is trivial. KMSW Chapter 5 explicitly verifies Hypothesis 3.6.3 and the intertwining relation of Theorem 2.6.2 for generic parameters on pure inner twists. Its Theorem 5.0.5 gives multiplicity one for this tuple, with the singleton split factors $\Pi_{K,w}$ and the unramified spherical members at inert places. The frame and link are unramified at those inert places, as required for these spherical factors. The even degree of F and splitting of the coefficient and ramified places also match the field and ramification conditions in the Labesse application in the proof of BLGGT II, Theorem 3.1. This packet argument uses Mok's quasi-split formula and parameter-refined stabilization as above. The actual unitary group · Unramified setting · Even tower degree · KMSW, §1.7 and Theorem 5.0.5 · Mok, Theorem 5.1.2

Weak descent in the release. The input Π_K is regular algebraic, conjugate-self-dual and cuspidal, with ramification confined to a finite set of split places containing those over ℓ. In $\mathrm{DescWeakOfRunTh}$ the RunTh dichotomy gives either a definite-unitary constituent spherical at every nonsplit finite place and matching Π_K at almost every split place, or regular block data D with the same cofinite split Satake data but excluding Π_K. The adapter $\mathrm{descWeak_of_runTh}$ rules out the second branch with PreOf and $\mathrm{NoOther}$. In the patching entry point the release supplies $\mathrm{SR.descWeak}$ to $\mathrm{HypothesesH'}$. It uses the proved runTh, whose split tests match by the determinant-twisted constant term and whose nonsplit unit tests match at hyperspecial level. In the second branch $\mathrm{partNorm}$ normalizes the blocks to unitary data, and $\mathrm{splitRigid}$ forces D to contain Π_K from the cofinite split Satake equalities. The generic packet calculation above gives the stronger local statements for the seed at the specified places. Transfer of the contributing constituents of the fixed ordinary modules uses the same-family base-change argument above. RunTh and weak descent Descent adapter Split rigidity Patching entry point

At split tame places the common kernel on inertia of the four linked residual characters acts trivially on R_K. The fixed-type deformation relation makes $\mathrm{tr}(R_K(h)(R_K(g)-1))=0$ for every h and every g in this kernel. Semisimplicity and the trace pairing on the characteristic-zero image algebra give the asserted triviality. The remaining finite abelian image has order prime to ℓ. Reducing and lifting its characteristic roots with multiplicities identifies its four characters as the Teichmüller lifts of the linked reductions. At dyadic places the unlinked list is $(\xi\psi,\psi,\xi,1)$ on inertia before the frame permutation and the common $\bar{\lambda}_K$ twist. This test uses the full restricted conductor and has zero monodromy. The prescribed integral type τ_0 supplies the invariant in its tensor with the split constituent via the source's fixed-type test. At each link's own scalar auxiliary place the residual character indexed by the trivial frame character gives $\bar{\lambda}_K=1$ on inertia, hence $\lambda_K=1$ there by its prime-to-ℓ inertia order. The parity factor is unramified there and the original frame kills this inertia. The auxiliary local factor is therefore spherical and supplies a vector at principal level one. At $\ell=3$ the auxiliary value test is $p_K\equiv 2 \bmod 3$ at the chosen degree-one place. Inertia comparison · Prescribed type · Auxiliary inertia · Scalar auxiliary condition

At a marked $w\mid\ell$ let η_i be the four flag characters of R_K in order. For the inverse-Artin coordinate u of g, the smooth characters are $\omega_i(u)=\iota(\eta_i(g)\prod_\tau \tau(u)^{h_{\tau,i}})$. The flagged representation becomes crystalline after restriction to a finite extension killing the smooth inertia parts. Indeed, for each successive extension, tensor the preceding crystalline subrepresentation by the inverse of the new quotient. The resulting representation has labeled Hodge–Tate weights at most -2, and its Tate dual has weights at least 1, so both H^0 groups vanish. The crystalline Bloch–Kato dimension formula, local duality and the Euler characteristic

then give $H^1_f = H^1$, so the next extension is crystalline there. Thus monodromy is zero. Temperedness makes the normalized inducing characters $\omega_i|\cdot|^{3/2}$ unitary and their ratios avoid $|\cdot|^{\pm 1}$. The local component is the corresponding irreducible principal series. The prescribed smooth torus eigenspace and the refinement in flag order occur at any admissible Iwahori depth. Evaluation at the identity of the induced model detects the simultaneous eigensystem. The unipotent index $q_w^{j(4-j)}$ and the half-modulus $q_w^{-j(4-j)/2}$, together with the common shift in the inducing characters, give raw eigenvalue $\prod_{i<j}\omega_i(\varpi_w)q_w^{-i}$ for step j. Local principal series · Ordered Iwahori vector · Bloch–Kato, §3 · Zelevinsky, Theorem 4.2

Put $a_i = \eta_i(g_\varpi)$ and $c_\varpi = \kappa_\ell(g_\varpi) = q_w^{-1}\prod_\tau \tau(\varpi_w)$ for the lift with inverse-Artin coordinate ϖ_w. The four ordered geometric Weil–Deligne roots are $\iota(a_i\prod_\tau \tau(\varpi_w)^{h_{\tau,i}})$, indexed by $i = 0,1,2,3$. Dividing the raw step by the lowest coefficient value $\iota(\prod_{i<j,\tau}\tau(\varpi_w)^{h_{\tau,i}-i})$ gives the four normalized Hecke eigenvalues in ℓ-adic coefficients,

$$\left(a_0,\ c_\varpi a_0a_1,\ c_\varpi^3 a_0a_1a_2,\ c_\varpi^6 a_0a_1a_2a_3\right).$$

Each is a unit and the last is the central step. Their reductions give the ordered profile even if local residual characters coincide at $\ell = 3$. The ratios of adjacent steps and the finite torus character on units recover each marked character by local reciprocity. Cyclotomic reciprocity · Torus and normalized step comparison · Classical ordinary normalization

Tensor these local vectors and contract with the real coefficient representation on each link. The result has that link's prescribed finite type, auxiliary level and simultaneous ordinary eigenvalues. The four good spherical coefficients from the determinant comparison and the four ordered step reductions place it in the same residual projector. With the fixed-level control and determinant interpolation for this projected module, the resulting Hecke character y_K satisfies $x_K = y_K \circ (B_K \to \mathbb{T}_K)$ for the parity-adjusted arithmetic point x_K. The determinant equality follows from the good split polynomials and Chebotarev, and the ordered character equality follows from the local ratios and torus law. Thus x_K kills J_K. Twist stability for $\beta_K^{n_{\mathrm{par}}}$ and its inverse identifies the ideal extended through the parity-adjusted comparison with that extended through the zero-exponent map. Applying this on both links makes the framed seed kill the sum of the extended ideals. This gives seed support at its evaluation. The prime-containment step uses the separate fixed-weight cotangent comparison. The following patch-level construction uses the transfer and Galois comparisons for all contributing summands on the selected levels. Placement and comparison · Twist stability · Two extended ideals

Detailed seed descent and source comparisons

Conditional support propagation on the two links

Fix the two fitted deep Count II links over L/L^+ with the same comparison character and distinct scalar auxiliary places. Let B_i be their base-level determinant rings and J_i the annihilators of their projected ordinary modules, for $i = 1, 2$. Write $X.B$ for the complete local coefficient ring generated by the determinant coefficients and ordered characters of the symplectic frame, and $X.A$ for the finite-type algebra with its normalized matrices and flags. The two maps $\phi_i : B_i \to X.B$ at the target normalization scale roots and ordered characters as above. Set $J_X = \phi_1(J_1)X.B + \phi_2(J_2)X.B$. All support tests use this sum. Two links · Framed space · Prime containment tests

Ordinary modules at the Taylor–Wiles levels. On either link keep the same complete coefficient DVR O, weight algebra A, torus character Θ, integral type τ_0 including the dyadic factors, and principal level one at its scalar place v_0. This retains the compact scalar action through τ_0 and Θ and the full central step in the ordinary Hecke action. Given a Taylor–Wiles realization at depth N, put $q_j = |\mathbf{k}_{w_j}|$ at its new split places, and $a_j = v_\ell(q_j - 1) \geq N$. At w_j use the hyperspecial $K_j = \mathrm{GL}_4(\mathcal{O}_{w_j})$, the Iwahori I_j, and its normal subgroup $I_{j,1}$ cut out by requiring each diagonal residue to have prime-to-ℓ order. Write $K = \prod_j K_j$, $I = \prod_j I_j$ and $I_1 = \prod_j I_{j,1}$. The quotient $G_N = I/I_1$ is $\prod_j (\mathbf{Z}/\ell^{a_j}\mathbf{Z})^4$, with all four diagonal coordinates. The fixed level outside the added places and the residual profile on the remaining spherical and ordinary operators are inherited from the base. At K the projected module is the same base module M_0 with annihilator J_i. The level-change argument uses equality of the hyperspecial spaces and density of the good Frobenius lifts still indexed after removing the added places to identify their projected parts. The map from the enlarged-level universal determinant ring to B_i pulls back the base determinant along the quotient of Galois groups and preserves the four ordered local characters. The determinant comparison at every level uses the preceding spectral transfer and matched Galois rows on separating classical specializations of these projected spaces. In particular it uses split local–global compatibility at the new places. Selected level and base action · Hyperspecial residual projector · Iwahori subgroups

Let M be the ordinary module at I_1 before residual projection. It is already finite free over A by ordinary level control and the trivial-stabilizer argument above, which applies with I and I_1 and the fixed principal level at v_0. Let e be the A-linear ordinary projector on the space of forms, obtained by taking factorial powers of the contracting ordinary operator in the limit. The level-I_1 double cosets with the coefficient-place factor at Iwahori depth zero have free G_N-orbits. Choose one in each orbit and let χ be their characteristic function, pulled back to the adelic arguments of the forms. Multiplication by χ preserves the prescribed coefficient equivariance, including the torus character. On the space of forms, the sum of the translates of χ under G_N is 1. On M set $b(m) = e(\chi m)$ and write gm for diamond translation. The factorial powers commute with the diamond action. For a given m and coefficient precision, take one factorial power large enough for m and for all $\chi g^{-1} m$. The sum of the translates then gives $\sum_g g b(g^{-1}m) = m$ modulo that precision. Separatedness gives the identity over A. Trivial stabilizers · Ordinary limit · Orbit representatives and trace · Trace and ordinary projection

Give $A[G_N] \otimes_A M$ the action on the first factor. The map $g \otimes m \mapsto gm$ onto M has the $A[G_N]$-linear section

$$m \longmapsto \sum_{g \in G_N} g \otimes b(g^{-1}m).$$

Indeed, translation by h reindexes the sum by $g = hk$, and the composite is the trace identity. The source of the split surjection is finite free over $A[G_N]$ by the underlying A-freeness of M. The residual Hecke idempotent cuts out the simultaneous topologically nilpotent part of the commuting operators shifted by their specified residual values. It commutes with the diamond action. On that summand the commuting auxiliary operators U_{ji} give another such projector, onto the chosen simultaneous residual Frobenius roots. Its idempotent also commutes with the diamond action, since the $1 + Y_{ji}$ act as the diamond generators in the same commutative operator algebra. The resulting M_N is thus finite projective over $A[G_N]$. This finite algebra over the local ring A is local, since its reduction over the residue field is the group algebra of a finite abelian ℓ-group in characteristic ℓ, with nilpotent augmentation ideal. Hence M_N is finite free over

$$A[G_N] \simeq A[[Y_{ji}]] / \left((1 + Y_{ji})^{\ell^{a_j}} - 1 \right)_{j,i}.$$

This proves freeness even when some labeled Frobenius root residues coincide. Trace criterion · Residual Hecke idempotent · Projection at the auxiliary eigenvalues · Local group ring

Augmentation at finite depth. The diamond norm $\sum_{g\in G_N} g$ identifies coinvariants with I-invariants in the projected ordinary representation. For a free group-ring module this is the map taking each basis vector to its group sum, and the kernel is generated by the differences $Y_{ji}m$. Taking a group-ring direct summand preserves the calculation, and the norm intertwines the auxiliary U_{ji}. Thus $M_N/(Y)M_N$ identifies with the selected residual Frobenius summand of the ordinary module at I. Write P_j for the base determinant polynomial of geometric Frobenius at w_j. The elementary symmetric polynomials in U_{ji} act on the hyperspecial vectors by the coefficients of P_j with the usual alternating signs. Sending $1\otimes m$ to the inclusion of $m\in M_0$ and letting the formal roots act by the U_{ji} therefore gives a map from the tensor product of the splitting algebra of the P_j with M_0 to the Iwahori module. The map back is obtained by applying the local spherical Hecke sums between the two families of polynomials dual for the divided-difference pairing, and summing with coefficients in the splitting algebra. For each j the Hecke sandwich identity makes the composite on the Iwahori module multiplication by $\prod_{i<i'}(U_{ji'}-q_jU_{ji})$. On the splitting extension the corresponding identity is the expansion by the same dual polynomials of the divided difference of that product times an input polynomial, evaluated on the hyperspecial vectors via the symmetric functions. Taking the product over j gives both composites. Multiplying the return map by the complementary polynomial factor then gives the exact factor used by the release. Restricting the maps to the chosen residual summands completes the base splitting algebra at the labeled root residues. Diamond norm · Polynomial pairing · Spherical comparison · Sandwich identities · Composite on the splitting algebra · Multiplying the return map · Passing to completed summands

In the resulting maps

$$S_N\otimes_{B_i} M_0 \;\rightleftarrows\; M_N/(Y)M_N,$$

S_N is the completed splitting algebra just described. The maps are linear relative to the ring map that sends Y_{ji} to zero, U_{ji} to the corresponding root and the enlarged-level determinant and ordered characters to their base pullbacks. Their composites are multiplication by F_N on the right and its image on the left, where the product is over ordered pairs of distinct indices,

$$F_N=\prod_j\prod_{i=0}^{3}\prod_{i'\neq i}(U_{ji}-U_{ji'})(U_{ji}-q_jU_{ji'}).$$

Inverting F_N identifies the augmentation with the base extension. The separate local determinant relations on M_N have the form $(F_N\,\mathrm{tr}(gQ))^{24}M_N=0$ for each group element g. Here Q is $\prod_i(\mathrm{Fr}_j-U_{ji})$ for geometric Frobenius, or $\big(\prod_{k\neq i}(\mathrm{Fr}_j-U_{jk})\big)(\tau_j-1-Y_{ji})$ for the chosen tame inertia generator τ_j. At classical evaluations the factor implements the distinct-root and norm-ratio tests in the local relations. The geometric lemma and Bernstein–Zelevinsky finiteness bound the length of the torus filtration by $4!$, so vanishing on its characters gives the stated power on generalized eigenspaces. Separation of the classical evaluations then gives the integral identities. Together with base weight freeness these constructions give $\mathrm{HasAut}(B_i,J_i)$ for every datum and realization on each of the two selected links. Exact factor · Augmentation and local relations · Local test · Generalized eigenspaces · Assembly for every datum · Two linked calls

Patch-level ordinary augmentation

The patched test on one link. Use the base setting $S' = S$, an odd ℓ, local class field data with ranks summing to $[L^+ : \mathbb{Q}]$, and a nontrivial residual cyclotomic value on the linked group. Use the $\mathrm{HasAut}(B_i, J_i)$ system constructed above and the paired cohomology input below. Let $f_i : B_i \to k[[u]]$ be a curve satisfying the source's NiceI conditions. They require the ordered ordinary characters to be distinct with $\psi_{w,j} \neq \psi_{w,k}\kappa_\ell$ for $j < k$, absolute irreducibility on the full cyclotomic kernel, dense separable characteristic polynomials there, and vanishing of the cyclotomic intertwiner at the scalar auxiliary decomposition group. Suppose $J_i \subseteq \ker f_i$. We show $J_i \subseteq \mathfrak{p}_i$ for every prime $\mathfrak{p}_i \subseteq \ker f_i$. Curve conditions · Patching hypotheses · Single-curve statement

Choose the Taylor–Wiles detection elements with distinct roots on the generic curve and use the corresponding prime families at arbitrarily large depth. Normalize the curve after finite coefficient enlargement. The resulting weight ring is $\Lambda' = O'[[s_0, X_1, \dots, X_D]]$, with s_0 specializing to the curve uniformizer and $D = 4[L^+ : \mathbb{Q}]$. For q_{TW} auxiliary places the tuple ring R_∞ has $r_{\mathrm{TW}} = 4q_{\mathrm{TW}}$ patch variables Y. Write I_∞ for the determinant, polarization, ordinary, killed-inertia and auxiliary Frobenius and inertia relations, and set $\Delta = F_N|_{q_j=1}$ with the limiting roots in R_∞. At each fixed precision the congruences $q_j \equiv 1 \bmod \ell^N$ give this limiting factor in the augmentation comparisons. It is $\prod_j \prod_{i<i'}(U_{ji} - U_{ji'})^4$, which is nonzero at the generic curve by the distinct-root choice there. The relations $(1 + Y_{ji})^{\ell^{a_j}} - 1$ vanish at any fixed weight and patch precision for all sufficiently large N, so the finite-level freeness gives freeness over the corresponding truncated power series ring. At $Y = 0$ the finite-precision maps to and from the base augmentation module X_0 have composites Δ. The compatible limits after inverting Δ give the finite patched module with quasi-regular Y and augmentation X_0, after the coefficient and splitting extensions. The auxiliary relations put I_∞ in the radical of its annihilator. Augmentation maps · Patched module · Auxiliary relations · Limiting discriminant factor

Complete at the lifted curve prime $\mathfrak{q}$ of R_∞, which contains $(Y) + I_\infty$ and avoids Δ. Write $H = \widehat{(R_\infty)_\mathfrak{q}}/I_\infty\widehat{(R_\infty)_\mathfrak{q}}$ and compute the module before quotienting by I_∞. Its localization and completion are nonzero by the hypothesis on J_i and the augmentation fiber test. The completed weight localization is regular of dimension $D + 1$. Indeed $\Lambda' \to k'[[u]]$ is onto by the coefficient residue map and $s_0 \mapsto u$, and the kernel has height $D + 1$. The augmentation is a local direct summand of the base-changed finite free augmentation, so it is free over this complete local weight ring. The patch variables followed by a regular system of weight parameters give a module regular sequence of length $D + 1 + r_{\mathrm{TW}}$. Every minimal support prime is associated, so its quotient dimension is at least this depth. Cutting by the r_{TW} patch parameters leaves the finite augmentation of dimension at most $D + 1$, giving the reverse bound. The same support dimensions hold in H. Weight localization · Dimension of patched support · Augmentation support

The paired cohomology estimates give the upper bound. At precision u^m the bound $\mathrm{B1}$ for paired first cocycles modulo coboundaries has exponent $m(D + r_{\mathrm{TW}}) + O(1)$ in the size of the curve residue field. The $\mathrm{B2}$ estimate gives the even lifting relation with bounded u-power loss among $g + 1$ arrays of g detection classes. Together with the auxiliary determinant relation giving the odd cocycle spanning test, these give the three limit conditions $\mathrm{K1}$–$\mathrm{K3}$. At a Type I point there is one full matrix block, so the off-diagonal term of the relative cotangent bound is zero. Over the complete coefficient DVR at the curve, with residue field $k'((u))$, the ring H is therefore a quotient of a power series ring in $D + r_{\mathrm{TW}}$ variables. A minimal support prime of H has quotient dimension $D + 1 + r_{\mathrm{TW}}$. Its inverse image in this regular local domain must be zero. Hence the presentation is an isomorphism and the entire localized spectrum is

supported. The discriminant test uses the nondegenerate trace pairing on the full matrix algebra, tested on two spanning lists of sixteen group matrices. Its determinant is nonzero generically and stays nonzero at the generizations, also at $\ell = 3$. Paired cohomology bounds · Cotangent estimate · Dimension comparison · Discriminant test

Lift any $\mathfrak{p}_i \subseteq \ker f_i$ below the chosen curve prime through the coefficient and splitting extensions, then through augmentation and the faithfully flat completion. The lifted prime contains Y and lies in the patched support. At such a prime the patched and augmentation residue fibers agree. Nakayama and contraction therefore give $J_i \subseteq \mathfrak{p}_i$. This proves the single-curve case of the patched support implication. Support descent · Single-curve application

Application to both links. Consider a curve $y : X.A \to k'[[u]]$ specializing to the closed coefficient point, killing J_X, with absolutely irreducible generic representation and distinct ordered characters on the designated $U_1(v)$ directions. The fitted link converts this into the tests above for each $f_i = y|_{X.B} \circ \phi_i$. On $U_1(v)$ the cyclotomic character and residual representation are trivial, so distinctness gives the ratio exclusion. The fitted irreducibility and separability argument also uses the absence of the specified residual cyclotomic quadratic self-twist and a residual ratio of order greater than four. At the scalar auxiliary place the reduced matrix at a witness is $\chi(g)I + N$, with $N^4 = 0$ and $\bar{\kappa}_\ell(g) \neq 1$. The reduced cyclotomic intertwiner operator has nonzero scalar term $(\bar{\kappa}_\ell(g) - 1)\chi(g)$ plus a nilpotent sum of commuting left and right multiplications by N. The intertwiner operator remains invertible over the curve ring by the unit determinant criterion. Hence there is no nonzero intertwiner over its fraction field. Each prime $\mathfrak{p} \subseteq \ker(y|_{X.B})$ pulls back below both curve kernels. Applying the patched test to the two pullbacks and extending the ideals gives $J_X \subseteq \mathfrak{p}$. Rank-four curve conversion · Cyclotomic kernel comparisons · Scalar auxiliary test · Two-link inclusion

From the marked seed to the target. The preceding descent places the linked characteristic-zero seed in fixed-level support. Its fixed-weight tangent injects into the prescribed polarized adjoint Selmer group. With the local tangent-to-finite comparison and adjoint vanishing, the full localized determinant ring has tangent dimension at most the weight dimension, and the supported finite torsion-free quotient gives the matching lower bound. The full ring is regular of this dimension, so the annihilator vanishes there. Apply the fixed-weight test on each link through its parity-adjusted map. The equality of extended ideals established above carries the two conclusions to the common framed ring. The resulting prime-containment statement at the seed gives at least one supported minimal component at its residual endpoint. Mere vanishing on the seed evaluation would only give containment at its evaluation kernel. Seed support · Fixed-weight tangent · Residual endpoint Parity invariance of the extended ideal

At a good residual endpoint, write $d_F = [F : \mathbb{Q}]$ for the totally real tower degree, s_ℓ for the number of places over ℓ, and $d_v = d_F / s_\ell$. The flat complete framed chart has a presentation over O with a variables and b' equations satisfying $a - b' \geq 2d_F + 2 - 2s_\ell$. Put $N_* = 2d_F + 1 - 2s_\ell$. Connectedness gives a path of minimal components with adjacent intersection dimensions at least $N_* + 1$. The nine block-support loci have dimensions at most $d_F + 20C + 1 < N_*$ under the tower bound $d_F \geq 30(C+1)^2$, with $s_\ell \leq C$ and the chart bounds on strict and even cocycles and on the odd global and local terms. On the supported characteristic-ℓ coefficient locus, a collision on $U_1(v)$ has dimension at most $2d_F + 1 - d_v$. This follows by pinning one symplectic pair at every place and the other away from v, then using one comparison at v and weight finiteness up to radical. On the irreducible part, trace products with the fixed diagonal element's projectors recover matrix entries up to two symplectic basis scalings, and each local

lower symplectic flag contributes at most four coordinates. The resulting bound $2d_F + 1 - d_v + 2 + 4s_\ell$ is below N_* since $d_v \geq 30(C+1) > 2 + 6s_\ell$. Chart comparison · Reducible dimensions · Collision dimensions · Framed fiber

For an edge P, Q with J_X already contained in P, intersect the edge with the special fiber. There is a prime of quotient dimension at least N_* over $P + Q + (\varpi)$, where ϖ is a coefficient uniformizer. It avoids the finite list of reducibility and irreducible collision loci. To obtain a one-dimensional quotient still avoiding the list, choose a nonzero product in the prime quotient from the defining ideals of those loci and from the maximal ideal, then take a maximal prime after inverting that product. Its contraction has quotient dimension one by the principal ideal theorem and prime avoidance. Normalizing gives a curve over a finite residue extension with the required irreducibility and distinct characters. It kills J_X because it lies on P. The two-link curve test applied to the contraction of Q puts Q on support. Induction on the connected graph, followed by flat contraction, shows that one supported component through a good residual point forces all minimal components there to be supported. Edge propagation · Prime avoidance · Curve normalization

The marked residual chain connects the target to either seed pattern, using the additional residual line for the pattern with the two paired entries. Each residual line has a prime kernel below both endpoint kernels, hence a shared minimal component. Starting at the seed, the chart implication gives support on all components through the first endpoint, and the shared component starts the next application. At the target, a minimal prime below the integral kernel also lies below its residual kernel. The all-components result puts J_X in that minimal prime, so the integral target kills both extended ideals. To pass from this to automorphy, apply the arithmetic point comparison on each link, use the common twist and distinct auxiliary places to glue the local conclusions, then apply the prescribed descent to the counting fields. At the target gaps $(1, j-1, 1)$ the split local test for potential cyclotomic linkage remains necessary. Residual chain · Target support · Gluing and descent · Local level witness

Inputs to the conditional argument. The patched test uses CohP for local and global finiteness, Tate duality, Euler characteristics, Greenberg–Wiles, normalized Poitou–Tate and the tame local descriptions, together with the local ranks, cocycle detection and auxiliary prime realization in the link. The two HasAut systems follow from the level construction above with the same base annihilators. The remaining automorphic premises are the projected spectral transfer at base and auxiliary levels, the matching Galois determinants and ordered characters on the separating arithmetic evaluations, and the local compatibility at the prescribed Taylor–Wiles places. The release packages these through its literature hypotheses and weak definite-unitary descent. Starting the propagation also uses the attached inducing cusp, its fixed-level realization, and the separate fixed-weight adjoint rigidity comparison on each link. The chart step uses the deep tower cohomology bounds and weight finiteness. Finishing uses arithmetic classicality and the stated local and global descent conditions. Thorne's residually Schur module and patching results give a precedent with their own generator count and local hypotheses. Allen–Newton–Thorne's lifting theorem assumes $\ell > 3$ and individually polarized residual constituents, so the paired calculation at three uses the stated link, spectral and cohomological hypotheses. Cohomology package · Ordinary patching input · First-order count · Obstruction bound

Detailed patching and propagation argument · Thorne, §§4 and 6

4. From one host to the points at a cutoff

Fix ε, J and let S consist of the real cutoffs admitting $\mathrm{WinDatum}$. The scale s records $s.x = x \in S$, $s.J = J$, $s.\varepsilon = \varepsilon$, $s.L = x^{\varepsilon}$ and $s.b = \exp(x^{\varepsilon})$. **Release input (admissible ranks and host).** Given a tuple (Y, q, d, e) at x, the rank theorem extends its frame to R with $\mathrm{RankVal}(s, R)$ and the host admissibility condition. In particular $R.t \le \lfloor x^{400} \rfloor$ eventually. The host theorem uses these ranks, the conductor bounds and the splitting conditions of §2. Its output is H with its arithmetic field, two auxiliary primes g_0, g_1, two deformation primes, an initial level prime and sets TT, T_2 of frame elements. It satisfies the size, local prime class and conductor requirements $\mathrm{HostGood}$. There is a separate simpler conductor-one host theorem with $J = 1$ and a trivial extra requirement. The construction consuming the cohomological windows uses the supplier with the simultaneous conditions below. Rank supply Host data Choice of tuple and common scale Basic host theorem

The extra host requirement G_ψ has two parts. First, $\mathrm{EAtH}(\mathrm{Level.cyc}(\psi), H, TT, x, \varepsilon, 4)$ asserts the kept and exceptional character conditions, accuracy $x^{-1+\varepsilon}|TT|$ in every coset count, and centered-ramp and log-derivative bounds for the bounded characters. Its ideal shares are $4/(3o+1)$ on the identity coset and $3/(3o+1)$ on the others, with $o = |\langle \psi \rangle|$. For $\mathbf{1}$ the single share equals one exactly. The nonempty quadratic twists still need the pair estimates. The release invokes its bounded-twist lemma on the same q, d used by the host. Second, G_ψ asserts $\mathrm{VVle}(H.E, R.t, \sqrt{s.b})$. **Proved here.** The host field satisfies $\mathrm{IsChiField}$ for the chosen e, q, d, so this part follows by restricting the pair's budget to $j < R.t$. A point with the same $s, R, H.E, TT$ and arithmetic host carries both assertions. Host analytic interface Proof of both host assertions Exact transport into point data

The primes and the level ideal on the selected field

Keep the field and the rank data at this cutoff. Write $N = QS\,qd$, $K = \mathbb{Q}(\sqrt{-d})$, $b = \exp(x^{\varepsilon})$. The field construction gives $d \mid N$, $d \equiv 3 \bmod 4$, and $\log N \ll x^{\varepsilon/5}(\log x)^2$. It takes the level ideal in K to be the unit ideal, whose norm is one. Paper II constructs the sets of thick places and tops on this frame. For each shift the thick places keep the same digits, and the weighted exceptional-orbit estimate uses the same tops as the character average. Level ideal and prime set · Thick digits and their translates · Weighted choice of the tops

Primes represented by a norm in a short interval. Suppose $d \ge 100$ is prime, $d \equiv 3 \bmod 4$, $M \ge 3$, and put $B = 4dM^2$. There are absolute $c, c' > 0$ such that, whenever $0 < h \le 1$, $\log B \le c\sqrt{\log X}$, and $h \ge e^{-c\sqrt{\log X}}$, at least

$$c' \frac{hX}{B^3 \log X}$$

rational primes $p \in (X, (1+h)X]$ have a generator of norm p congruent to $1 \bmod M\mathcal{O}_K$. Writing that generator as $(a + b_0\sqrt{-d})/2$ gives $a^2 + db_0^2 = 4p$, $M \mid b_0$, and $2M \mid a - b_0 - 2$. Indeed $1, (1+\sqrt{-d})/2$ is an integral basis. These are the congruences required of a norm prime by the release. Norm and congruence conditions · Short-interval input

To prove the estimate, take the ray class field of K of modulus $M\mathcal{O}_K$, and let h_M be its degree over K. The ray exact sequence gives $h_M \le h_K M^2$. The imaginary quadratic Minkowski bound gives an integral ideal of norm $O(\sqrt{d})$ in every class. At norm j there are at most the divisor function $\tau(j)$ such ideals. Thus

$h_K \ll \sqrt{d}\log(2d)$ and $h_M \ll B$. Every Hecke character in this extension has conductor dividing $M\mathcal{O}_K$, so its analytic conductor parameter in Thorner–Zaman's Theorem 1.4 is bounded by $4dM^2$. The base degree is two. For $X \le u \le 2X$, that theorem counts the prime ideals with trivial Artin symbol by

$$\frac{F(u)}{h_M}\bigl(1 + O(e^{-c_1\sqrt{\log X}})\bigr).$$

Here $F(u) = \mathrm{Li}(u)$ in the absence of an exceptional zero and $F(u) = \mathrm{Li}(u) - \mathrm{Li}(u^\beta)$ otherwise. The displayed error follows from the two terms in the theorem, since $\log X/\log B \ge \sqrt{\log X}/c$. One can take $c_1 > 0$ fixed for $c \le 1$, and then decrease c to meet the theorem's size threshold. Thorner–Zaman, Theorems 1.4 and 3.3

In the exceptional case Stark's bound in their Theorem 3.3 gives $(1-\beta)\log B \gg B^{-2}$. Consequently $1 - u^{\beta-1} \gg B^{-2}$ on this interval. This follows from $1 - e^{-s} \gg \min(1, s)$ for $s \ge 0$ and $\log X \ge \log B$. Integration of $F'(u) = (1-u^{\beta-1})/\log u$ gives $F((1+h)X) - F(X) \gg hXB^{-2}/\log X$. The same bound holds for $F = \mathrm{Li}$. Each endpoint error before division by h_M is $O(Xe^{-c_1\sqrt{\log X}}/\log X)$, since $F(u) \ll X/\log X$. Choose $3c < c_1$. Then $hB^{-2} \ge e^{-3c\sqrt{\log X}}$ absorbs both errors, uniformly once c is small enough. This gives the lower bound with denominator B^3 after division by h_M. Prime ideals of residue degree two account for $O(\sqrt{X})$ terms, and primes dividing Md for $O(\log B)$. Their total is absorbed because $hB^{-3} \ge e^{-4c\sqrt{\log X}}$. Shrinking c forces the requisite lower bound on X for all permitted B. Each remaining rational prime contributes at most two ideals. Trivial ray class means the ideal is generated by an element congruent to one at the modulus. Since the ideal is integral, its generator is integral at every finite place. This proves the estimate. Ray class version in the release · Thorner–Zaman

The two primes with prescribed types. Put $Y_0 = (kb)^8$. For $M = N$ the preceding bound with $h = 1$ gives at least $Y_0^{9/10}$ norm primes in $(Y_0, 2Y_0]$ eventually. Indeed $\log(4dN^2) = o(\sqrt{\log Y_0})$ and $3\log(4dN^2) + \log\log Y_0 = o(\log Y_0)$. For a candidate g, its excluded coefficient primes form the set

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

For fixed a at most three primes qualify, since the product of their exact prime powers divides $g^a - 1$. Thus $|\mathcal{B}_r(g)| \le 3r + 1$. Fix g_0 and $\ell \in \mathcal{B}_r(g_0)$. For a given a, let w be the least positive integer with $\ell^w > Y_0^{a/4}$. Every g for which a witnesses exclusion satisfies $g^a \equiv 1 \bmod \ell^w$. There are at most $2a$ residue classes, by the cyclic unit group at odd primes and the decomposition of dyadic units into a sign and a cyclic group (the moduli 2 and 4 satisfy the same bound). Each class occurs at most $1 + Y_0^{3/4}$ times in this interval. Including $g = \ell$, the number of excluded partners for g_0 is at most

$$(3r+1)\bigl(1 + 2r^2(1 + Y_0^{3/4})\bigr) < Y_0^{9/10}$$

eventually, since $r \le x^5$ and $Y_0 \ge e^{8x^\varepsilon}$. A second norm prime g_1 therefore has disjoint excluded set. The norm congruence gives $g_i \equiv 1 \bmod N$. For the later level norm prime at size T_s, use $M = N(g_0g_1)^2$, $h = 1/V$. Here $1 \le V \le x^{12}$, $\log M = O(x^\varepsilon)$, and $\log T_s \ge c_s m x^\varepsilon$ with fixed $c_s > 0$ and $m \ge x^{1/3}$. Both short-interval hypotheses hold eventually, and the positive lower bound supplies a prime in the required interval. Excluded prime definition and count · Fixed-coefficient count · Two primes and the later level norm prime Level interval

The deformation primes and the canonical extra prime. Linnik's theorem gives an absolute $A \ge 1$ such that a prime congruent to $1 \bmod D$ occurs below or at D^A for every $D \ge 2$, absorbing its fixed

multiplicative constant into A. Apply it successively to $15N$, $15Np_{\mathrm{de}}$, and $Np_{\mathrm{de}}p'_{\mathrm{de}}$. This gives distinct deformation primes $p_{\mathrm{de}}, p'_{\mathrm{de}} \equiv 1 \bmod 15N$ and a provisional level prime in $1 \bmod N$, all bounded by N^{C_8} for an absolute C_8. They lie below g_0, g_1 eventually. For the final set let r_0 be the least prime avoiding $D = N(p_{\mathrm{de}} - 1)(p'_{\mathrm{de}} - 1)$. A prime divisor of $D + 1$ shows $r_0 \le 2D$. Since $2 \le N$ and the level ideal has norm one,

$$r_0 N \le N^{3+2C_8}, \qquad q_0 = \min\{p \text{ prime} : p \equiv 1 \bmod r_0 N\} \le N^{(3+2C_8)A}.$$

In particular q_0 differs from both deformation primes, as r_0 divides neither $p_{\mathrm{de}} - 1$ nor $p'_{\mathrm{de}} - 1$. It lies below the type primes for large x, because $\log N = o(x^{\varepsilon})$. The five primes split in the abelian host field whose conductor divides N, and the resulting level set is $\{g_0, g_1, p_{\mathrm{de}}, p'_{\mathrm{de}}, q_0\}$. This gives the enlarged host constant at every sufficiently large supported cutoff. Least-prime theorem, Xylouris · Host level selection · Canonical prime bound and releveling

Common weights on the chosen host

Conditional sampling theorem. Fix the host at a supported cutoff, including TT, T_2, and fix the constants in the order required by the uniform base-level projection estimate. At all sufficiently late cutoffs, that estimate implies the existence of a single pair a, b supported on TT satisfying both weight windows and WOK. The cutoff threshold is uniform over eligible hosts and indices. The arithmetic proof uses only the local **mean** loss at large primes. The proof below gives the integral transfers at the serving depth from the release's segment estimates and cleaned paired-column certificate. Local projection and all size and depth checks · Base-level projection statement

Write $a_* = x\exp(x^{\varepsilon})$, $X' = \exp(x^{\varepsilon/3})$, $b_* = a_*/X'$, $\delta = x^{-\varepsilon/8}$, and $Y_0 = x^{13}$ in this sampling argument. Let $N_1 = O_{\varepsilon}(\log x)$ be the bound in Paper II §4 on the number of tops in an orbit window. With the fixed width constant $K.\mathrm{La}$, put $\mathcal{L} = K.\mathrm{La}\, N_1^2 \log a_*$ and $W_b = 2n\delta x^{\varepsilon}\mathcal{L}b_*/k$. The two intervals are

$$I_a = [a_*, a_* + a_*/(kY_0)] \cap \mathbb{Z}, \qquad I_b = [b_*, b_* + W_b] \cap \mathbb{Z}.$$

Both widths exceed one eventually. In the first, $k \le x^6$. In the second, $2n\delta \ge 1$, $\mathcal{L} \ge 1$ and b_*/k grows faster than any power of x. Sample uniformly from the finite nonempty product $I_a^{TT} \times I_b^{TT}$ and extend by zero. These intervals impose WinOK. The constants in the projection estimate are chosen before using the boxes, so the sampling distribution is fixed for all primes on this host. Box and width definitions · Weight windows

At a candidate prime $p < X'$, the cyclotomic coordinate of Frobenius partitions the frame into orbits. The candidates are in the prescribed inert class with orbit order at least $x^{\varepsilon/4}$. In the range $p \ge k^2Y_0^2$ the candidate condition excludes a power of p in a short relative band about X'. The signal S_p is m times the sum of the rank losses over the good orbits of Paper II §4. For an orbit with N_T tops, each assignment of the local weights has exactly $(|I_a||I_b|)^{|TT|-N_T}$ extensions. Its marginal is therefore uniform on the prescribed local boxes. The tops in the good orbits are disjoint, with sum of their cardinalities at most $|TT|$. Good orbits · Candidate exclusions and marginals

For $p < k^2Y_0^2$ the defined loss is bounded by $CN_Tkx^{-26} \le N_Tx^{-12}$ eventually, with C fixed once the orbit constants are fixed. At large primes the base-level projection estimate gives mean loss at most $C_{\mathrm{pr}}N_TN_1/Y_0$. Its certificate is uniform in the bounded interior weight coordinates, the local coefficient

ring and the target lattice. It holds at the required depth and bounds the least admissible rank loss from above. In transferring the estimate, the orbit test gives $N_T \le 40n'\delta$ for orbit order n', and the band exclusion and height bounds apply to the prescribed boxes. The supporting proof checks the signed odd-top term in the serving depth and the dimension-weighted change of boxes in the local estimate. Since $C_{\mathrm{pr}}N_1 \le x$ eventually, the large-prime loss has mean at most $N_T x^{-12}$ as well. Local means and quantitative transfer · Serving-depth and numerical bounds

It follows on the single global product that $\mathbb{E}S_p \le m|TT|x^{-12}$ for each candidate and $\Pr(S_p > x^{-\varepsilon}) \le x^{\varepsilon}\mathbb{E}S_p$. Write $u = \log X' = x^{\varepsilon/3}$. Summing over all integers up to X' gives $\sum_p (\log p)/p \le u(2+u)$ for the candidate sum. Thus

$$\begin{aligned}\mathbb{E}\sum_{p\ \text{candidate}} \frac{\log p}{p}\left(21_{\{S_p > x^{-\varepsilon}\}} + 4S_p\right) &\le 4(x^{\varepsilon}+1)m|TT|x^{-12}\,u(2+u)\\ &\le 48x^{-19/2+(9+2/3+1/50)\varepsilon}.\end{aligned}$$

For the last inequality use $m \le x^{1/2+\varepsilon/100}$, $|TT| \le 2x^{2+8\varepsilon+\varepsilon/100}$ and $u(2+u) \le 3u^2$. The excess saving $\gamma = 1/2 - (9+2/3+1/50)\varepsilon$ is positive for $\varepsilon < 1/52$. Taking $x^{\gamma} \ge 48$ makes the bound at most x^{-9}. A point in the finite product attains at most the mean and supplies WOK. Linearity of expectation allows the restrictions for different primes to overlap arbitrarily. Independence is used only for the coordinates of the weight box. Finite averaging and uniformity

A determinant estimate on the original interval of weights

The large-prime argument varies one terminal weight while keeping the other weights of its segment fixed. Its Gram matrix $M(z)$ has size s, with entries of degree at most J in the integral binomial basis over $\mathbb{Z}_p$. Take W consecutive integers with $1 \le W \le p^h$ and $J < p^h$. Assume $\det M$ is a nonzero polynomial with leading coefficient l and degree d, and $v_p(l) + d/(p-1) \le sL$. For a positive integer E, let $\beta(z)$ count the Smith factors divisible by p^E, including zero factors. Then

$$\frac{1}{W}\sum_z \beta(z) \le s\left(\frac{J}{W} + \frac{L + Jh/W}{E}\right).$$

This inequality also counts a weight at which the specialized Gram matrix is singular. Terminal Gram window and application

For the proof, expand $M(z+p^hU)$ in U. Its coefficients are integral by the binomial addition formula and $v_p(i) < h$ for $1 \le i \le J$. In Smith coordinates at z, a determinant coefficient of degree i uses a nonconstant term from at most i rows. Its valuation is at least $E\max(0, \beta(z) - i)$. Factor $\det M(Z) = l\prod_\zeta (Z-\zeta)$ over a finite splitting field, counting roots with multiplicity. Let d_z be the number of roots in $v_p(z-\zeta) \ge h$. Divide each factor of $\det M(z+p^hU)$ by a coefficient of least valuation. After reduction the product has degree exactly d_z, since each close root contributes degree one and each distant root degree zero. The coefficient of U^{d_z} therefore has valuation $v_p(l) + \sum_\zeta \min(h, v_p(z-\zeta))$. Each root belongs to at most one such close disc. For $0 < u \le h$ the number of weights in $v_p(z-\zeta) \ge u$ is at most $W/p^{\lceil u\rceil} + 1$. Integrating in u, summing over roots, and using $d \le sJ$ proves the bound. Integral change of variable and Smith estimate

Counting the initial segment and pairing its columns

Fix one eligible Frobenius orbit of length n' with N tops, let $q = p^{n'}$, and write $h = k - 1$ in this paragraph. A top u contributes symmetric degrees A_u, B_u. Put $S_k(d) = \binom{d+k-1}{k-1}$ for $d \geq 0$, zero for $d < 0$, and $d_{\mathrm{bi}} = \prod_u S_k(A_u) S_k(B_u)$. Cut the cyclic orbit at its even tops. A segment of length G_j has index $Q_j = p^{G_j}$, so $\prod_j Q_j = q$. On an internal letter of degree d and digit modulus P, divide each exponent of a monomial by P. Its remainder vector $b \in \{0, \dots, P-1\}^k$ satisfies $|b| \leq d$ and $|b| \equiv d \bmod P$. Conversely a composition of $(d - |b|)/P$ and the vector b recover the monomial. Thus the sum of the reduced symmetric dimensions over all admissible remainders is exactly $S_k(d)$. At a terminal degree B, keeping ties $0 \leq l \leq L = \lfloor \mathcal{L} B / k \rfloor$ gives $\sum_l S_{k-1}(B - l) = S_k(B) - S_k(B - L - 1)$. Both endpoints enter this count.

For a fixed digit class, interlacing in successive symmetric and dual Pieri products counts the intermediate paths with multiplicity equal to the dimension of the terminal highest weight. Restricting the initial degree from H_0 to H gives the exact factor $S_k(H)/S_k(H_0)$. The path exclusions for excessive loss in an internal strip have total dimension at most $e^{-\mathcal{L}/8}$ times the unrestricted dimension at H. These three factors give a lower bound for each class. Summing over digits and multiplying over segments gives $d_{\mathrm{bi}} q^{-h} (1 - \kappa_{\mathrm{margin}})$ as a lower bound for the model rank. Here κ_{margin} bounds the total loss from initial, terminal and internal cutoffs. The exponent h counts the free coordinates of a composition with its total prescribed. The edge normalization counts the hyperplanes in the residue space and has denominator $1 + q + \cdots + q^h$. Fixed-digit lower bound · Segment product · Digit multiplicities, endpoints and mean

Let κ_j denote the fraction of model directions in segment j whose integral Gram factor has valuation exceeding its allocated budget D_j, counting zero factors among the discarded directions. Products of good ranks give at least $d_{\mathrm{bi}} q^{-h} [1 - \kappa_{\mathrm{margin}} - \sum_j \kappa_j]_+$ directions. In passing from the enlarged to the original integer boxes, hold all lower weights fixed and restrict the upper weights first. With N_1 the orbit-window bound, s_j the number of internal tops and σ_j the number preceding the excluded prime-power band, the mean of the total loss is at most

$$\frac{e^{0.015} N_1}{Y} \left(C_s N + C_A \sum_j (s_j - \sigma_j + 1) \right).$$

This uses the internal Gram estimates, the terminal Gram coefficient bounds and the integral comparison from level lattices to blocks. At a fixed lower weight the enlarged upper box has volume ratio at most $(1 + \Delta/|I_a|)^N \leq e^{0.015}$. All losses are nonnegative, so restriction costs at most this ratio. Averaging over the original lower box gives the displayed estimate. Original-box marginals

To identify the pairing on an actual edge $e = (S, L)$, choose a basis b with $S = bW^k$, $L = bD_p^{-1} W^k$, where W is the unramified valuation ring and $D_p = \mathrm{diag}(p, 1, \dots, 1)$. Put $w^2 = p$ and $D_w = \mathrm{diag}(w^{-1}, 1, \dots, 1)$. The reversed dual edge has basis $b' = b^{-T} D_p$. At every Frobenius embedding ι_s the adapted matrices satisfy

$$\iota_s(b') D_w = (\iota_s(b) D_w)^{-T}.$$

In the polynomial model the Fischer form pairs equal monomials of multi-index α with value $\prod \alpha_i!$. It is invariant under inverse-transpose substitutions on the two arguments. If d_ℓ are the individual letter degrees, multiply this form by $Z = p^{k \sum_\ell \lfloor \log_p d_\ell \rfloor} / \prod_\ell d_\ell!$, assigning a zero logarithm to degree zero. The

resulting form is integral on the coefficient lattices since the multinomial valuation for a degree d_ℓ is at most $k\lfloor \log_p d_\ell \rfloor$. It maps the cleaned lattice on the reversed edge into the dual test lattice on e. Indeed, duality takes every competing incoming edge to a competing outgoing edge, and multiplication by w^c commutes with the map. Lattice and edge definitions · Two-sided lift

In identity coordinates the polynomial lift of each retained block symbol has the prescribed symbol modulo w^c on the selected line and vanishes to precision c on every other line. At precision n subtract lifts of each error term $w^i\tau_i$ computed at precision $n-i$. Multiplication by w^i restores its denominator allowance. The localizing polynomials preserve the off-line vanishing throughout this induction. Apply these lifts to the source in basis b and to the test in basis b'. Their vertex denominators are at most the stated serving depth $\lfloor t_0 \rfloor_+$. The two powers of w in the rescaled pairing contribute exactly the required power of p. The comparison on e is therefore Z times the Fischer pairing of the two identity symbols modulo w^c. Distinct labels have disjoint monomial supports. For a common label β, the block factorial comparison gives a scalar κ_β times the product of the segment forms, with $0 \le v_p(\kappa_\beta) \le 8Nk\log_p a$. This scalar can have positive valuation. On the tensor of the good Smith frames the matrix is diagonal modulo w^c. Its diagonal entries have valuations at most $D_0 = \sum_j D_j + 8Nk\log_p a$, and $c > 4D_0$. Dividing the rows by those entries gives a matrix congruent to the identity modulo w, hence invertible over the coefficient ring. This realizes the product rank on the two cleaned lattices. The construction uses the identity-coordinate localizing polynomials and their factorial estimates. The internal Gram and level-lattice bounds above control how many directions reach this step. Identity-edge induction · Normalized block forms · Coordinate maps, scalars and depth

The determinant on a capped prefix

Fix an internal tie, the coefficient prime p, and the degrees of the preceding letters. Write $K+1$ for the length of the prefix ending at its standard letter, r_0 for the number of coordinates, and S for the sum of its unsigned degrees d_j. On the monomial lattice A over $\mathbb{Z}_p$ use the form $B(x^\alpha, x^\beta) = \mathbf{1}_{\alpha=\beta}\alpha!/\prod_j d_j!$. Impose the caps on all prefixes through this letter and take highest vectors of fixed weight for the subgroup fixing its last coordinate x. For order e in x, let L_e be the image obtained by taking the coefficient of order e, with its factor x^e, on integral vectors divisible by x^e. Write $U_e = \mathbb{Q}_p L_e$, $A_e = A \cap U_e$ and $h_e = \dim U_e$.

Two integral estimates govern the answer for $K \ge 1$. Put

$$c_1 = \binom{r_0}{2}\lfloor \log_p(S + r_0 - 1)\rfloor, \qquad c_0 = 2\binom{r_0}{2}\lfloor \log_p(4S + r_0 - 1)\rfloor, \qquad b_0 = c_0 + c_1, \quad D_0 = (K+1)c_1.$$

Then $p^{b_0}A_e \subseteq L_e \subseteq A_e$. Indeed division by x^e shifts the last degree and its cap down by e. Restriction at $x = 0$ in this reduced degree has an integral lift after scaling by p^{c_0}. Its proof uses the auxiliary $\mathfrak{gl}_{r_0}$ action on the sum of the restricted degrees. Its highest vectors have last degree zero and hence lift. Type projection and divided-power generation each cost at most $c_0/2$, and the actual restriction image is stable under those divided powers. Returning the lift to the prescribed highest weight costs c_1. Next take the Fischer-orthogonal projection P_e to U_e. The nested cap projectors, followed by the highest-weight projector on the chosen weight, give $p^{D_0}P_eA \subseteq A_e$. The length-zero and single-letter caps each act by zero or the identity on this multidegree. This accounts for the number of denominators in D_0. Bounded restriction lift · Projection on a positive-length prefix

For a lattice L write $L^{\#}$ for its dual under B in its rational span. In the monomial basis the entries of B are reciprocals of multinomial integers, so $A^{\#} \subset A$. Since A_e is saturated in A, every integral functional on A_e extends to A. Orthogonal projection of its representing vector gives $A_e^{\#} = P_e A^{\#}$. Thus $p^{D_0} A_e^{\#} \subseteq A_e$. Dualizing the inclusion for L_e gives $p^{b_0} L_e^{\#} \subseteq A_e^{\#}$. A second multiplication by p^{b_0} returns from A_e to L_e. Every Smith valuation of B on L_e is therefore at most $D_0 + 2b_0$. Equivalently, changing from A_e to L_e adds $2\ \mathrm{length}(A_e/L_e)$ to the Gram determinant valuation. Hence

$$v_p(\det B|_{L_e}) \le h_e((K+3)c_1 + 2c_0) \le h_e \frac{(K+7)r_0^2}{2} \log_p(4S + r_0).$$

For $K = 0$ each nonzero graded space has one monomial with Gram entry $1/\binom{o}{e}$ for last degree o. The rank-zero determinant is one. Smith argument and integral partner · Graded lattice, cap action and denominator bounds

Internal and terminal averages on the chosen weights

In a segment with s internal ties, let σ be the first late tie. At a late tie i, fixing the digits and all other coordinates makes the compatible upper weights an arithmetic progression of step P_i. The free dual degree $n_i = (A_i - R_i)/P_i - j_i$ runs through each integer in one interval. Here R_i is the sum of the fixed remainders and j_i is the fixed tie level. The caps on the prefix through its standard letter apply throughout this interval. In a basis adapted to order e, permutation terms mixing different orders have smaller degree in the frozen Gram determinant. Its leading coefficient is the product of the graded determinants just bounded and their nonzero leading scalar factors. The width of the interval and the leading coefficient can thus be used together in the determinant-window estimate. Leading coefficient of the internal determinant

More precisely, let b_x^{int} count the Smith factors above the internal threshold at level x, and let $Y_x(A)$ be the capped polynomial space there. For a fixed compatible digit class $\mathcal{C}$, write $A^{(i)}$ for A with its i-th coordinate at the last compatible value. The continuation multiplicity of a discarded highest vector is bounded by its multiplicity at this endpoint. The integral extension of the highest-weight partner operators places all the retained directions in a saturated sublattice. The internal determinant estimate, with prefix degree at most $7/Y$ times the progression length, gives

$$\sum_{A \in \mathcal{C}} b_x^{\mathrm{int}}(A) \le \frac{31}{Y} \sum_{i=\sigma}^{s-1} \sum_{A \in \mathcal{C}} \dim Y_x(A^{(i)}).$$

At the terminal tie the space is fixed as the scalar exponent varies. Write d_x for the free dual letter's dimension, $r_x = d_x \dim Y_x(A)$ and W for the number of terminal exponents. The monomial with exponents ν has terminal pairing $\nu!/n!$, where $n = |\nu|$. Multiplying its partner by $n!/\nu!$ gives a unit. The terminal polynomial after binomial clearing has this internal tensor Gram as its leading matrix up to a specified scalar. Its coefficient bound and the same window estimate give at most $W(d_x b_x^{\mathrm{int}}(A) + 5r_x/Y)$ discarded directions, summed over the terminal exponents. This gives the segment constant $C_M = 1$ on the compatible progressions. Passing to the paired columns uses the integral comparison from the separately normalized level lattices to the monomial block. The highest-weight partner extension, coefficient integrality and this last comparison are the remaining inputs to this deduction. Highest-weight partner extension · Terminal coefficient estimate · Passage from levels to the block · Internal and terminal counts with their thresholds

Integral harmonic splitting on the flat rows

At a large eligible prime p, work over the admitted coefficient DVR O with $w^2 = p$ and $v(p) = 1$. For one numbered Frobenius orbit let T be the set of tops and $N_T = |T|$. The row rank is $k \geq 3$, and write $m_{\mathrm{loc}} = k - 1$ for the local weight index. At a top s the two degrees (P_s, Q_s) are (A_s, B_s) for even s and (B_s, A_s) for odd s. A shift subtracts j_s from both, with $0 \leq j_s \leq \min(P_s, Q_s)$. The flat row then has weight $(P_s - j_s, 0, \dots, 0, -Q_s + j_s)$ and the off-top factors are trivial. Use the rounded radius a with $P_s + Q_s + 1 \leq 8a$. Alternating degrees · Bidegree model

For one top write $S_{P,Q}$ for the bidegree space in the standard and dual variables, $h = \sum_i x_i y_i$, and $F = \sum_i \partial_{x_i} \partial_{y_i}$. On degree (P, Q), $Fh - hF = P + Q + k$. Thus for $Fv_0 = 0$ in degree (a', b') one has $F(h^i v_0) = i(a' + b' + k + i - 1)h^{i-1} v_0$. The top harmonic projector is

$$\mathcal{H}_{P,Q} = \sum_{\ell'=0}^{\min(P,Q)} \frac{(-1)^{\ell'}}{\binom{P+Q+k-2}{\ell'}} h^{(\ell')} F^{(\ell')}, \qquad h^{(\ell')} = h^{\ell'}/\ell'!, \quad F^{(\ell')} = F^{\ell'}/\ell'!.$$

Applying F cancels consecutive terms, and the projector is the identity modulo h. Induction on the smaller degree gives the direct sum of the strings $h^j \ker F|_{S_{P-j,Q-j}}$. The coefficient map for $a' = P - j$, $b' = Q - j$ is

$$\mathcal{D}_j = \frac{\mathcal{H}_{a',b'} F^{(j)}}{j! \binom{a'+b'+k+j-1}{j}}.$$

The string formula shows that $h^j \mathcal{D}_j$ is the projection onto the j-th summand. To bound its denominators in the coefficient lattice, expand the ℓ'-term on a monomial $x^\mu y^\nu$. For multi-indices $|\alpha| = \ell'$ and $|\beta| = j + \ell'$ the factorial ratio from the divided powers and $1/j!$ is

$$\frac{\binom{j+\ell'}{j} \beta!}{j! \alpha!} = \binom{j + \ell'}{j}^2 \frac{\binom{\ell'}{\alpha}}{\binom{j+\ell'}{\beta}}.$$

The remaining monomial coefficients in the numerator are integral binomials. The denominators are one k-part multinomial and two binomials with upper arguments at most $P + Q + k$. Legendre's formula bounds their loss by $(k + 1)\lfloor \log_p(P + Q + k) \rfloor$. The same bound works for $\mathcal{H}$, and multiplication by h^j is integral. Both operations are equivariant for a change of row frame. Harmonic identities · Coefficient map · Gauss lattice estimate

The identification with the actual flat lattices uses the first-column entries and the signed $(k - 1)$-minors of the first $k - 1$ columns. Send x to the first column of Z^+ and y to the signed cofactor vector of the independent matrix Z^-. Call the resulting map ι_{pol}. It is injective because independent matrices realize arbitrary first columns and any nonzero cofactor vector. Let β be its specialization to $Z^+ = Z^- = Z$. The two-letter lattice and the flat lattice are the integral images under ι_{pol} and β, respectively, by the integral flag-minor basis theorem for shapes $(P, 0, \dots, 0)$, $(Q, \dots, Q, 0)$ and $(a' + b', b', \dots, b', 0)$. The image of the tied matrix is dense in $x \cdot y = 0$: when $x, y \neq 0$ and $y(x) = 0$, complete x to a basis of $\ker y$ and scale another basis vector to obtain the prescribed cofactors. Thus $\ker \beta$ in each bidegree is the corresponding multiple of h. For a lattice frame g, the transported polynomial for weight λ is $\det(g)^{-\lambda_k} f(g^{-1} Z)$. The cofactor identity $C(g^{-1} Z) = \det(g)^{-1} g^t C(Z)$ gives exactly this determinant twist on the dual letter and on the flat row. In particular the integral flat image has a coefficient-integral lift in each frame, with its

transported two-letter coordinates. Integral Schur modules · Flag minors · Cofactor transformation · Lattice and incidence comparison

On the actual rows use the projection $\beta\mathcal{D}_j\iota_{\mathrm{pol}}^{-1}$ and inclusion $\beta(f)\mapsto\iota_{\mathrm{pol}}(h^j\mathcal{H}_{a',b'}f)$. The latter is well defined by the incidence kernel. Tensoring gives $r^0_{\mathbf{j}}, i^0_{\mathbf{j}}$ with $\sum_{\mathbf{j}} i^0_{\mathbf{j}} r^0_{\mathbf{j}} = 1$. Each is integral on its row lattices after multiplication by p^u, where

$$u = N_T\left\lceil (m_{\mathrm{loc}}+2)(\log_p(8a+m_{\mathrm{loc}})+1)\right\rceil.$$

This holds on the Frobenius-twisted source, intermediate edge and target lattices and on the neighboring edges in the intersections. In adapted frames these three lattices use g_s, $g_s\,\mathrm{diag}(w^{-1},1,\dots,1)$ and $g_s\,\mathrm{diag}(w^{-2},1,\dots,1)$. Tensor comparison · Adapted edge bases

Put $B=\sum_s P_s$, $A_{\mathbf{j}}=\sum_s(P_s-j_s)$ and $d_{\mathbf{j}}=B-A_{\mathbf{j}}$. These are the edge exponents. The two pairing conventions are $b_{\mathrm{bi}}(x,\Phi)=\Phi(p^Bx)$ and $b_{\mathbf{j}}(y,\psi)=\psi(p^{A_{\mathbf{j}}}y)$. Set $r_{\mathbf{j}}=w^{d_{\mathbf{j}}}r^0_{\mathbf{j}}$ and $i_{\mathbf{j}}=w^{-d_{\mathbf{j}}}i^0_{\mathbf{j}}$. On the two sides the maps are $x\mapsto p^u r_{\mathbf{j}}x$ and $\Phi\mapsto p^{2u+d_{\mathbf{j}}}(\Phi\circ i_{\mathbf{j}})$. For an integral row lattice or its integral dual scaled by w^E, the corresponding map lands in the flat lattice or flat dual scaled by $w^{E+d_{\mathbf{j}}}$. On the dual side this uses $p^{2u+d_{\mathbf{j}}}w^{-d_{\mathbf{j}}}=p^{2u}w^{d_{\mathbf{j}}}$ and integrality of $p^u i^0_{\mathbf{j}}$. The exponent $-t$ for the source and target constraints becomes $-t+d_{\mathbf{j}}\ge -t$. The exponents $-B, c-B$ become $-A_{\mathbf{j}}, c-A_{\mathbf{j}}$ in the edge and adjacent-edge constraints. The cleaned inclusions therefore hold at the same depth t and conductor c. The pairing identity is

$$b_{\mathbf{j}}\big(p^u r_{\mathbf{j}}x,\ p^{2u+d_{\mathbf{j}}}(\Phi\circ i_{\mathbf{j}})\big)=p^{3u}b_{\mathrm{bi}}(i_{\mathbf{j}}r_{\mathbf{j}}x,\Phi).$$

Consequently a paired bidegree certificate with scalar z and size r yields integral flat pairing matrices whose sum is $p^{3u}zI_r$, after using the original integral change of columns. For each matrix keep the Smith factors dividing $p^{3u}z$. After division by $p^{3u}z$, the omitted parts vanish modulo w, so the sum of the retained ranks is at least r. Integral basis changes and rescaling of the retained factors give flat certificates with scalar $p^{3u}z$. The loss in its valuation is $3u$ and the release allocates $D_1=D_0+4u$ when $v(z)\le D_0$. Cleaned lattice inclusions · Smith extraction · Pairing and budgets

The initial paired-column input is $\mathrm{admB2}$ at depth $\lfloor t_0\rfloor_+$, conductor $c_4=c_*+2N_Tc_{\mathrm{cl}}+2n'c_{\mathrm{cl}}$ and budget $D_0=\mathrm{totalDegreeBudget}$ for orbit length n'. It uses the segment and digit-class estimates on both cleaned bidegree lattices. The split itself preserves depth. For the shifts retained in the mean, $j_s\le J_0=\mathcal{L}b'/k$, each odd-top minimum in t_0 decreases by at most J_0. The shifted certificate is served at $\lceil t_0(A-\mathbf{j},B-\mathbf{j})+N_TJ_0\rceil_+$. The sum of ranks uses all shifts and the edge rank cap at D_1 and c_4. The serving inequalities give $4D_1<c_4$, also required in the middle-row comparison. Paired-column input · Shift and edge rank cap

Contraction on the bounded middle rows

Put $c=c_{\mathrm{cl}}$ and $r_{\mathrm{mid}}=k-2$. At a top with base weight $(a_s,\mu_s,-b_s)$, take $(P,Q,\eta)=(a_s-c,b_s-c,\mu_s)$ at an even slot and $(b_s-c,a_s-c,-\operatorname{rev}\mu_s)$ at an odd slot. The lower, actual and upper operative weights are $\lambda_-=(P,0^{r_{\mathrm{mid}}},-Q)$, $\lambda=(P+c,\eta,-Q-c)$ and $\lambda_+=(P+2c,0^{r_{\mathrm{mid}}},-Q-2c)$. Write $\theta=\lambda-\lambda_-=(c,\eta,-c)$, $\theta^\vee=(c,-\operatorname{rev}\eta,-c)$ and $\zeta=\lambda_+-\lambda_-$. The upper and lower translations here are made in the original sampling coordinates. Operative weights in the two parities · Contraction statement

The middle pairing. Write $\mathscr{D}_A(\nu)$ for the flag-minor module of shape ν over a coefficient ring A. Over the fraction field it consists of polynomials with right upper triangular character $\prod_j b_{jj}^{\nu_j}$ under $Z \mapsto Zb$. The integral flag-minor basis theorem identifies the module over a DVR with the coefficient-integral part of this field space, since the column multidegree and the polynomial identities for column additions descend coefficientwise. Products add shapes. Put $\beta = c + \eta$, $\alpha = Q + \beta$, $M = Q + 2c$, and $\alpha' = M - \operatorname{rev}\alpha = c - \operatorname{rev}\eta$. These are the middle shapes of the auxiliary row and the two contraction inputs, with output middle shape $(M, \dots, M)$. Set $e_0 = \alpha_{r_{\mathrm{mid}}}$, $\gamma = \alpha - e_0\mathbf{1}$, $D = M - e_0$ and $\delta(Y) = \det Y$. Deleting full-height minors identifies $\mathscr{D}_A(\alpha)$ with $\mathscr{D}_A(\gamma)$ by division by δ^{e_0}. The complementary map

$$\tau f'(Y) = \delta(Y)^D f'(Y^{-t}J)$$

uses the reversing permutation matrix J. Jacobi's minor identity replaces each first-column minor by a signed complementary-row first-column minor divided by δ. There are $D = \alpha'_1$ factors, and complementing their heights gives shape γ. Complementing back, with full-height columns for empty minors, gives surjectivity over the integral ring. The two shapes have equal Weyl dimensions by reversal and a determinant shift, so this is a lattice isomorphism. Integral flag minors · Fulton, §§8–9 · Complementary substitution

On $\mathscr{D}_{\mathbb{Q}}(\gamma)$ let $\langle\ ,\ \rangle_F$ be the Fischer form with monomial norms $\prod_{ij} u_{ij}!$ on Y^u. Normalize by $\mathfrak{n} = \langle h, h\rangle_F$, where $h = \prod_j \Delta_j^{\gamma_j - \gamma_{j+1}}$ is the leading-minor highest polynomial, $\gamma_{r_{\mathrm{mid}}+1} = 0$. The normalized form B_0 is nondegenerate by positivity over $\mathbb{R}$ and contravariant for transpose by the differential formula for the Fischer pairing. Over $\mathbb{Z}_{(p)}$ let $\mathscr{D} = \mathscr{D}_{\mathbb{Z}_{(p)}}(\gamma)$ and let $\mathcal{U} \subset \mathscr{D}$ be spanned by coefficients of $h(gY)$ for g a product of elementary unipotents in independent parameters. The highest row-weight line is spanned by h, with $h(I) = 1$. Evaluation at I kills the other row weights, which are Fischer orthogonal to h. Thus $B_0(F, h(gY)) = F(g^t)$. A polynomial integral against $\mathcal{U}$ therefore takes integral polynomial values on a generic unit lower triangular matrix, by triangular elimination of its transpose into elementary matrices. Gaussian elimination and the right triangular character put F in the localization of the coefficient ring at the product of the leading minors. Intersecting with $\mathbb{Q}[Y]$ gives $\mathbb{Z}_{(p)}[Y]$, since reduction modulo p shows that multiplying a primitive polynomial by powers of that product preserves primitivity. The converse follows from $F(g^t)$ having integral coefficients for integral F. The same evaluation identity shows $\mathcal{U}$ spans over $\mathbb{Q}$, since vanishing on the unit lower triangular matrices implies vanishing on the Gaussian open cell by the right character. Hence $\mathcal{U}^{\#} = \mathscr{D}$ and $\mathscr{D}^{\#} = \mathcal{U} \subset \mathscr{D}$, with $\#$ taken for B_0. Normalized Fischer form

The highest norm is

$$\mathfrak{n} = \frac{\prod_{i=1}^{r_{\mathrm{mid}}} (\gamma_i + r_{\mathrm{mid}} - i)!}{\prod_{i<j} (\gamma_i - \gamma_j + j - i)}.$$

Indeed the Fischer norm is the complex Gaussian mean of $|h(Y)|^2$, with independent entries of density $\pi^{-1} e^{-|z|^2}$. The polynomial h is the matrix coefficient on the unit highest vector in the tensor product of the exterior powers specified by the column heights of γ. In the factorization $Y = UR$ by Gram–Schmidt, U is Haar unitary and independent of the upper triangular R by left unitary invariance. The squared diagonal entries of R are independent gamma variables of shapes $r_{\mathrm{mid}} - i + 1$, by conditioning on the preceding columns. Their γ_i-th moments give the factors $(\gamma_i + r_{\mathrm{mid}} - i)!/(r_{\mathrm{mid}} - i)!$. On the irreducible generated by the highest vector, Schur averaging gives $1/\dim V_\gamma$ for the mean squared unitary coefficient. The Weyl dimension formula gives the displayed quotient. Fulton and Harris, §24

For a monomial of column degrees γ_i and a power $b = p^j$ with $j \ge 1$, the contribution of column i to the valuation of the numerator divided by its monomial factorial is

$$\left\lfloor \frac{\gamma_i + r_{\mathrm{mid}} - i}{b} \right\rfloor - \sum_l \left\lfloor \frac{u_{li}}{b} \right\rfloor = \left\lfloor \frac{\sum_l (u_{li} \bmod b) + r_{\mathrm{mid}} - i}{b} \right\rfloor \le r_{\mathrm{mid}} - 1.$$

The last column contributes zero and powers above $\gamma_1 + r_{\mathrm{mid}} - 1$ contribute zero. The denominator in $\mathfrak{n}$ can only improve this bound. Thus $B = p^E B_0$ is integral for $E = (r_{\mathrm{mid}} - 1)^2 \lfloor \log_p(\gamma_1 + r_{\mathrm{mid}} - 1) \rfloor$ when $\gamma_1 > 0$. For the constant line use $E = 0$. Since $\mathscr{D}^{\#} \subset \mathscr{D}$, the image of B contains $p^E \mathscr{D}^*$, where $\mathscr{D}^*$ is the ordinary linear dual. Smith form over $\mathbb{Z}_{(p)}$ gives diagonal powers p^{e_i} with $0 \le e_i \le E$, and the bases extend to O. In particular $\gamma_1 \le 2c \le 8a$ gives

$$K_0 = \left\lceil m_{\mathrm{loc}}^2 (\log_p(8a + m_{\mathrm{loc}}) + 1) \right\rceil \;\ge\; e_i.$$

The middle pairing is $\mathcal{L}(f, f') = B(f/\delta^{e_0}, \tau f')$. For $A_g f(Y) = f(gY)$ the identity $A_{g^t} \tau A_g = \det(g)^D \tau$ and transpose contravariance show that simultaneous left substitution multiplies $\mathcal{L}$ by $\det(g)^M$. The pairing is integral on the two middle lattices and has $\dim V_\eta$ diagonal Smith pairs, by the two lattice isomorphisms. Twisting its target by the inverse determinant character of degree M gives the invariant functional used in the release's QPh. Invariant pairing and its integral bases · Release Smith bound

Row maps. Put $D(Y) = \mathrm{diag}(1, Y, 1)$. Apply $\mathcal{L}$ coefficientwise in Z to $v(ZD(Y))s'(ZD(Y'))$, with independent Y, Y', to define $C(v, s')$ on the shapes $\nu(\lambda), \nu(\theta^\vee)$, where $\nu(\omega) = \omega - \omega_k \mathbf{1}$. The coefficients in Y, Y' have shapes α, α' by the right triangular identities of the inputs and are integral on the integral lattices. The output transforms under right middle triangular substitution by $\det^M$. Additions from the first column into the middle columns commute through $D(Y)$ as combinations of such additions, and the last column has degree zero. The first column has degree $P + Q + 4c$. Thus the output belongs to $\mathscr{D}_O(\nu(\lambda_+))$ on integral inputs, and left substitution in Z commutes with C. For the middle Smith vectors u_i, u'_j, lift u_i/δ^Q and u'_j by adjoining the first row to every middle minor and inserting enough Z_{11} factors to give first row degree $2c$. The lifts s_i, s'_j have last row degree zero and specialize at $D(Y)$ to the prescribed polynomials. Their shapes are $\nu(\theta), \nu(\theta^\vee)$. Coefficient construction · Output character

Every lower-flat F satisfies $F(ZD(Y)) = \delta(Y)^Q F(Z)$ by its height-one and height-$(k-1)$ minors. Consequently $C(Fs_i, s'_j) = FR_{ij}$ with R_{ij} independent of F. Use $F_0 = \Delta_{k-1}^Q$ to compute R_{ij}. The result $F_0 R_{ij}$ has shape $(Q + 4c, (Q + 2c)^{r_{\mathrm{mid}}}, 0)$ and first and last row degrees $Q + 4c$ and 0. Every contributing minor product must use row one in every factor and omit the last row, so it is a multiple of $Z_{11}^{2c} \Delta_{k-1}^{Q+2c}$. Evaluate at I to read off the Smith scalar, and cancel F_0. This yields

$$C(Fs_i, s'_j) = \delta_{ij} p^{e_i} FH, \qquad H = Z_{11}^{2c} \Delta_{k-1}^{2c}.$$

Here H is the highest polynomial of weight ζ, of bidegree $(2c, 2c)$ in the two letters. In the frame g the determinant twists add for multiplication by s_i and H, and for contraction because $\lambda_k + \theta_k^\vee = (\lambda_+)_k$. Thus all three maps are integral on their common row lattices. Diagonal contraction · Frame and determinant factors

Cleaned lattices and primitive inclusion. At a chosen edge $e = (S, L)$ use the adapted source frames described above and tensor the three maps, with trivial factors off T. Call the lower, actual and upper tensor weights H_-, X, H_+, and let $d = N_T c$. For any tensor weight W put $A_W = \sum_s W_{s,1}$ and write

$H_W(S), H_W(L), H_{W,e}$ for its integral tensor lattices. The normalized edge lattices $R_W = w^{-A_W} H_{W,e}$ and $R'_W = w^{-A_W} H^\vee_{W,e}$ are perfect for $B_W(v, G) = G(p^{A_W} v)$. Source cleaning at (t, c_4) imposes membership in $w^{-t} H_W(S)$, $R_W(e)$ and $w^{c_4} R_W(f)$ on other outgoing edges at S. For tests use $w^{-t} H_W(L)^\vee$, $R'_W(e)$ and $w^{c_4} R'_W(f)$ on other incoming edges at L. A row of weight ω and first row degree j scales at the intermediate and target lattices by $w^{j+\omega_k}$ and $p^{j+\omega_k}$, with $\omega_k \le j + \omega_k \le \omega_1$ on the tableau basis. Thus $H_W(S) \subseteq R_W(e)$, and $H_W(L) \subseteq R_W(e)$ when $\sum_s W_{s,k} = -A_W$. Normalized pairing · Cleaned lattices

Let i multiply by the tensor of transported H's, $U_{\mathbf{i}}$ multiply by the tensor of transported s_{i_s}'s, and $L_{\mathbf{j}}$ contract with the tensor of transported s'_{j_s}'s. Thus $L_{\mathbf{j}} U_{\mathbf{i}} = \delta_{\mathbf{ij}} p^{e_{\mathbf{i}}} i$ with $e_{\mathbf{i}} = \sum_s e_{i_s}$. The first two maps preserve source cleaning at (t, c_4) because the multipliers lie in the auxiliary source lattices, hence in every normalized outgoing lattice, and the leading sums add. In the normalized edge coordinates i is multiplication by a tensor of flag polynomials taking value one at the identity. The auxiliary intermediate scale is w^{2c} per top, which cancels against the normalization. Reduction modulo w shows that a primitive multiplier preserves the least coefficient valuation. Thus each row inclusion is saturated in the coefficient lattice, with free quotient over the DVR. Tensoring the resulting integral row retractions and extending scalars gives a retraction π of i carrying $R_{H_+}(e)$ into $R_{H_-}(e)$. Primitive inclusion · Tensor comparison

The fixed second input of $L_{\mathbf{j}}$ has adapted exponent c per top. It lies in its auxiliary normalized edge lattice and in p^{-d} times its target lattice, hence in p^{-d} times its normalized lattice at every other incoming edge. Here both auxiliary endpoint sums are $d, -d$. With $A_{H_+} - A_X = d$, the test map is $\widehat{L}_{\mathbf{j}} G = p^d G \circ L_{\mathbf{j}}$ and satisfies $B_X(v, \widehat{L}_{\mathbf{j}} G) = B_{H_+}(L_{\mathbf{j}} v, G)$. At L the factor p^d cancels the p^{-d} in the input under integral contraction, giving the same depth t. On e normalized pairing gives the test inclusion with no loss. On another incoming edge f, the contraction sends $R_X(f)$ into $p^{-d} R_{H_+}(f)$. The other-edge condition on the new test is therefore $w^{c_4 - 2d} R'_X(f)$. The resulting source and dual conductor is $c_4 - 2d$. Source and dual inclusion estimates · Normalized pullback identity

Cross pairing and rank. Take the two flat certificates on e at common depth t, conductor c_4 and valuation budget D_1, with ranks $n_\pm$ and pairings $z_\pm I$. Form the integral cross matrix of the lower sources after i against the upper tests. Keep the u nonzero Smith entries dividing $z_- z_+$, extracting $z_- z_+ I_u$ by integral operations. The other $s_0 = n_- - u$ lower sources pair into $w z_- z_+ O$ against all upper tests. Use the lower tests pulled back through π with multiplier $p^{A_{H_-} - A_{H_+}}$, changing them contragrediently to the lower source basis. These are in $R'_{H_+}(e)$ by perfectness, and pair with the s_0 sources as $z_- I_{s_0}$. Combining with the n_+ upper vectors and tests gives the block matrix

$$\begin{pmatrix} z_- I & T_{\mathrm{bad}} \\ B & z_+ I \end{pmatrix}, \qquad T_{\mathrm{bad}} \in w z_- z_+ M_{s_0, n_+}(O).$$

Right multiplication by $\left(\begin{smallmatrix} z_+ I & 0 \\ -B & z_- I \end{smallmatrix}\right)$ gives $z_- z_+$ times a matrix congruent to identity modulo w, so integral column operations give the scalar identity. All the sources are cleaned and the tests are integral at the chosen normalized edge. Transport to the $g_q = (q^k - 1)/(q - 1)$ outgoing edges of the **same source**, with $q = p^{n'}$. These edges correspond to lines in $p^{-1} S / S$ and its stabilizer acts transitively. The cross-edge entries lie in $w^{c_4} O$ by source cleaning alone. As $v_w(z_- z_+) \le 4 D_1 < c_4$, the combined matrix has full rank, whence $s_0 + n_+ \le d_+ = \dim V_{H_+} / g_q$. In particular $u \ge n_- + n_+ - d_+$. Cross-pairing argument · One-sided edge rank cap

Apply $U_{\mathbf{i}}$ and $\widehat{L}_{\mathbf{j}}$ to the retained pairs. The result has rank uD_{mid} for $D_{\mathrm{mid}} = \prod_{s\in T} \dim V_{\eta_s}$. Its block scalars are $p^{e_{\mathbf{i}}} z_- z_+$, which can be made $p^{N_T K_0} z_- z_+$ by integral test rescaling. Thus the valuation budget is $2D_1 + N_T K_0$. Write $d_- = \dim V_{H_-}/g_q$, $d_X = \dim V_X/g_q$ and let $\kappa_\pm$ be the flat losses at the lower and upper translations. The rank is at least

$$D_{\mathrm{mid}}\big((1-\kappa_-)d_- - \kappa_+ d_+\big).$$

After removing the middle factors in the Weyl product for X, the other $2k-3$ factors per row lie between the corresponding flat factors. Their numerators grow by at most $2c$ between the flat weights, except the first-to-last numerator, which grows by $4c$. Since $P, Q \ge b'/3$ in the interior, all ratios are bounded by $1 + 6c/b'$. With $\epsilon_0 = 12 N_T k c/b' \le 1/9$ this gives $(1-\epsilon_0)d_X \le D_{\mathrm{mid}} d_- \le d_X$ and $d_+ \le 2d_-$. The normalized actual loss is thus bounded by $\kappa_- + 2\kappa_+ + \epsilon_0$. The boundary uses the trivial loss when the lower translation leaves the original box. Integral rank budget · Weyl comparison

The serving depth is the upper flat depth $t = \lceil t_0(a+c, b+c) + N_T J_0 \rceil_+$. Its weight-dependent term is the sum over odd tops s of $\min(b_s + c, (a_{u(s)} + c)/p^{d(s)})$, with $u(s)$ the preceding even top and $d(s)$ its cyclic distance to s. Both terms use the first operative weight at their respective parity. The minimum increases by at most c from the untranslated weights, and monotonicity serves the lower pair at the same depth. After the top transfer the conductor is $c_4 - 2N_T c$. The off-top argument below uses the remaining $2n' c_{\mathrm{cl}}$ of conductor and at most that much extra depth. The serving inequalities give the full certificate at $t \ge t_I$, with conductor c_* and pairing budget D_p. Odd-top depth term · The two translated pairs · Serving inequalities · Final serving depth

Full bounded middle-row proof

Restoring the bounded weights outside the tops

Fix the orbit of length n' and its subset T of tops. For $s \notin T$ write λ_s for the dominant weight in its operative order. At an odd position this reverses and negates the original row. Put

$$A_o = \sum_{s\notin T} \lambda_{s,1}, \qquad B_o = \sum_{s\notin T} \lambda_{s,k}, \qquad J = A_o - B_o \le 2(n' - |T|)c_{\mathrm{cl}}.$$

Let $e = (S, L)$ be a lattice edge with $pL \subseteq S \subseteq L$ such that L/S has residue dimension one. On the top factors take source vectors and dual tests at depth t and conductor $c \ge J$ with pairing matrix P. Assume $PM = zI_h$ for an integral matrix M and a nonzero scalar z. Restoring all the other factors gives the same scalar z, rank $h\, d_o$, depth $t + J$ and conductor $c - J$, where d_o is the dimension of their tensor product. In particular the valuation allowance on z and the proportion of the available dimension both persist. Edge and weight conventions · Intermediate and endpoint lattices · Statement used in the projection estimate

Proof. Work over the coefficient DVR with uniformizer w and $w^2 = p$. At a lattice vertex write $H(S)$ and $H(L)$ for the integral tensor lattices of the factors outside T, and let H_e be their lattice at the intermediate edge. The intermediate lattice over this DVR is $S + wL = S + w^{-1}pL$. In a basis for S adapted to $S \subset L$, its first basis vector is w^{-1} times the source vector, and the first target basis vector is p^{-1} times that vector. This description holds after every Frobenius twist because the twist fixes p. In a factor of weight λ_s, the standard integral tableau basis uses shape $\lambda_s - \lambda_{s,k}\mathbf{1}$ and the determinant twist. In such a basis let a_s count the occurrences of the first letter. A column is strictly increasing, so $0 \le a_s \le \lambda_{s,1} - \lambda_{s,k}$. Under the intermediate change of lattice the inverse coordinate substitution multiplies a basis polynomial by w^{a_s}, and

the factor $(\det g)^{-\lambda_{s,k}}$ for a basis matrix g contributes $w^{\lambda_{s,k}}$. Tensoring gives a basis (v_i) of $H(S)$ in which H_e and $H(L)$ have bases $(w^{j_i}v_i)$ and $(p^{j_i}v_i)$ with $B_o \le j_i \le A_o$. The top and complementary factors use disjoint sets of variables. Products of their bases form an integral basis at every vertex and edge. Integral tableau basis · Adapted generators · Tensor basis at the three lattices

The normalized edge lattices are $R_e = w^{-A_o}H_e$ and $R'_e = w^{-A_o}H_e^\vee$. Their respective bases are $r_i = w^{j_i-A_o}v_i$ and $\rho_i = w^{-j_i-A_o}v_i^\vee$. They pair to δ_{ij} under $(r,\rho) \mapsto \rho(p^{A_o}r)$ because $p^{A_o}w^{-2A_o} = 1$. Relative to the vertex basis $p^{-j_i}v_i^\vee$ of $H(L)^\vee$, the second multiplier is again $w^{j_i-A_o}$. Hence

$$H(S) \subseteq R_e \subseteq w^{-J}H(S), \qquad H(L)^\vee \subseteq R'_e \subseteq w^{-J}H(L)^\vee.$$

The vertex lattices are independent of the adapted frame, since a change in $\mathrm{GL}_k(O)$ and its inverse preserve the integral polynomial lattice. The following comparisons at neighboring edges use their inclusions of these same vertex lattices. For any other edge f with source S, these inclusions give $r_i \in w^{-J}R_f$. For any other edge with target L, they give $\rho_i \in w^{-J}R'_f$. Normalized edge pairing · Frame independence · Comparison at adjacent edges

Tensor the given source vectors with every r_i and the given tests with every ρ_i. At the chosen edge these tensors are in the normalized lattices because the leading weights add. At the source and dual target vertices their depth is at most $t + J$. At competing edges the top conditions supply w^c and the added basis costs at most w^{-J}, which gives the required divisibility by w^{c-J}. The pairing on tensors is the product of the pairings on the two factors. The new matrix is the original pairing matrix tensored with I_{d_o}, and tensoring the integral column operation with I_{d_o} puts it in the form zI_{hd_o}. Both the certified rank and the ambient dimension multiply by d_o. The neighboring edges are indexed by the same lattices after restoring the weights. Under the release hypothesis $c \ge 2n'c_{\mathrm{cl}}$, inclusion gives the stated uniform depth $t + 2n'c_{\mathrm{cl}}$ and conductor $c - 2n'c_{\mathrm{cl}}$. For weights constant along each complementary row, $J = 0$ even if the determinant twists themselves are nontrivial. Depth and competing-edge conditions · Product of pairing matrices and rank normalization

The remaining integral estimates in the initial rank bound concern the internal Gram ties, the terminal Gram coefficients and passage from the untied levels to the block lattices. The pairing step uses $\mathrm{admB2}$ for integral cleaned source and test vectors and $\mathrm{admB3}$ for the margin loss. The determinant estimate, harmonic splitting and the two weight transfers give the mean bound $C_{\mathrm{pr}}N_TN_1/Y_0$ from these inputs. Returning from the enlarged intervals to the original uniform measure costs a factor at most two. Segment estimates · Digit-class comparison · Paired-column and margin inputs Recovery of the original mean

Integral harmonic splitting

Realization with prescribed weights and level

Word supply. Write k, m, f for the corresponding sizes in R, and put $\Gamma = \exp(C_\Gamma m^5 f x^\varepsilon)$ and $D_{\mathrm{word}} = \lceil 3\Gamma^2 x^{C_F} \rceil$. The rank estimates give $m^5 f x^\varepsilon \le 128 x^{9/2+9\varepsilon}$ and $k \ge x^{9/2+10\varepsilon}$. For fixed $C_\Gamma, C_F \ge 0$, $\log D_{\mathrm{word}} = O(x^{9/2+9\varepsilon} + \log x) = o(k)$. In particular $D_{\mathrm{word}} \le 2^{k-3}$ eventually. The rank-count argument weakens this integer bound to 2^{k-2}, and Paper III uses its logarithm only to absorb the additive loss into Q_0. Thus there is an injection $\mathrm{Fin}(D_{\mathrm{word}}) \to \{0,1\}^{k-2}$ eventually. Choose it once for the kit. The point provenance gives a bijection from its own words to $\mathrm{Fin}(D_{\mathrm{word}})$, with bits equal to those of the common code followed by zeros. Hence identical word indices yield identical bits across points. The

additional word tests $\mathrm{ChCode_at}$ come from the object and numerical hypotheses on the same kit. Rank estimates and injection Word bijections and agreement Further word tests · Chosen word count · Rank-count application · Peer check of the slope loss

For the geometric construction fix the weights in WinOK, the word code, and a shift $u_0 < n$. Use $F = E^{+}$, $f = 2n$, $r = k + m - 1$, $\Delta = \lceil x^{c_D} \rceil$, and $h = \lceil x^{c_h} \rceil$, with c_D, c_h fixed and $c_h \geq 402$. In the ordinary field, finite inertia at q, d can act only on the respective quadratic factor, and inertia at 2 can act only on the real dyadic factor. No finite inertia contains complex conjugation, which changes both imaginary radicals. Hence E/F is unramified at finite places. Arithmetic and realization theorem

At each real place take $v_i = 4\Delta(k - 2 - i)$, $0 \leq i \leq k - 2$, and $u_{i+1} = v_i - \frac{1}{2} - \Delta c_i$ for the internal coordinates of the word, with $c_i \in \{0, 1\}$. The two endpoints at the top t_0 are $a(t_0) + (k-1)/2$ and $-b(t_0) - (k-1)/2$. Away from the tops they are $4\Delta(2k - 2) + 1/2$ and $-1/2$. For $0 \leq j < m$ set $b_j = B_\tau + 2jh$, $e_j = b_j + h$, where the base B_τ is $b(t_0) + \lfloor b_* / x \rfloor$, $\lceil \exp(x^{\varepsilon}) \rceil$, or $\lceil s.w''(n) \rceil$ at a top, thick or thin place, respectively, with $s.w''(n) = b_*/(xn)$. A shift of B_τ at one thick place by less than $|\mu(E)|$ cancels the infinity-type character on CM torsion. Strict interlacing holds since each internal u_{i+1} lies between $v_i - 4\Delta$ and v_i, and the endpoint inequalities use $4k\Delta \leq a_*$, $\lfloor b_* / x \rfloor \geq k$. The b, e gaps are $h \geq (t+3)/2$ since $t \leq x^{400}$ eventually. The compact rank-k dimension is the Weyl product $\prod_{i<j} (u_i - u_j)/(j - i)$. The row parameter uses v_i and $-b_j$, with alternating signs on the latter entries. Its normalized polynomial is $\prod_{i<j} |\lambda_i - \lambda_j|/(j-i)$, with $\lambda = (v_i, -b_j)$. Weight formulas and torsion shift · Row formal-degree polynomial · Prescribed family of weights

At either completely split type prime g, choosing one member of a split pair identifies the unitary group with GL on that member. The other component is the inverse adjoint. Reduction gives the column upper unitriangular subgroup in rank k. In the row group of rank r, use the upper unitriangular matrices with the central diagonal block of size $k+1$ fixing its last basis vector. This imposes k extra zero entries above the diagonal in that block. The central diagonal block is multiplicative on upper block triangular matrices and respects inversion, so these matrices form a subgroup. Counting free entries and dividing $|\mathrm{GL}_s(\mathbb{F}_g)| = g^{s(s-1)/2} \prod_{i=1}^{s} (g^i - 1)$ by the subgroup orders gives

$$I_{\mathrm{col}}(g) = \prod_{i=1}^{k} (g^i - 1), \qquad I_{\mathrm{row}}(g) = g^k \prod_{i=1}^{r} (g^i - 1).$$

Each g_i contributes f such indices. The signed tame list at g_i consists of the exponents $0, \pm 2, \ldots, \pm 2m$ of a primitive residue unit. Outside $\mathcal{B}_r(g_i)$ its prime-to-ℓ order is at least $(g_i - 1)/g_i^{1/4} > 4m$ eventually, so the list stays distinct under the prime-to-ℓ projection. The disjoint excluded sets ensure that one of the two choices works at each coefficient prime. Exact subgroup and tame calculation · Local level definitions

The signed row covolume

Let $F = E^{+}$, $f = [F : \mathbb{Q}]$, and let $J = I_{r-m/2, m/2}$ be the Gram matrix of the selected row Y_0. Give its finite integral stabilizer volume one and use compact-dual Haar at each real place. Then

$$\mathrm{vol}\big(U(J)(F) \backslash U(J)(\mathbb{A}_F)\big) = \sum_{[L] \in \mathrm{gen}(I_r)} \frac{1}{|\operatorname{Aut}(L)|} = M_r.$$

The ordinary construction installs the adelic quotient volume for this signature on the datum used by the row count. Choice of quotient measure · Row size

At a finite place the extension is split or unramified quadratic. There is a unit u of norm -1. In the split case take $(-1, 1)$. In the unramified case lift a residue-field solution through the principal unit filtration, using surjectivity of the residue trace at each step, including residue characteristic two. The diagonal matrix with entries 1 on the positive coordinates and u elsewhere gives an integral isometry with the identity lattice. To compare rational gauges take an F-rational top form ξ on the skew-Hermitian matrices and set $(\omega_H)_1 = (\det H)^{-r}(X \mapsto HX)^*\xi$ for Gram matrix H. Under $H = A^* H' A$ the skew-coordinate change $Y \mapsto A^{-*} Y A^{-1}$ has determinant $N(\det A)^{-r}$, as one sees by splitting the quadratic algebra. The determinant prefactor cancels it. Thus the finite stabilizers have equal gauge volume.

At a real place take imaginary diagonal and complex upper triangular coordinates on the Lie algebra and let dC_H be their coordinate density. For the compact form, the last-column map $U(s) \to S^{2s-1}$ gives successive factors $2\pi^s/(s-1)!$ in its total coordinate Haar. Write $V_r = \prod_{s=1}^{r} 2\pi^s/(s-1)!$. The compact-dual convention equates the two real densities in the complexified Lie algebra. The map $X \mapsto JX$ preserves the absolute coordinate density, and $|\det J| = 1$. Hence the compact-dual identity density is dC_H/V_r for both forms and its ratio to the rational gauge agrees at each real place. The analytic convention multiplies the group measure by V_r and divides its discrete-series formal degree by V_r. The factors cancel over the f real places. Compact-dual convention · Change of Haar

The Tamagawa convergence factors are the same because the two groups have the same character module and dimension. Both unitary groups have Tamagawa number two. The equal local comparisons thus give equal adelic quotient volumes in the specified measures. For the positive form the real group has volume one. The finite double cosets classify the lattices in the genus, and the stabilizer in each class contributes the reciprocal of its finite order. This gives the displayed sum. Spectral occurrence of a row with the required finite type and rational periods is a further assertion about this quotient. Kottwitz, *Tamagawa numbers* · Definite double cosets · Gauge and compact-dual density calculation

The mass ratio and the word-level prime

Let M_s denote the mass of the genus of the positive identity lattice in rank s. Write I_{col} and I_{row} for the products of the corresponding local indices over all $2f$ type places. For a strictly decreasing list y of length s, put $V(y) = \prod_{i<j}(y_i - y_j)$ and $W(y) = V(y)/\prod_{j=0}^{s-1} j!$. At a real place τ the two lists are $u = u(w, \tau)$ and $\lambda = (v_0, \dots, v_{k-2}, -b_0, \dots, -b_{m-1})$. The compact Weyl dimension and the normalized discrete-series formal-degree polynomial in the release give the two explicit sizes

$$d_w = M_k I_{\mathrm{col}} \prod_{\tau} W(u), \qquad D_* = M_r I_{\mathrm{row}} \prod_{\tau} W(\lambda).$$

The formal-degree normalization follows by inserting the decreasing parameters in the product of absolute root differences divided by $\prod_{i<j}(j-i)$. The argument also works for any D_* at least $(2/3)^f$ times the displayed row expression. These sizes are defined before the prime for the word and before the selection of the automorphic row. Definitions of both sizes · Row polynomial and the lower comparison

The column size is a positive integer. On a lattice in the genus, the choices of complete residue flags with a basis of each successive line form a set of size I_{col}. The finite group of lattice automorphisms acts freely on this set. Indeed, at a place above g_i a stabilizer embeds into a conjugate of the inverse image of the upper

unitriangular group. This inverse image is pro-g_i, since each principal congruence quotient has g_i-power order and the reduction is a g_i-group. A finite subgroup of a pro-g_i group has g_i-power order. The local representation is faithful by scalar extension, and $g_0 \neq g_1$. Thus the stabilizer has order one. Summing the orbit counts over the genus gives $I_{\mathrm{col}} M_k$. Finally u is strictly decreasing in a single translate of $\mathbb{Z}$ with the compact Harish-Chandra shift. Subtracting the half sum of positive roots gives a dominant integral weight, so each $W(u)$ is a positive integer by the Weyl dimension formula. Column level

Put $v_E = |\operatorname{disc} F|^{1/f}/2$, so $\pi \le v_E$ and $\log v_E \le L = x^{\varepsilon}$ on the selected host. Let η be the quadratic character of E/F. Paper V's norm-one sphere calculation gives, for $k < s \le r$,

$$\frac{M_s}{M_{s-1}} = ((s-1)!)^f (v_E/\pi)^{sf} (2v_E)^{-f/2} L(s, \eta^s).$$

At residue norm q, the Euler factor is $(1 - \eta(\mathfrak{p})^s q^{-s})^{-1}$. Its absolute logarithm is at most $-\log(1 - q^{-s})$. Grouping over rational primes bounds $|\log L(s, \eta^s)|$ by $f \log \zeta(s)$. The factorials in the telescoping product cancel exactly with the Weyl denominators in $\prod_\tau W(\lambda)/W(u)$. Mass multiplier in the release

For the remaining real factors, compare every interior pair $1 \le i < j \le k-2$. Interlacing gives $u_i - u_j \le v_{i-1} - v_j$. The pairs $(i-1, j)$ run once through all nonadjacent pairs of the v's. Each of the other $k-2$ differences in $V(v)$ is 4Δ. The $2k-3$ pairs of u touching an endpoint have difference at most $8a_*$ eventually, uniformly in the word and place. At a top this follows from $u_0 \le 2a_*$ and $u_{k-1} \ge -a_*$, and elsewhere $k\Delta$ grows polynomially while $a_* = x \exp L$. Appending the negative entries to v contributes at least $\prod_j b_j^{k-1}$, since $v_i \ge 0$ and the b_j have gaps at least one. Hence

$$\frac{V(\lambda)}{V(u)} \;\ge\; \frac{(4\Delta)^{k-2}}{(8a_*)^{2k-3}} \prod_{j=0}^{m-1} b_j^{k-1}.$$

The chosen weights satisfy $b_j < e_j \le b_j(1 + 1/x)$ and $1 \le e_j \le 2 \exp L$. Dividing the logarithm by k therefore gives at least $\sum_j \log e_j - O_\varepsilon(L + (m/k)L + m/x)$, using $\log x = O_\varepsilon(L)$. Selected place bounds · Interlacing hypotheses

At a type place the index ratio is $g^k \prod_{s=k+1}^{r} (g^s - 1)$. Its power of g is $mk + m(m-1)/2$. The logarithmic cost from the factors $1 - g^{-s}$ is bounded by an absolute constant, since $g \ge 3$. In the mass multiplier, $\sum_{s=k+1}^{r} s = (m-1)k + m(m-1)/2$. After division by fk, the field power contributes at least $m \log(v_E/\pi) - L$. The factors $(2v_E)^{-f/2}$ and the Euler products together cost $O((m/k)(1+L))$, and the optional row factor costs $\log(3/2)/k$. Combining these inequalities, using $m \le x$ eventually, proves uniformly in w

$$\frac{\log(D_*/d_w)}{fk} \;\ge\; m(\log g_0 + \log g_1) + \frac{1}{f} \sum_{\tau, j} \log(e_j(\tau) v_E/\pi) - C_\varepsilon\big(L + m^2 L/k\big).$$

The summands in the average are nonnegative. Both $\log g_i \ge 8L$, and $k \ge m^2$, $m \ge 2C_\varepsilon$ eventually. The quotient is at least $15mL$. Also $2 \log \Gamma/(fk) = 2C_\Gamma m^5 L/k \le mL$ eventually. Thus $S_w = (D_*/d_w)^{1/(fk)} \Gamma^{-2/(fk)}$ satisfies $\log S_w \ge 14mL$. The signed row covolume calculation above identifies the mass factor at the chosen stage. The automorphic interpretation uses the specified level and real weights. Numerical window and penalty · Comparison on realized data

The norm-prime estimate above applies with $X = S_w$, $M = N(g_0 g_1)^2$, $h = 1/(fk)$, since $fk \le x^{12}$, $m \ge x^{1/3}$ and $\log M = O(x^{\varepsilon})$. It gives

$$S_w \le T(w) \le (1+1/(fk))S_w, \qquad \pi_T(w)\overline{\pi_T(w)} = T(w), \qquad \pi_T(w) \equiv 1 \pmod{N(g_0g_1)^2\mathcal{O}_K}.$$

The primes split completely in E and are larger than X' and the five initial level primes eventually. The host level set adds the $T(w)$. The choice is per word, with repetition permitted. Mass normalization, window bounds and their hypotheses · Pinned numerical window proof

At each of the f places above $T = T(w)$ the word level specifies a nonzero covector on the residue space. There are $T^k - 1$ possibilities since $\mathrm{GL}_k(\mathbb{F}_T)$ acts transitively on its nonzero covectors. The type flags have already removed all stabilizers. Thus adjoining the word-level data multiplies the class count by $(T^k-1)^f$. A coefficient-valued function on the framed classes can take an arbitrary value of the weight space on each class. Under the identification with that function space its dimension is $d_w(T^k-1)^f$. The resulting ratio in the final dimension condition is

$$\frac{\Gamma^2 d_w (T^k-1)^f}{D_*} = \left(\frac{T}{S_w}\right)^{fk} (1-T^{-k})^f.$$

The first factor lies in $[1, e]$. The second lies in $[1 - fT^{-k}, 1]$, and $fT^{-k} \le 1/2$ eventually. The ratio therefore lies in $[1/2, e]$. This is where the pre-level mass estimate is used in fitting the sizes after the prime has been chosen. Covector index, class count and space of coefficients · Final dimension condition

Write $N_j = k + j(t+1)$, $R_j = r + jt$ and $q_j = \lfloor (m-j)/2 \rfloor$. Take E-spaces with diagonal Gram matrix having the indicated number of plus signs first, and give them the self-dual lattices $\mathcal{O}_E^{\mathrm{rank}}$.

Space	Rank	Signature at each real place
$X_j, 0 \le j \le m$	N_j	$(N_j, 0)$
X'	R_m	$(R_m, 0)$
$Y_j, 0 \le j \le m$	R_j	$(R_j - q_j, q_j)$
$V_j, 1 \le j \le m$	$r + t$	$(r+t, 0)$
$W_j, 1 \le j \le m$	r	$(k + j - 1, m - j)$

Each form J defines its group by $\bar{A}^{\mathsf{t}} J A = J$. The positive forms give compact real factors. Successive X_j add $t+1$ positive lines. For $j < m$, embed X_j into the first N_j coordinates of Y_j. The remaining $m - 1 - j$ coordinates admit $\lfloor (m-1-j)/2 \rfloor$ hyperbolic pairs $e_+ + e_-, (e_+ - e_-)/2$. A spare negative coordinate occurs exactly when j is even, and the vector $\pi_T e_-$ there has norm $-T$. Here T is the norm prime of the distinguished word. At the top $Y_m = X'$, and the last line of X_m contains the perpendicular vector $\pi_T e$ of norm T. The coordinate lattice is self-dual since J is unimodular. Forms, embeddings and hyperbolic frames · Pinned signed model

At T the auxiliary column lattice $M_b = \mathcal{O}_E^{k-1} \oplus \pi_T \mathcal{O}_E$ satisfies $M_b + \bar{\pi}_T M_b^{\#} = \mathcal{O}_E^k$. The self-dual split move is given on the two factors of the selected coordinate by (T, T^{-1}). On writing $O = \mathcal{O}_{F,v}$ at $v \mid T$, $O \oplus O$ and $TO \oplus T^{-1}O$ meet in $TO \oplus O$, of index T on each side. At a such places the index on each side is T^a, in particular T^f if every place over T is used. Negative or positive orthogonal lines and their fractional ideals extend the lattice models, with the dual ideal on a unit-norm line equal to the inverse conjugate ideal. For a finite adelic isometry a_f the transported lattice is $E^{\mathrm{rank}} \cap a_f \widehat{M}$. Clear denominators of a_f and a_f^{-1} by a single integer D. Chinese remaindering in $D^{-1}M/DM$ realizes the prescribed completions and bounds the index on each side of the intersection by $D^{[E:\mathbb{Q}]\,\mathrm{rank}\,M}$. Left rational isometries

transport this intersection and the decorated right stabilizer preserves its reductions. For a sequence of finite adelic moves the product of the successive denominator bounds suffices. The shift u_0 resets the distinguished host index, with TT, T_2, weights and code carried by provenance. Local moves and integral gluing

The two automorphic passes and the point ledger

Release inputs (the automorphic realization). The first global step after the signed and decorated lattice models is the saturated automorphic seed. It is a row reading at the prescribed real weights and finite types, compatible with the column models and their intertwining maps. In the untwisted pass the active coefficient places, selected by the local required content of the lifts, must be a union of fibres over the fixed row coefficient field L_0. A whole fibre is selected whenever one of its places is. The first state has empty localization and zero local depth. Its pinned construction invokes the fine adelic reading and its extension in the ambient coefficient field, the lower-unitary multiplicity equality and the weight-specific rational row test. That test uses the specified local Weil-vector matching, the theta integral comparison and its recurrences for adapted plane changes. The finite-type trace argument and the construction of the separating periods below give the two spectral requirements under their respective hypotheses. The rationality step uses the normalized matched theta comparison on the same grid, including both adapted-plane recurrences. Saturated seed and precise spectral hypotheses

In the first pass choose the fixed fields, split-place orientations and true columns. The twist characters initially equal 1 and the central tuples include the two distinguished characters as well as 1. Keeping the prescribed Galois level, install the rational row lattice, cusp data, defining forms, induced series and moduli data. The strict-henselian and finite spectral steps give a finite-level state with the same support. The discrete-series replacement carries the actual row evaluation and its norm identity, the real volume normalization, the local types and the column models. Join the row, column, division and theta data using the matched characters and intertwiners. This gives the first core and its export. Construction through the spectral replacements · First and second pass

In the second pass use that export to select finite-order twists at the active places with conjugation, agreement over the fixed coefficient field and the local avoidance conditions. Enlarge the ambient coefficient field for their values while preserving the fixed embedded fields and the rational prime list. Run the row, spectral and kernel constructions again with these twists. The resulting core supplies good local twists for every required exponent prime to the relevant residue characteristic. The normalized volume and column-count comparisons identify its actual pre-level sizes with D_*, d_w. These comparisons, and the column denominator choice $\mathrm{MinkDiv}$, enter the recursive object theorem. Second pass and the required count and volume equalities · Actual two-pass assembly

The resulting RecOK records Δ, the error parameter $6x^{-C_F}$, finite active support, vanishing lift corrections off that support, fullness of the ramified division module, the column denominator and tame conditions, and local comparison at $T(w)$. For any valid size record at the same scale and weights, LedOK supplies the local containment and denominator bounds for the canonical parameters. After refilling the parameters, glue the pulled-back kernel lattice at served places, the prescribed comparison lattice above $T(w)$, and the ramified division lattice at all other places. The exceptional set is finite since the served primes are at most X'. The local saturation and comparison bounds give the stated localizations and allow the ledger to record their lengths divided by ramification indices. The division-lattice replacement preserves QD, which requires used tuples, good twists with conjugation and the specified prime-to-residue-

characteristic powers, and stability of the row lattice under the full coefficient integer ring. This gives the object at the original scale and weights, with provenance $\mathrm{Pt.From}$ and a canonical instance. Exact ledger gluing and QD · Extraction of the point

Saturation of the reference support

At the first pass fix the arithmetic host, the frames, and the prime norms. In the CM field E, there are only finitely many integral generators of a prescribed norm $\pi\bar{\pi} = T$. In every complex embedding the generator has modulus $\sqrt{T}$, and the Minkowski lattice meets this compact set in finitely many points. Choose a fine automorphic realization for each required generator tuple, assuming these realizations exist. Their reference fields lie in one finite coefficient field L_0. Each pair of full lattices in a common vector space is commensurable, so the selected comparisons have finite prime support. Reference choices and the two index pairs

Let $A \subseteq L_0 \subseteq L$ be the chosen reference field, fixed coefficient field and enlarged coefficient field in their distinguished embeddings. Write U_A for the support of the two signed reference indices, including their conjugates. For either pair M, N, scalar extension carries $M \cap N$ to the intersection after extension by flatness. At $w \mid v$ each finite quotient length is multiplied by $e(w/v)$, as one sees on a residue-field composition factor. Thus the signed difference also scales by $e(w/v)$. The Hermitian dual commutes with this extension in a basis by the inverse Gram matrix and its conjugate scalar action. Given the source's simultaneous base change of the two reference pairs, the active set in L is the inverse image of U_A together with the primes over rational residue characteristics where the prescribed content is nonzero. Both sets are unions of fibres over primes of L_0. Taking this restriction and setting all initial division depths to zero proves the required saturation. This gives finite support for each fixed host. A conductor exponent at a supported prime requires its own estimate. Restriction over the fixed field

How the finite types exclude proper blocks

Here is the cuspidal multiplicity argument for a row with the prescribed good Green types. Write $r = k + m - 1$ and $s = m/2$. The two unitary groups have signatures $(k + s - 1, s)$ and $(k + i, m - i - 1)$ at every real place for $0 \le i < m$. Their finite adelic isomorphism carries the chosen frames and levels. In the image of the finite projector on a passing irreducible constituent choose a single tensor of finite types $\bigotimes_v \tau_v$ appearing there. The polynomial condition on that tensor gives, for each $1 \le a < r$, a split place v where the degree-r polynomial labeling τ_v has no divisor of degree a. For $2a > r$ use complementary divisors. The place can depend on a. Polynomial condition on the selected tuple · Extraction from the actual projector

The projector in this argument can be identified coefficient by coefficient on the fitted finite quotient. Write $E_{v,f}$ for the character projector of the Green series labelled by f at v. Then $E_0 = \sum_{\boldsymbol{f}\ \mathrm{good}} \bigotimes_v E_{v,f_v}$. The embedding of the prescribed coefficient field sends the central idempotent of the reference class lattice to this very sum. For $0 < a < r$, average over the product U_a of the finite $(a, r-a)$ radicals at the selected places above g_0, and denote its idempotent by u_a. In each tuple there is a place whose polynomial has no degree-a divisor. The finite parabolic restriction there vanishes, so every tensor summand kills this average. Thus $E_0 u_a = 0$. Equivalently, $|G|[1](E_0 u_a) = \sum_{\tau:\, E_0|_\tau = 1} \dim(\tau) \dim(\tau^{U_a}) = 0$, by taking the regular trace. Every passing irreducible tuple therefore supplies a place for each cut. The level-character projection

restricts the image of E_0 further. This proves the finite support assertion on a row carrying the specified type data. Character descent and the trace test

At this place the radical of the finite parabolic with block sizes $a, r-a$ has zero invariants in τ_v. Indeed the inducing finite cuspidal blocks have the degrees of the irreducible factors of the Green polynomial with multiplicity. In the double-coset formula a block crossing the cut has invariants under a proper radical and contributes zero. The surviving blocks on the first side would give a divisor of degree a. In the local group, taking principal-congruence invariants of a representation induced from this parabolic gives finite parabolic induction. This follows by restricting its functions to $\mathrm{GL}_r(\mathcal{O}_v)$ and then reducing modulo the maximal ideal. Compact invariants are exact over $\mathbb{C}$, so every subquotient of the local induction also fails the type test. In particular a representation induced from a proper rational parabolic has zero contribution at the split witness for its first isotropic step. This excludes the residual spectrum for the chosen tuple. Finite and local parabolic argument

For the trace comparison assume stabilization of the discrete trace for these unitary inner forms and the archimedean pseudo-coefficient identities, with matching measures. Compress each finite Hecke function on both sides by the idempotent of the chosen tuple. At the split witness for a, averaging the idempotent over the integral radical gives zero. The parabolic descent of every compressed function there consequently vanishes. The proper elliptic endoscopic datum of sizes $a, r-a$ at that split place has exactly this transfer, up to the matching character factors. Every proper term of the stabilized comparison thus vanishes. At infinity the two prescribed discrete series have the same decreasing regular infinitesimal character. The stable elliptic orbital integral of the trace-one pseudo-coefficient is the signed Weyl character expression for that list. At signature (p,q) the sign is $(-1)^{pq}=1$ since r is odd, and the orbital integral vanishes on nonelliptic regular tori. The strong gap assumption gives the trace selector on the unitary spectrum. The identity datum therefore has equal stable traces on the two forms. Real orbital signs and spectral hypotheses · Stabilization in Kaletha–Mínguez–Shin–White, §3.1

At fixed level and infinity type the contributing cuspidal list is finite. In the common finite Hecke algebra, the nonzero spaces of level invariants of its distinct irreducibles are inequivalent simple modules. The density theorem on their direct sum provides an element acting as the identity on the specified constituent and as zero on all the others. The equality of the two traces gives the difference of the two multiplicities times the rank of the type projector on this constituent. Its rank is positive. This proves equality for the individual passing constituent, including the case of multiplicity zero. The argument uses the polynomial support condition on the same row chosen for the starting datum. Multiplicity separation and scope of the trace argument

A common normalization of the row periods

Work on the fine datum restricted to the active coefficient places, given by the support of the two reference indices and their conjugates, together with the places forced by the lift content. The restriction keeps the framed group, core and good-type idempotent and uses its restricted tuple dictionary, including the reference tuple μ_0. In the source, the period calculation takes place on the copy with zero division depths and the same group. The same-group comparison carries the tests back to this datum. Let V be the whole complex row, and write $F_\mu : A_\mu \to V$ for the map of the rational tuple source into that row. The Hilbert model H_{row} of V identifies with the good cuspidal multiplicity space $\mathrm{EEat}(1)$ via an injective map onto it. In the Hilbert calculation we use $F_\mu(x) = \mathrm{Th}_\mu(j_{X,\mu}x)$ and $K_{\mathrm{H}} = \mathrm{embLL}(L)$ for the image of the

ambient coefficient field. The scalar twist at the end gives the corresponding maps and field in V. Restriction and weight step · Same-group row construction · Good Hilbert model · Complex row

Cuspidal transfer on the restricted datum. For each $i < m$ and each irreducible admissible finite-adelic σ with a nonzero flag-level vector fixed by the good projector, the required equality is

$$m_{\mathrm{cusp},Y_0}(\lambda,\sigma) = m_{\mathrm{cusp},W_i}(\lambda_i,\sigma\circ\Phi_i^{-1}).$$

Here Φ_i carries the finite frames, and λ, λ_i are the prescribed real models at all real places. The preceding stable trace argument gives this equality under the stated transfer and pseudo-coefficient hypotheses. The reference reading uses the equality in the lower-unitary comparison. The weight comparison uses it on the restriction alongside rationality of the period maps. Good constituent and transported multiplicities · Separate spectral inputs

The separating periods. Fix the grid and its Weil data and Haar measures on the actual row. The constructor takes a nonzero smooth tensor w by descending a top test vector along the chain of intertwiners. For $u \in H_{\mathrm{row}}$ let b be its cuspidal row map from the smooth real tensor model, and set $f = b(w)$. With Φ a finite grid datum, ρ real Schwartz data, Ψ a definite column, and ξ a complex-linear functional on its coefficient space, use

$$P_{\rho,\xi;\Psi,\Phi}(u) = B^{+}\big(\mathrm{pairVec}(\xi,\Psi),\,\theta_{\chi}(\rho,\Phi)f\big).$$

Here θ_{χ} is the upward row lift on the grid. The period B^{+} first evaluates the lift at $n\iota(h)$ and integrates against the conjugate of the prescribed Bessel character over the unipotent flag. It then integrates against the scalarized column $\mathrm{pairVec}(\xi,\Psi)$ on the definite stabilizer. The conjugation acts on the fixed character. Evaluation at w and the convergent integrals are complex-linear in u. The row-tower and Bessel convergence results give this linearity on the entire good target. Test vector and period construction · Row lift and Bessel integral · Linearity · Convergence

Take any $u \neq 0$. Its map b is nonzero and belongs to the cuspidal row at the principal level above g_0g_1, since that subgroup lies in the framed row level with trivial character there. The bottom-period argument takes a real test at which the cusp form is nonzero and an adelic point where its modulus exceeds half its supremum. The real period estimate and the grid depth make the distinguished term of a translated coefficient Poincare series dominate the other terms. The series gives a form of the bottom core group whose pairing with the row on the embedded $U(Y_0)$ is nonzero. The real translates of the given nonzero w detect this same pairing by irreducibility of the real model and continuity. Changing variables moves the translate to the core form, and density gives a compact-finite core vector. Thus some bottom core form evaluated at a compact-finite vector pairs nontrivially with $b(w)$ in the pinned quotient measure. Principal level · Bottom evaluation and principal depth · Poincare series and its period · Transfer to the fixed tensor

At the specified even-rank core levels this core evaluation is in the span of descending theta chains from definite columns. The release proves surjectivity on these levels by checking the upward adjoints on the spanning Hecke components. It uses transferred parameters, nonvanishing of the real and semilocal zeta factors, and the normalized doubling formula with its nonzero L-value. Pairing with $b(w)$ is linear on the span, so some single descending chain gives a nonzero bottom period. The depth conditions and spanning by the admissible conjugated Fock vectors transfer nonvanishing first to a Fourier–Jacobi period, then alternately to Bessel and Fourier–Jacobi periods. Since $m \geq 2$ is even, the final Bessel period is $P_{\rho,\xi;\Psi,\Phi}(u) \neq 0$ for some choice of its data. These steps use the prescribed real branching and gap estimates and the split and unramified Weil comparisons, with the convergent Siegel–Weil and doubling

identities on the auxiliary core tower. Separation inputs · Surjectivity at the even ranks · Spanning by descending columns · Passage to one chain · Depth conditions and Fock tests · Last Bessel period

The terminal Bessel forms have dimension at most one at each real place in the prescribed gap. Taking their tensor product gives a factorization $P_{\rho,\xi;\Psi,\Phi}(u) = a(\rho,\xi)\lambda(u,\Psi,\Phi)$ for the fixed w. If the row is nonzero, take one nonzero period and fix its real pair (ρ_*,ξ_*), so $a(\rho_*,\xi_*) \neq 0$. For each nonzero u, a period found by the preceding argument has $\lambda(u,\Psi,\Phi) \neq 0$ and remains nonzero at this fixed pair. For each Φ, the columns indexed by a definite double coset and a weight basis vector span all definite columns. At the neat level their values at the representatives are the basis vector on the indicated class and zero elsewhere, and equivariance recovers the full column from those values. A period linear in the column and nonzero on this span must be nonzero on a basis column. The finite tests ψ_a are precisely these periods at (ρ_*,ξ_*), indexed by Φ, class and weight coordinate. Their common kernel on H_{row} is zero. The argument applies to arbitrary linear combinations of cusp forms and multiplicity copies, including row vectors outside the span of the tuple images. The good idempotent alone acts as the identity on this whole target. Terminal real line · Fixed pair and whole-row separation · Class and weight values · Span at fixed datum · Choice of the actual tests

The matched integral and relative scalar. The good projector e acts over the coefficient field on A_μ. It fixes the target, so equivariance gives $F_\mu(ex) = eF_\mu(x) = F_\mu(x)$. For the projected sources and the same tests just constructed the release uses

$$\zeta_\mu C_\mu\, \psi_a(F_\mu(ex)) = \kappa\, v_{\mu,a}(ex), \qquad C_\mu = \prod_{i,j} c_{ij,\mu}.$$

Each $c_{ij,\mu} \neq 0$ compares the two orders of the theta operations around a square. Flipping the squares exchanges the downward tuple path and the upward testing path and contributes C_μ and a finite phase. Unfolding the resulting path into finite cosets gives a second phase. Their product $\zeta_\mu \in K_{\mathrm{H}}^\times$ is a root of unity. The orthogonal-sum jet comparison gives a single κ for all terms, tuples and tests of this choice. The finite sum $v_{\mu,a}$ is over normalized branch pairings on the definite classes. With the prescribed character values, local phases, jets, and coefficient pairings in K_{H}, its rational volume factors and reciprocals of stabilizer orders show $v_{\mu,a}(ex) \in K_{\mathrm{H}}$. Projection of the sources · Simultaneous identity · Square flips · Unfolding and jet comparison · Rational finite sum

Matching in the local quotients. Write $I(\Phi;h)$ and $I'(\Phi';h)$ for the theta integrals over the two source quotients, of ranks a and a', with $h \in U(W)$ and $\dim_E W = a + a'$. Use the same trace character, the prescribed Weil splittings and $\nu = \alpha/\beta$. Right translation by the integrated group kills the finite augmentation differences. At a good split place these differences are the entire kernel of the local quotient map, and data of full source rank surject onto the primed quotient. Use such a representative at the distinguished good split place v_1. Match the sums of tensor images in the polarized real quotient and the finite quotients, where the spherical images agree off $T = S_0 \sqcup Q$ and the two split matrix zeta maps give the images at T. The pinned factorization uses vanishing against cusp forms to descend the polarized comparison to this common quotient. Its spherical, ramified and real Hom calculations determine the local comparisons used in descent. Under the stated rank and real weight gaps, absolute convergence and integrable Schwartz bounds, using full source rank on the primed side, extend the polarized comparison to the allowed Schwartz limits. Source note · Local maps and support · Full-rank partners · Quotient comparison and limits

Definite endpoint. Use the classical anisotropic Siegel–Weil value for the standard evaluation section at $s = (M-N)/2$, where M is the integrated dimension and N is the rank being doubled. It gives the theta

average multiplied by 1 for $M > N$ and by 2 for $M \leq N$. At the definite endpoint, double each source in turn and use $M = a + a'$ and $N = a$ or a'. For the box tensor of a test and its conjugate, vanishing of its theta lift forces the doubled zeta integral to vanish at this point. Unfolding with the pinned measures and clearing the good denominators before evaluation gives a product of nonzero real and exceptional zeta factors and continued Hecke values at $M/2 - i$, $0 \leq i < N$, for $\nu^{\pm 1}$. The real exponent gap is used for the values outside absolute convergence. Thus neither lift map vanishes identically. Gan–Qiu–Takeda, Theorem 7.1 · Yamana, Theorem 2 · Normalized unfolding · Endpoint nonvanishing

On the compact quotient of $U(W)$ the invariant mixed pairing on the common domain determines a scalar c for which $D = I - cI'$ is orthogonal to the primed image. Its Gram form is a scalar multiple of the primed Gram form. For a matched test with nonzero primed lift, the reverse theta integrals of D and I' are constant, with respective averages zero and a nonzero multiple of the primed norm squared. Double W with integrated rank a' and continue the unfolded Gram comparison with its moderate-growth bound. At $s = -a/2$ the factor-2 Siegel–Weil value computes both pairings from those reverse integrals. The Gram scalar of D is therefore zero. This gives $I(\Phi; h) = cI'(\Phi'; h)$ for all matched pairs and every h, and the nonvanishing witnesses give uniqueness and $c \neq 0$. Mixed pairing and reverse lift · Endpoint continuation · Conditional comparison constructor

Adapted-plane recurrences. Remove a hyperbolic plane from W and from the changing source. With $d = (a - a')/2$ at the upper stage, the β-move lowers a' and the α-move lowers a after the primed side is definite. On the adapted line parabolic the two Levi powers separate the constant-term components. In the β-component the unprimed term is evaluation and the primed term is a cone integral. In the α-component their roles are reversed. Unramified cone shells give the partial L^{S_0} factor, and ramified Fourier transforms give the epsilon factors with their conductor powers and uniformizer values. In the pinned witnesses the two line-ideal character values have exponents $(-1, -1)$ for β and $(1, 1)$ for α. Write $D_\beta(\nu)$ and $D_\alpha(\nu)$ for the products of these ideal values and the epsilon factors at $-d$ and $1 - d$, respectively. The local maps match evaluation to cone integration on the lower quotient. Lifting a nonzero lower comparison by surjective evaluation makes cancellation in the chosen component legitimate. With c_1 the lower scalar this gives

$$c\, G_\beta D_\beta(\nu) L^{S_0}(d + 1, \nu) = c_1, \qquad c = G_\alpha D_\alpha(\nu) L^{S_0}(d, \nu) c_1.$$

The geometric factors G_β, G_α are nonzero and fixed before the characters vary when the geometry and real maps are fixed. Starting at the definite endpoint, these recurrences determine the upper scalar. The pinned induction handles the remaining constant-term components and uses its cuspidality and orthogonality arguments to force the residual to vanish. Adapted planes · Line ideals · Cone factors · Constant terms · Induction

Compatibility on the selected grid. The square constructor transports both orders into the comparison triangle, including their Haar measures. Its nonzero matching scalar multiplies c to give the square scalar. The argument takes the classical theta and doubling facts and the pinned constructors as inputs. It still assumes that the selected grid identifies the unfolded quotient, compact and additive measures, the normalized real and finite maps on the actual tensors, and the required data in its chosen coefficient field. In particular, the stage transports must put the tuple and reference on common real maps before comparing their geometric factors. Rescaling the two quotient measures by tuple-dependent factors would change the square ratio by the corresponding relative Haar factor. Square comparison · Haar transport · Common real maps · Grid assembly

The definite comparison and the local and special-value ratios below give $C_\mu / C_{\mu_0} \in K_{\mathrm{H}}^{\times}$ when the indicated field conditions hold. The release chooses the ambient field to contain the required roots, square roots, jets, ratios and definite comparison data. If $\kappa \neq 0$, put $\widetilde{\psi}_a = (C_{\mu_0}/\kappa)\psi_a$. For every tuple and source vector,

$$\widetilde{\psi}_a(F_\mu(x)) = \frac{v_{\mu,a}(ex)}{\zeta_\mu \, C_\mu / C_{\mu_0}} \in K_{\mathrm{H}}.$$

The one nonzero rescaling preserves separation, and the reference ratio is 1. Dividing each phase and κ by ζ_{μ_0} sets the reference phase to 1. When a period image is nonzero, some $\psi_a(F_\mu(ex))$ is nonzero, so the simultaneous identity forces $\kappa \neq 0$ and fixes it from that same test once the phases of the matched comparison are fixed. This works even if the reference map vanishes. When $\kappa = 0$ all period images vanish by separation, and the unscaled tests work. A zero target admits the empty family. A zero good part of any source contributes zero by equivariance. Field data on this family · Rationality of the square ratios · Rescaling and zero case · Reference phase convention

In the complex row V, the source twists the scalar action on the Hilbert model by the identity or complex conjugation, according to its relation to the distinguished embedding at v_0. Apply the same conjugation convention to the output of each test. In the conjugate case a complex scalar is conjugated once on the input and once on the output, so the transported test is complex-linear. Its values on the tuple sources are in $K = v_0(L)$, with L_0 used separately for support saturation. Equivalently, the kernel of the joint map on $\left(\bigoplus_\mu A_\mu\right) \otimes_K \mathbb{C}$ descends to K. Composing with the separating tests gives a matrix over K with the same kernel, and field extension is exact. Conversely, a descended kernel gives a K-basis of the quotient whose images form a complex basis of the joint image. Extending this to a basis of V and taking dual coordinates also separates the complement of the joint image. Scalar twist of tests · Transfer to the complex row of the datum

The special values in the common row normalization

Consider a square (i, j) of the theta construction, with $0 \leq i, j < m$ and zero-based indices. Write $c_{ij,\mu}$ for its nonzero comparison scalar at tuple μ and μ_0 for the reference tuple. The associated unitary Hecke character of E is $\nu_{ij,\mu} = \mu_j^0 \xi_\mu / \chi_{m-i}$, where ξ_μ is a finite-order twist and $\xi_{\mu_0} = 1$. Use the finite Hecke L-function with its primitive finite Euler factors, and set

$$R_{ij,\mu}(s) = \frac{L_E(s, \nu_{ij,\mu})}{L_E(s, \nu_{ij,\mu_0})}, \qquad E_0 = \frac{t_{\mathrm{rk}} + r + 6}{2}, \qquad s_* = E_0 + 1.$$

Here t_{rk} is the odd rank increment denoted by t in the Hermitian construction, and $r = k + m - 1$ is odd. Indeed the rank at (i, j) is $r + j t_{\mathrm{rk}} + i(r + 6)$, so the half increment to $(i + 1, j + 1)$ is the same E_0 for every square. Character and tuple convention · Grid ranks

Nonvanishing at the integers in the comparison. At each place τ of the CM type the angular exponent of ν is $2h_\tau$ with $h_\tau > H_{ij}$, where H_{ij} is half the sum of the two diagonal ranks. To see the strict inequality, the weight construction writes $h_\tau = s + \kappa_{ij}(s, \tau)$ with $\kappa_{ij}(s, \tau) > 0$ up to $s = H_{ij}$. This exponent is common to every finite-order twist. Let e_{ij} be half the difference of the two bottom ranks. The rank formulas give $0 \leq e_{ij} \leq H_{ij}$. Then $L_E(n, \nu) \neq 0$ for every integer $1 - e_{ij} \leq n \leq H_{ij}$, and also for every positive integer for either auxiliary finite-order twist of the reference character. Angular weights and ranks · Hecke characters on the selected squares

For $n > 1$ this follows from the absolutely convergent Euler product. At 1, Hecke's continuation theorem makes both $L_E(s, \nu)$ and $L_E(s, \nu^2)$ holomorphic because the angular type of each is nontrivial. For real $\sigma > 1$ the Euler logarithm of $\zeta_E(\sigma)^3 |L_E(\sigma, \nu)|^4 |L_E(\sigma, \nu^2)|$ is nonnegative. At each unramified prime power its numerator is $3 + 4\cos\theta + \cos 2\theta = 2(1 + \cos\theta)^2$, and at a ramified prime the missing factors leave a numerator at least two. If $L_E(1, \nu)$ vanished, its fourth power would overwhelm the pole of order three, contradicting the lower bound one. The same proof applies to the inverse and the indicated twists. For $1 - e_{ij} \le n \le 0$ use Hecke's functional equation. The finite value of the inverse at $1 - n$ is nonzero. The gamma arguments on the two sides are $n + h_\tau$ and $1 - n + h_\tau$, positive integers at every place. All gamma, conductor and root-number factors are finite and nonzero. This proves nonvanishing of the finite value at n. Finite Euler-product convention · Reflection at integers

For the joint rational row test let $K \subset \mathbb{C}$ be the image of its coefficient field. On the compatible grid the upper half-rank difference starts at E_0. Put $q = \lfloor (m - j)/2 \rfloor$, $p = \lfloor (m - j - 1)/2 \rfloor$ and write $R = R_{ij,\mu}$. At the integers in the recurrences, L^{S_0} equals the primitive finite L_E because the omitted factors above S_0 are already 1 by ramification. Compare each move at μ with the same move at μ_0 after the common real-map transport. The β-moves increase d by one and put $R(E_0 + 1), \dots, R(E_0 + q)$ on the left. The α-moves decrease d by one and put the p ratios starting at $R(E_0 + q)$ in descending order on the right. Division leaves reference-over-tuple local factors for β and tuple-over-reference factors for α. The remaining factor $B_{ij,\mu}$ is the product of the step-factor ratios, the stage and endpoint transport units, and the prescribed definite bottom ratio. The latter comes from evaluating a common nonzero test in the polynomial quotient by finite orbit sums with the prescribed coefficients and volumes. Primitive Euler factors · Tower transports · Definite bottom ratio

The resulting identity is

$$\frac{c_{ij,\mu}}{c_{ij,\mu_0}} \prod_{a=1}^{q} R(E_0 + a) = B_{ij,\mu} \prod_{a=q-p+1}^{q} R(E_0 + a).$$

The empty product is one. Every cancelled factor is nonzero already in the half-plane of absolute convergence. Since m is even, $q - p = 0$ for odd j and $q - p = 1$ for even j. The square ratio is therefore $B_{ij,\mu}$ for odd j and $B_{ij,\mu}/R_{ij,\mu}(s_*)$ for even j. The remaining factor is inverted because it occurs on the left. The prescribed field data must place the local and transport factors, the bottom ratio and the required surviving L-ratios in $K^\times$ for the individual square conclusion. Multiplication gives the ratio C_μ / C_{μ_0} used in the common rescaling. Step-factor arithmetic · The two lengths · Square ratio and field implication

In particular, once the definite and local comparisons give $B_{ij,\mu} \in K^\times$, the single product

$$\mathcal{R}_\mu = \prod_{i=0}^{m-1} \prod_{\substack{0 \le j < m \\ j \text{ even}}} R_{ij,\mu}(s_*)$$

belongs to $K^\times$ exactly when $C_\mu / C_{\mu_0} \in K^\times$, where $C_\mu = \prod_{i,j} c_{ij,\mu}$. This supplies a sufficient field condition for the common rescaling of the separating tests in the row argument. There is a uniform archimedean estimate at the chosen embedding. Taking the logarithm from the convergent Euler series gives $|\log R_{ij,\mu}(s)| \le 2[E : \mathbb{Q}] \log \zeta(s)$. For $s \ge 2$ the sum after the term 2^{-s} is at most $\int_2^\infty u^{-s} du$, hence $|\log \mathcal{R}_\mu| \le 3[E : \mathbb{Q}] m^2 2^{-s_*}$. The following field calculation determines the extension generated by these values. Later valuation estimates in the release also use reflected values and bottom theta norms. Field comparison on the bottom triangle · Finite-place valuation argument

The field generated by the critical ratios

Write $K_0 = \mathbb{Q}(\sqrt{-d}) \subset E$ for the imaginary quadratic subfield, reserving B for an embedded coefficient field containing all embeddings of E and the finite character values on the grid. For a fixed tuple μ, let θ_μ be the finite character of $G_{K_0}^{\mathrm{ab}}$ corresponding by class field theory to $\xi_\mu|_{\mathbb{A}_{K_0}^\times}$. If $H_\mu^{(n)}$ is the fixed field of $\ker \theta_\mu^n$ in K_0^{ab}, then

$$B(R_{ij,\mu}(s_*)) = BH_\mu^{(1)}, \qquad B(\mathcal{R}_\mu) = BH_\mu^{(m^2/2)}.$$

The first equality also holds for every integer $1 \le s \le \mathrm{bigN}(D) + 1$ required by the individual-ratio clause. Here $\mathrm{bigN}(D) + 1$ is the release's common upper bound on that list of arguments. Field clause · Angular exponents throughout the list

Indeed both Hecke characters are algebraic of weight zero. Their angular exponents $2h_\tau$ agree at every complex place and $h_\tau > s$ on this range. The arguments $s + h_\tau$ and $1 - s + h_\tau$ of the archimedean gamma factors are positive. The finite-order twist contributes a rank-one Artin character. Blasius's critical-value theorem for CM Hecke characters and his twist formula over a totally complex field give the ratio, up to $B^\times$, as one nonzero Artin period. The common powers of $2\pi i$ cancel. The de Rham quotient selects one of the two Hodge entries $-h_\tau, h_\tau$ at each place. This half is independent of the square and of s, since there is no middle Hodge component. The embeddings in the chosen CM type all restrict to the same embedding of K_0. The transfer defining the Artin period is consequently the usual transfer from G_{K_0} to G_E^{ab}, or its conjugate for the other choice of half. On ideles it sends a K_0-idele diagonally into the E-ideles. Conjugation inverts the restricted twist because the twist is trivial on $\mathbb{A}_F^\times$. Blasius, §§M.11–M.17 and L.6–L.8 · Restriction of the selected CM type

For $g \in G_B$ the multiplier $g(R_{ij,\mu}(s))/R_{ij,\mu}(s)$ is thus $\theta_\mu(g)^{\pm 1}$, with a common sign. All these nonzero ratios for a fixed tuple span the same B-line. Fixing a single ratio amounts to killing the character, whereas fixing the product amounts to killing its $m^2/2$-th power. If $d_\mu = |\theta_\mu(G_B)|$, the two extension degrees are respectively d_μ and $d_\mu/\gcd(d_\mu, m^2/2)$. In particular the full field condition for the primitive ratios in one tuple can be tested at one square and one integer. The row rescaling uses only the product, given the local and definite bottom comparisons. · Primitive conductors and Blasius periods

In the ordinary existence construction the ambient row field can include all the individual class fields. At fixed preliminary datum and frame the moving seed supplies a finite set Ξ of twist characters before the final ambient field is chosen. Add generators for $H_\xi^{(1)}$ for each $\xi \in \Xi$ to the finite list of algebraic numbers prescribed in that ambient field. Generators for the host embeddings and character values are included as well. For the latter one may use the existing character-value lists on the representative grids for the finitely many keys (ξ, Q), where Q records oriented ramified places. Pullback along enlargement is surjective on coefficient places and preserves their residue characteristic. It therefore gives exactly the same twist set and the same avoided norms on E. The construction preserves the embedded small field $FO(D)$, whose degree and root discriminant occur in the size estimate. Consequently this prescription accommodates the ratios at the original cutoff and field-size budget. It supplies a field of definition. The local valuation inequalities for the squares, the definite bottom comparison and the identities for the chosen theta kernels enter separately. The field choices and order of quantifiers · Pullback along enlargement · Finite prescription and ring class fields

Conjugates and the signed finite degree

Let B contain the values of the characters on finite ideles and the images of the host under all its embeddings. For $\eta : B \hookrightarrow \mathbb{C}$, form $R_{ij,\eta}$ at s_* by applying η to the Euler coefficients at each prime of E. Blasius's theorem is an identity in a tensor product. With Ψ the selected CM type and φ the Artin character of the twist, it gives

$$(R_{ij,\eta})_\eta = (b_{ij} \otimes 1)c(\Psi, \varphi), \qquad b_{ij} \in B^\times.$$

If $g \in G_{\mathbb{Q}}$ and $\eta = g|_B$, the η-component of $1 \otimes g$ on a tuple (z_λ) is $g(z_{\mathrm{id}})$. The reciprocity formula for $c(\Psi, \varphi)$ therefore gives $g(R_{ij,\mathrm{id}}) = \omega(g) R_{ij,\eta}$ with $\omega(g)$ a root of unity common to all squares. For $g \in G_{K_0}$ the induced type gives $\omega(g) = g(\theta_\mu(g)^{\pm 1})$ with the same sign as in the field comparison. This application uses the full tensor identity, including the conjugates $\eta(b_{ij})$. Blasius, M.13, M.17 and L.6–L.8

For either Hecke character χ in the ratio, take a ray-principal power of an unramified prime $\mathfrak{p}$. Its chosen generator a is 1 modulo the conductor, so a power of $\chi(\mathfrak{p})$ is a product of integral powers of $\tau(a)/\bar{\tau}(a)$. The embeddings of a CM field preserve these conjugate pairs. Thus every conjugate Euler coefficient has modulus one. The identity $\chi(\bar{a}) = \chi(a)^{-1}$ follows from triviality on the real ideles and persists at every coefficient embedding. Distinct conjugate primes give positive paired factors in the Euler product. An unramified prime fixed by conjugation has coefficient 1 or -1, again giving a positive factor. Thus $R_{ij,\eta} > 0$. For $s \geq 2$, $\log \zeta(s) \leq 2^{-s} + \int_2^\infty u^{-s} du \leq 3 \cdot 2^{-s}$. Applying the Euler logarithm bound to both values in each of the $n = m^2/2$ ratios gives

$$\left|\log |g(\mathcal{R}_\mu)|\right| \leq 3[E : \mathbb{Q}] m^2 \, 2^{-s_*} \qquad (g \in G_{\mathbb{Q}}).$$

Consequently for any number field L containing $\mathcal{R}_\mu$ the same upper bound holds for $\left|\sum_{\mathfrak{P}} \mathrm{ord}_{\mathfrak{P}}(\mathcal{R}_\mu) \log N\mathfrak{P}\right| / [L : \mathbb{Q}]$, by taking the absolute norm. The finite prime contributions enter this sum with their signs. Angular characters and their restrictions

There is also an exact exponent for total positivity. Put $d = |\theta_\mu(G_{K_0})|$. For g fixing K_0, the phase of $g(\mathcal{R}_\mu)$ is $\omega(g)^n$, of order dividing $d/\gcd(d, n)$. Every other g is hc for h fixing K_0 and complex conjugation c, and c fixes the distinguished positive real value. Choose h with $\theta_\mu(h)$ of order d. Its image under h, and also its inverse, have order d. The phase of $h(\mathcal{R}_\mu)$ thus has order exactly $d/\gcd(d, n)$. This proves that $d/\gcd(d, n)$ is the least power giving a positive real number at every embedding. That power belongs to B by the field calculation. The relative extension degree over B uses $|\theta_\mu(G_B)|$, which may be smaller. Transfer and field of the Artin period · All conjugates and their phases

The valuation at the prime selecting the twist

Fix a coefficient place w over ℓ. Off the active set its tuple is the reference and every ratio is one. On the active set its character is $\Delta\psi_w$. The deformation Δ is the first deformation character except when ℓ equals its ramification prime, in which case it is the second. It is unramified above ℓ. The moving character ψ_w is trivial on the full local group $E_P^\times$ for every $P \mid \ell$, including the uniformizer. Restriction to $K_{0,\mathfrak{p}}^\times$ multiplies the local characters over $P \mid \mathfrak{p}$. Thus at every $\mathfrak{p} \mid \ell$ the class field of the restricted seed splits completely and the fields $H_{\mu(w)}^{(a)}$ are unramified for every positive power a. These statements persist over any completion B_b at $b \mid \ell$. For $v_\ell(\ell) = 1$ and $e_b = e(b/\ell)$, every individual ratio has valuation in $e_b^{-1}\mathbb{Z}$. Dividing it by the corresponding power of a uniformizer of B_b gives a unit whose residue transforms by the prime-to-ℓ part of

the restricted deformation character, up to the same inversion. Selection and switch · Triviality on the local group · The restricted class field at the selected prime

At an even column put $\nu = \nu_{ij,\mu_0}$, $R = L_E(s_*, \nu\Delta\psi_w)/L_E(s_*, \nu)$ and $R_d = L_E(s_*, \nu\Delta)/L_E(s_*, \nu)$. The quotient $Z = R/R_d$ has local multiplier that of ψ_w, so $Z \in B_b^\times$ at the compatible embedding. Its primitive numerator and denominator have identical Euler factors above ℓ. In the toric evaluation the denominator uses the deformation at the base order, and the numerator uses the enlarged order at which the moving primes are ramified. This gives the primitive Z in the square identity. Euler factors at the orders used in the comparison · The square at the selected place

The fitted theta estimate in the release has the following numerical consequence. Its local comparison gives $Z\,S_\lambda = uT_\lambda(a)$, with u a unit, S_λ the base sum of theta coordinates, and $T_\lambda(a)$ the twisted sum for an exponent prime to ℓ. The estimate for all these powers gives $v_\ell(T_\lambda(a)) \le v_\ell(S_\lambda) + 1/\ell^2$. Taking valuations in the identity and using $Z \in B_b^\times$ gives

$$v_\ell(R) - v_\ell(R_d) \le \frac{\lfloor e_b/\ell^2 \rfloor}{e_b}.$$

In particular the error is at most zero when $e_b < \ell^2$. The bound on the smaller field's root discriminant allows ramification at the deformation primes, so this last criterion is local to the specified base. Fitted theta hypothesis · Passage to the valuation of the ratio

The lower estimate uses the reference value $L_E(s_*, \nu)$. Let β_0 be its twelfth power with the prescribed CM-period normalization, P_0 its monic minimal polynomial over the row field, $e_w = \mathrm{ord}_w(\ell)$, and D_ℓ the loss per twelfth power in the toric denominator theorem. That theorem controls the normalized twisted value $\beta_t = R^{12}\beta_0$ at every place above ℓ. Summing the lower bound $v_\ell(\beta_t) \ge -12D_\ell$ over conjugates fixing the row field gives

$$v_\ell(Z) \;\ge\; -D_\ell - \frac{\mathrm{ord}_w(P_0(0))}{12 e_w \deg P_0} - v_\ell(R_d).$$

Here R is in the row field. The local degree, conductor and partial-zeta denominators enter D_ℓ through integrality of the toric lattice values. Rounding this lower endpoint up in $e_b^{-1}\mathbb{Z}$ is valid. These two estimates identify the arithmetic needed to bound the scalar at b. Reciprocity fixes its eigenline, and multiplication by an element of $B_b^\times$ preserves that eigenline. Denominator theorem · Reference normalization and its valuation · Both valuation bounds at the surviving integer

One character for all the theta evaluations

The estimate for $T_\lambda(a)$ above comes from choosing the order of the twist after the finite set of theta, row and Selmer tests is known. The numerical argument permits every exponent a prime to ℓ at the chosen order. Write π for a uniformizer in the coefficient ring of a test and $e = \mathrm{ord}_\pi(\ell)$. If $h(X) = \sum h_j X^j$ is integral with one known nonzero coefficient h_w, put $d_j = \mathrm{ord}_\pi(h_j)$ for nonzero coefficients. At a primitive ℓ^k-th root ζ, the valuations of the terms of $h(\zeta - 1)$ are

$$d_j/e + j/\varphi(\ell^k).$$

Indeed $(1-\zeta^a)/(1-\zeta)$ is a unit for every coprime a, and the product of $1-\zeta^a$ over the primitive roots is ℓ. At most $w + d_w$ positive depths k have two minimizing indices. At such a depth choose minimizers

$i_k < j_k$. Since $1/\varphi(\ell^k)$ strictly decreases with k, the two minimizing inequalities give $j_k \le i_{k'}$ when $k < k'$. The depths with $i_k < w$ therefore number at most w. For the others $j_k > w$, hence $d_{j_k} < d_w$ by comparison with the w-th term. These nonnegative orders strictly decrease through the remaining exceptional depths. At every other depth the unique minimum is the exact valuation for every primitive root. Exceptional depths · Valuation for all primitive roots

To apply this count, divide the split auxiliary prime pairs into left and right lists with their prescribed Frobenius coordinates. Clear the toric coefficients by ℓ^{N_0}. The group ring, written in variables $1 + X_q$, gives a polynomial whose specialization is the cleared theta sum. Dropping the right list with its character set equal to one gives the polynomial on the left list multiplied by the pair Euler corrections. The first nonzero residual term of each correction has degree zero or two in its own left variable. The constant term of the left-list polynomial has order $e(N_0 + v_\ell(S_\lambda))$, since the corrections for deleting its pairs at the trivial character are units. Substitute $X_q = (1+T)^{\ell^{e_q}} - 1$ with fixed widely spaced exponents e_q. Reduction gives distinct powers of T for the designated lowest terms of the Selmer polynomials. The right exponents can be taken large enough that the group relations contribute zero modulo π^{b+1} through the required degree, with b the order bound on the left constant. If the reduced Euler product first has a unit in degree d, its product with a polynomial of constant order $\le b$ has a coefficient of order $\le b$ by degree $(b+1)d$. Otherwise reducing the product first modulo π and then successively modulo π^{b+1} forces the constant to be divisible by π^{b+1}. These observations give one bounded-degree, bounded-order coefficient for each test. The bounds hold at arbitrarily high common auxiliary levels using the toric clearing, pair deletion, row determinant and auxiliary-prime hypotheses on the same data. Low-coefficient implication · Simultaneous substitutions

Sum $w + d_w$ over the finite test list, using the uniform bounds for the chosen coefficients. Choose the auxiliary level large enough to contain a consecutive interval of depths longer than this sum, with $w/\varphi(\ell^k)$ below the allocated error for every test. One depth then gives unique minima for them all. Its character has the prescribed nontrivial inertia on each chosen pair and is trivial above ℓ. For a toric label divide the selected value by ℓ^{N_0}. Its normalized valuation lies between $-N_0$ and $v_\ell(S_\lambda) + \ell^{-2}$ and is the same for every coprime power at that label. The second pass identifies the new character with a coprime power of the first, with the same test coefficients and a possible inversion. This is why the Fourier estimate applies to the character in the local ratio Z. Common depth and error · All-powers comparison · Full proof of the Fourier bound for the selected seed

The provisional and final kits on the same host

Choose the constants before the cutoff varies. For fixed sufficiently small ε and the host constants, choose the lower bounds for the standard record C, including the local projection thresholds. Choose c_D above its bound and $c_h \ge 402$ above the bounds that then apply. Take one onset satisfying the host, construction, weight and code requirements. These thresholds are uniform in x, the eligible hosts, and the shifts and words at each cutoff. In particular the word bound and the local projection estimates do not call for a separate onset for each word or prime. Recursive quantifiers · Uniform projection thresholds

Use a pair in the nonempty integer windows and one code to build a provisional kit at each supported cutoff. Each kit consists of a point for every shift $0 \le u_0 < n$. Apply the weight result in its family form, $\mathrm{WClosed}$, to a member at that cutoff. Provenance identifies its scale, rank, TT, prime classes and Frobenius coordinates with those of the originating host. The point and host signals agree with the same

numerical parameters and orbit representatives. Hence the recovered pair satisfies both weight requirements on that same host. Reapply object existence with this pair, the same host and code, once for every shift. The conditional arithmetic sampling proof above already takes place on the host, and the provisional-kit route is how the release carries it through to the final objects. Both passes of the automorphic construction take place within each application of object existence. Transport from the provisional point · Provisional and final choices · Final kit

Proved here (one family). Store the final kit at every $x \in S$ above one threshold x_1, with its objects and sizes as $F.D(x,i)$ and $F.L(x,i)$. For formal indexing the release chooses one fixed supported cutoff $y \geq x_1$ as a fallback. It uses the fallback at $x \leq x_1$, and the index subtype is empty for $x > x_1$ outside S. Every supported $x > x_1$ uses its own kit. Within a kit, $\mathrm{Pt.From}$ and $\mathrm{TieSame}$ preserve the ranks, q, d, QS, the Artin data, the torus and auxiliary and deformation primes, TT, T_2, the weights and code. Consequently every member has G_ψ at its actual scale. The design theorem requests a standard record large enough for the prescribed lower bounds, takes the common onset with the numerical requirements, carries the thick set, tame, QD, TT and word tests into the design, and sets $CTu = c_T(\varepsilon, 9)$ with the lower bounds adjusted for that assignment. Its cutoff set is still S. Indexing and same-cutoff data Actual scale and retained cutoff set Design construction

5. Fitting the comparison and consuming the family

The geometric existence route fixes a class of families and a selection on that class. **Release input (representation and geometric rows).** The input from local Langlands, the unitary group and theta correspondence includes the two theta-integral identities, matching local factors, and the definite-side covolume bound. The source obtains a Fourier coefficient fit and an evaluation grid. It takes the larger of the two grid thresholds for the row functions and for the evaluation theorem and uses that grid for both. The linked row bundles and unramified exponent condition give witnesses for the definite-side hypotheses. $\mathrm{classSel}$ combines the row-existence statements with the lattice-volume identities and the base and retwisting existence statements to obtain one class and one selection satisfying the row requirements together. Its existence proof is recS, using the host-to-design construction of §4 with the recursive input obtained from these rows. In particular the release calls $\mathrm{classSel}$ on a conjunction of the required rows. Choosing a separate family for each row would give a weaker existence assertion. Theta, grid and linked bundles Representation and covolume premises in the final call Common class and selection Simultaneous rows and existence on that class

Release inputs for the comparison, with the divisor and column constructions supplied through the interfaces of Papers II–V. On this class the closed row requirements include the common map on the column lattices, the Bessel pairing identification, the rank and fullness assertions, the twisted-place comparison LinkN, and the remaining norm and irreducibility estimates used in $\mathrm{LinkGLOBAL}$. “Closed” here specifies quantifiers. There is a positive cap for ε, then a required bound on the constants for each ε, then a conclusion for every family of the class satisfying that bound. Eventual pointwise conclusions may have an onset depending on that family. The rank input identified as $\mathrm{SmallLevelRankEq}$ equates each selected lift rank $d_{Z,0}$ to D_0. The divisibility input $\mathrm{LinkDIVt}$ takes $\mathrm{LinkEta}$ and DIVinN and supplies $\mathrm{LinkDIV}$. That output includes $B_k(u)(\nu_u) \subseteq CC$ for every word u, where ν_u is its column lattice and CC is the common target lattice, the Bessel comparison by unitary maps and a nonzero scalar, full row lattice, and the stated bounds on the heights and local errors. The common-map theorem and the

quantitative divisor, column and norm statements are substantial inputs. Closed quantifiers Common-map definition Bessel, rank and divisibility interfaces Actual comparison inputs

Proved here (using the length budget). At each point, VVle and the identification $\mathrm{sha_eq}$ of its record lengths with the localization lengths give finiteness and bound the nonnegative finite partial sums of vvTerm by $\sqrt{b}$. Their supremum therefore bounds the infinite sum by $\sqrt{b}$. Let $L = x^{\varepsilon}$, $B_0 = fmbL$, and $N_c \geq 0$ be the size budget of the point. With the selected lift rank identified by $\mathrm{SmallLevelRankEq}$, LinkN bounds the local sum etaSum by

$$D_0 \left[C\left(fmbL^2 + (\log x)^2 N_c \right) + \textstyle\sum_{\ell} \mathrm{vvTerm}(\ell) \right].$$

At a sufficiently late cutoff, $L \geq 1,\ fm \geq 1,\ (\log x)^2 \leq L$ and $\sqrt{b} \leq b \leq B_0$. The first term in the bracket is at most $|C| L (B_0 + N_c)$, and its other term is at most $L(B_0 + N_c)$. This is DIVinN with constant $|C| + 1$. A maximum of the thresholds for the three inputs and the logarithm estimate makes the deduction uniform over the late members. Thus the localization lengths supply the error allowance in the divisibility bridge. The independent place coverage, column and lattice assertions still enter the bridge. Identification and summability Input inequality and desired bound Pointwise and family deductions

Proved here (order of choices). Intersect the cap from design existence with the caps from the closed assertions. After fixing ε below this positive minimum, join the finitely many required lower bounds on the parameters. Request the single design family at this joined bound. Apply its closed assertions, including the preceding cohomological deduction for divisibility, to this witness F. The output $\mathrm{LinkGLOBAL}$ takes $\mathrm{ChosenE}, \mathrm{LinkG}, \mathrm{LinkS}, \mathrm{LinkIRR}, \mathrm{LinkDIV}, \mathrm{LinkMu}$ on F, and returns constants and one onset with $\mathrm{ChainTied}$ at every later $x \in S$. Its proof combines the onsets of its input estimates. A tied chain has $D_d.x = x$, chooses all its points from I_x, and identifies its arithmetic data, top set, weight windows, and word types with those of F. The actual shape call supplies $e_D = 2, e_L = 0$, independent of ε. This proves the assembly theorem of §1 relative to its stated inputs. Closed join and same-family application Actual exponents Global premises and onset Common size threshold Chain ties

Supplied by Papers II–V (scalar comparison). The scalar comparison uses the chain and the arithmetic condition at the same member i and the same TT. It yields, after enlarging the threshold to cover the chain, arithmetic and scalar onsets,

$$G_z(x) = \sum_{n \leq x} \Lambda_{\mathrm{vM}}(n)(x/n - 1) - x \log x + (1 + \gamma_E)x \ \leq\ x^{1/2 + 9\varepsilon} \qquad (x \in S,\ x \text{ sufficiently large}).$$

Here $\mathrm{Level.cyc}(\mathbf{1})$ keeps only the trivial character, so its correction sum over the other kept characters is empty. Given $1/2 < \theta < 1$, choose ε below the common cap with $9\varepsilon \leq \theta - 1/2$ and then take the family.
Proved for ordinary zeta in Paper I, Theorem 1. The released window criterion uses the resulting bound by x^{θ} on arbitrarily late $[Z, Z^{3/2}]$. Its analytic premise, FormB, turns this into zero-freeness to the right of θ for the kept characters in the release's pole convention. Ordinary zeta is the conductor-one case. The classical zeta functional equation and the given nonvanishing statements give the asserted critical-line conclusion. The full theorem takes this step for general kept characters and then passes through the imprimitive extension. Conductor-one kept set Scalar input uses the same member Window estimate and analytic quantifiers Analytic window premise Final pass to all characters Ordinary critical-line reduction

The favorable-prime refinement of §3.3 can also give a real tail of cutoffs, with the other inputs of this construction. **Proved here from that refinement.** Given a sufficiently large Z, use $X = Z^2$ in the conductor-one selection and take $W = Z$. Since $c = k = 1$, §2 supplies $[Z, Z^{3/2}]$, including its left

endpoint, so every sufficiently large x belongs to S. The releveling, provisional-family and final-family constructions preserve that set. The same-family comparison therefore holds on a full tail and the ordinary all-cutoff Landau consumer applies. The proof of the exported theorem already uses the weaker window assertion. This refinement still needs the pointwise host cohomological estimates and the common scalar, divisor and column inputs. Preserved cutoff set in design All-cutoff analytic interface Prior tail trace

For the favorable ordinary pair, §3.3 supplies the dyadic term on the entire prescribed interval. The general pair-independent route in §3.1 still uses the exceptional-twist H^2 estimate. The classical proof in §3.4 supplies the odd minus term for every host in this specialization. In the odd plus term the remaining steps include the ordinary model and special-fibre hypotheses of the Count I cotangent comparison in §3.6, the rank-three attachments in I and V, and the rank-four automorphic construction in II. The weight selection and signed Hermitian spaces in §4 are on the same arithmetic host. The internal Gram and integral block-lattice estimates in §4, the occurrence and multiplicities of the selected representations, the saturated automorphic seed with the prescribed types, and the theta comparisons in §§4–5 still enter the geometric construction. Cohomological assembly Existence and comparison assembly