Validated Dirichlet approximation and weighted arithmetic moments

GRH research working note

6 November 2026

Abstract

Adaptive outward dyadic arithmetic and theta-integral quadrature give a requested-error construction for three fixed completed Dirichlet functions. Fresh samples and executed rational readout training certify the supplied uniform 10^{-24} examples on the stated compact; no practical extreme-precision runtime bound is asserted. Separately, a fixed-twist sextic Hecke-family moment improves an original masked remainder estimate, while retaining a positive cutoff loss. We give the precise external analytic inputs and distinguish these results from zero-location and spectral determinant claims. A classical quadratic pivot improves one row sector; joint Gauss phases have an exact masked identity. An ordinary all-character conductor comparison improves the unsigned raw envelope to $H + DH^{1/6} + (DH)^{2/3}$, while an executed finite arithmetic Fourier readout certifies correlation and signed-observable errors without asserting a varying-family signed saving. A moving common-profile construction aligns the actual coefficients on nonprincipal sixth-power rows, and an exact joint example rules out variance-only centering; the fixed-profile signed target remains open. An exact conductor-12 character twist and gcd–Möbius inversion give a separate bound for the signed coprime slice, with the former 7/6 comparison retained as history. A signed common-divisor zeta expansion and a complete-column exact laboratory retain every raw mask; weighted divisor certificates do not close the required asymptotic saving. An exact rough-prime support identity gives an explicit target-budget mask correction; the remaining coherent signed main term is still unestimated. A pinned fixed-order preprint theorem, applied to the physical sixth-power-free radicands, improves the global unsigned critical power to 9/8 and the $A_6 = 1$ signed slice to 13/12. Mean-retaining coherent reductions and conditional scalar-input diagnostics keep the fixed-profile signed target open. No proof of GRH or complete Lean formalization is asserted. Primitive root normalization, all-component finite dual prefixes, actual-cutoff Schur transfer and a full equal-norm localization retain the coherent unequal-norm target as an open problem. An actual Gauss-weighted short-divisor tail admits a physical cubic-matrix estimate. A ramified theta construction now covers the actual fixed finite-character family, and a whole-factor partition bounds the original pure finite dual prefix with normalized power 101/144, retaining every ray, raw zero and finite cutoff. The infinite-frequency complement and full signed assembly remain open; the whole signed and unsigned exponents are unchanged. The paired cutoff analysis gives target-size estimates on stated pure and mixed selections of the signed equality–complement cross, with an exact ordered divisor identity.

1 A finite, validated approximation result

Let $\chi_3 = (0, 1, -1)$, $\chi_4 = (0, 1, 0, -1)$ and $\chi_5 = (0, 1, -1, -1, 1)$, extended periodically, with parities $a = 1, 1, 0$. We approximate the completed functions

$$\Lambda_q(s) = (q/\pi)^{(s+a)/2}\Gamma((s + a)/2)L(s, \chi_q).$$

The primitive theta/Mellin identity is the explicit classical analytic input; no hypothesis about zeros is used. These three examples do not assert a uniform result for all conductors, general Hecke functions or Artin holomorphy.

Theorem 1 (Scope of the supplied finite certificate)**.** *For each $q = 3, 4, 5$, the supplied exact-rational calculation constructs a complex polynomial P_q of degree 28 satisfying*

$$\sup_{|s-1/2|\le 7/10} |P_q(s) - \Lambda_q(s)| < 5 \cdot 10^{-15}.$$

The computation includes enclosures of all 64 true Fourier-circle nodes, completed-function sample error, analytic truncation, Fourier aliasing, the exact closed-form 180-step gradient-descent iterate and coefficient midpoint rounding. The Python implementation and informal analytic proof have independent checks; this is not a kernel-checked Lean theorem.

Proof. Write $F(c+z) = \sum_{n\ge 0} a_n z^n$, with $c = 1/2$, $R = 7/10$, $\rho = 6/5$, $r = 12/5$. The independently recomputed rational theta majorants are strictly below $M_q = 4, 10, 6$, respectively, on $|z| \le r$. Use true roots $\zeta_j = e^{2\pi i j/64}$ and rational labels y_j with proved error $|y_j - F(c + \rho\zeta_j)| \le \eta$. Use the real Euclidean gradient on complex coefficient coordinates for

$$\mathcal{L}(b) = \frac{1}{128} \sum_{j=0}^{63} \left| \sum_{n=0}^{28} b_n \zeta_j^n - y_j \right|^2.$$

Discrete Fourier orthogonality gives the least-squares optimum $b_n^* = 64^{-1} \sum_j y_j \bar{\zeta}_j^n$, $0 \le n \le 28$, and gradient $\nabla\mathcal{L}(b) = b - b^*$. For step $1/4$ and zero initialization, the exact 180-step iterate is $b_n^{180} = (1 - (3/4)^{180}) b_n^*$. The code evaluates this closed form; it does not claim an explicit180-update optimizer run.

Set $u = R/\rho = 7/12$, $v = R/r = 7/24$ and $w = (\rho/r)^{64} = 2^{-64}$. Let $\widehat{b}_n$ be an outward rectangle enclosing b_n^* and ϵ_n the sum of the coordinate halfwidths of the rectangle for b_n^{180}. Choose each actual polynomial coefficient as that rectangle's exact rational midpoint c_n. The normalized basis is explicit:

$$P_q(s) = \sum_{n=0}^{28} c_n \left(\frac{s - 1/2}{6/5} \right)^n.$$

Thus the exported coefficient with index n multiplies the normalized variable, rather than the raw power $(s - 1/2)^n$. Its error on the compact disk is bounded by

$$E_{\text{opt}} = (3/4)^{180} \sum_{n=0}^{28} |\widehat{b}_n|_1 u^n, \qquad E_{\text{arithmetic}} = \sum_{n=0}^{28} \epsilon_n u^n,$$

$$E_{\text{oracle}} = \eta \frac{1 - u^{29}}{1 - u}, \qquad E_{\text{alias}} = M_q \frac{w}{1 - w} \frac{1 - v^{29}}{1 - v},$$

$$E_{\text{Taylor}} = M_q \frac{v^{29}}{1 - v}.$$

Here $|\widehat{b}_n|_1$ is the sum of the maximum absolute real and imaginary coordinates, a safe upper bound for every enclosed complex modulus. Cauchy's estimate $|a_n| \le M_q r^{-n}$ and the Fourier identity $b_n^* = a_n \rho^n + \sum_{\ell \ge 1} a_{n+64\ell} \rho^{n+64\ell}$ for exact labels prove the alias and Taylor bounds. Label perturbations contribute at most η per Fourier coefficient. The remaining two terms follow directly from the exact iterate and coefficient rectangles. Each term is rounded upward to a 180-bit dyadic rational; the reported total is the exact sum of those reported terms. The certified totals below are strictly below the asserted threshold. Detailed sample enclosures are proved in the extended edition and encoded with exact integer endpoints. □

Conductor	Maximum sample radius	Uniform polynomial bound
3	3.026×10^{-17}	1.786×10^{-15}
4	7.206×10^{-17}	4.455×10^{-15}
5	1.404×10^{-16}	2.906×10^{-15}

The displayed values are rounded upward; exact rational records establish the strict inequalities.

2 What this does and does not certify

The final JSON records exact rational centers, interval endpoints, source hashes and all five error terms. Float values are display summaries only. A high-precision comparison is an independent diagnostic, not the reason an interval contains its mathematical target. Arithmetic uses outward integer operations, and invalid float/order inputs are rejected, including through typed caches. Earlier pilots and failed validation probes are retained.

This is a fixed finite tolerance/configuration. Increasing accuracy needs new truncation and precision choices; the fixed $N = 11, T = 128$ interface is not an arbitrary-accuracy oracle. The broader approximation existence theorem is separate. The certificate says nothing about the location of zeros: P_q is not required to be real-rooted or a self-adjoint spectral determinant. The independent arithmetic determinant identity, positivity and original signed short-row/covariance estimates remain open. No GRH proof or complete new Lean4 formalization is claimed.

3 Validated primitive operations

Real intervals are closed dyadic intervals $[l2^{-224}, h2^{-224}]$ with integer endpoints; complex intervals are Cartesian products. Addition is exact at this scale. For multiplication take the minimum/maximum of all four integer endpoint products, divide by 2^{224} and round down/up. For a reciprocal whose interval omits zero, reverse the endpoint order and round $2^{448}/h$ down and $2^{448}/l$ up. Negative denominators and scalar division use the same endpoint rules. Each operation encloses all represented real values by monotonicity; induction proves inclusion for every expression built from them. Rectangular complex multiplication uses these real rules. The sum of coordinate halfwidths bounds the complex distance from its exact rational center. Binary floating arguments are never silently converted.

For $0 < x < 1$, the alternating arctangent sum and the endpoint obtained by adding the next signed term enclose $\arctan x$. Sixty terms at $1/5$ and $1/239$, together with Machin's identity $\pi = 16\arctan(1/5) - 4\arctan(1/239)$, give the interval for π. The negative coefficient reverses the second interval. The identity follows from tangent addition and the range of the two angles; no decimal value of π is a trusted input.

For rational $x > 0$, range-reduce $x = 2^k v$, $1 \le v < 2$. With $z = (v-1)/(v+1) \le 1/3$,

$$\log v = 2\sum_{j=0}^{J-1} \frac{z^{2j+1}}{2j+1} + \mathcal{R}, \qquad 0 \le \mathcal{R} \le \frac{2z^{2J+1}}{(2J+1)(1-z^2)}.$$

The geometric tail bounds each remaining denominator below by $2J+1$. Use $J = 92$, the analogous expression for $\log 2$, and sign-aware interval multiplication by k. This covers negative k as well.

For complex z, repeatedly halve until its interval ℓ^1 magnitude bound is $A \le 1/2$. The degree 60 Taylor polynomial for e^z then has remainder at most

$$\frac{A^{61}}{61!}\frac{1}{1-A/62}.$$

Successive absolute terms have ratio at most $A/62$, so geometric comparison proves this bound. Add its symmetric interval to each coordinate, then square outward the required number of times. The identity $e^{2z} = (e^z)^2$ and interval inclusion prove the final enclosure. Caches are type-sensitive: otherwise a cached exact argument can accidentally bypass rejection of an equal-valued float. No implementation change is allowed during the recorded run; the producer verifies source bytes again before emitting its certificate.

4 Recomputing the compact growth constants

For the larger growth disk, $c = 1/2$, $r = 12/5$, take $k = 2$ when $a = 1$ and $k = 1$ when $a = 0$. Both real Mellin exponents on that disk are at most k. Define

$$Q_k(x) = \sum_{j=0}^{k} \frac{k!}{(k-j)!x^{j+1}}, \qquad U_{16}(x) = (1 + x/16)^{-16}, \qquad \delta = 3/q.$$

Integration by parts gives $\int_1^\infty t^k e^{-xt}\,dt = e^{-x}Q_k(x)$. The same Gaussian split and geometric-tail comparison used below, now with cutoff $n_0 = 8$, imply the exact rational majorant

$$\widehat{M}_q = 2\left[\sum_{n=1}^{7} n^a U_{16}(\delta n^2)Q_k(\delta n^2) + C_a Q_k(64\delta)\frac{U_{16}(32\delta)}{1 - U_{16}(8\delta)}\right], \quad C_0 = 1, \quad C_1 = (1 + \delta^{-1})/2.$$

Replacing the actual $d = \pi/q$ by the smaller positive δ increases the integral majorant. Exact fraction comparison gives $\widehat{M}_3 < 4$, $\widehat{M}_4 < 10$, $\widehat{M}_5 < 6$. The producer recomputes these fractions, asserts the strict comparisons and exports them; the constants in the readout proof are thus checked rather than supplied by a floating growth estimate.

5 The completed-function integral and its tails

Let $d = \pi/q$ and $\theta_q(t) = \sum_{n\ge1} \chi_q(n)n^a e^{-dn^2 t}$. The standard primitive theta transformation, followed by splitting the Mellin integral at1, gives

$$\Lambda_q(s) = \int_1^\infty \theta_q(t)t^{(s+a)/2-1}\,dt + \int_1^\infty \theta_q(t)t^{(1-s+a)/2-1}\,dt.$$

The second root-number coefficient is1 for these three characters: their finite Gauss sums are $i\sqrt{3}$, $2i$ and $\sqrt{5}$, respectively. For the last, $2\cos(2\pi/5) - 2\cos(4\pi/5) = \sqrt{5}$; the first two are direct conjugate-root differences. Division by $i^a\sqrt{q}$ gives the claimed normalization. Exponential decay makes the two integrals entire. This classical theta identity, not an assumption of Artin holomorphy or desired positivity, is the analytic input of the implementation.

The actual circle nodes are enclosed by computing $e^{2\pi ij/64}$ with the validated primitives. All input boxes fit $D = [-1, 2] + i[-4/3, 4/3]$. Thus both exponents $p = (s+a)/2-1$ and $p' = (1-s+a)/2-1$ satisfy $|\Re p| \le 2$, $|\Im p| \le 2/3$. The domain guard precedes all term discards. Retain $1 \le n \le 11$ and integrate $1 \le t \le 128$.

Put $\delta = 3/q < d$, $U(x) = (1+x/128)^{-128}$ and $Q(x) = x^{-1} + 2x^{-2} + 2x^{-3}$. Since $(1+x/128)^{128} \le e^x$, $e^{-x} \le U(x) < 1$ for $x > 0$. Integration by parts gives $\int_1^\infty t^2 e^{-xt}\,dt = e^{-x}Q(x)$. With $C_0 = 1$, $C_1 = (1 + \delta^{-1})/2$ the omitted-n tail of both integrals is

$$2C_a Q(\delta 12^2)\frac{U(\delta 12^2/2)}{1 - U(12\delta)}.$$

Indeed $n^a e^{-\delta n^2/2} \le C_a$, the function Q decreases, and $(12+j)^2 \ge 12^2 + 24j$. These inequalities yield a geometric bound after splitting the exponent in halves. The retained-n time tail is at most

$$\sum_{n=1}^{11} 2n^a 128^3 U(128\delta n^2) Q(128\delta n^2).$$

It follows by substituting $t = 128u$ and majorizing $t^{\Re p}$ by t^2. The two tails are disjoint: omitted n use the full integral, whereas retained n contribute only $t > 128$.

6 High-order integration with a certified remainder

Split each dyadic interval $[2^j, 2^{j+1}]$, $j = 0, \dots, 6$, into four equal panels. For its midpoint m and halfwidth h, $h/m \le 1/9$. For $f(z) = e^{-\lambda z} z^p$, $\lambda = \pi n^2/q$, the principal logarithm is analytic on $|z - m| \le \rho = m/2$. On that circle,

$$|f(z)| \le M_\rho = 3 \max\{(3m/2)^2, (m/2)^{-2}, 1\} U(\delta n^2 m/2).$$

Here $\Re z \ge m/2$, $|\arg z| \le 1$ and $e^{|\Im p|} \le e^{2/3} \le 3$. The maximum handles both signs of $\Re p$. Cauchy's coefficient estimate and cancellation of odd powers on $[-h, h]$ therefore bound the error after degree $2K$ by

$$\frac{2hM_\rho (h/\rho)^{2K+2}}{(2K+3)(1-(h/\rho)^2)}.$$

Indeed integrate each remaining even term, replace its denominator by the first omitted denominator, then sum a geometric series. Here $K = 12$. Multiply by $2n^a$ for both Mellin terms and sum all 28 panels and retained n. Keeping this majorant even for later discarded terms only overcounts error.

To evaluate the polynomial stably, set $u = h/m$, $q_{-1} = 0$, $q_0 = 1$ and use

$$q_{j+1} = \frac{u(p - \lambda m - j) q_j - (\lambda m) u^2 q_{j-1}}{j+1}.$$

Multiplying $f' = (-\lambda + p/t) f$ by $t = m + x$ and matching coefficients proves this recurrence, with $q_j = f^{(j)}(m) h^j / (j! f(m))$. The integrated polynomial is $2he^{-\lambda m} m^p \sum_{j=0}^{K} q_{2j}/(2j+1)$; all its arithmetic is outward.

Fixed absolute exponential-rounding errors can be amplified by high-order coefficients in late panels. The code therefore first bounds a complete (n, panel) pair, including both powers, by

$$4hn^a (m+h)^2 U(\delta n^2 (m-h)).$$

For real $t \ge 1$ its phase has modulus1 and $t^{\Re p} \le (m+h)^2$, proving the bound. If it is below 2^{-90}, skip that pair entirely and add its bound to the final radius. This is controlled truncation, not numerical underflow. The final sample radius is the interval-sum halfwidth plus the omitted-n, time, Cauchy and discard bounds. All are recorded separately and checked as nonnegative. No high-precision special-function routine is trusted here.

7 Reproduction and remaining interfaces

Run the preserved source on the existing Linux Python environment with no new packages: `python3` on `examples/analytic_oracle/theta_oracle.py`, with argument `-output NEW`. It refuses to overwrite an existing experiment directory. The final-v5 experiment contains 192 samples, three rational

degree 28 polynomials, copied source and SHA-256 hashes. Independent validation checks rational endpoint arithmetic, transcendental enclosures, 1160 exact coefficient identities, high- precision completed-function comparisons and the exported component sums.

The fixed-parameter oracle solves a concrete evaluation/certification gap. It does not identify a positive arithmetic operator, force real-rooted polynomials, prove the signed original-family short-row bound or supply complete Lean4 verification. Surface/order and neural architectures may use these certificates as independent targets, but a small learned loss is not a substitute for the remaining arithmetic and categorical theorems.

8 Requested-error readouts from preserved certified labels

The accepted final-v5 sample certificate is preserved unchanged. For each $q = 3, 4, 5$, it supplies $m = 64$ rational complex labels y_j, bounds $|y_j - \Lambda_q(c + \rho\zeta_j)| \le \eta_j$, and outward rectangles enclosing the true roots $\zeta_j = e^{2\pi ij/m}$, where $c = 1/2$, $\rho = 6/5$. The classical primitive theta/Mellin identity and accepted growth bounds $M_q = 4, 10, 6$ on $|s - c| \le r = 12/5$ remain the analytic inputs. This extension accepts an exact rational error request $\varepsilon > 0$ on the fixed compact $|s - c| \le R = 7/10$. It changes the finite readout, while retaining the certified sample errors and the 64-node alias budget. It is not an arbitrary-accuracy theta-sampling interface.

Theorem 2 (Executed rational readout with explicit refusal)**.** *For each supported conductor and rational request, the finite-data selection below either returns a normalized rational polynomial P and an executed, exactly replayable gradient-descent trace with*

$$\sup_{|s-c|\le R} |P(s) - \Lambda_q(s)| \le \varepsilon,$$

or refuses the request with the exact minimum base budget over all 64 *allowed degrees. Refusal concerns attainability by this particular fixed-data majorant; it is not an impossibility theorem for approximation.*

Proof. Let B_n be an outward rectangle enclosing $b_n^* = m^{-1} \sum_j y_j \bar{\zeta}_j^n$, $0 \le n < 64$, computed from the preserved root rectangles. Put $t_n = \mathrm{mid}(B_n)$ and let e_n be the sum of its coordinate halfwidths. These are exact rationals, and $|t_n - b_n^*| \le e_n$. For a selected $d < 64$ execute, using exact rational operations,

$$b_n^0 = 0, \qquad b_n^{k+1} = \tfrac{3}{4} b_n^k + \tfrac{1}{4} t_n, \qquad 0 \le n \le d.$$

This is genuine gradient descent with step $1/4$ for the explicitly specified real Euclidean objective $J(b) = \frac{1}{2} \sum_{n=0}^{d} |b_n - t_n|^2$, whose gradient is $b - t$. All iterates and exact objective values are retained. The output is literally

$$P_k(s) = \sum_{n=0}^{d} b_n^k \left(\frac{s - c}{\rho} \right)^n .$$

Set $u = R/\rho = 7/12$, $v = R/r = 7/24$, $w = (\rho/r)^{64} = 2^{-64}$, $S_d(x) = \sum_{n=0}^{d} x^n$ and $\eta = \max_j \eta_j$. The five budgets are

$$E_{\mathrm{opt}}(k) = \sum_{n=0}^{d} |b_n^k - t_n|_1 u^n, \qquad E_{\mathrm{oracle}} = \eta S_d(u),$$

$$E_{\mathrm{alias}} = M_q \frac{w}{1 - w} S_d(v), \qquad E_{\mathrm{Taylor}} = M_q \frac{v^{d+1}}{1 - v},$$

$$E_{\mathrm{arithmetic}} = \sum_{n=0}^{d} e_n u^n.$$

Here $|x+iy|_1 = |x| + |y|$. Cauchy's estimate and true-root Fourier orthogonality give the alias and Taylor terms; certified label errors give the oracle term. Triangle inequality applied to $b_n^k - b_n^* = (b_n^k - t_n) + (t_n - b_n^*)$ gives the other two. Thus their exact sum bounds the compact-uniform error.

Define $C_d = E_{\text{oracle}} + E_{\text{alias}} + E_{\text{Taylor}} + E_{\text{arithmetic}}$. Select the smallest $d < 64$ with $C_d < \varepsilon$; if none exists, report $\min_{d<64} C_d$ and refuse. For a selected degree, $E_{\text{opt}}(k) = (3/4)^k \sum_n |t_n|_1 u^n \longrightarrow 0$. Therefore executing updates until $E_{\text{opt}}(k) \le \varepsilon - C_d$ terminates after finitely many steps. The serialized coefficients, budgets, targets and trace use exact rational endpoints; displayed decimal approximations are not proof inputs. □

9 Objective semantics, precision floors and replay

The original true-circle sample-label least-squares objective is

$$\mathcal{L}_y(b) = \frac{1}{2m} \sum_{j=0}^{m-1} \left| \sum_{n=0}^{d} b_n \zeta_j^n - y_j \right|^2 .$$

With $d < m$, Fourier orthogonality gives $A^*A = mI$ and the real Euclidean coordinate gradient $\nabla \mathcal{L}_y(b) = b - b^*$. The executed objective J instead has the exact rational midpoint target t. The difference $\nabla \mathcal{L}_y - \nabla J = t - b^*$ is enclosed coefficientwise by e_n. Consequently this is an actual rational GD run for J and a certified perturbed-gradient implementation for $\mathcal{L}_y$. It is not the exact optimizer trajectory of the ideal-root objective, nor a replay of neural or quantum training. An interval enclosure of an ideal-root GD trajectory would be a different explicitly labelled object.

Each actual update contracts $b^k - t$ by $3/4$, so the exact stored losses satisfy $J(b^{k+1}) = (9/16)J(b^k)$. A replay recomputes every rational coordinate update, loss and weighted residual. The closed form $b^k = (1 - (3/4)^k)t$ is an independent check on the executed trace, not its producer. Exact rational loop coordinates introduce no further midpoint or floating export error. The arithmetic budget covers the difference between the certified Fourier rectangles and their midpoint targets, including the inherited root enclosure and outward interval operations. A subsequent float conversion requires its own exact conversion-error bound.

The normalized polynomial convention is essential: a stored coefficient with index n multiplies $((s - 1/2)/(6/5))^n$, not $(s - 1/2)^n$. The input certificate is pinned by SHA-256, including its source hashes, all 192 sample enclosures, root rectangles and analytic radius components. Producing source is captured at entry, checked unchanged before emission, and copied into a fresh output directory. An existing output is never overwritten. An independent replay should also check the fixed input hash, all dimensions, degree choice, the first stopping iterate, all five exact terms and their literal sum against the requested rational tolerance.

The attainable boundary of this majorant is

$$C_{\min,q} = \min_{0 \le d < 64} C_{q,d}.$$

Every request $\varepsilon > C_{\min,q}$ has strict budget slack for some degree and the proved finite stop rule. At or below that boundary this interface refuses. The degree minimizing C_d need not be 63: Taylor error decreases with degree while the oracle and arithmetic sums increase. In particular $C_d \ge \eta$ and $C_d \ge M_q 2^{-64}/(1 - 2^{-64})$. Increasing the number of optimization steps does not lower these fixed-data budgets. Approximate boundary values on the stated compact are $7.2929 \cdot 10^{-17}$, $1.7370 \cdot 10^{-16}$ and $3.3741 \cdot 10^{-16}$ for $q = 3, 4, 5$; only the exact exported rational minima are used for acceptance. These are certificate floors, not lower bounds on the actual errors.

A full adaptive theta oracle would separately need proof-preserving choices of omitted-n cutoff, time cutoff, quadrature order, whole-panel discard budget, primitive series lengths, interval precision and serialized-bound precision. The present $N = 11$, $T = 128$, 224-bit arithmetic, 60/92/60 primitive series lengths, 2^{-90} discard threshold and 180-bit rounded budgets cannot support every positive request. Increasing a nominal bit count alone leaves the fixed primitive remainders and analytic tails in place. The preserved labels are not recomputed here. The classical identity, finite arithmetic and informal proof are not new Lean declarations. No zero-location, all-conductor, positive arithmetic determinant, quantum-advantage or GRH conclusion follows.

The interactive replay is available at the certified readout laboratory. All decimal graphics are illustrative views of the linked exact records.

10 Requested-error theta samples for three fixed conductors

For the primitive characters $\chi_3 = (0, 1, -1)$, $\chi_4 = (0, 1, 0, -1)$, $\chi_5 = (0, 1, -1, -1, 1)$ with parities $a = 1, 1, 0$, retain the classical primitive theta/Mellin identity and root number one already stated above. Put $D = [-1, 2] + i[-4/3, 4/3]$ and $\delta = 3/q$. The following adaptive extension is implemented with independent interval and sample checks; it is an informal analytic theorem, not a Lean result.

Theorem 3 (Finite adaptive sample specification)**.** *For every supported conductor, exact rational $s \in D$, and rational $\eta > 0$, the choices below terminate and give a rational complex midpoint y and exact rational radius $R \le \eta$ with $|y - \Lambda_q(s)| \le R$. The same conclusion holds for a specified true circle node when its root and node enclosures are refined on every precision retry.*

Proof. For $x \ge 0$ define the exact rational bound $E(x) = 2^{-\lfloor x \rfloor}$. Since $e > 2$, $e^{-x} \le E(x)$; this bound is nonincreasing and tends to zero. Let $Q(x) = x^{-1} + 2x^{-2} + 2x^{-3}$ and $C_0 = 1$, $C_1 = (1 + \delta^{-1})/2$. For $N \ge 1$, $n_0 = N + 1$, $J \ge 1$, $T = 2^J$, use

$$B_n = 2C_a Q(\delta n_0^2) \frac{E(\delta n_0^2/2)}{1 - E(\delta n_0)},$$

$$B_t = \sum_{n=1}^{N} 2n^a T^3 E(\delta n^2 T) Q(\delta n^2 T).$$

Here $\delta n_0 \ge 6/5$ ensures a positive denominator; $E(x)$ equals one for $0 \le x < 1$, so this safeguard is essential. On the four panels of each $[2^j, 2^{j+1}]$, write midpoint m, halfwidth h, $\rho = m/2$, $v = h/\rho \le 2/9$, and $B(m) = \max\{(m + \rho)^2, (m - \rho)^{-2}, 1\}$. Use the Cauchy remainder

$$B_Q = \sum_{\text{panels}} \sum_{n=1}^{N} \frac{2n^a \, 2h \, 3B(m) E(\delta n^2 m/2) v^{2K+2}}{(2K + 3)(1 - v^2)}.$$

Choose in turn the first $N \ge 1$, first $J \ge 1$, and first $K \ge 0$ for which their respective budgets are at most $\eta/8$. Gaussian decay, then exponential time decay, then $v \le 2/9$ prove termination. These are the accepted tail and Cauchy arguments with each true exponential separately replaced by its upper bound E.

There are $P = 4JN$ potential pairs. A nonzero-character pair may be omitted only if its entire two-power contribution has bound

$$b = 4hn^a (m + h)^2 E(\delta n^2 (m - h)) \le \eta/(8P).$$

Add every omitted bound to B_{discard}, giving $B_{\text{discard}} \le \eta/8$. Evaluate all other integrated Taylor polynomials by the stated exact coefficient recurrence using outward rectangles. Refine both interval

precision and all primitive series remainders until the final rectangle's sum of coordinate halfwidths satisfies $R_{\mathrm{arith}} \le \eta/4$. For this fixed finite expression all input and primitive intervals shrink; continuity and outward arithmetic imply $R_{\mathrm{arith}} \to 0$. Precision doubling therefore terminates, and the sum of the five raw components is at most $3\eta/4$. Choose an integer output scale B such that $5\,2^{-B} \le \eta/8$ and round each component upward by $\mathrm{up}_B(x) = 2^{-B}\lceil 2^B x\rceil$. The reported radius is their literal exact sum. With $\mathbf{b} = (R_{\mathrm{arith}}, B_n, B_t, B_Q, B_{\mathrm{discard}})$, it obeys

$$|y - \Lambda_q(s)| \le \sum_{i=1}^{5} \mathrm{up}_B(b_i) \le \frac{3\eta}{4} + 5\,2^{-B} \le \frac{7\eta}{8} < \eta.$$

The output retains exact rational midpoints, rounded components, B, raw interval radius, parameters, input metadata and source hashes. The raw analytic bounds are exactly reconstructible; they need not be serialized as potentially enormous fractions. The output precision varies with the request and introduces no fixed reporting floor. □

This fixed-conductor sample result uses no hypothesis on zeros. It does not remove a downstream fixed-node Fourier alias floor, supply an all-conductor theorem or a positive arithmetic determinant, or prove GRH.

11 Adaptive sampling parameters and precision termination

The primitive theta/Mellin identity with root number one yields

$$\Lambda_q(s) = \sum_{n\ge 1} \chi_q(n) n^a \int_1^\infty e^{-\pi n^2 t/q}\big(t^{(s+a)/2-1} + t^{(1-s+a)/2-1}\big)\,dt.$$

On $D = [-1,2] + i[-4/3, 4/3]$, both exponents p satisfy $|\Re p| \le 2$ and $|\Im p| \le 2/3$. For exact $x \ge 0$, put $E(x) = 2^{-\lfloor x\rfloor}$. The elementary inequality $e > 2$ gives $e^{-x} \le e^{-\lfloor x\rfloor} \le E(x)$ and $E(x) \le 2e^{-x\log 2}$. Use this solely as a pointwise nonincreasing upper bound; it is not multiplicative and has no asserted exponential integral identity.

Integration by parts gives $\int_1^\infty t^2 e^{-xt}\,dt = e^{-x}Q(x)$ for $Q(x) = x^{-1} + 2x^{-2} + 2x^{-3}$. For $a = 0$ the bound $n^a e^{-\delta n^2/2} \le C_0 = 1$ is immediate. For $a = 1$, $n \le (1+n^2)/2$ and $n^2 e^{-\delta n^2/2} \le 2/(e\delta) \le 1/\delta$ give $C_1 = (1+1/\delta)/2$. For $n = n_0 + j$, the inequality $n^2 \ge n_0^2 + 2n_0 j$ therefore bounds the omitted-n sum by

$$2C_a Q(\delta n_0^2)\frac{e^{-\delta n_0^2/2}}{1 - e^{-\delta n_0}} \le 2C_a Q(\delta n_0^2)\frac{E(\delta n_0^2/2)}{1 - E(\delta n_0)} = B_n.$$

Take $n_0 = N + 1 \ge 2$; then $\delta n_0 \ge 6/5$ and the latter denominator is positive. Substituting $t = Tu$ in each retained-n time tail gives

$$B_t = \sum_{n=1}^{N} 2n^a T^3 E(\delta n^2 T) Q(\delta n^2 T).$$

The two tails are disjoint because the omitted-n tail uses the whole integral and the time tail uses only retained n.

For $j = 0, \dots, J-1$ and $\ell = 0, \dots, 3$ take

$$m = 2^j\left(1 + \frac{2\ell+1}{8}\right), \qquad h = \frac{2^j}{8}, \qquad T = 2^J.$$

These panels cover $[1,T]$ and have $h/m \le 1/9$. On $|z-m| = \rho = m/2$, principal log is analytic, $\Re z \ge m/2$, and $|\arg z| \le 1$. For $f(z) = e^{-\pi n^2 z/q} z^p$, the angular multiplier is at most $e^{2/3} \le 3$, whence

$$|f(z)| \le 3B(m)E(\delta n^2 m/2), \qquad B(m) = \max\{(3m/2)^2, (m/2)^{-2}, 1\}.$$

Cauchy's coefficient estimate, cancellation of odd terms, and geometric summation yield the B_Q in Theorem 3. Including zero-character or later discarded terms in that sum is harmless. Gaussian, time, and geometric decay prove termination of the sequential first-success searches with $B_n, B_t, B_Q \le \eta/8$. For fixed q and $L = \max\{1, \log(1/\eta)\}$ their asymptotic scales can be chosen as $N = O(\sqrt{L})$, $T = O(L)$, $J = O(\log L)$, and $K = O(L)$. Polynomial prefactors in N, T, J contribute only $O(\log L)$ to the logarithm of the error estimates; this is a truncation-size statement, not a rational-arithmetic bit-complexity bound.

For real $t \in [m-h, m+h]$, $t \ge 1$ and $t^{\Re p} \le (m+h)^2$. The bound for a whole pair, including both powers, is therefore $4hn^a(m+h)^2 E(\delta n^2(m-h))$. There are $P = 4JN$ potential pairs; summing only bounds at most $\eta/(8P)$ gives $B_{\text{discard}} \le \eta/8$. No partial recurrence or isolated exponential is discarded. For active pairs retain the coefficient recurrence

$$r_{-1} = 0, \quad r_0 = 1, \quad r_{j+1} = \frac{(h/m)(p - \lambda m - j)r_j - \lambda m (h/m)^2 r_{j-1}}{j+1}, \qquad \lambda = \pi n^2/q.$$

Coefficient comparison in $tf' = (-\lambda t + p)f$ proves that the integrated Taylor polynomial is

$$2he^{p\log m - \lambda m} \sum_{j=0}^{K} \frac{r_{2j}}{2j+1}.$$

Outward arithmetic encloses its exact value, including the signed character weight and both powers.

Use an isolated p-bit dyadic context, with all primitive remainders refined as p increases. Put $\delta_p = 2^{-p-16}$. Each arctangent series separately stops when its next-term bound is at most δ_p, giving total weighted analytic tail width at most $20\delta_p$ in Machin's formula, before outward dyadic rounding. Each reduced logarithm series separately has tail at most δ_p; for $x = 2^k v$, the combined bound is $(1+|k|)\delta_p$, before rounding. For the complex exponential, after reduction to $A \le 1/2$, choose the Taylor order so the bound

$$\frac{A^{H+1}}{(H+1)!} \frac{1}{1 - A/(H+2)}.$$

is at most δ_p. Choose H by finite exact search; its denominator is positive. Add a symmetric exponential error enclosure to each coordinate before outward squaring. The fixed guard16 is not itself an acceptance bound: the final complete rectangle radius must satisfy $R_{\text{arith}} \le \eta/4$. All interval caches must include p and preserve exact input-type guards. The successful p and retry precisions/radii are recorded. Individual series lengths are determined by these, the exact inputs and the pinned source and are reconstructible by replay; they are not explicitly stored.

For fixed N,J,K and exact target point, the expression has finitely many operations and only nonzero fixed integer denominators. Each primitive input is bounded for sufficiently large p; exponential halving counts are consequently bounded. Exact primitive tails and dyadic rounding widths tend to zero. Induction through the finite continuous expression, including finitely many outward squarings, gives $R_{\text{arith}} \to 0$. Doubling p after a failed comparison terminates. The argument also covers true roots and nodes regenerated at every p, since their widths tend to zero. A fixed nondegenerate s rectangle has no corresponding shrinking-input claim.

A growing-order primitive estimate. To justify convergence despite a growing exponential Taylor order, write $\epsilon_p = 2^{-p}$, $\delta_p = 2^{-p-16}$, and let $d(X)$ be the sum of a rectangle's coordinate widths and $m(X)$ its ℓ^1 magnitude bound. The rounded reduced input W satisfies $a = m(W) \le 1/2 + 2\epsilon_p < 3/5$ and $d(W) \le 2^{-k} d(Z) + 4\epsilon_p$; the remainder is at most 2^{-H}, so the first accepted order satisfies $H \le p + 18$. For the computed term boxes, outward multiplication and integer division give $m_j \le a m_{j-1} + 6\epsilon_p$ and $d_j \le a d_{j-1} + m_{j-1} d(W) + 12\epsilon_p$. Thus $m_j \le a^j + 15\epsilon_p$ and $\sum_{j=0}^{H} d_j \le (25/4) d(W) + 75 H \epsilon_p$; adding the remainder increases width by at most $4(\delta_p + \epsilon_p)$. The pre-squaring width is therefore $O(d(Z) + p2^{-p})$ and its magnitude is bounded (indeed below 3 for $p \ge 32$). Each outward square has width at most $2m(X)d(X) + 8\epsilon_p$; the bounded number of squarings consequently preserves convergence to zero. Finally, exact rational partial sums and endpoint rounding give Machin width at most $20\delta_p + 40\epsilon_p$, and logarithm width at most $(1 + |k|)(\delta_p + 2\epsilon_p)$ for fixed $x = 2^k v$. These estimates justify the primitive-width step without assuming a fixed Taylor order or asserting a global arithmetic complexity bound.

For serialization, initialize

$$B = \max\{32, \operatorname{bitlen}(\operatorname{den}\eta) - \operatorname{bitlen}(\operatorname{num}\eta) + 8\}$$

and increase B until $5\,2^{-B} \le \eta/8$. Write $\widehat{x} = 2^{-B}\lceil 2^B x \rceil$ for each of the five raw components. Their exact reported sum is at most $3\eta/4 + 5\,2^{-B} \le 7\eta/8$, and every component remains an upper bound. The output records B, the raw interval radius and complete rectangle, N/J/K, discard threshold/pair counts, exact target metadata and source hashes. These reconstruct all raw analytic bounds and pair sets without serializing huge exact tail denominators. For the practical requests below B is87 and114, respectively. Exact endpoints remain rational; the mathematical specification assumes an integer representation sufficient for the requested size. Dynamic rounding removes the old fixed reporting floor. For rational s, validate its exact coordinates in closed D before conversion: an outward dyadic enclosure of $\Im s = 4/3$ extends beyond D, but the analytic bounds concern the true point. Circle nodes satisfy the strict exact bounds $-7/10 \le \Re s \le 17/10$, $|\Im s| \le 6/5$.

The exact parameter-design checker gives

η	q	N	J	T	K	active pairs
10^{-24}	3	12	7	128	19	71
10^{-24}	4	14	8	256	20	65
10^{-24}	5	16	8	256	20	111
10^{-32}	3	14	7	128	25	84
10^{-32}	4	16	8	256	26	72
10^{-32}	5	18	8	256	26	129

with exact rational budgets and no special-function oracle. The arithmetic precision is selected from the actual interval radius, not inferred from this table. This extension retains the existing classical theta identity; it adds no zero-location, positive determinant, all-conductor, Lean or GRH claim. The following end-to-end adaptive readout supplies its own node-count, alias, degree, Fourier and optimization choices.

12 An end-to-end adaptive analytic readout

The new driver changes the sample truncations, interval precision, circle size and degree; it does not reuse the frozen labels to assert a smaller error. Its scope is still the three stated primitive quadratic characters and the fixed compact $|s - 1/2| \le 7/10$.

Theorem 4 (Requested-error construction in the stated finite family)**.** *Given a positive rational ε and $q \in \{3, 4, 5\}$, the specified exact-arithmetic algorithm terminates in the mathematical unbounded-resource model and returns a normalized rational polynomial P with*

$$\sup_{|s-1/2|\le 7/10} |P(s) - \Lambda_q(s)| \le \varepsilon.$$

Its sample enclosures, rational optimizer updates and all error budgets are explicit. This is an approximation theorem, with the stated classical theta identity as an input; it supplies no zero-location or positive determinant conclusion and is not a Lean theorem or a practical runtime bound.

Proof. Put $u = 7/12$, $v = 7/24$. Choose $d \ge 0$ so $10v^{d+1}/(1-v) \le \varepsilon/8$, and choose $m > d$ so

$$10\frac{2^{-m}}{1-2^{-m}}\frac{1}{1-v} \le \varepsilon/8.$$

The driver rounds m upward to a multiple of eight, which only decreases this bound. Set the requested sample error $\eta = \varepsilon(1-u)/8$. The adaptive sampler gives all m true circle nodes and labels with error at most η. Form their outward Fourier rectangles through degree d, and refine the nodes and finite evaluation arithmetic if their weighted midpoint radii exceed $\varepsilon/8$. This refinement terminates: for these fixed finite d, m, node uncertainty and every finite-expression arithmetic width tend to zero with precision. Merely lifting fixed boxes would not suffice, so the driver regenerates the nodes and samples on this retry.

Choose rational midpoint targets t_n and execute the previously specified quadratic-objective updates until $\sum_{n=0}^{d} |b_n^k - t_n|_1 u^n \le \varepsilon/4$. Their geometric contraction proves finite termination. The oracle, aliasing, Taylor and Fourier-midpoint terms are each at most $\varepsilon/8$, and the optimization term is at most $\varepsilon/4$. Hence the raw five-term sum is at most $3\varepsilon/4$. Choose a serialization precision B with $5\,2^{-B} \le \varepsilon/8$, round each term upward, then add the five reported rationals exactly. The printed bound is at most $7\varepsilon/8 \le \varepsilon$. The polynomial uses $((s-1/2)/(6/5))^n$, and the preceding Fourier/Cauchy proof gives its asserted compact-uniform bound. □

The trace is losslessly encoded by the fixed target vector and a rational scale extracted from each *executed* coefficient state, together with its objective, weighted residual and full-state hash. Every coefficient can be recovered exactly; a separate replay expands the states and checks every update. The producer does not generate its trace from a closed form. The saved objective remains the rational Fourier-midpoint quadratic, with its perturbation from ideal-root least squares explicitly budgeted.

13 Termination, compression and the boundary of the guarantee

The degree and circle-size choices use only rational comparisons. Since $0 < v < 1$, the Taylor test succeeds at finite d; since $2^{-m} \to 0$, the alias test succeeds at finite $m > d$. The common growth bound10 is safe for all three families; the certificate also recomputes the previously proved exact conductor-specific majorants below $4, 10, 6$.

At fixed d, m, η, the analytic truncation parameters are fixed while precision is doubled. The computed sample centers stay bounded: their underlying finite expression is fixed and the true completed-function values are bounded on the sample circle. Its arithmetic widths tend to zero, as do the regenerated true-root enclosures. Fourier computation is a finite composition of additions, products and division by the positive integer m. Its midpoint-radius sum therefore tends to zero.

This proves the driver's additional precision retry, without pretending that changing a box's nominal precision shrinks its uncertainty.

The accepted sampler error may remain above its actual arithmetic width because its analytic tails and output bounds are fixed for this η. That causes no problem: the oracle allocation uses the total certified sample radius, whereas the Fourier-midpoint allocation uses only finite coefficient-rectangle width. No analytic sample error is silently moved into the latter term or removed by numerical convergence.

For the stored training trace, the exact loop gives $b^{k+1} = 3b^k/4 + t/4$ from $b^0 = 0$. Inductively every state is a rational scalar multiple of the same target vector. The producer first executes the actual vector updates, then extracts this scalar from a nonzero target coordinate, checks all coordinates, and records a canonical full-vector SHA-256 hash. If the entire target vector is zero, its zero state needs only scale0. The compressed record is thus an exact representation of the executed states, not a floating-point sketch. Hashes are integrity aids; exact vector reconstruction and the independently checked recurrence establish the mathematical claim.

Coefficient export uses exact rational numerators and denominators. A chunked decimal integer encoder/parser handles endpoints beyond Python's usual digit-conversion guard without changing global interpreter settings. No fixed 180-bit serialization floor remains. The decimal graphics are separate approximate projections, not polynomial coefficients to which the exact certificate may be applied without a conversion-error bound.

All producing source bytes are captured, copied into fresh output paths and checked unchanged. Historical fixed producers are pinned and untouched. For the new modules the run also checks the import-time source snapshots, so a stale imported module cannot be attributed to later source bytes. After output creation, a failed evaluation, final provenance check or serialization emits uncertified failure metadata. Invalid requests and stale import-time sources are refused before creating output artifacts. Old outputs are never overwritten. Extreme requests can be computationally impractical: the termination proof gives no machine-time or memory guarantee, and finite test runs establish only their reported cases.

Nothing requires the output polynomials to be real-rooted, to equal an arithmetic Fredholm determinant, or to have a self-adjoint realization. Consequently arbitrary compact approximation of these known functions does not imply GRH. Uniformity in changing conductor, general Hecke/Dedekind/elliptic/Artin scopes, the independent determinant comparison, and actual analytic/category Lean objects remain separate obligations.

14 An arithmetic-family moment with every original row retained

Let $K = \mathbb{Q}(\sqrt{-3})$, $\mathcal{O} = \mathbb{Z}[\omega]$, and write N for the ideal or element norm. Fix a finite-order Hecke character ν and a finite set S of prime ideals containing those above 6 and those dividing its conductor. For $0 \neq u \in \mathcal{O}$, let κ_u be the character of $K(u^{1/6})/K$ defined using arithmetic Frobenius, and put $\psi_u = \mathrm{prim}(\nu\kappa_u)$. Here prim means the underlying primitive Hecke character: it is formed before introducing any divisor variable or evaluating conductor-dependent zero extensions. The rows below are all actual Eisenstein elements, including units and prime powers supported in S.

The external inputs are abelian Kummer reciprocity and character uniqueness, the primitive Hecke functional equation and its usual vertical growth, Stirling's formula, scalar Dedekind-zeta convexity, and Blomer–Goldmakher–Louvel's large sieve for their specified sixth-order family. We use the pinned arXiv:1112.1650v1, Theorem 1.3, Lemma 2.1 and Appendix II; its Section 5 supplies the all-ideal decomposition whose fixed-twist proof is redone here. No statement about all primitive characters of comparable conductor, Artin holomorphy, Lindelöf bounds or zero locations is assumed.

Theorem 5 (The original-family numerator moment)**.** *For every $\varepsilon > 0$, uniformly for $H \geq 1$ and $t \in \mathbb{R}$,*

$$\sum_{0<Nu\leq H} |L(1/2+it,\psi_u)|^2 \ll_{\nu,S,\varepsilon} H^{1+\varepsilon}(1+|t|)^{1+\varepsilon}. \tag{1}$$

For every fixed $0 < \delta < 1/2$, the same estimate holds with $1/2+it$ replaced by $1/2+\delta+it$, with the constant also depending on δ. If the central bound has constant C_0 at an exponent $0 < \varepsilon_0 \leq 1/4$, an admissible off-central constant is

$$200C_0\left(\frac{3/2+\delta}{1/2-\delta}\right)^2. \tag{2}$$

Proof. For the canonical BGL character $\vartheta_{\mathfrak{a}}$, Appendix II gives a radicand $\alpha_{\mathfrak{a}}$ with $(\alpha_{\mathfrak{a}})_S = \mathfrak{a}\mathfrak{g}^{-6}$. Since $\mathcal{O}_S$ is principal, comparison with any literal u whose outside-S ideal is $\mathfrak{a}$ gives $\kappa_u = \eta\vartheta_{\mathfrak{a}}$ for one of finitely many fixed S-unit Kummer twists η. This is an equality of primitive characters, established outside a finite set and extended by uniqueness. For each outside-S ideal $\mathfrak{a}$ and S-supported integral ideal $\mathfrak{b}$, at most six elements generate $\mathfrak{a}\mathfrak{b}$. The resulting S-supported multiplicity sum converges at exponent $1+\varepsilon$.

For each fixed $\xi = \nu\eta$, split the canonical parameter as $\mathfrak{a} = \mathfrak{a}_0\mathfrak{a}_1\mathfrak{a}_2^6$, where $\mathfrak{a}_0$ is squarefree, $\mathfrak{a}_1$ is squarefull and sixth-power-free, and $(\mathfrak{a}_0, \mathfrak{a}_1) = 1$. The primitive conductor is at most a fixed multiple of $N(\mathfrak{a}_0\mathfrak{a}_1)$. The common-gamma approximate functional equation, with $X = H/(N\mathfrak{a}_2)^6$ and $T = 1+|t|$, has common maximal length $B = X^{1/2}T(HT)^{\eta}$ for any sufficiently small fixed $\eta > 0$. Mellin inversion removes its row-dependent conductor weight. Split a column as $\mathfrak{b}_S\mathfrak{b}_1\mathfrak{b}_2$, with $\mathfrak{b}_S$ supported in S, $\mathfrak{b}_1$ squarefree and $\mathfrak{b}_2$ squarefull, the last two coprime and prime to S. The original primitive values at $\mathfrak{b}_S$ remain bounded row multipliers. The S-supported and squarefull column weights are summable at $1/2+\eta$. Multiplying squarefree column coefficients by the fixed ξ preserves their squared norm, whatever its finite order. BGL's large sieve therefore bounds the energy for fixed $\mathfrak{a}_1, \mathfrak{a}_2$ by

$$(HT)^{O(\eta)}\{M + B + (MB)^{2/3}\}, \qquad M = X/N\mathfrak{a}_1.$$

Squarefull counting and summation give $X + B\sqrt{X} + (XB)^{2/3} \ll (HT)^{O(\eta)}XT$. The $\mathfrak{a}_2$ sum converges at exponent 6. Principal rows are separated before the pole-free functional equation and bounded by scalar zeta convexity. This proves (1).

For the off-central statement divide each $L(s,\psi)$ by $(s+1)^{5/4}$ and, on principal rows, also multiply by $(s-1)/(s+1)$. These are bounded analytic functions on $\operatorname{Re} s > 1/2$, continuous on the boundary. Poisson representation and Jensen's inequality transfer the central squared moment. The convolution of $P_\delta(y) = \delta/[\pi(\delta^2+y^2)]$ with $(1+|t|)^{-3/2+\varepsilon_0}$ is at most eleven times that weight when $0 < \varepsilon_0 \leq 1/4$. Restoring the denominator and the principal pole-removal factor gives (2). The extended edition gives the detailed derivation with all primitive local values. □

Separate fixed-unitary extension. The moment theorem also holds for any one fixed unitary Hecke character ν. For precision choose $\nu(\mathfrak{a}) = \nu_0(\mathfrak{a})(N\mathfrak{a})^{-i\tau}$, with angular parameter $m \in \mathbb{Z}$ and completed gamma factor

$$\gamma_\nu(s) = 2(2\pi)^{-(s+|m|/2+i\tau)}\Gamma(s+|m|/2+i\tau).$$

All finite Kummer twists have this same infinity type. Use $T_\nu = 1+|t+\tau|+|m|$, comparable to $1+|t|$ with ν-dependent constants. Pure norm rows have $L(s,\psi) = \zeta_K(s+i\tau)$ and are separated before the pole-free functional equation; their pole at $1-i\tau$ is removed by $(s-1+i\tau)/(s+1+i\tau)$. No uniformity in m, τ or the conductor of ν is asserted. This extension of the moment is separate from the following finite-order type-I residue formula.

14.1 The exact masked remainder and the remaining loss

Retain the original zero-extended symbol $\chi_d(u) = (u/d)_6$ on good squarefree ideals. Let μ be the ideal Möbius function and set

$$\beta_U(d) = \sum_{\substack{ab=d \\ Na, Nb \le U}} \mu(a)\mu(b), \quad T(u,d) = S \cup \operatorname{supp}(u) \cup \operatorname{supp}(d).$$

Let $\psi_{u,2} = \operatorname{prim}(\psi_u^2)$, formed independently of squaring any zero sentinel. With $\mathfrak{f}_u$ and $\mathfrak{f}_{u,2}$ the primitive conductors, the full-mask series is

$$F_{u,d}(s) = \frac{L(s,\psi_u)}{L(2s,\psi_{u,2})} C_{T(u,d)}(s), \tag{3}$$
$$C_T(s) = \frac{\prod_{p \in T,\ p \nmid \mathfrak{f}_u} (1 - \psi_u(p)(Np)^{-s})}{\prod_{p \in T,\ p \nmid \mathfrak{f}_{u,2}} (1 - \psi_{u,2}(p)(Np)^{-2s})}.$$

Initially it equals $\sum_{c \text{ squarefree},\ (c,T)=1} \psi_u(c)(Nc)^{-s}$ for $\operatorname{Re} s > 1$. Every prime of T remains deleted, including those forgotten by either primitive conductor.

For finite-order ν, $D \ge 1$, $U \ge 1$, $U^2 \le D$, and $W \in C_c^\infty((0,\infty))$ supported in $[1,2]$, put

$$\widehat{W}(s) = \int_0^\infty W(x) x^{s-1}\, dx, \qquad \mathcal{M}_4(W) = \sum_{j=0}^{4} \sup_x |(x\partial_x)^j W(x)|.$$

On $\sigma = 1/2 + \delta$, the exact residue-subtracted type-I remainder is

$$E_u = \frac{D^\delta}{2\pi} \int_{\mathbb{R}} \widehat{W}(\sigma + it) D^{it} \sum_{\substack{d \text{ squarefree} \\ (d,S)=1}} \frac{\beta_U(d)\nu(d)\chi_d(u)}{(Nd)^{\sigma+it}} F_{u,d}(\sigma + it)\, dt. \tag{4}$$

This definition retains the actual β_U, all phases, d-masks, conductors and the independently primitive square.

Theorem 6 (An improved masked residual estimate)**.** *Under the preceding finite-order hypotheses, for every $\varepsilon > 0$,*

$$\sum_{0 < Nu \le H} |E_u|^2 \ll_{\nu,S,\delta,\varepsilon} \mathcal{M}_4(W)^2 D^{2\delta} U^{2-4\delta} H (HD)^\varepsilon. \tag{5}$$

Proof. The same deleted set in both factors of (3) gives $|F_{u,d}(\sigma + it)| \ll (Nu\, Nd)^\eta |L(\sigma + it, \psi_u)|$: the reciprocal incomplete Euler product at $2\sigma > 1$ is bounded by $\zeta_K(2\sigma)$, and its remaining finite numerator is bounded by a small power of its radical. Apply Minkowski to (4) in the exact row space and use Theorem 5. The actual coefficient sum is at most

$$\sum_d \frac{|\beta_U(d)|}{(Nd)^{\sigma-\eta}} \le \left(\sum_{Na \le U} (Na)^{-\sigma+\eta} \right)^2 \ll U^{1-2\delta+2\eta}.$$

Four Mellin integrations by parts bound the remaining height integral by a constant times $\mathcal{M}_4(W)$. Square the result and absorb the small losses using $U^2 \le D$. □

At $H = D^{3/4}$ and $U = D^a$ with fixed $0 < a \le 1/2$, the upper bound's excess over H is $2\delta + (2 - 4\delta)a$, up to small losses; as $\delta \downarrow 0$ it still leaves $2a$. This is a limitation of the upper bound, not an energy lower bound. The extracted principal residue, the original short term at smaller

scales and the signed type-II minus type-I covariance are not controlled by (5). The next arithmetic task is to estimate the complete d-weighted moment before its absolute coefficient sum, or establish sufficient signed covariance. No positive arithmetic determinant, zero-location theorem or GRH proof follows. The analytic argument uses the external classical inputs above; finite diagnostics and selected existing Lean local identities do not formalize the Hecke functions, approximate functional equation, Poisson transfer or these moment bounds.

15 Family identification, the approximate functional equation and masks

We prove the analytic statements of Section 14 from the explicit classical inputs stated there. Write $I(S)$ for the nonzero integral ideals prime to S, and $\mathcal{O}_S$ for the ring of S-integers. Constants may depend on the fixed field, S, ν and any declared positive small exponent. No such constant is asserted uniform in the conductor or infinity type of a varying ν.

15.1 A finite comparison with the literal Kummer rows

Use Appendix II of BGL, with $n = 6$. Its fixed local modulus and representatives may be chosen once, since S contains the primes above 6 and $\mathcal{O}_S$ has class number one. For each $\mathfrak{a} \in I(S)$ the construction writes

$$\mathfrak{a} = (x)E\mathfrak{g}^6, \qquad \alpha_{\mathfrak{a}} = xm_E, \qquad (m_E)_S = E,$$

so $(\alpha_{\mathfrak{a}})_S = \mathfrak{a}\mathfrak{g}^{-6}$. Choose $y \in K^\times$ generating $\mathfrak{g}$ as an $\mathcal{O}_S$-ideal. If $(u)_S = \mathfrak{a}$, then

$$e = \alpha_{\mathfrak{a}} y^6/u \in \mathcal{O}_S^\times, \qquad [\alpha_{\mathfrak{a}}] = [eu] \quad \text{in } K^\times/K^{\times 6}.$$

There is a precision here: Appendix II allows choosing $\mathfrak{g}$ relative to a tested denominator. Fix one admissible global decomposition first. Its radicand gives the canonical character $\vartheta_{\mathfrak{a}}$ on all primes outside the finite set $S \cup \operatorname{supp}(\mathfrak{a}\mathfrak{g})$ where that choice applies. Uniqueness of Hecke characters, equivalently Chebotarev in the finite abelian quotient, then identifies the underlying primitive characters everywhere. With matching arithmetic-Frobenius conventions,

$$\kappa_u = \kappa_{e^{-1}}\vartheta_{\mathfrak{a}}. \tag{6}$$

If the residue-symbol convention is conjugated, conjugate the canonical family and its coefficient sequence; the same large sieve follows.

The group $\mathcal{O}_S^\times/(\mathcal{O}_S^\times)^6$ is finite. Thus (6) uses one of a finite set of fixed Kummer twists η, all unramified outside S and of bounded S-conductor. This is a comparison of primitive characters, not a pointwise multiplication rule for independently zero-extended arrays at bad primes. In particular, a conductor drop must not be replaced by a product of artificial zeros.

Since $\mathcal{O}$ is principal with six units, for fixed $\mathfrak{a} \in I(S)$ and an S-supported integral ideal $\mathfrak{b}$ there are at most six u satisfying $(u) = \mathfrak{a}\mathfrak{b}$. For any nonnegative row quantity, the literal moment is consequently at most

$$6 \sum_{\substack{\mathfrak{b}\ S\text{-supported} \\ N\mathfrak{b} \le H}} \sum_{\eta} \sum_{\substack{\mathfrak{a} \in I(S) \\ N\mathfrak{a} \le H/N\mathfrak{b}}} \left|L\left(s, \operatorname{prim}(\nu\eta\vartheta_{\mathfrak{a}})\right)\right|^2. \tag{7}$$

An ideal bound $X^{1+\varepsilon}T^{1+\varepsilon}$ therefore transfers to all literal rows, since

$$\sum_{\mathfrak{b}\ S\text{-supported}} (N\mathfrak{b})^{-1-\varepsilon} = \prod_{p \in S} (1 - (Np)^{-1-\varepsilon})^{-1} < \infty.$$

No quotient of the original unit rows or deletion of S-prime powers has occurred.

15.2 Fixed twists, exact tame conductors and the squarefree sieve

Fix $\xi = \nu\eta$ and write $\rho_{\mathfrak{a}} = \mathrm{prim}(\xi\vartheta_{\mathfrak{a}})$. Its fixed conductor is supported in S before the Kummer twist. On columns prime to S, $|\xi(\mathfrak{b})| = 1$. BGL Theorem 1.3 thus gives, for any fixed $\mathfrak{a}_1$,

$$\sum_{\substack{\mathfrak{a}_0 \in I(S) \text{ squarefree} \\ N\mathfrak{a}_0 \le M}} \left| \sum_{\substack{\mathfrak{b}_1 \in I(S) \text{ squarefree} \\ N\mathfrak{b}_1 \le B}} \lambda_{\mathfrak{b}_1} \xi(\mathfrak{b}_1) \vartheta_{\mathfrak{a}_1}(\mathfrak{b}_1) \vartheta_{\mathfrak{a}_0}(\mathfrak{b}_1) \right|^2 \ll (MB)^{\varepsilon}\{M + B + (MB)^{2/3}\} \sum_{\mathfrak{b}_1} |\lambda_{\mathfrak{b}_1}|^2. \tag{8}$$

The first fixed multiplier preserves the coefficient norm; the second has modulus at most one. Restricting the positive outer row sum is harmless. A coprimality condition with another fixed ideal belongs to λ. In particular, ν need not have order dividing six and is not asserted to belong to BGL's finite ray-class character set.

There is a unique valuation decomposition

$$\mathfrak{a} = \mathfrak{a}_0\mathfrak{a}_1\mathfrak{a}_2^6, \quad \mathfrak{a}_0 \text{ squarefree}, \quad \mathfrak{a}_1 \text{ squarefull and sixth-power-free}, \quad (\mathfrak{a}_0, \mathfrak{a}_1) = 1.$$

Indeed, for $v_p(\mathfrak{a}) = 6q + r$, $0 \le r \le 5$, put the residual prime in $\mathfrak{a}_0$ if $r = 1$, and its exponent r in $\mathfrak{a}_1$ if $2 \le r \le 5$. BGL Lemma 2.1 gives $\vartheta_{\mathfrak{a}\mathfrak{b}^6} = \vartheta_{\mathfrak{a}}$ and $\vartheta_{\mathfrak{a}_0\mathfrak{a}_1} = \vartheta_{\mathfrak{a}_0}\vartheta_{\mathfrak{a}_1}$ on ideals prime to S in this sixth-power-free case. Those are the identities needed here; no arbitrary multiplicativity assertion is substituted.

At a prime outside S, tame Kummer inertia at residual exponent r has order $6/\gcd(6, r)$. It is nontrivial for $r = 1, \dots, 5$ and has conductor exponent one. The fixed ξ is unramified there, so cannot cancel that inertia. At primes in S the local sixth-power classes are finite, hence only finitely many primitive conductor types occur after multiplication by the fixed ξ. Consequently

$$Q_{\rho_{\mathfrak{a}}} = |D_K| N\mathfrak{f}_{\rho_{\mathfrak{a}}} \le C_{\xi,S} N(\mathfrak{a}_0\mathfrak{a}_1), \tag{9}$$

independently of $\mathfrak{a}_2$. The equality of the conductor's outside-S part with $\mathrm{rad}(\mathfrak{a}_0\mathfrak{a}_1)$ does not identify that conductor with the original larger deleted set.

15.3 Removing the row-dependent functional-equation weights

First take finite-order ν. All characters then have trivial infinity type at the complex place. The primitive completion may be normalized as

$$\Lambda(s, \rho) = Q_\rho^{s/2} \gamma_K(s) L(s, \rho), \qquad \gamma_K(s) = 2(2\pi)^{-s}\Gamma(s), \qquad \Lambda(s, \rho) = \epsilon_\rho \Lambda(1 - s, \bar{\rho}), \quad |\epsilon_\rho| = 1.$$

For nonprincipal rows the completed function is entire. Apply this functional equation to the Mellin integral with even auxiliary factor $G(z) = e^{z^2}$ and denominator z, shifting across $z = 0$. At $s = 1/2 + it$ this gives two weighted Dirichlet series, truncated below to polynomials; the dual prefactor has modulus one. Principal rows are excluded from this pole-free step and treated separately below.

For fixed $\mathfrak{a}_2$, let $X = H/(N\mathfrak{a}_2)^6 \ge 1$ and $T = 1 + |t|$. An explicit normalized kernel is

$$V_t(x) = \frac{1}{2\pi i} \int_{(\eta)} e^{z^2} \frac{\Gamma(s + z)}{\Gamma(s)} T^{-z} x^{-z} \frac{dz}{z}, \qquad x = \frac{2\pi N\mathfrak{b}}{\sqrt{Q_\rho} T}, \tag{10}$$

with $\eta > 0$ small. Its factors combine to $(\sqrt{Q_\rho}/(2\pi N\mathfrak{b}))^z$ in the actual completed-L kernel; there is no extra factor of T. The denominator z supplies the residue at zero. On $z = \eta + iy$, Stirling and $|t + y| \le T(1 + |y|)$ give

$$\left|\frac{\Gamma(1/2 + \eta + i(t + y))}{\Gamma(1/2 + it)T^{\eta+iy}}\right| \ll_\eta (1 + |y|)^{C_\eta} e^{\pi|y|/2}.$$

The factor e^{-y^2} makes this an integrable majorant uniformly in t. The analogous bound holds on every fixed positive shifted line and for the dual gamma kernel. Moving far enough to the right gives arbitrary power decay beyond a common total-ideal cutoff

$$B' = X^{1/2}T(HT)^\eta, \tag{11}$$

up to fixed ξ, S, K constants. For any requested $A > 0$, a sufficiently large positive kernel shift, depending also on η, makes the truncation error $O((HT)^{-A})$ uniformly in all rows in this block. Ideal counting $O_K(Y)$ bounds the absolute tails. Such errors are negligible after the finite row sums by taking A larger if needed.

Insert (10) on $\operatorname{Re} z = \eta$. The factor $(\sqrt{Q_\rho}T)^z$ is independent of the column and is a row multiplier of modulus at most $(CXT^2)^{\eta/2}$. Its imaginary part is a phase. The remaining column coefficient is $(N\mathfrak{b})^{-1/2-\eta-it-iy}$, apart from fixed constants. Minkowski or Cauchy with the integrable gamma majorant therefore removes the row-dependent weight at cost $(HT)^{O(\eta)}$ in the squared norm. No independence of the root numbers is assumed or needed.

For each column use the unique decomposition

$$\mathfrak{b} = \mathfrak{b}_S\mathfrak{b}_1\mathfrak{b}_2, \quad \mathfrak{b}_S \text{ } S\text{-supported}, \quad \mathfrak{b}_1 \text{ squarefree}, \quad \mathfrak{b}_2 \text{ squarefull}, \quad (\mathfrak{b}_1, \mathfrak{b}_2) = 1,$$

where $\mathfrak{b}_1, \mathfrak{b}_2 \in I(S)$; a prime of outside-$S$ valuation one goes in $\mathfrak{b}_1$, and a prime of valuation at least two goes with its full exponent in $\mathfrak{b}_2$. For fixed $\mathfrak{b}_S, \mathfrak{b}_2$ retain $\rho_{\mathfrak{a}}(\mathfrak{b}_S)$ as the original primitive row value. On the good factors the exact canonical identity applies. The remaining factors involving $\mathfrak{b}_2$, including $\vartheta_{\mathfrak{a}_0}(\mathfrak{b}_2)$, are bounded row multipliers. Their moduli are at most one; only then are they discarded in an inequality. In particular the value at $\mathfrak{b}_S$ is not replaced by the product of zero-extended primitive arrays.

For the $\mathfrak{b}_1$ polynomial the coefficient sequence is independent of $\mathfrak{a}_0$; it includes the indicator $(\mathfrak{b}_1, \mathfrak{b}_2) = 1$ and the common cutoff $N\mathfrak{b}_1 \le B'/N(\mathfrak{b}_S\mathfrak{b}_2)$. Padding it by zeros up to B' is legitimate. Its squared norm is at most $\zeta_K(1 + 2\eta)$. Minkowski over $\mathfrak{b}_S, \mathfrak{b}_2$ is legitimate because

$$\sum_{\mathfrak{b}_S, \mathfrak{b}_2} N(\mathfrak{b}_S\mathfrak{b}_2)^{-1/2-\eta} < \infty.$$

The finite Euler product handles $\mathfrak{b}_S$; the squarefull series converges above $1/2$. Apply (8) to each such polynomial and to the conjugate dual family. For fixed $\mathfrak{a}_1, \mathfrak{a}_2$ the resulting energy is

$$(HT)^{O(\eta)}\{M + B' + (MB')^{2/3}\}, \qquad M = X/N\mathfrak{a}_1. \tag{12}$$

The additional condition $(\mathfrak{a}_0, \mathfrak{a}_1) = 1$ simply restricts the positive row sum. The large-sieve small exponent can be chosen sufficiently small along with η and absorbed in the indicated small loss.

15.4 All-ideal summation and principal rows

Every squarefull ideal $\mathfrak{c}$ has a unique expression $\mathfrak{d}^2\mathfrak{e}^3$ with $\mathfrak{e}$ squarefree. This follows primewise by using exponent zero or one in $\mathfrak{e}$ according to the parity of an exponent at least two. Therefore

$$\#\{\mathfrak{c} \text{ squarefull} : N\mathfrak{c} \le X\} \ll_K X^{1/2}, \qquad \sum_{\mathfrak{c} \text{ squarefull}} (N\mathfrak{c})^{-s} < \infty \quad (s > 1/2).$$

For the count, sum $O((X/(N\mathfrak{e})^3)^{1/2})$ and use $\zeta_K(3/2) < \infty$. Summing (12) over $\mathfrak{a}_1$ gives

$$(HT)^{O(\eta)}\{X + B'\sqrt{X} + (XB')^{2/3}\} \ll (HT)^{O(\eta)}XT,$$

since the squarefull series converges at 1 and 2/3, $B' = \sqrt{X}\,T(HT)^{\eta}$, and $T^{2/3} \le T$ for $T \ge 1$. The $\mathfrak{a}_2$ sum is bounded by $H\sum_{\mathfrak{a}_2}(N\mathfrak{a}_2)^{-6} = H\zeta_K(6)$. Choosing all internal losses small for the requested final ε proves the fixed-twist all-ideal central moment.

There are $O_K(H)$ ideal rows in total. On any row where the primitive character is principal its L-function is ζ_K, and ordinary scalar convexity gives $|\zeta_K(1/2+it)|^2 \ll_\varepsilon T^{1+\varepsilon}$. This treats those rows without a pole-free approximate functional equation. Finally (7) proves (1) for every actual nonzero Eisenstein element.

15.5 Poisson transfer with an explicit constant

For a nonprincipal row define

$$f_\psi(s) = \frac{L(s,\psi)}{(s+1)^{5/4}},$$

using the holomorphic logarithm on $\mathrm{Re}(s+1) > 0$. On a principal row additionally multiply by $(s-1)/(s+1)$. Analytic continuation and ordinary degree-two convexity, used separately at each fixed conductor, show that f_ψ is bounded analytic on the right half-plane $\mathrm{Re}\, s > 1/2$, with continuous boundary and decay at infinity. The principal pole is removed before applying Poisson representation to its real and imaginary parts. Jensen gives

$$|f_\psi(1/2+\delta+it)|^2 \le \int_{\mathbb{R}} P_\delta(t-v)|f_\psi(1/2+iv)|^2\,dv.$$

The exact row family is finite for fixed H, so its sum interchanges with the integral without a limiting-family issue. The boundary pole-removal factor has modulus at most one.

Take $0 < \varepsilon_0 \le 1/4$ and $\alpha = 3/2 - \varepsilon_0 \in [5/4, 3/2)$. Since $|3/2 + iv| \ge (1+|v|)/\sqrt{2}$, the central squared moment of the f_ψ is bounded by $3C_0H^{1+\varepsilon_0}w_\alpha(v)$, where $w_\alpha(v) = (1+|v|)^{-\alpha}$. For $|t| \ge 2$, split $|v-t| \le |t|/2$ from its complement. The near contribution to $P_\delta * w_\alpha$ is at most $2^\alpha w_\alpha(t) < 3w_\alpha(t)$. In the far region use $P_\delta(t-v) \le 4\delta/(\pi t^2)$ and $\int w_\alpha = 2/(\alpha-1)$, giving at most

$$\frac{2\delta 3^\alpha}{\pi(\alpha-1)}w_\alpha(t) \le 8w_\alpha(t).$$

Here $(1+|t|)^\alpha/t^2 \le 3^\alpha/4$ for $|t| \ge 2$. For $|t| \le 2$, the convolution is at most one and hence at most $3^\alpha w_\alpha(t) < 6w_\alpha(t)$. Thus $P_\delta * w_\alpha \le 11w_\alpha$ uniformly for $0 < \delta < 1/2$. Restoring the denominator costs less than $6(1+|t|)^{5/2}$, and the inverse principal factor is bounded by $(3/2+\delta)/(1/2-\delta)$. The product $11 \cdot 3 \cdot 6 = 198 < 200$ proves (2). The constant is uniform as $\delta \downarrow 0$, but not as the contour approaches the principal pole line $\mathrm{Re}\, s = 1$. No zero-free hypothesis appears.

15.6 The separately scoped unitary infinity type

Let $\nu(\mathfrak{a}) = \nu_0(\mathfrak{a})(N\mathfrak{a})^{-i\tau}$ be fixed and unitary, with angular parameter m and gamma factor γ_ν stated in Section 14. The external primitive functional equation is explicitly

$$Q_\rho^{s/2}\gamma_\nu(s)L(s,\rho) = \epsilon_\rho Q_\rho^{(1-s)/2}\gamma_{\bar\nu}(1-s)L(1-s,\bar\rho), \qquad |\epsilon_\rho| = 1.$$

A constant norm-twist phase in another completion convention is absorbed in ϵ_ρ. All finite Kummer twists share this infinity type. Replace T in the kernel by $T_\nu = 1 + |t+\tau| + |m|$. The numerator

gamma ratio now has real base $1/2 + |m|/2 > 0$ and vertical parameter $t + \tau$; its conjugate dual has parameter $-(t + \tau)$. Stirling gives the same integrable Gaussian-majorized kernel on each fixed shifted line. Constants can depend on m, τ; T_ν is comparable to T for this one fixed ν. The finite-prime conductor analysis, modulus-one coefficient twist, all-ideal decomposition and all original primitive row values remain unchanged.

A pure norm row has $L(s, \rho) = \zeta_K(s + i\tau)$; exclude it from the pole-free approximate functional equation and use scalar convexity at $t + \tau$. In the Poisson step cancel its actual pole by

$$h_\tau(s) = \frac{s - 1 + i\tau}{s + 1 + i\tau}.$$

This factor is analytic in the relevant right half-plane. On its boundary $|h_\tau| \le 1$, while inside $|h_\tau^{-1}| \le (3/2 + \delta)/(1/2 - \delta)$. Hence the same transfer factor 200 applies to the now ν-dependent central constant. If the opposite norm convention is chosen, replace τ by $-\tau$ in the gamma factor and pole alike. This proves the separate fixed-unitary moment assertion; it does not claim uniformity in a family of infinity types or replace the finite-order principal-residue expression below.

15.7 Keeping the full type-I mask during transfer

For $\sigma = 1/2 + \delta$, define L^T by removing every prime of $T(u, d)$, not just conductor primes. The identity (3) is equivalently

$$F_{u,d}(s) = \frac{L(s, \psi_u)}{L^{T(u,d)}(2s, \psi_{u,2})} \prod_{p \in T(u,d),\ p \nmid \mathfrak{f}_u} (1 - \psi_u(p)(Np)^{-s}).$$

Absolute convergence at $2\sigma > 1$ bounds the reciprocal incomplete product by $\zeta_K(2\sigma)$, without requiring any zero-free theorem in the critical strip. For every fixed $\eta > 0$,

$$\prod_{p \in T} (1 + (Np)^{-\sigma}) \ll_{\sigma,\eta} (N \operatorname{rad} T)^\eta, \qquad N \operatorname{rad} T \le C_S Nu\, Nd.$$

Indeed $1 + q^{-\sigma} \le q^\eta$ beyond a fixed threshold, and only finitely many prime ideals lie below it. Therefore

$$|F_{u,d}(\sigma + it)| \ll_{\nu,S,\delta,\eta} (Nu\, Nd)^\eta |L(\sigma + it, \psi_u)|. \tag{13}$$

This is a bounded inequality after the exact mask identity, not an identity that removes original zeros. For example, with principal ν and $u = \pi^3$ or π^6 at $N\pi = 7$, the primitive local ratio, restoration and full masked factor at $s = 1$ are respectively

$v_\pi(u)$	primitive ratio	restoration	masked factor
3	48/49	49/48	1
6	8/7	7/8	1

The squared primitive character in the first row is unramified at π although the original primitive character is ramified there. Squaring a zero sentinel would give the wrong denominator.

For finite-order ν, the normalized type-I piece is the retained d sum of the smoothed squarefree series at $Y = D/Nd$, divided by $\sqrt{D}$. Its possible pole is only at $s = 1$, on principal rows. Contour shifting with the full mask gives $I_u = P_u + E_u$, with E_u exactly (4) and

$$P_u = \mathbf{1}_{\psi_u = 1} \sqrt{D}\, \widehat{W}(1) \sum_{\substack{d \text{ squarefree} \\ (d,S)=1}} \frac{\beta_U(d) \nu(d) \chi_d(u)}{Nd} \rho_{T(u,d)},$$

$$\rho_T = \frac{\operatorname{res}_{s=1} \zeta_K(s)}{\zeta_K(2)} \prod_{p \in T} (1 + (Np)^{-1})^{-1}.$$

This does not impose positivity on P_u. Since β_U is supported on $Nd \le U^2 \le D$, every inner scale Y is at least one. Four integrations by parts give $|\widehat{W}(\sigma+it)| \ll_\sigma \mathcal{M}_4(W)(1+|t|)^{-4}$, which justifies the shifted integrals and their horizontal limits with the classical Hecke growth estimates.

Apply Minkowski to (4) in ℓ^2 of the actual u rows, then use (13) and the off-central moment. For sufficiently small internal losses this bounds $\|E\|_2$ by

$$CD^{\delta}H^{1/2+O(\eta)}\sum_{d}\frac{|\beta_U(d)|}{(Nd)^{\sigma-\eta}}\int_{\mathbb{R}}|\widehat{W}(\sigma+it)|(1+|t|)^{(1+O(\eta))/2}\,dt.$$

All original $\nu(d)\chi_d(u)$ have modulus at most one and remain inside the exact expression before this inequality. The actual truncated factorization gives, without replacing β_U by arbitrary coefficients,

$$\sum_{d}\frac{|\beta_U(d)|}{(Nd)^{\sigma-\eta}} \le \left(\sum_{Na\le U}(Na)^{-\sigma+\eta}\right)^2 \ll_{\delta,\eta} U^{1-2\delta+2\eta}.$$

Choose $\eta < \delta$ if necessary and use the fixed-field ideal count. The height integral is bounded by $C\mathcal{M}_4(W)$. Squaring and absorbing the small losses in $(HD)^{\varepsilon}$, using $U^2 \le D$, proves (5). Thus the formerly conditional numerator moment input in this residual transfer is discharged, while its positive-power cutoff loss is retained explicitly.

15.8 What remains open and what has been verified

The new remainder envelope improves the old scalar-convexity H power $3/2-\delta+\varepsilon$ to $1+\varepsilon$. At fixed U one may choose δ for a requested small-loss exponent; for $U = D^a$ with fixed $a > 0$, the remaining excess tends to $2a$ as $\delta \downarrow 0$. Neither statement controls P_u, the other signed piece, nor their covariance. In the original Vaughan identity the short $2a_U$ term must be restored at smaller scales where its support can meet the annulus. Uniform $\mathcal{M}_4$ bounds do not justify interchanging a profile supremum with the row moment or assuming $U^2 \le D'$ after a scale D' is reduced.

The independent prose audit is recorded in `planning/analytic-oracle/arithmetic-review/adversarial-afm-review-20261104.md`. The source-hashed finite checker tests valuation/support decompositions, squarefull coordinates, height exponents and rational Poisson envelopes. Its finite success does not prove reciprocity, analytic continuation, the large sieve, the approximate functional equation or the moments. No new Lean declarations accompany these analytic theorems. Existing Lean local Euler identities are a separately limited verification layer. The independent positive arithmetic determinant comparison, original signed low-gcd short-row estimate, all required contour uniformities and GRH remain open; Dirichlet, Hecke, Dedekind and general Artin scopes are not identified with one another.

Pinned primary sources and external inputs. The principal analytic family input is Blomer–Goldmakher–Louvel, arXiv:1112.1650v1, Theorem 1.3, Lemma 2.1, Appendix II and Section 5. Our fixed-ν proof changes column coefficients and redoes the functional-equation reduction; it does not claim the published corollary directly states every fixed twist. For comparison, Gao–Zhao, arXiv:2201.01885v2, Lemma 2.9 and Section 3.2, uses the sextic large-sieve/moment interface. Its GRH-conditional central-density theorem is not invoked. Abelian Kummer identification uses the classical reciprocity theorem, with the project's pinned Milne *Class Field Theory*, version 4.03 (6 August 2020), Theorem 0.8 and Chapter I, Section 1, a reference: Milne's notes. The primitive degree-one Hecke continuation and completed functional equation are explicit classical external inputs; the project's pinned Kudla notes on Tate's thesis, Theorem 4.3 and Corollary 4.4, supply a reference.

Their sketch and incomplete displayed epsilon formulas are not used as independently certified constants. Standard Stirling, Mellin inversion, Poisson representation and scalar convexity complete the external analytic dependencies. None is replaced by a finite diagnostic or a GRH-equivalent assumption.

16 Recombining the actual weighted cutoff

Keep the finite-order Eisenstein family and original masks of Section 14. For $U \geq 1$ and a squarefree ideal m, put

$$A_U(m) = \sum_{d|m} \beta_U(d) = \sum_{\substack{ab|m \\ (a,b)=1,\ Na,Nb\leq U}} \mu(a)\mu(b).$$

Every summation variable here is an integral ideal. The last formula uses squarefreeness of m; an unconstrained square of a Möbius polynomial would also include forbidden overlapping supports.

Theorem 7 (Exact recombination and an absolute-half-plane rate)**.** *Write $T_0 = S \cup \mathrm{supp}(u)$ and $J_{u,U}(s) = \sum_{d \text{ squarefree},\ (d,S)=1} \beta_U(d)\nu(d)\chi_d(u)(Nd)^{-s}F_{u,d}(s)$. For $\mathrm{Re}\, s > 1$,*

$$J_{u,U}(s) = \sum_{m \text{ squarefree},\ (m,T_0)=1} A_U(m)\psi_u(m)(Nm)^{-s},$$

$$|A_U(m)| \leq 3^{\omega(m)}, \qquad A_U(m) = \mu(m) \quad (Nm \leq U).$$

For every $\kappa > 0$ and $0 < \epsilon < \kappa$, uniformly for $\mathrm{Re}\, s \geq 1 + \kappa$, all original rows and all $U \geq 1$,

$$\left| J_{u,U}(s) - \frac{1}{L^{T_0}(s,\psi_u)} \right| \leq U^{-\kappa+\epsilon}\{\zeta_K(1+\epsilon)^3 + \zeta_K(1+\epsilon)\}.$$

Here L^{T_0} removes every prime of T_0, including conductor-drop primes; the bound is in the absolute-convergence half-plane only.

Proof. On a surviving term $m = dc$, the two squarefree ideals c, d are coprime, and both avoid T_0. The original product $\nu(d)\chi_d(u)\psi_u(c)$ is then $\psi_u(m)$. Grouping by m gives A_U. Each prime of m has three possible positions in $(a, b, m/(ab))$, so the absolute sum is at most $3^{\omega(m)}$. When $Nm \leq U$ all divisors fit both cutoffs; the primewise sum is $1 - 1 - 1 = -1$, giving $\mu(m)$.

The same bound justifies absolute convergence and regrouping for $\mathrm{Re}\, s > 1$, because on squarefree ideals $3^{\omega(m)} = \tau_3(m)$ and $\sum_m \tau_3(m)(Nm)^{-v} = \zeta_K(v)^3$ for $v > 1$, where this latter sum runs over all nonzero integral ideals. Subtract the absolutely convergent Möbius series for $1/L^{T_0}$. Only $Nm > U$ remain. Rankin's inequality with exponent $\eta = \kappa - \epsilon$ bounds their absolute sum by

$$U^{-\eta} \sum_m \{\tau_3(m) + 1\}(Nm)^{-1-\epsilon},$$

which proves the stated constant independently of deleted primes, row multiplicities or character phases. □

For W supported in $[1, 2]$, the same exact regrouping gives the original normalized type-I polynomial

$$I_u = D^{-1/2} \sum_{m \text{ squarefree},\ (m,S)=1} A_U(m)\nu(m)\chi_m(u)W(Nm/D).$$

This formula retains all original zero sentinels. It permits a direct family large sieve before any absolute d sum; its quantitative moment is developed separately below. The displayed absolute-half-plane rate alone is not that critical-strip moment.

17 Euler packaging, coprimality and the analytic boundary

For $p \notin T_0$, set $z_p = \psi_u(p)(Np)^{-s}$ and $t_p = z_p/(1+z_p)$. On $\operatorname{Re} s > 1$, deleting an additional surviving prime divides the squarefree Euler product by $1+z_p$. Consequently the exact retained expression is

$$J_{u,U}(s) = F_{u,1}(s)B_{u,U}(s), \qquad B_{u,U}(s) = \sum_{d \text{ squarefree},\ (d,T_0)=1} \beta_U(d) \prod_{p \mid d} t_p.$$

Original zeros enforce $(d,u)=1$ even when a sixth power drops from its primitive conductor. This factorization cannot restore a forbidden d row by replacing a zero sentinel with an unramified primitive value. The equality continues wherever both sides are holomorphic.

Without the cutoffs, $\beta_\infty(d) = \mu(d)2^{\omega(d)}$ on squarefree ideals. Absolute convergence justifies

$$B_{u,\infty} = \prod_{p \notin T_0} (1-2t_p) = \prod_{p \notin T_0} \frac{1-z_p}{1+z_p}, \qquad F_{u,1}B_{u,\infty} = \prod_{p \notin T_0} (1-z_p) = \frac{1}{L^{T_0}(s,\psi_u)}.$$

Theorem 7 gives a quantitative rate rather than an unsupported exchange of this limit with a critical-strip contour.

A local counterexample to the naive square. In a one-prime projection at a good Eisenstein prime of norm 7, take the principal local phase and $s=1$. Then $z=1/7$, $t=1/8$ and, for $U \geq 7$, the exact polynomial is $B_U = 1-2t = 3/4$. The unconstrained square $(1-t)^2 = 49/64$ differs by $1/64$: it admits the forbidden pair $a=b=p$. After multiplying by $F = 1+z = 8/7$ it gives $7/8$ instead of the correct $6/7$. This is an actual local Euler calculation, not an infinite-series or mean-square counterexample. It identifies the exact coprimality condition that a proposed mollifier-square simplification must preserve.

The exact intermediate coefficient band. For a squarefree m with $U < Nm \leq U^2$, one has $A_U(m) = -\mu(m)$. Indeed the unrestricted disjoint-assignment sum is $\mu(m)$. If $Na > U$, every $b \mid m/a$ automatically has $Nb \leq Nm/Na < U$, and $\sum_{b \mid m/a} \mu(b)$ vanishes unless $a = m$. Hence the excluded a assignments contribute exactly $\mu(m)$; the excluded b assignments contribute another $\mu(m)$. They cannot overlap because $NaNb > U^2 \geq Nm$ contradicts $ab \mid m$. Subtraction proves the claim. Together with the small-norm band, this gives $A_U = \mu$ for $Nm \leq U$ and $A_U = -\mu$ for $U < Nm \leq U^2$. The original smoothed columns have $Nm \geq D \geq U^2$, so the new band does not justify replacing their coefficients by $-\mu(m)$: generally they lie beyond it.

Theorem 8 (A stronger cutoff convergence would require zero-freeness)**.** *Fix a nonprincipal primitive ψ_u. If these actual $J_{u,U}$ converge locally uniformly on $\operatorname{Re} s > 1/2$ as $U \longrightarrow \infty$, then $L(s,\psi_u)$ has no zero there.*

Proof. Classical abelian Hecke continuation makes the numerator of $F_{u,1} = L^{T_0}(s,\psi_u)/L^{T_0}(2s,\psi_{u,2})$ holomorphic. Its denominator is nonzero on $\operatorname{Re} s > 1/2$ by the absolute Euler product at $2s$. Each $B_{u,U}$ is a finite sum with nonzero denominators $1+z_p$ on $\operatorname{Re} s > 0$. Thus every $J_{u,U}$ is holomorphic on the asserted half-plane. A locally uniform limit G is holomorphic there. On $\operatorname{Re} s > 1$, the already proved convergence gives $L^{T_0}(s,\psi_u)G(s) = 1$. The identity theorem extends this equality to the connected half-plane, precluding a zero of L^{T_0}. The finitely deleted factors $1-\psi_u(p)(Np)^{-s}$ do not vanish when $\operatorname{Re} s > 0$, so the same conclusion holds for the primitive L-function. □

This is an implication, not a proof that the stronger convergence holds, and it does not assert a converse. In particular, a family-averaged energy bound is not silently promoted to locally uniform convergence for each individual character. Principal rows have a pole at 1 and were excluded from this proposition. General Artin holomorphy is not an input.

Reproducible exact diagnostics. The new source `examples/arithmetic_cutoff/exact_recombination.py` uses actual split/inert prime ideals and raw sextic symbols in $\mathbb{Z}[\omega]$, with exact rational arithmetic. It checks the direct masked sum, coefficient regrouping, Euler packaging, both finite full-cutoff limits and original zeros, including third/sixth-power rows. Its overlapping- support negative controls are not tests of an infinite large sieve or of zero location. The analytic convergence argument above uses ordinary ideal Euler products and the explicitly stated Hecke continuation theorem; no new Lean declaration or complete GRH proof is claimed.

18 A cutoff-uniform envelope for the original masked remainder

Fix $K = \mathbb{Q}(\omega)$, $\mathcal{O} = \mathbb{Z}[\omega]$, a finite-order Hecke character ν, and a finite set S containing the primes above 6 and the conductor of ν. All unspecified ideal sums below are over nonzero integral ideals. Let $\chi_m(u) = (u/m)_6$ be the original symbol, zero when $(m, u) > 1$, and put

$$\beta_U(d) = \sum_{\substack{ab=d \\ Na, Nb \le U}} \mu(a)\mu(b), \qquad A_U(m) = \sum_{d \mid m} \beta_U(d).$$

The columns d, c, m are squarefree and prime to S where indicated. Retain the primitive characters $\psi_u = \mathrm{prim}(\nu\kappa_u)$, $\psi_{u,2} = \mathrm{prim}(\psi_u^2)$, and the full deleted set $T(u, d) = S \cup \mathrm{supp}(u) \cup \mathrm{supp}(d)$. For $H, D, U \ge 1$, $U^2 \le D$, and $W \in C_c^\infty((0, \infty))$ supported in $[1, 2]$, the original normalized type-I piece is

$$I_u = D^{-1/2} \sum_{\substack{m \text{ squarefree} \\ (m,S)=1}} A_U(m)\nu(m)\chi_m(u) W(Nm/D). \tag{14}$$

Its exact principal residue is

$$P_u = \mathbf{1}_{\psi_u=1} \sqrt{D}\, \widehat{W}(1) \sum_{\substack{d \text{ squarefree} \\ (d,S)=1}} \frac{\beta_U(d)\nu(d)\chi_d(u)}{Nd} \rho_{T(u,d)},$$

$$\rho_T = \frac{\mathrm{res}_{s=1}\, \zeta_K(s)}{\zeta_K(2)} \prod_{p \in T} (1 + (Np)^{-1})^{-1}. \tag{15}$$

Define $E_u = I_u - P_u$; this is the same residue-subtracted remainder as the fully masked Mellin contour, with no cutoff limit or profile supremum.

Theorem 9 (Uniformity in the actual truncated cutoff)**.** *For every $\varepsilon > 0$,*

$$\sum_{0 < Nu \le H} |E_u|^2 \ll_{\nu,S,\varepsilon} \|W\|_\infty^2 (HD)^\varepsilon \{H + DH^{1/3} + (HD)^{2/3}\}. \tag{16}$$

The sum is over every actual nonzero Eisenstein element, including units, S-supported powers and sixth-power rows. The constant is independent of U under $U^2 \le D$. This is an application of the classical sextic large sieve with exact original coefficients, not a zero-location theorem.

Proof. On an absolutely convergent line, expand the restored squarefree series $F_{u,d}$. Its condition $(c, T(u,d)) = 1$ includes $(c,d) = 1$, and the original symbols multiply. Grouping $m = cd$ gives (14) by Mellin inversion. Each prime of m may belong to the first truncated Möbius divisor, to the second, or to neither. Consequently $|A_U(m)| \le 3^{\omega(m)} \ll_\eta (Nm)^\eta$, uniformly in U. The length-D coefficient sequence in (14) therefore has squared norm $\ll_\eta \|W\|_\infty^2 D^\eta$.

The all-ideal row norm in Section 3 of Blomer–Goldmakher–Louvel, arXiv:1112.1650v1 gives

$$B_2^1(X,D) \ll_\eta (XD)^\eta \{X + DX^{1/3} + (XD)^{2/3}\}.$$

For completeness, split $a = a_1 a_2^2 \cdots a_5^5 a_6^6$ dyadically; $a_1, \dots, a_5$ are squarefree and pairwise coprime. If their dyadic sizes are A_j, put $R = \prod_{j=1}^6 A_j$ and $M = \max(A_1, A_2)$. Pivot on the first or second variable, using the squarefree powered sieve proved simultaneously in that paper's Section 3.2. The other variables are bounded column multipliers. The block energy is at most

$$(XD)^\eta \{R + DR/M + D^{2/3}R/M^{1/3}\} \sum_m |\lambda_m|^2.$$

The inequalities $R \le X$, $R^3 \le XM^3$ and $R^3 \le X^2 M$ follow from $\prod_j A_j^j \le X$ and the pivot choice. Summing the logarithmically many blocks proves the displayed bound.

Finite S-unit and ray-class partitions convert the original literal symbol into the canonical column value $\vartheta_m(a)$ with fixed modulus-one coefficient multipliers. On noncoprime pairs both values are zero: squarefree m has outside-S conductor exactly m. This orientation preserves sixth-power zeros. At most six elements generate each outside-S ideal times an S-supported ideal; the S-supported sums converge at powers 1, 1/3 and 2/3. Thus (16) holds first with I_u in place of E_u.

A principal ψ_u requires every outside-S valuation of u to be divisible by 6. There are $O_{\nu,S}(H^{1/6})$ such literal rows. Equation (15), $|\rho_T| \ll_K 1$, and the actual coefficient sum

$$\sum_d \frac{|\beta_U(d)|}{Nd} \le \left(\sum_{Na \le U} \frac{1}{Na} \right)^2 \ll \log^2(2U)$$

give $\sum |P_u|^2 \ll DH^{1/6} \|W\|_\infty^2 \log^4(2U)$. This is absorbed by the right side of (16), using $U^2 \le D$ and a smaller internal exponent. Finally $|E_u|^2 \le 2|I_u|^2 + 2|P_u|^2$. The extended proof specifies the canonical-family transfer and the powered-sieve input precisely. □

Combining this with the previously proved fixed-twist AFM transfer yields

$$\sum_{0 < Nu \le H} |E_u|^2 \ll_{\nu,S,\delta,\varepsilon} \mathcal{M}_4(W)^2 (HD)^\varepsilon \min\{D^{2\delta} U^{2-4\delta} H,\ H + DH^{1/3} + (HD)^{2/3}\}. \tag{17}$$

At $H = D^{3/4}$ and $U = D^a$, $0 \le a \le 1/2$, the excess over H is at most $\min\{2\delta + (2-4\delta)a, 1/2\}$, up to arbitrarily small losses. The new envelope improves the old one for $a > 1/4$ as $\delta \downarrow 0$, but still leaves an exponent loss of 1/2. It does not settle the original signed short-row bound, type-II covariance, positive arithmetic determinant comparison or GRH. The principal-residue bound used here is sufficient for this envelope, not the stronger target needed to reassemble the full signed argument.

19 All-ideal reciprocity and the surviving block obstruction

The proof uses no new zero-free input. Its external analytic ingredient is the pinned BGL large sieve and the powered estimates established in its proof; elementary ideal counting and the already restored Mellin identity supply the remaining interfaces. All constants may depend on fixed K, S, ν and declared small exponents, never on U.

19.1 Exact recombination retains the primitive corrections

On $\operatorname{Re} s > 1$, start from the full defining series, with its original zero masks:

$$F_{u,d}(s) = \sum_{\substack{c \text{ squarefree} \\ (c,T(u,d))=1}} \nu(c)\chi_c(u)(Nc)^{-s}.$$

Its primitive ratio involves $L(s,\psi_u)$, the independently primitive $L(2s,\psi_{u,2})$ and every deleted Euler factor. Expanding this *restored* series gives the exact equality

$$\begin{aligned}\sum_{\substack{d \text{ squarefree} \\ (d,S)=1}} &\beta_U(d)\nu(d)\chi_d(u)(Nd)^{-s}F_{u,d}(s) \\ &= \sum_{\substack{m \text{ squarefree} \\ (m,S)=1}} A_U(m)\nu(m)\chi_m(u)(Nm)^{-s}. \end{aligned} \tag{18}$$

Indeed $(c,d)=1$ makes $m=cd$ squarefree, with $d \mid m$. If m meets u, both original symbols retain their zero. Only on the absolutely convergent line is this infinite series used. Mellin inversion there gives the finite smoothed polynomial (14). Nothing evaluates (18) as a Dirichlet series on $\operatorname{Re} s = 1/2+\delta$ or interchanges a cutoff limit with continuation across unknown zeros.

For squarefree m, a factorization $ab = d \mid m$ assigns each prime to a, to b, or to neither. The actual inequalities $Na, Nb \le U$ select a subset of these assignments. Thus

$$|A_U(m)| \le 3^{\omega(m)} \ll_\eta (Nm)^\eta.$$

The last divisor bound follows by choosing a fixed threshold so $3 \le (Np)^\eta$ beyond it, then absorbing the finitely many smaller prime ideals in a constant. When $Nm \le U$ all assignments occur, so $A_U(m) = \prod_{p\mid m}(1-1-1) = \mu(m)$; no such simplification is made for the length-D columns in general. With $\lambda_m = D^{-1/2}A_U(m)\nu(m)W(Nm/D)$, ideal counting and a smaller internal η give

$$\sum_m |\lambda_m|^2 \ll_\eta \|W\|_\infty^2 D^\eta, \qquad \operatorname{supp}\lambda \subseteq \{D \le Nm \le 2D\}. \tag{19}$$

19.2 The powered squarefree input and all-ideal pivot

Let ϑ_a be BGL's canonical primitive sixth-order family. Their Section 3 defines B_1^j using squarefree row/column ideals and kernel $\vartheta_a(m)^j$, while B_2^j uses *all* ideal rows and kernel $\vartheta_m(a)^j$. Theorem 1.3 states the squarefree bound for $j=1$; its Section 3.2 proves the estimates simultaneously for $1 \le j \le 5$, because the recursion in Lemma 3.4 uses $j\ell$. In particular the two inputs needed here are

$$B_1^j(M,D) \ll_\eta (MD)^\eta\{M + D + (MD)^{2/3}\}, \qquad j = 1,2. \tag{20}$$

To see the final step in the published recursion explicitly, its $E_{4/3}$ estimate is $M^{4/3}+D+(MD)^{2/3}$. If $M \le D$, its first term is at most $(MD)^{2/3}$; if $M > D$, use reciprocity/duality (3.2) and the estimate with M, D reversed. The finite set of powers is treated together; no inference from a fixed conductor count to an arbitrary primitive-character sieve occurs.

Write each ideal row uniquely as

$$a = a_1 a_2^2 a_3^3 a_4^4 a_5^5 a_6^6,$$

where $a_1, \dots, a_5$ are pairwise-coprime squarefree ideals and a_6 is arbitrary. For valuation $6q+r$, $0 \le r \le 5$, place the prime in a_r if $r > 0$ and its power q in a_6. Dyadically let $A_j \le Na_j < 2A_j$,

with $A_j \geq 1$, and put $R = \prod_{j=1}^{6} A_j$. A nonempty block satisfies $\prod_j A_j^j \leq X$ when $Na \leq X$. Choose $\ell = 1$ if $A_1 \geq A_2$, and $\ell = 2$ otherwise; put $M = A_\ell$. Keep every other ideal fixed. Evaluation of the single *column* character gives

$$\vartheta_m(a) = \vartheta_m(a_\ell)^\ell \vartheta_m(a/a_\ell^\ell).$$

The second factor is a bounded column multiplier, including its original zero if m meets any fixed factor. It may depend on all the fixed ideals but not on the pivot row. Pairwise-coprimality and the exact row norm restriction merely select a subset of the positive pivot-row sum. Apply (20) to the first factor. There are $O_K(R/M)$ choices of the other ideals by ideal counting. The block energy is bounded by

$$(XD)^\eta \{R + DR/M + D^{2/3}R/M^{1/3}\} \sum_m |\lambda_m|^2. \tag{21}$$

All three monomial bounds used to sum these blocks are elementary. First $R \leq \prod_j A_j^j \leq X$. If $M = A_1$, then

$$(R/M)^3 = A_2^3 \cdots A_6^3 \leq A_1 A_2^2 A_3^3 A_4^4 A_5^5 A_6^6 \leq X,$$

since $A_2 \leq A_1$ and $j \geq 3$ for the remaining factors. If $M = A_2$, use $A_1 \leq A_2$ in the same comparison. Finally,

$$\frac{R^3}{M(\prod_j A_j^j)^2} = \frac{A_1}{M A_2 A_3^3 A_4^5 A_5^7 A_6^9} \leq 1.$$

Thus $R/M \leq X^{1/3}$ and $R/M^{1/3} \leq X^{2/3}$. Insert these in (21) and sum the logarithmically many blocks. This proves the all-ideal estimate used in Theorem 9. It is precisely the mechanism of BGL Lemma 3.4 and (3.6)–(3.9), with the final powered bound inserted. The simpler fixed-squarefull-row argument would leave the weaker $D\sqrt{X}$ term.

19.3 Literal rows, reciprocal orientation and principal count

Let $a = (u)_S$ denote the outside-S ideal. The finite S-unit comparison gives $\kappa_u = \eta_u \vartheta_a$ as primitive characters, with η_u in a finite fixed set. For $(a, m) = 1$, BGL reciprocity therefore gives

$$\chi_m(u) = \eta_u(m) C([a], [m]) \vartheta_m(a). \tag{22}$$

The notation is deliberate: the left side is the project's original $(u/m)_6$, not BGL's canonical character indexed by m. When $(a, m) > 1$, the left side is zero. The right side is also zero, because squarefree m has outside-S primitive conductor exactly m. Equation (22) hence extends to every pair. Using $\vartheta_a(m)$ on noncoprime pairs instead would be false: primitive ϑ_{b^6} has restored primes of b, whereas the original symbol must still vanish there.

Partition rows by η_u and $[a]$, both finite. On each part the factor $\eta_u(m)C([a], [m])$ is a fixed modulus-one column multiplier. A positive row sum over that part is bounded by the full B_2^1 row sum. For each S-supported ideal b and outside-S ideal a, at most six actual elements generate ab. Thus the required literal energy is bounded by the all-ideal norm summed over b, with $X = H/Nb$. The geometric products

$$\sum_{b \ S\text{-supported}} (Nb)^{-q} = \prod_{p \in S} (1 - (Np)^{-q})^{-1} < \infty, \qquad q = 1,\ 1/3,\ 2/3,$$

transfer the three terms without losing any unit or S-power row. Together with (19), this proves the claimed I energy.

If ψ_u is principal, ν is unramified outside S, and tame Kummer inertia at $v_p(u)$ has order $6/\gcd(6, v_p(u))$. It must be trivial, so each such valuation is divisible by 6. Writing $(u) = ba_6^6$ with b supported in S, ideal counting gives

$$\#\{0 < Nu \le H : \psi_u = 1\} \ll_K H^{1/6} \prod_{p \in S} (1 - (Np)^{-1/6})^{-1}.$$

This is only a necessary condition, sufficient for an upper count. The exact residue formula (15) and $|\widehat{W}(1)| \le \|W\|_\infty$ then give the principal energy bound used above. No sign or positivity of that residue is assumed.

19.4 What the envelope still fails to prove

The $DH^{1/3}$ term is still larger than $H(HD)^\varepsilon$ at $H = D^{3/4}$. The monomial comparison itself is saturated by $A_1 = A_2 = X^{1/3}$ and $A_3 = \cdots = A_6 = 1$. This proves sharpness of that *particular comparison*, not an energy lower bound for actual arithmetic coefficients. In particular, no conclusion that the signed A_U cannot cancel is justified. The next load-bearing input must exploit those actual truncated Möbius coefficients inside the mixed moment, or establish the required signed type-II minus type-I covariance. The smaller-scale short term, principal contribution, full arithmetic determinant identification and GRH remain separate obligations.

Exact finite checks use the split Eisenstein primes generated by $3 + \omega$, $4 + \omega$, $5 + 2\omega$, $6 + \omega$, of norms $7, 13, 19, 31$. Their actual residue symbols are computed in $\mathbb{Q}(\zeta_6)$, retaining sixth-power zeros. Coefficient recombination and dyadic inequalities pass; they check finite algebra and bookkeeping, not the external analytic sieve or the contour theorem. No new Lean declarations accompany this argument.

20 High cutoffs, actual bilinear factors and a signed covariance test

Retain the original Eisenstein family, fixed finite-order ν, fixed S, common $W \in C_c^\infty((0,\infty))$ supported in $[1, 2]$, and every raw shared-prime zero. Put

$$C_U(m) = \sum_{\substack{abc=m \\ Na>U,\, Nb>U}} \mu(a)\mu(b), \qquad r_U(m) = \#\{(a, b) : ab = m,\ Na, Nb > U\}.$$

All factors are ideals, squarefree and pairwise coprime; on the active annulus $A_U = C_U - \mu$. A coefficient identity is different from a row-energy bound.

Proposition 10 (The high-cutoff window)**.** *If $D > 1$ and $2D/7 \le U^2 \le D$, then for every surviving squarefree column $D \le Nm \le 2D$,*

$$C_U(m) = \mu(m) r_U(m), \qquad A_U(m) = \mu(m)(r_U(m) - 1), \qquad |A_U(m)| \ge 1.$$

For each fixed nonzero W, the normalized squared coefficient norm satisfies

$$D^{-1} \sum_{\substack{m \text{ sf} \\ (m,S)=1}} |A_U(m) W(Nm/D)|^2 \gg_{W,S} 1 \tag{23}$$

for large D, uniformly in this window. This is not a lower bound for row energy.

Proof. The least norm of a nonunit ideal outside S is at least 7. A long-long term with $c \neq 1$ would have $Nm > 7U^2 \geq 2D$. Thus $ab = m$, including both endpoints. The fixed-point-free swap $(a, b) \mapsto (b, a)$ makes r_U even, so A_U is an odd nonzero integer. The last claim follows by counting squarefree ideals on a fixed interval where $|W|$ is bounded below; the constant and error are proved in the extended proof. □

The exact remaining polynomial is genuinely bilinear, with $(a, b) = 1$, $U < Na, Nb < 2D/U$, $D \leq Na\,Nb \leq 2D$, and $x_{ab}(u) = x_a(u)x_b(u)$, including zeros. Both factors see the same row $u = \varepsilon\alpha\beta^2$. Cauchy followed by the accepted powered BGL bound gives only

$$Q_{B_X}(C_U, C_U) \ll_{\nu,S,\epsilon} (DX)^{\epsilon}\|W\|_{\infty}^2\{\sqrt{D}\,X^2 + DX + D^{5/6}X^{5/3}\}. \tag{24}$$

At $X \asymp D^{1/4}$ the exponents are $1, 5/4, 5/4$; the earlier regrouped bound has $1/2, 5/4, 13/12$ and is stronger. No BT_σ saving is obtained. Multiplying two marginal operator norms does not supply a coupled fourth moment.

A finite test in the original family. There is no automatic nonnegative signed covariance, even when $C_U(m)\mu(m) \geq 0$. Take $\nu = 1$, S the primes above $2, 3$, $D = 400$, $U = 17$, and all 72 literal rows of B_7. Choose the common nonnegative smooth profile constructed in the extended proof, with values 10 at 439/400, 1 at 475/400, and zero at every other integer norm divided by 400. There are exactly four active squarefree ideal columns, of norms $439, 439, 475, 475$. In that order,

$$A = (1, 1, 1, 1), \quad C = (0, 0, 2, 2), \quad \mu = (-1, -1, 1, 1), \qquad \Re G_{B_7} = \begin{pmatrix} 72 & 0 & 0 & 18 \\ 0 & 72 & 18 & 0 \\ 0 & 18 & 72 & 0 \\ 18 & 0 & 0 & 72 \end{pmatrix}.$$

Exact raw-symbol enumeration, independently reproducible from the ideal and residue-field definitions, gives

$$Q(C, \mu) = -\frac{27}{25}, \qquad Q(C, C) = \frac{36}{25}, \quad Q(\mu, \mu) = \frac{864}{25}, \quad Q(A, A) = \frac{954}{25}. \tag{25}$$

The real Gram matrix is positive definite, with eigenvalues $54, 90$; its positivity does not force a cross covariance to be positive. This fixed finite example refutes that shortcut, not the asymptotic BT_σ target or the original signed mean-square estimate.

The two-cutoff contour remains coupled. For $T_0 = S \cup \mathrm{supp}(u)$, $\psi = \psi_u$ and independently primitive $\psi_2 = \mathrm{prim}(\psi^2)$, put $q_p = \psi(p)(Np)^{-s}$, $\alpha_p = (Np)^{-z}$, $\beta_p = (Np)^{-w}$ outside T_0. The actual three-placement Euler factor is $1 + q_p - q_p\alpha_p - q_p\beta_p$. Its exact factorization is

$$\mathcal{F}(s, z, w) = \frac{L^{T_0}(s, \psi)}{L^{T_0}(s + z, \psi)L^{T_0}(s + w, \psi)} \prod_{p \notin T_0} \left(1 - \frac{q_p^2(1 - \alpha_p)(1 - \beta_p)}{(1 - q_p\alpha_p)(1 - q_p\beta_p)}\right). \tag{26}$$

Exact double Perron inversion is valid on positive auxiliary lines $c, d > 1 - \Re s$ for $1/2 < \Re s < 1$, retaining the original principal pole exclusion. These safe lines reproduce the previous $U^{2-4\delta}$ energy loss at $\Re s = 1/2 + \delta$, with the same common-profile derivative and fixed-margin dependencies. The correction product alone does not justify moving those lines through reciprocal-L poles. In fact $\mathcal{F}(s, z, 0) = 1/L^{T_0}(s + z, \psi)$ as a meromorphic identity. No zero-free region, reciprocal moment, analytic saving, new Lean formalization or GRH proof is supplied by this construction.

21 Proof details for the high-window and two-cutoff tests

21.1 The coefficient lower bound and its precise limitation

The element lattice $\mathbb{Z}[\omega] \subset \mathbb{C}$ has covolume $\sqrt{3}/2$. Counting disks with boundary error $O(\sqrt{x})$ and dividing by the six units, using that this ring is a PID, gives

$$N_K(x) = \kappa x + O(\sqrt{x}), \qquad \kappa = \frac{\pi}{3\sqrt{3}}.$$

Finite inclusion-exclusion at S gives $N_{K,S}(y) = \kappa y \prod_{p\in S}(1-(Np)^{-1}) + O_S(\sqrt{y})$. Insert $1_{\mathrm{sf}}(m) = \sum_{d^2|m}\mu(d)$. Absolute convergence at 2, the tail $\sum_{Nd>\sqrt{x}}(Nd)^{-2} = O(x^{-1/2})$, and $\sum_{Nd\le\sqrt{x}}(Nd)^{-1} = O(\log(2x))$ yield

$$N_{\mathrm{sf},S}(x) = c_S x + O_S(\sqrt{x}\log(2x)), \qquad c_S = \frac{\kappa}{\zeta_K(2)}\prod_{p\in S}(1+(Np)^{-1})^{-1} > 0.$$

These tail estimates follow by partial summation from the displayed lattice count. For a fixed nonzero W, choose $[a,b] \subset (1,2)$ with $|W| \ge \eta > 0$. Proposition 10 gives, for large D,

$$D^{-1}\sum |A_U(m)W(Nm/D)|^2 \ge \tfrac{1}{2}\eta^2 c_S(b-a).$$

There is no uniform lower constant over all profiles with one specified supremum norm. Nor can a coefficient lower bound be combined with an operator upper bound to produce a row-energy lower bound: no lower singular value has been established on the actual vector.

21.2 What legal factor Cauchy actually proves

For dyadic $Na \asymp A, Nb \asymp B$, retain the exact $(a,b)=1$ mask, hyperbolic support and $W(Na\,Nb/D)$ in

$$T_a(u) = \sum_b \mu(b)\nu(b)W(Na\,Nb/D)x_b(u).$$

Cauchy gives

$$|\sum_a \mu(a)\nu(a)x_a(u)T_a(u)|^2 \ll A\sum_a |T_a(u)|^2.$$

For fixed a, fixed second row factor and fixed unit/ray classes, the b coefficients are independent of the pivot row and have squared norm $\ll B\|W\|_\infty^2$. The prior classical power-one/two BGL input has $\Delta(X,B) = X + B + (XB)^{2/3}$. Summing the $O(A)$ fixed columns and $O(X)$ other row factors gives

$$\mathcal{E}_{A,B} \ll_{\nu,S,\epsilon} (DX)^\epsilon \|W\|_\infty^2 \frac{A^2B}{D} X\Delta(X,B) \ll (DX)^\epsilon\|W\|_\infty^2\{AX^2 + DX + D^{2/3}A^{1/3}X^{5/3}\}.$$

In the high window $A, B \asymp \sqrt{D}$, and only boundedly many dyadic pairs occur. This proves (24); exchanging factors or row pivots leaves its DX term. A product of marginal second-moment bounds supplies no estimate for the two transforms on the same (r,s). Any useful saving must control that coupled fourth moment or the signed actual-coefficient covariance, with all masks retained.

21.3 A smooth profile and the exact covariance computation

Let

$$\phi(t)=\begin{cases}\exp(1-1/(1-t^2)), & |t|<1,\\ 0, & |t|\ge 1,\end{cases}\qquad W_x(y)=x\phi(4(400y-439))+\phi(4(400y-475)).$$

For $x\ge 0$ this is nonnegative and smooth; its disjoint supports lie strictly inside $(1,2)$. For every integer n, $W_x(n/400)$ equals x at $n=439$, 1 at $n=475$ and zero otherwise. Thus the computation uses exact values of a genuine common C_c^∞ profile, rather than claiming a finite rational weight is smooth. At $x=10$, $\|W_x\|_\infty=10$.

The primary ideal generators, in the order used above, are

$$(-23,-18),\quad (-5,18),\quad (10,-15),\quad (25,15),$$

where (a,b) denotes $a+b\omega$. The first two have norm 439. The last two factor as $(-2,3)(-5,0)$ and $(-5,-3)(-5,0)$, with norms $19\cdot 25$. Since $19,25>17$, the long-factor count is 2 there and zero on the prime columns. All squarefree ideals at both selected norms are included. The complete B_7 uses the four good primes of norms $7,7,13,13$ for distinct ordered row factors and all six units. It has 72 distinct rows, maximum norm 2197; each is nonprincipal by its outside-S Kummer valuation 1 or 2 modulo 6, whose nontrivial tame character has conductor exponent 1.

For the raw Gram $G_{ij}=\sum_{u\in B_7} x_{m_i}(u)\overline{x_{m_j}(u)}$, exact finite-field exponentiation and lattice enumeration give the displayed real matrix. The weighted real vectors are $\ell=(0,0,2,2)$, $v=(-x,-x,1,1)$, and $a=(x,x,1,1)$. Therefore

$$Q(C_U,\mu)=D^{-1}\ell^t(\Re G)v=\frac{36-9x}{50}.$$

Its diagonal contribution is $36/50>0$ and its off-diagonal contribution is $-9x/50$. For $x=10$ this is $-27/25$ and gives (25). The identity $(\Re G-72I)^2=18^2I$ proves the real eigenvalues $54,90$; these are not being asserted as the eigenvalues of the full complex Gram. The exact source is `examples/arithmetic_cutoff/signed_covariance.py`. A second test retains raw zeros on the 246 rows $0<Nu\le 89$ whose outside-S part is nonunit. Its covariance at $x=10$ is $-69/50$. This subset is guaranteed nonprincipal; it is not called the entire nonprincipal family. Here $89=\lfloor 400^{3/4}\rfloor$. These are finite counterexamples to automatic covariance positivity, not asymptotic counterexamples to BT_σ or GRH.

21.4 Double Perron without a divergent critical-strip series

For $T_0=S\cup\operatorname{supp}(u)$ retain $F_0(s)=L^{T_0}(s,\psi)/L^{T_0}(2s,\psi_2)$, with the independently primitive square character. Set $t_p=q_p/(1+q_p)$ outside T_0. The finite actual cutoff is

$$J_U(s,u)=F_0(s)B_U(s),\qquad B_U(s)=\sum_{\substack{a,b\ \mathrm{sf},\ (a,b)=1\\ (ab,T_0)=1,\ Na,Nb\le U}}\mu(a)\mu(b)\prod_{p\mid ab}t_p.$$

For $\sigma=\Re s>1/2$, the series

$$\mathcal{G}(s,z,w)=\sum_{\substack{a,b\ \mathrm{sf},\ (a,b)=1\\ (ab,T_0)=1}}\mu(a)\mu(b)\prod_{p\mid ab}t_p\,(Na)^{-z}(Nb)^{-w}=\prod_{p\notin T_0}\left(1-t_p(Np)^{-z}-t_p(Np)^{-w}\right)$$

converges absolutely on $\sigma+\Re z>1$, $\sigma+\Re w>1$, uniformly in all vertical parts at fixed positive margins. This follows from $|t_p|\le (Np)^{-\sigma}/(1-(Np)^{-\sigma})$. There $\mathcal{F}=F_0\mathcal{G}$, without evaluating an infinite c series in the critical strip. Initial absolute convergence also proves (26) by the local identity

$$(1-q)(1+q-q\alpha-q\beta)=(1-q\alpha)(1-q\beta)-q^2(1-\alpha)(1-\beta).$$

The correction product is holomorphic on

$$\sigma+\Re z>0,\quad \sigma+\Re w>0,\quad 2\sigma+\min(\Re z,0)+\min(\Re w,0)>1.$$

On margins $\sigma+\Re z,\sigma+\Re w\ge\eta>0$ and last exponent $\ge 1+\epsilon$, its local deviation from 1 is at most $4(1-2^{-\eta})^{-2}(Np)^{-1-\epsilon}$. This proves normal convergence and a phase/deletion-uniform vertical upper bound. Nonvanishing of the correction is not asserted.

For $1/2<\sigma<1$, positive lines $c,d>1-\sigma$, and $V=\lfloor U\rfloor+1/2$, define

$$\mathcal{I}_{T_1,T_2}=\frac{1}{(2\pi i)^2}\int_{c-iT_1}^{c+iT_1}\int_{d-iT_2}^{d+iT_2}\mathcal{F}(s,z,w)\frac{V^{z+w}}{zw}\,dw\,dz.$$

The half-integer V includes the actual equality cutoff, since $Na\le U$ is exactly $Na<V$. Put $\Delta_V=\inf_{n\in\mathbb{Z}_{\ge1}}|\log(V/n)|\ge 1/(2V+1)$, and let $\mathcal{A}=\sum_{a,b}|\mu(a)\mu(b)\prod_{p\mid ab}t_p|(Na)^{-c}(Nb)^{-d}$ with the same disjoint support. For $c,T>0$ and $v>0$, $v\ne1$, define

$$K_{c,T}(v)=\frac{1}{2\pi i}\int_{c-iT}^{c+iT}\frac{v^z}{z}\,dz.$$

Finite-height termwise integration of the absolutely convergent $\mathcal{G}$ and the one-variable tail estimate $|K_{c,T}(v)-1_{v>1}|\le 2v^c/(\pi T|\log v|)$ give

$$|\mathcal{I}_{T_1,T_2}-J_U|\le|F_0|\mathcal{A}V^{c+d}\left\{\frac{2}{\pi T_1\Delta_V}+\frac{2}{\pi T_2\Delta_V}+\frac{4}{\pi^2T_1T_2\Delta_V^2}\right\}. \tag{27}$$

Thus $\mathcal{I}_{T_1,T_2}\to J_U$ as both heights tend independently to infinity, away from the original principal pole. This is an oscillatory limit, not an absolutely integrable infinite vertical integral.

Taking $T_1=T_2=V^2$ gives the safe upper bound $|J_U|\ll(Nu)^\eta|L(s,\psi)|V^{c+d}\log^2(2V)$, with constants retaining fixed margins, ν,S, and σ. Choosing c,d close to $1-\sigma$ reproduces $U^{2-2\sigma+\eta}$ in amplitude and $U^{2-4\delta+\eta}$ in energy on $\sigma=1/2+\delta$. Insertion into the original common-profile Mellin argument retains $M_4(W)^2$ and all previously required height/conductor uniformities. No derivative-free improvement follows from this Perron estimate.

The correction converges even on some negative auxiliary lines, but moving there can cross $z=\rho-s$ or $w=\rho-s$, from zeros of the primitive L-function. The exact meromorphic slice $\mathcal{F}(s,z,0)=1/L^{T_0}(s+z,\psi)$ forbids assuming universal cancellation of these poles. At a numerator zero the slice must be understood meromorphically, not by an unjustified pointwise substitution. Sequential shifts have single boundary integrals and crossed-zero residues as well as the kernel double residue; away from s zeros or poles that double residue is $1/L^{T_0}(s,\psi)$. A principal L pole supplies a reciprocal zero, not a reciprocal pole. Uniform control of actual reciprocal integrals/residues or the signed covariance remains a new input. Numerator AFM alone supplies neither.

22 A quantitative covariance budget without a new exponent

Retain the original Eisenstein sextic symbols, fixed finite-order ν, fixed S, one possibly complex common $W\in C_c^\infty$ supported in $[1,2]$, all squarefree columns prime to S, and every raw shared-prime

zero. In the accepted window $2D/7 \le U^2 \le D$, put $r_m = r_U(m)$, $\lambda_m = \mu(m)\nu(m)W(Nm/D)$ and $v_m = \overline{\lambda_m}$. Thus $C_U = \mu r$ and $A_U = \mu(r-1)$ on the active annulus. For a finite literal row block B, let

$$G_{mn} = \sum_{u\in B} x_m(u)\overline{x_n(u)}, \qquad R = \operatorname{diag}(r_m), \qquad M_j = \sum_m r_m^j |W(Nm/D)|^2 \quad (0 \le j \le 2).$$

The primitive $\psi_u, \psi_{u,2}$ and full deleted sets retain their previous meanings; no conductor-drop value replaces a raw zero.

Proposition 11 (An actual-vector lower budget)**.** *If $0 \le G \le \Lambda_B I$, then*

$$\Re Q_B(C_U, \mu) \ge -\frac{\Lambda_B}{2D}\{\sqrt{M_0 M_2} - M_1\}. \tag{28}$$

In this high window, uniformly in U,

$$\sum_{D\le Nm\le 2D} r_m \ll_K D, \qquad \sum_{D\le Nm\le 2D} r_m^2 \ll_K D\log(2D). \tag{29}$$

Consequently

$$\Re Q_B(C_U, \mu) \ge -C_K \|W\|_\infty^2 \Lambda_B \sqrt{\log(2D)}.$$

The vector inequality is sharp among positive contractions with the specified v, Rv; no such sharpness is asserted for the arithmetic Gram. The moment bound counts actual balanced ideal factorizations, rather than assuming independence of the two row factors.

On B_X, inserting the previously accepted powered-BGL norm gives only

$$\Re Q_{B_X}(C_U, \mu) \ge -C_{\nu,S,\epsilon}\|W\|_\infty^2 (DX)^\epsilon \{X^2 + DX + D^{2/3}X^{5/3}\}.$$

At $X \asymp D^{1/4}$ the exponents remain $1/2, 5/4, 13/12$. At $H = D^{3/4}$ this budget still has the old $D^{1/2}$ excess over H; it supplies no BT_σ saving. A negative-part budget is weaker than $2\Re Q(C_U, \mu) \ge Q(C_U, C_U) + Q(\mu, \mu) - R_\sigma$.

For $S_R = \sqrt{R}$, the exact identity

$$\tfrac{1}{2}(RG + GR) = S_R G S_R + \tfrac{1}{2}[S_R, [S_R, G]]$$

localizes the commutator to unequal factor-count levels, with entry $(\sqrt{r_m} - \sqrt{r_n})^2 G_{mn}$. If B retains all six literal units, the actual unit sum further gives $G_{mn} = 0$ unless $Nm \equiv Nn \pmod{36}$. These are exact deletions, not a power saving. Even zero different-level pairings leave the nonnegative within-level energies $\sum_k (k-1)^2 Q_B(\mu \mathbf{1}_{r=k}, \mu \mathbf{1}_{r=k})$. The unequal-product coupled fourth moment and its mixed Möbius term remain unestimated. No new analytic exponent, energy lower bound, Lean formalization or GRH conclusion follows.

23 Proof and arithmetic boundary of the covariance budget

23.1 The positive-contraction calculation

Since G is the exact raw Gram, it is positive semidefinite. If $\Lambda_B > 0$, put $T = 2G/\Lambda_B - I$, a self-adjoint contraction. Then

$$\Re(v^* R G v) = \frac{\Lambda_B}{2}\{M_1 + \Re(v^* R T v)\} \ge \frac{\Lambda_B}{2}\{M_1 - \|Rv\|_2 \|v\|_2\},$$

which is (28); $\Lambda_B = 0$ is immediate. The Hermitian rank-two matrix $\{(Rv)v^* + v(Rv)^*\}/2$ has possible nonzero eigenvalues $(M_1 \pm \sqrt{M_0 M_2})/2$. Its negative spectral projection realizes the lower bound among all $0 \le G \le \Lambda_B I$, a linear-algebra sharpness statement only. If $a \le r_m \le b$ on the weighted support, $a + b > 0$, the pointwise inequality $r_m^2 \le (a+b)r_m - ab$ also gives

$$\Re Q_B(C_U, \mu) \ge -\frac{\Lambda_B M_0 (b-a)^2}{8D(a+b)}.$$

For $a = b = 0$ the covariance is zero. The constant follows by maximizing $\sqrt{(a+b)t - ab} - t$ on $[a, b]$.

23.2 The balanced-divisor second moment

Set $T_0 = \sqrt{D}$. The high-window factor restrictions imply

$$\sqrt{2/7}\, T_0 < Na, Nb < \sqrt{14}\, T_0.$$

The elementary ideal count $\#\{a : Na \le Y\} \ll_K Y$ proves the first assertion of (29). For the second, two ordered factorizations $ab = cd = m$ have the unique form

$$a = gh, \quad c = gk, \quad b = k\ell, \quad d = h\ell, \qquad g = (a, c).$$

Squarefreeness makes g, h, k, ℓ pairwise coprime and squarefree. To upper-bound this positive count, enlarge to arbitrary ideals. For $M = \max(Nh, Nk)$ the individual bounds are

$$Ng \le \frac{\sqrt{14}\, T_0}{M}, \qquad N\ell \le \frac{\sqrt{14}\, T_0}{M}, \qquad 1 \le M \le \sqrt{14}\, T_0.$$

On $2^j \le M < 2^{j+1}$ there are $O_K(2^{2j})$ choices of (h, k) and $O_K((T_0/2^j)^2)$ choices of (g, ℓ), with bounded terminal ranges absorbed in the constant. Each dyadic range contributes $O_K(D)$; there are $O(\log(2D))$ ranges. This proves the second moment with no extra logarithm from counting only $NgN\ell$. Restricting to primes outside fixed S decreases the count. Hence

$$M_0, M_1 \ll_K \|W\|_\infty^2 D, \qquad M_2 \ll_K \|W\|_\infty^2 D \log(2D).$$

These constants are independent of the cutoff in the stated window. Only the common supremum norm is used, including for complex W.

23.3 What the commutator actually controls

Write

$$t_{mn} = D^{-1}\Re\{\lambda_m \overline{\lambda_n} G_{mn}\}, \qquad \delta_{mn} = (\sqrt{r_m} - \sqrt{r_n})^2.$$

The exact commutator identity proves

$$\Re Q_B(C_U, \mu) \ge Q_B(\sqrt{r}\,\mu, \sqrt{r}\,\mu) - \tfrac{1}{2}\sum_{m,n} \delta_{mn}|t_{mn}|.$$

The displayed loss retains the actual character and profile phases before taking absolute values; it is at most the phase-blind loss obtained from $|t_{mn}| \le D^{-1}|\lambda_m \lambda_n G_{mn}|$. On the accepted four-column smooth B_7 fixture at amplitude $x \ge 0$, the positive term is $18/25$, this coefficient-adapted loss is $9x/50$, and the lower bound is attained. Purely imaginary entries of the full complex Gram make the phase-blind loss larger. This calibrates the new budget and does not give an asymptotic obstruction.

For the unit deletion let $\zeta_6 = 1+\omega$. Every good prime norm q is $1 \pmod 6$, and direct residue-field exponentiation gives $x_p(\zeta_6) = \zeta_6^{(q-1)/6}$. Multiplying over $p \mid m$ and reducing $Nm = \prod_{p|m}(1+6e_p)$ modulo 36 gives

$$x_m(\zeta_6) = \zeta_6^{(Nm-1)/6 \bmod 6}.$$

Multiplication of all literal rows by ζ_6 preserves their raw zeros and proves the norm-congruence deletion in the core. The vector budget may be applied separately on these six classes. No D-dependent number of classes or conductor saving is produced.

The moment proof above counts the coefficient diagonal $ab = cd$. It has not estimated the unequal-product character quadruples in the coupled row fourth moment. Replacing the inherited DX norm by $D^{1-\sigma}X$ merely because equal-level and different-class entries vanish is the unsupported step. Small negative commutator loss alone also leaves the within-level energies in the core. Principal residues are unchanged; this covariance is used on the nonprincipal blocks where the original residue is zero, not as a new centering prescription.

23.4 The complete near-saturating face

For a six-factor row block let $R_0 = \prod_{j=1}^{6} A_j$, $M = \max(A_1, A_2)$ and $\prod_j A_j^j \le H$. The accepted exact-mask block norm is

$$\Lambda_A \ll_{\nu,S,\epsilon} (HD)^{\epsilon}\{R_0 + DR_0/M + D^{2/3}R_0/M^{1/3}\}.$$

The column moment and covariance proofs apply to every such block. Write $x_j = \log A_j$, $h = \log H$, $c_0 = \sum_j jx_j$ and $r_0 = \sum_j x_j$. Then, for $x_1 \ge x_2$,

$$h - 3(r_0 - \max(x_1, x_2)) = h - c_0 + (x_1 - x_2) + x_4 + 2x_5 + 3x_6;$$

for $x_2 \ge x_1$ replace $x_1 - x_2$ by $2(x_2 - x_1)$. All terms are nonnegative. Thus the middle monomial is saturated on

$$A_1 = A_2, \qquad A_1A_3 = H^{1/3}, \qquad A_4 = A_5 = A_6 = 1,$$

with full capacity. The balanced $A_3 = 1$ case is only an endpoint. The a_3 direction remains nonprincipal by its outside-S valuation $3 \pmod 6$. If $R_0/M \ge H^{1/3}D^{-\eta}$, the total displayed slack is at most $3\eta \log D$. Therefore a fixed power gain must be uniform on this full neighboring region. The proved budget still costs DR_0/M; it gives no global improvement over the accepted envelope.

The independent standard-library diagnostic checks nine ideal moment cases, eight raw Gram cases with 27 000 phase values and 5 184 raw zeros, exact radical commutator/vector budgets, and 227 exponent cases. Finite tables use common norm profiles realizable by smooth disjoint bumps; they are not tests of an infinite analytic estimate. Full derivation, source bindings and preserved checks are under `planning/covariance-budget-20261105/attempt/` and `outputs/checks/2026-11-05-covariance-budget/attempt/`.

24 The fourth-moment conductor and its full raw zero set

Retain the original Eisenstein sextic symbols $x_m(u) = (u/m)_6$, fixed finite-order ν, fixed S, and one possibly complex common smooth W supported in $[1,2]$. Put $\lambda_m = \mu(m)\nu(m)W(Nm/D)$ and $r_m = r_U(m)$. Columns remain squarefree and prime to S. In $2D/7 \le U^2 \le D$, the actual polynomial is

$$\mathcal{C}_U(u) = \sum_{\substack{(a,b)=1,\ Na,Nb>U \\ D \le Na\,Nb \le 2D}} \mu(a)\mu(b)\nu(ab)W(Nab/D)x_a(u)x_b(u).$$

Its squared row energy has four factor indices a, b, c, d, with $(a, b) = (c, d) = 1$. These internal masks are essential. At each prime,

$$e_p = \mathbf{1}_{p|a} + \mathbf{1}_{p|b} - \mathbf{1}_{p|c} - \mathbf{1}_{p|d} \quad \text{belongs to } \{0, \pm 1\}.$$

Thus, for the two squarefree products $m = ab, n = cd$, the primitive oscillatory conductor outside fixed S is $q = mn/(m, n)^2$. The full raw zero-preserving period is $L = mn/(m, n)$:

$$x_m(u)\overline{x_n(u)} = \mathbf{1}_{(u,(m,n))=1}\theta_{m/(m,n),\, n/(m,n)}(u).$$

No common-prime zero is restored by taking the primitive character.

A generic fourth moment $|\sum_m \lambda_m x_m(u)|^4$ has four individually squarefree columns $m_1, \dots, m_4$ with no such internal coprimality. Its signed local exponent $e_p = \sum_{i=1}^{2} \mathbf{1}_{p|m_i} - \sum_{i=3}^{4} \mathbf{1}_{p|m_i}$ can also be ± 2, giving a nontrivial cubic component. For $L = \mathrm{rad}(m_1 m_2 m_3 m_4)$, the exact row-power interface is

$$Q_j = \prod_{\substack{p|L \\ je_p \not\equiv 0 \ (6)}} p, \qquad Z_j = L/Q_j, \qquad \prod_{i=1}^{2} x_{m_i}(v^j)\overline{\prod_{i=3}^{4} x_{m_i}(v^j)} = \theta_{j,e}(v)\mathbf{1}_{(v,Z_j)=1}.$$

Here $\theta_{j,e}$ is primitive on Q_j, with its own zero extension. At $j = 3$ an exponent ± 2 becomes primitive-principal but its raw zero remains. For the original two-squarefree-product expansion, $Q_j = q$ for all $1 \le j \le 5$, including the quadratic $j = 3$ direction. For the accepted row blocks $Nr_j \asymp A_j$, put $R_0 = \prod_j A_j$, $P = \max(A_1, \dots, A_5)$ and retain $\prod_j A_j^j \le H$. The full A_3 saturation face and all row-factor masks remain.

The small-conductor pair count recovers an already controlled high-gcd sector; it is not a new analytic advance. On the remaining original sector $Nq > P^2$, favorable smooth completion costs $\min(A_j, \sqrt{Nq}) = A_j$ for $1 \le j \le 5$. The literal sharp cutoff is not replaced by this smooth trial. The signed large-conductor fourth-moment and mixed Möbius sums remain open; no improved full exponent, BT_σ, new Lean theorem, zero-location result or GRH conclusion follows.

25 Scope of the raw fourth-moment reduction

Write $c_{a,b} = \mu(a)\mu(b)\nu(ab)W(Nab/D)$ on the exact factor support in the core. Before absolute Cauchy, the normalized energy is

$$\frac{1}{D} \sum_{(a,b),(c,d) \text{ on that support}} c_{a,b}\overline{c_{c,d}} \sum_{u \in B} x_a(u)x_b(u)\overline{x_c(u)x_d(u)}.$$

Squarefreeness of ab and cd forces each positive and negative prime multiplicity to be at most one, proving $e_p \in \{0, \pm 1\}$. For the mixed term against the actual μ polynomial the exponent is $\mathbf{1}_{p|a} + \mathbf{1}_{p|b} - \mathbf{1}_{p|n}$, with the same conclusion and the same full union mask. The common possibly complex W, ν, hyperbolic support and all raw zeros stay in these identities.

In contrast, unrestricted four-column memberships $(1, 1, 0, 0)$ give $e_p = 2$, a cubic component; $(1, 1, 1, 0)$ give $e_p = 1$ although both product supports contain p. Neither the parity symmetric difference of four supports nor the symmetric difference of two nonsquarefree product supports gives the general conductor. For the interface in the core, its local character order is

j	$1, 5$	$2, 4$	3	6
$e_p = \pm 1$	6	3	2	1
$e_p = \pm 2$	3	3	1	1

Nontrivial entries have tame conductor exponent one; trivial entries remain deleted at $p \mid v$. Multiplying the interface over all six row factors, with the fixed unit/S phase, preserves every original zero, including those at sixth-power factors. All inter-factor, squarefree, ray and capacity restrictions remain external restrictions on the same row sum; this identity asserts no independence.

At a residue field of size Q, let χ have order six and ρ be nontrivial additive. The zero-extended complete sum $G_k(t) = \sum_{z \neq 0} \chi(z)^k \rho(tz)$ satisfies

	$t = 0$	$t \neq 0$
$k \not\equiv 0$ (6)	0	$\lvert G_k(t)\rvert = \sqrt{Q}$
$k \equiv 0$ (6)	$Q - 1$	-1.

For $t \neq 0$, substitution gives $G_k(t) = \overline{\chi(t)^k} G_k(1)$; for nontrivial k, expanding the norm square gives $(Q-1) - \sum_{a \neq 0,1} \chi(a)^k = Q$. The proof applies at split and inert primes. By CRT the full period-L transform vanishes if any local frequency on Q_j is zero; otherwise its magnitude is

$$(NQ_j)^{1/2} \prod_{p \mid Z_j} \left| -1 + Np\, \mathbf{1}_{\text{local frequency at } p \text{ is } 0} \right|.$$

Canceled common primes thus supply Ramanujan factors.

Exact deletion, with all remaining restrictions kept in F, gives

$$\sum_r \theta(r) F(r) \mathbf{1}_{(r,Z)=1} = \sum_{d \mid Z} \mu(d) \theta(d) \sum_z \theta(z) F(dz).$$

For a squarefree pivot let E be the radical of its fixed excluded primes, including Z and the relevant other row factors, with conductor primes omitted since θ already vanishes there. The additional squarefree identity is

$$\mu(r)^2 \mathbf{1}_{(r,E)=1} = \sum_{\substack{h^2 \mid r \\ (h,E)=1}} \mu(h) \sum_{\substack{d \mid E \\ d \mid r}} \mu(d).$$

Since $(d,h) = 1$, its lifted row sum has terms $\mu(d)\mu(h)\theta(d)\theta(h)^2 \sum_z \theta(z) F(dh^2 z)$. The scaled F retains the exact window, primary/ray restrictions and other masks. Divisor triangles cost at most $(HD)^\epsilon$; they do not restore a forbidden nonzero.

The previous fixed-radial-profile high-gcd proof (`research/short-row-high-gcd.md`, independently reviewed in `planning/reviews/book/short-row-formal.md`) already bounds its off-diagonal tail by $O_{\Phi,\epsilon}(D^{2+\epsilon}/G_0^2)$ for $Nh \geq G_0$. For annular $m = hr, n = ht$, $Nq \leq Y$ implies $Nh \geq D/\sqrt{Y}$, so the old smooth-scope bound is $O_{\Phi,\epsilon}(D^\epsilon Y)$. The present elementary count $O(D\sqrt{Y})$, factor multiplicities $r_U(m) r_U(n) \ll_\epsilon D^\epsilon$, and $\#B_A = O(R_0)$ give only $O(\|W\|_\infty^2 D^\epsilon R_0 \sqrt{Y})$ on a literal sharp block. This is a weaker conductor-form recovery in its own bookkeeping scope, not a new high-gcd theorem; the smooth estimate is not silently transferred to a selected sharp off-diagonal piece. The full A_3 face and neighborhood slack are also inherited.

Finally, primitive smooth Poisson completion has bound $O_V(\min(A_j, \sqrt{Nq}))$ on the original support. For $Nq > P^2$ and $1 \leq j \leq 5$, every scaled length $A_j/(Nd\,Nh^2)$ is below $\sqrt{Nq}$, so even this favorable trial takes its trivial length term. Its constant uses fixed smooth seminorms and need not stay bounded when a sharp boundary is approached. In the admissible subsector $(m,n) = 1$, $Nq = Nm\,Nn \asymp D^2$. A complete zero mean does not control the nonzero frequencies of the short original masked row window. The unresolved target is their joint signed sum with weights $r_U(m) r_U(n) \lambda_m \overline{\lambda_n}$ and $r_U(m) \lambda_m \overline{\lambda_n}$, before per-conductor absolute estimates. No new global analytic conclusion is claimed.

26 A quadratic pivot on the cube factor

Retain the original zero-extended sextic symbols $\chi_m(u) = (u/m)_6$ over $K = \mathbb{Q}(\omega)$, the fixed S containing all primes above 6 and comparison moduli, and the squarefree good columns $D \le Nm \le 2D$. Write an outside-S row ideal uniquely as $r_1 r_2^2 r_3^3 r_4^4 r_5^5 r_6^6$: only $r_1, \dots, r_5$ are pairwise-coprime squarefree ideals; r_6 is arbitrary and may overlap them. Use primary generators. Fix one literal unit and one S-supported row factor, retaining their original values. Let $A_j \le Nr_j < 2A_j$, $A_j \ge 1$, $R_0 = \prod_j A_j$ and $\prod_j A_j^j \le H$; the actual capacity restriction on rows is retained. Coefficients $b_m \in \mathbb{C}$ are arbitrary and fixed independently of every variable row.

Proposition 12 (The cube-factor block norm)**.** *For every $\varepsilon > 0$, the raw Gram on this actual block obeys*

$$\sum_{u \text{ in the block}} \left| \sum_m b_m \chi_m(u) \right|^2 \ll_{\nu,S,\varepsilon} (HD)^{\varepsilon} \left\{ R_0 + \frac{DR_0}{A_3} \right\} \sum_m |b_m|^2.$$

The constant is uniform in the fixed unit/S-factor, frozen good factors, H, D and the coefficient sequence. There is no rational-norm-squarefree restriction. Consequently, with $M = \max(A_1, A_2)$, the old powered-BGL bound and this bound give

$$\Lambda_A \ll_{\nu,S,\varepsilon} (HD)^{\varepsilon} \left\{ R_0 + \min\left(\frac{DR_0}{M} + \frac{D^{2/3} R_0}{M^{1/3}}, \frac{DR_0}{A_3} \right) \right\}. \tag{30}$$

The external analytic input is Goldmakher–Louvel, arXiv:1112.1642v2, Corollary 1.2 with $n = 6$, already an explicit classical cube-power sieve. The refinement is this application, not a new large-sieve theorem. For the original normalized coefficients $b_m = D^{-1/2} A_U(m) \nu(m) W(Nm/D)$, the accepted divisor bound gives $\sum_m |b_m|^2 \ll_\eta \|W\|_\infty^2 D^\eta$, uniformly in $U \ge 1$ with $U^2 \le D$ (and the common W may be complex). Choosing smaller internal exponents yields the corresponding uniform block energy envelope. General cutoff or factor-count weights remain inside the arbitrary coefficient norm; that norm is not silently replaced by one. Summing S-supported row factors uses the reduced capacity $H/N\mathfrak{b}$ and the established convergent geometric sums.

On the complete old saturated face $A_1 = A_2 = A$, $A_3 = B$, $A_4 = A_5 = A_6 = 1$, $AB = H^{1/3}$, the new bound is $A^2 B + DA^2$ in place of the old middle monomial DAB. For $H = D^{3/4}$ put $A = D^a$, $B = D^{1/4-a}$, $0 \le a \le 1/4$. When $a < 1/8$, its leading exponent is $1 + 2a < 5/4$; the gain requires B/A to grow by a power for a power improvement. For $a \ge 1/8$, the combined bound still has leading exponent $5/4$. Thus the A_3-dominant part improves, while no new global exponent, BT_σ estimate, signed covariance theorem, Lean theorem or GRH conclusion follows.

Inspecting the surviving sector. The quadratic-pivot inspector plots the complete face, the five monomial powers and the remaining target gap. These are analytic bound exponents, not measured operator norms or trained-network predictions. The refined norm can be inserted in the earlier covariance budget; it does not remove its within-level energies.

A necessary near-face condition. Put $\mathcal{M}_A = DR_0/\max(M, A_3)$. The combined envelope is at most $(HD)^{\varepsilon}\{R_0 + \mathcal{M}_A + D^{2/3} R_0/M^{1/3}\}$. If $\mathcal{M}_A \ge DH^{1/3} D^{-\eta}$, then the earlier near-face slack is at most $3\eta \log D$ and, additionally, $A_3 \le MD^\eta$. Thus a uniform remaining leading-term estimate may focus on this smaller neighborhood; it must still cover the entire $B \le A$ part of the face. This is a necessary condition on explicit upper-bound monomials, not a lower bound for actual arithmetic energy.

27 Proof and scope of the quadratic pivot

Freeze all factors except r_3, and write the actual literal row $u = Fr_3^3$. Multiplicativity of the original numerator gives the exact zero-preserving identity

$$\chi_m(u) = \chi_m(F)\chi_m(r_3)^3.$$

Put $c_m = b_m\chi_m(F)$, a coefficient sequence independent of r_3 with $\sum |c_m|^2 \le \sum |b_m|^2$. In particular the arbitrary sixth-power factor is the coefficient mask

$$\chi_m(r_6^6) = \mathbf{1}_{(m,r_6)=1},$$

not the constant one. No condition $(r_3, r_6) = 1$ is introduced. The actual capacity and coprimality of r_3 with r_1, r_2, r_4, r_5 only restrict this positive outer row sum; after freezing, it may be enlarged to all squarefree good pivot ideals in the dyad. The variable collision mask (m, r_3) remains in the original symbol.

Apply the previously proved literal comparison (22) to the primary generator of r_3 itself:

$$\chi_m(r_3)^3 = \eta_{r_3}(m)^3 C([r_3], [m])^3 \vartheta_m(r_3)^3.$$

The twist η_{r_3} ranges over a fixed finite set of S-unit Kummer characters, and $[r_3]$ over the fixed finite reciprocity group. Partition pivot rows by $(\eta_{r_3}, [r_3])$. On each part the first two factors are one fixed modulus-one column multiplier. Thus its coefficient sequence is independent of the pivot row and has squared norm at most $\sum |b_m|^2$. This finite partition, rather than an assertion that arbitrary row-dependent phases may be absorbed, resolves the fixed-S phase issue.

Goldmakher–Louvel Definition 1 and Remark 3 identify the cubic powers of a sixth-order Hecke family as a quadratic Hecke family; their Corollary 1.2 gives directly

$$\sum_{\substack{a \text{ sf,good} \\ Na \le 2A_3}} \left| \sum_{\substack{m \text{ sf,good} \\ Nm \le 2D}} z_m \vartheta_m(a)^3 \right|^2 \ll_{K,S,\eta} (A_3 D)^\eta (A_3 + D) \sum_m |z_m|^2.$$

This is the required column-indexed orientation and allows arbitrary complex coefficients and squarefree ideals of nonsquarefree rational norm. The canonical family has primitive trivial-infinity-type characters, a fixed finite reciprocity group, and the same-class primitive-product property required by Definition 1. At each $p \mid m$ outside S, the local tame sextic character has order six, so its cube has order two and conductor exponent one. Thus its outside-S primitive conductor is exactly m; the S-part has only finitely many fixed conductor types. Accordingly both sides of the comparison are zero at $(m, a) > 1$. No conductor-drop value replaces the raw zero, and no global formula for the full conductor is required. The fixed ν stays in coefficients.

Enlarge the positive row sum separately on each finite part and apply this sieve. Ideal counting gives $O_K(A_j)$ possibilities for each frozen r_j, including arbitrary r_6, hence $O_K(R_0/A_3)$ frozen tuples. Summing gives $(A_3D)^\eta(R_0 + DR_0/A_3)\sum |b_m|^2$. Since $A_3 \le H$, adjust η to obtain Proposition 12; literal unit multiplicity is at most six. Together with the prior block bound this proves (30). It does not change the actual coefficient normalization or principal residues.

On the face in the core, the old bound has exponents $a + 1/4$, $5/4$, $11/12 + 2a/3$, whereas the quadratic bound has $a + 1/4$ and $1 + 2a$. Its improvement for $a < 1/8$ leaves the balanced point and $a \ge 1/8$ with the old $5/4$ leading exponent. Even the improved endpoint $a = 0$ costs D, above the target $D^{3/4}$. The pinned GL PDF has SHA256 `5cc5f03388d3996b4333f07be51f989f`

60779fe172f8bd24c7aa6ddad64dd304. This proof invokes that classical result; finite exponent checks do not reprove it, and no new formalization is claimed.

For the last condition, use $R_0/M \le H^{1/3}$ from the exact capacity inequalities. The hypothesis implies $R_0/A_3 \ge H^{1/3}D^{-\eta}$, so $A_3/M = (R_0/M)/(R_0/A_3) \le D^{\eta}$. It also implies $R_0/M \ge H^{1/3}D^{-\eta}$, which is exactly the previous slack hypothesis. Finally $\min(a+b,c) \le \min(a,c)+b$ for nonnegative a,b,c proves the separated envelope, with no replacement of a signed moment by a positive lower bound.

28 Joint completion before conductor absolute values

In the original C–C expansion write $m = hr, n = ht$ with h,r,t pairwise coprime squarefree, $q = rt \ne 1$, $L = hq$, $Q = Nq$, $\ell = NL$. Set $\theta = x_r\overline{x_t}$ and $\xi_q = x_q^3$. The original internal masks give $\theta^3 = \xi_q$ and primitive θ^j on q for $1 \le j \le 5$, while $\mu(hr)\mu(ht) = \mu(q)$. Fix generators α of h, β of q, choosing $\alpha = 1$ for $h = 1$. With $\delta = 1 + 2\omega$, put $e_\beta(z) = \exp(2\pi i \operatorname{Tr}(z/(\delta\beta)))$, $\tau_\beta(\eta) = \sum_{z \bmod q} \eta(z)e_\beta(z)$, and $c_h(k) = \prod_{p\mid h}(Np\mathbf{1}_{k\in p} - 1)$, $c_1 = 1$. For $f_j(v) = \theta(v)^j\mathbf{1}_{(v,h)=1}$, $1 \le j \le 5$, $f_6(v) = \mathbf{1}_{(v,L)=1}$ and $G_j(k) = \sum_{v \bmod L} f_j(v)e_{\alpha\beta}(kv)$, joint completion gives

$$\prod_{j=1}^{6} G_j(k_j) = Q^2\xi_q(-\alpha)\tau_\beta(\xi_q)\overline{\theta\Big(\prod_{j=1}^{5} k_j^j\Big)} \prod_{j=1}^{5} c_h(k_j)c_L(k_6). \tag{31}$$

This holds when the first five frequencies are q-units; otherwise the left side is zero. The mixed Gauss phase is orientation-independent, with every common-prime zero retained. The exact row Fourier weight still contains all squarefree, ray, inter-factor, capacity and profile restrictions, including the full A_3 face. Its weighted joint conductor/frequency estimate remains open. This identity and its finite Parseval budget give no improved global exponent or GRH conclusion.

29 The surviving joint identity and its missing estimate

CRT gives $G_j(k) = \theta(\alpha)^j\tau_\beta(\theta^j)\overline{\theta(k)^j}c_h(k)$ for q-unit k, and $G_6 = c_L$. Pairing Gauss powers $1,5$ and $2,4$ yields $\prod_{j=1}^5 \tau_\beta(\theta^j) = Q^2\xi_q(-1)\tau_\beta(\xi_q)$; the CRT factors multiply to $\theta(\alpha)^{15} = \xi_q(\alpha)$. This proves (31), with all additive-generator phases explicit.

On a fixed literal class write $u = u_0\prod_{j=1}^6 v_j^j$ and let $F(v)$ be its exact finite row weight. Only $v_1,\dots,v_5$ are squarefree and pairwise coprime; v_6 is arbitrary with its permitted overlaps. Retain every original mask and the sharp capacity boundary in F. Define $\widehat{F}_L(k) = \sum_v F(v)e_{\alpha\beta}(-\sum_j k_jv_j)$ and, with $t = q/r$, put

$$\begin{aligned} a(r) &= r_U(hr)r_U(ht)\nu(r)\overline{\nu(t)} \\ &\qquad \times W(Nh\,Nr/D)\overline{W(Nh\,Nt/D)}\theta_{r,t}(u_0), \\ A(K) &= \sum_{r\mid q} a(r)\overline{\theta_{r,q/r}(K)}, \\ B(K) &= \sum_{\substack{k \bmod L,\ (k_1\cdots k_5,q)=1 \\ \prod_{j=1}^5 k_j^j = K \bmod q}} \widehat{F}_L(k)\prod_{j=1}^{5} c_h(k_j)c_L(k_6). \end{aligned}$$

Set $a(r) = 0$ at excluded orientations; the actual factor counts and common possibly complex W

remain. Fourier inversion gives exactly

$$E_{\mathcal{T}} = \frac{1}{D} \sum_{(h,q)\in\mathcal{T}} \mu(q) \frac{Q^2 \xi_q(-\alpha)\tau_\beta(\xi_q)}{\ell^6} \sum_{K\in(\mathcal{O}/q)^\times} B(K)A(K). \tag{32}$$

Sum the original u_0 classes and blocks. This retains low gcd, including $h = 1$, before conductor absolute values. For C–μ replace the second $r_U(ht)$ by 1. A column-orientation-dependent restriction cannot be hidden in a common F when applying the next identity.

Distinct orientations have nontrivial cubic quotient, so

$$\sum_K |A(K)|^2 = \varphi(q) \sum_{r|q} |a(r)|^2, \qquad |E_{h,q}| \le \frac{Q^{5/2}}{D\ell^6} \left(\varphi(q) \sum_{r|q} |a(r)|^2 \right)^{1/2} \|B\|_2.$$

The product-frequency fibers have size $\varphi(q)^4$ under uniform averaging. The actual B is not uniform: all orientations agree on the coherent set $x_p(K)^2 = 1$ for every $p \mid q$, of size $\varphi(q)/3^{\omega(q)}$. Centered representatives have $|k_j| \le \sqrt{\ell}$. Restricting to any frequency subset Ω has error bounded by the same displayed budget with $\|B - B_\Omega\|_2$; no small lower-height tail is proved.

The full face $A_1 = A_2 = A$, $AA_3 = H^{1/3}$, $A_4 = A_5 = A_6 = 1$ remains. At its pure A_3 endpoint, inversion in the fixed factors leaves $(\sum_r a(r)) \sum_v F_3(v)\xi_q(v)\mathbf{1}_{(v,h)=1}$: the actual row character there is orientation-independent. Literal length-one Fourier factors can undo complete phase balance. A quantitative signed estimate for the actual h, q, K sum is still required. Exact split/inert Gauss, full-mask CRT and orientation tests pass 991 assertions; receipt: `outputs/checks/2026-11-05-joint-dual/research/joint-gauss-check-v1.json`. These diagnostics do not prove an incomplete-moment saving or GRH.

30 A radical-conductor refinement of the raw envelope

Retain the original literal symbols $\chi_m(u) = (u/m)_6$, extended by zero, the fixed good set S, and the squarefree ideal columns $D \le Nm \le 2D$. On an actual dyadic row block write

$$u = u_0 r_1 r_2^2 r_3^3 r_4^4 r_5^5 r_6^6, \qquad A_j \le N r_j < 2A_j, \qquad \prod_{j=1}^{6} A_j^j \le H.$$

Only $r_1, \dots, r_5$ are pairwise-coprime squarefree ideals; r_6 is arbitrary and may overlap them. Fix primary generators, a literal unit and an S-supported factor in u_0. Put $R = \prod_j A_j$, $M = \max(A_1, A_2)$ and $Q = R/A_6$. For arbitrary complex coefficients b_m, independent of every variable row, define

$$\mathcal{E}_A(b) = \sum_{u \text{ in the actual block}} \left| \sum_m b_m \chi_m(u) \right|^2, \qquad \Lambda_A = \sup_{b\ne 0} \frac{\mathcal{E}_A(b)}{\sum_m |b_m|^2}.$$

Proposition 13 (An ordinary conductor comparison)**.** *For every $\varepsilon > 0$,*

$$\Lambda_A \ll_{S,\varepsilon} (HD)^\varepsilon A_6 \{D + Q^2\}. \tag{33}$$

Consequently the original residue-subtracted remainder satisfies, uniformly for $H, D, U \ge 1$ and $U^2 \le D$,

$$\sum_{0<Nu\le H} |E_u|^2 \ll_{\nu,S,\varepsilon} \|W\|_\infty^2 (HD)^\varepsilon \{H + DH^{1/6} + (DH)^{2/3}\}. \tag{34}$$

Here ν, S are fixed and $W \in C_c^\infty((0,\infty))$ is one common, possibly complex profile supported in $[1,2]$; there is no profile supremum.

The input to (33) is the classical ordinary sieve for *all* primitive characters of conductor norm at most Q, with constant $D+Q^2$, not a higher-order-family theorem for arbitrary conductor labels. The following section gives its standard Gauss-sum and two-dimensional lattice proof, including prime-power conductors at S. For fixed r_6, the first five factors have primitive outside-S conductor $q = r_1 \cdots r_5$. The literal primitive character agrees with the raw symbol even at collisions with q; the sixth-power factor instead remains the coefficient mask $\mathbf{1}_{(m,r_6)=1}$. At most $5^{\omega(q)} \ll_\eta (Nq)^\eta$ tuples give the same primitive character. Applying the ordinary sieve and then counting r_6 proves (33). No rational-norm-squarefree restriction is imposed.

To obtain (34), combine this with the accepted powered-BGL block bound from Section 18. With $S_A = R/M$ and $T_A = D^{2/3}R/M^{1/3}$, capacity gives $R \le H$, $A_6 \le H^{1/6}$, $RS_A \le H$ and $T_A \le (DH)^{2/3}$. Hence

$$\begin{aligned}\min(R + DS_A + T_A, DA_6 + R^2/A_6) &\le R + T_A + DA_6 + \min(DS_A, H^2/S_A^2)\\ &\le H + DH^{1/6} + 2(DH)^{2/3}.\end{aligned}$$

Dyadic and fixed-S summation, the uniform actual coefficient norm, and the principal-residue estimate give the assertion for $E_u = I_u - P_u$.

At $H = D^{3/4}$ the new envelope has exponent $7/6$, improving the old $5/4$ by $1/12$; its excess above the target H is still $D^{5/12}$. On the *entire* old face $A_1 = A_2 = A$, $A_3 = B$, $A_4 = A_5 = A_6 = 1$, $AB = H^{1/3}$, one has $Q = A^2B = AH^{1/3} \le H^{2/3}$, so (33) costs $D^{1+\varepsilon}$ there. This includes $B \le A$. The new combined upper bounds reach $7/6$ at $(\log_D A_1, \log_D A_2) = (5/12, 1/6)$ with other coordinates zero; the pure point $\log_D A_1 = 3/4$ also retains the old BGL cross term. These are upper-envelope witnesses, not arithmetic energy lower bounds or counterexamples to signed cancellation.

The simpler proposed replacement of the radicand ab^2 by its radical ab inside BGL remains unjustified: after the existing fixed reciprocity phases, the mixed reciprocity quotient is $\vartheta_b(a')/\vartheta_a(b')$, which varies outside S. The present application does not require that replacement. No target signed moment, type-II covariance estimate, positive arithmetic determinant comparison, zero-location result, new Lean theorem or GRH conclusion follows from (34).

31 Proof of the ordinary-conductor refinement

31.1 Primitive Gauss sums at every conductor

Let $\mathcal{O} = \mathbb{Z}[\omega]$, $\delta = 1 + 2\omega$, and choose one generator f of each nonzero ideal $\mathfrak{f}$. Write $e_f(x) = \exp(2\pi i \operatorname{Tr}(x/(\delta f)))$, $U(\mathfrak{f}) = (\mathcal{O}/\mathfrak{f})^\times$, and $\varphi(\mathfrak{f}) = |U(\mathfrak{f})|$. This additive character is primitive, since the different is (δ). For a primitive multiplicative character ρ modulo $\mathfrak{f}$, extended by zero at nonunits, set $\tau_f(\rho) = \sum_{x \bmod \mathfrak{f}} \rho(x) e_f(x)$. We verify the full-modulus identity

$$|\tau_f(\rho)|^2 = N\mathfrak{f}, \qquad \sum_{x \bmod \mathfrak{f}} \rho(x) e_f(zx) = \overline{\rho(z)}\tau_f(\rho) \quad (z \in \mathcal{O}). \tag{35}$$

No squarefree-conductor hypothesis is used, so this includes wild prime powers in the fixed S-part.

The Ramanujan sum is

$$c_{\mathfrak{f}}(y) = \sum_{x \in U(\mathfrak{f})} e_f(xy) = \sum_{\mathfrak{d} \mid (\mathfrak{f}, (y))} N\mathfrak{d}\, \mu(\mathfrak{f}/\mathfrak{d}).$$

Indeed insert $\mathbf{1}_{(x,\mathfrak{f})=1} = \sum_{\mathfrak{b}\mid(x,\mathfrak{f})} \mu(\mathfrak{b})$; the complete additive sum with $\mathfrak{b} \mid x$ is $N(\mathfrak{f}/\mathfrak{b})$ when $\mathfrak{f}/\mathfrak{b} \mid (y)$ and zero otherwise. Changing variables $x = ay$ in the Gauss square now gives

$$\begin{aligned}|\tau_f(\rho)|^2 &= \sum_{a\in U(\mathfrak{f})} \rho(a) c_{\mathfrak{f}}(a-1) \\ &= \sum_{\mathfrak{d}\mid\mathfrak{f}} N\mathfrak{d}\,\mu(\mathfrak{f}/\mathfrak{d}) \sum_{\substack{a\in U(\mathfrak{f}) \\ a\equiv 1 \pmod{\mathfrak{d}}}} \rho(a) = N\mathfrak{f}.\end{aligned}$$

For every proper $\mathfrak{d} \mid \mathfrak{f}$, the reduction $U(\mathfrak{f}) \to U(\mathfrak{d})$ is surjective by local lifting and CRT. Primitivity says ρ is nontrivial on its kernel, so that inner sum vanishes. For $\mathfrak{d} = \mathfrak{f}$ it equals one. For unit z, substitution proves the second identity in (35). For all frequencies together, finite additive Parseval gives

$$\sum_{z \bmod \mathfrak{f}} \left|\sum_x \rho(x) e_f(zx)\right|^2 = N\mathfrak{f}\,\varphi(\mathfrak{f}).$$

The unit frequencies already exhaust this quantity by the first identity; every nonunit-frequency sum is therefore zero. This proves the asserted zero extension. For $\mathfrak{f} = \mathcal{O}$ there is one residue class and one character; take $\tau_1 = 1$ and the single frequency zero.

31.2 The classical ordinary sieve in these lattice conventions

For $D, Q \geq 1$, coefficients z_α supported on any subset of $\{\alpha \in \mathcal{O} : |\alpha|^2 \leq C_0 D\}$ satisfy

$$\sum_{N\mathfrak{f}\leq Q} \sum_{\rho \text{ primitive mod } \mathfrak{f}} \left|\sum_\alpha z_\alpha \rho(\alpha)\right|^2 \ll_{\mathcal{O},C_0} (D+Q^2)\sum_\alpha |z_\alpha|^2. \tag{36}$$

Here the primitive characters may be restricted to the unit-trivial subfamily. We give the standard proof to make the conductor comparison explicit.

First the reduced fractions x/f, with $x \in U(\mathfrak{f})$ and $N\mathfrak{f} \leq Q$, are distinct modulo $\mathcal{O}$. Their denominator ideal is intrinsic: $\{a \in \mathcal{O} : a(x/f) \in \mathcal{O}\} = \mathfrak{f}$. Thus equality modulo $\mathcal{O}$ forces the same denominator ideal and residue numerator. For different fractions and every $k \in \mathcal{O}$,

$$\left|\frac{x}{f} - \frac{y}{g} - k\right| = \frac{|xg - yf - kfg|}{|fg|} \geq \frac{1}{Q},$$

because the numerator is a nonzero Eisenstein integer and $|fg| = (N\mathfrak{f}N\mathfrak{g})^{1/2} \leq Q$. The fixed codifferent factor $1/\delta$ changes only the lattice constants. The modulus-one frequency zero obeys the same separation from all others.

For completeness, a ρ_0-separated set t_r in a fixed two-dimensional torus obeys the additive bound

$$\sum_r \left|\sum_{|\alpha|\ll\sqrt{D}} z_\alpha \exp(2\pi i \operatorname{Tr}(\alpha t_r))\right|^2 \ll (D + \rho_0^{-2})\sum_\alpha |z_\alpha|^2. \tag{37}$$

Use fixed linear coordinates to write the lattice as $\mathbb{Z}^2$ and its additive pairing as $e(n \cdot t)$. The support lies in $|n_i| \leq T$ with $T \asymp \max(1, \sqrt{D})$. To prove the dual inequality, majorize this box indicator by $e^{2\pi}\exp(-\pi\|n\|^2/T^2)$, expand the square, and apply Gaussian Poisson summation. Up to fixed constants the resulting kernel is

$$T^2 \sum_{\ell\in\mathbb{Z}^2} \exp(-cT^2\|t_r - t_s + \ell\|^2).$$

Take ρ_0 at most a fixed torus constant. The lifted points are $c'\rho_0$-separated, so disjoint disks bound their count in a disk of radius L by $O(1+(L/\rho_0)^2)$. Gaussian annular summation bounds every absolute kernel row sum by $O(T^2+\rho_0^{-2})$. Schur's inequality and finite-dimensional duality prove (37), with constants independent of the number of frequencies.

Put $\mathcal{S}(x/f)=\sum_\alpha z_\alpha e_f(\alpha x)$. By (35),

$$\rho(\alpha)=\tau_f(\overline{\rho})^{-1}\sum_{x\in U(\mathfrak{f})}\overline{\rho(x)}e_f(\alpha x).$$

After taking squared absolute values, enlarging the primitive-character sum to all characters of $U(\mathfrak{f})$ and using orthogonality yields

$$\sum_{\rho \bmod \mathfrak{f}\ \text{primitive}}\left|\sum_\alpha z_\alpha\rho(\alpha)\right|^2\le\frac{\varphi(\mathfrak{f})}{N\mathfrak{f}}\sum_{x\in U(\mathfrak{f})}|\mathcal{S}(x/f)|^2\le\sum_{x\in U(\mathfrak{f})}|\mathcal{S}(x/f)|^2.$$

Apply (37) with $\rho_0\gg Q^{-1}$ to prove (36).

Finally $\mathcal{O}$ has class number one. Choose one generator α_m of each ideal column, with $|\alpha_m|^2=Nm$; these are an injective subset of the lattice disk. A primitive trivial-infinity-type Hecke character is a unit-trivial primitive residue character on $\mathcal{O}/\mathfrak{f}$, and its ideal value equals its value on α_m. Descent of the residue character to a proper divisor would preserve unit-triviality and contradict primitive Hecke conductor. Thus (36) applies to arbitrary ideal coefficients of norm at most $2D$, including squarefree ideals of nonsquarefree rational norm.

31.3 Raw zeros, tuple multiplicity and block transfer

Freeze r_6 and the literal unit/S-factor, putting

$$c_m=b_m\chi_m(u_0)\mathbf{1}_{(m,r_6)=1},\qquad\sum_m|c_m|^2\le\sum_m|b_m|^2.$$

This is independent of every first-five variable. In particular the arbitrary overlapping sixth-power factor is not replaced by one. Let $v=\prod_{j=1}^5 r_j^j$ denote the chosen literal product. Every outside-S valuation is between one and five. BGL, arXiv:1112.1650v1, Lemma 2.1, together with the established finite S-unit comparison in (22), gives its primitive literal Kummer character κ_v of trivial infinity type, with

$$q=r_1\cdots r_5\mid\operatorname{cond}(\kappa_v)\mid\mathfrak{c}_S q,\qquad\kappa_v(m)=\chi_m(v)\quad((m,S)=1).$$

The equality includes collisions with q: both sides are zero. The fixed $\mathfrak{c}_S$ includes all wild conductor types and literal comparison twists. No row-dependent phase enters the coefficients.

A fixed primitive character determines its outside-S conductor q. Each prime of q has at most five assignments to the pairwise-coprime $r_1,\dots,r_5$, so tuple multiplicity is at most $5^{\omega(q)}\ll_\eta(Nq)^\eta$. All conductors in the dyadic block have norm at most C_SQ. Enlarge the positive row sum, apply (36), and count $O(A_6)$ possibilities for the frozen arbitrary r_6. This proves (33); the actual capacity, squarefree, ray and inter-factor restrictions only shrink the positive sum. Equivalently $\mathcal{E}_A(b)\le\Lambda_A\sum_m|b_m|^2$; there is no coefficient norm factor in the definition of Λ_A.

31.4 Uniform coefficient and principal-residue assembly

The other block estimate is exactly the accepted powered-BGL input (21), whose squarefree powers are justified in BGL Section 3.2. The capacity comparisons and elementary minimum in Section 30 give

$$\Lambda_A \ll (HD)^{\varepsilon}\{H + DH^{1/6} + (DH)^{2/3}\}.$$

Sum logarithmically many blocks with smaller internal exponents. An S-supported literal factor $\mathfrak{b}$ reduces capacity to $H/N\mathfrak{b}$; the required sums over $\mathfrak{b}$ converge at powers 1, 1/6, 2/3. Literal units have multiplicity at most six. This proves the raw bound with arbitrary fixed b_m.

For the actual coefficients use (19): $|A_U(m)| \le 3^{\omega(m)} \ll_\eta (Nm)^\eta$ gives

$$b_m = D^{-1/2}A_U(m)\nu(m)W(Nm/D), \qquad \sum_m |b_m|^2 \ll_\eta \|W\|_\infty^2 D^\eta,$$

uniformly for $U \ge 1$, $U^2 \le D$. The common W is not optimized separately for different rows. If the primitive ψ_u is principal, ν is unramified outside S, so tame Kummer inertia forces every outside-S valuation of u to be divisible by six. Ideal counting and the convergent S-supported sum at power 1/6 give $O_{\nu,S}(H^{1/6})$ such literal rows. In the exact formula (15), $|\rho_T| \ll_K 1$, $|\widehat{W}(1)| \le \|W\|_\infty$, and

$$\sum_d \frac{|\beta_U(d)|}{Nd} \le \left(\sum_{Na\le U}\frac{1}{Na}\right)^2 \ll \log^2(2U).$$

Thus $|P_u| \ll \sqrt{D}\|W\|_\infty \log^2(2U)$ on principal rows, and

$$\sum_{0<Nu\le H} |P_u|^2 \ll DH^{1/6}\|W\|_\infty^2 \log^4(2U).$$

This fits the new envelope. The exact identity $E_u = I_u - P_u$ and $|E_u|^2 \le 2|I_u|^2 + 2|P_u|^2$, with smaller internal exponents, prove (34). No cutoff limit or profile supremum occurs.

The ordinary comparison is an application refinement. A direct BGL radical relabelling would additionally require reciprocity of its mixed family, not just its conductor count: for pairwise-coprime a, b, a', b', the quotient after the original fixed ray phases is $\vartheta_b(a')/\vartheta_a(b')$, which varies outside S. GL, arXiv:1112.1642v2, Definition 1(2), is therefore not supplied by that shortcut. This does not rule out a different mixed-family theorem; it is simply unneeded for the proved weaker ordinary comparison. The remaining exponent gap and every independent signed, spectral, formal and GRH obligation stated in Section 30 remain open.

32 Learning the actual signed arithmetic correlations

The unsigned refinement in Section 30 leaves a signed problem. Here is a separate, executed finite learning objective on the original correlations, with a certificate for every coordinate and observable. It is not a uniform estimate for an expanding family. For ordered distinct squarefree columns $m = hr, n = ht$, put $q = rt$ and

$$\Gamma_i = \sum_{u\in\mathcal{R}} \chi_r(u)\overline{\chi_t(u)}\mathbf{1}_{(u,h)=1}, \qquad i = (h, q, r).$$

The symbols vanish at every collision with q; the common-prime mask is retained separately. With principal ν, actual complementary divisor counts $r_U(m)$, and one common complex profile $W_m = W(Nm/D)$, define

$$d_i^{CC} = r_U(m)r_U(n)W_m\overline{W_n}, \qquad d_i^{C\mu} = r_U(m)W_m\overline{W_n}, \qquad \lambda_i = 1 + |d_i^{CC}|^2 + |d_i^{C\mu}|^2.$$

The exact diagonal Δ_X and these offdiagonal coordinates give

$$E_X = \Delta_X + D^{-1} \sum_i \mu(q_i) d_i^X \Gamma_i, \qquad X \in \{CC, C\mu\}. \tag{38}$$

Thus $E_{CC} = \sum_{u \in \mathcal{R}} |C_u|^2$ and $E_{C\mu} = \sum_{u \in \mathcal{R}} C_u \overline{M_u}$ for the actual normalized row sums. Neither the Möbius sign nor a complex conjugate is discarded.

Proposition 14 (A finite supervised certificate)**.** *Let g_i be complex readout coordinates and $L(g) = \sum_i \lambda_i |g_i - \Gamma_i|^2$. For $K_X = \sum_i |d_i^X|^2/\lambda_i$ and $\lambda_{\min} = \min_i \lambda_i$, the substituted readouts satisfy*

$$|\widehat{E}_X - E_X|^2 \le D^{-2} K_X L(g),$$
$$\max_i |g_i - \Gamma_i|^2 \le L(g)/\lambda_{\min}.$$

The executed metric-gradient update $g^{(t+1)} = (g^{(t)} + \Gamma)/2$ gives $L(g^{(t)}) = 4^{-t} L(g^{(0)})$. Therefore every positive rational requested tolerance on these three squared bounds is attained after finitely many updates on this fixed exact instance, with unrestricted resources.

The implementation's resource guard refuses requests beyond its update budget; termination in this proposition is not a practical runtime guarantee. The labels Γ_i are computed exactly before training, so no speedup over exact label evaluation or generalization to unseen conductors is claimed.

The supplied instance uses $D = 6000, U = 60, H = 625$, S above $2, 3$, principal ν, and all 15 good squarefree ideal columns at activated norms $7507, 7519, 8281, 10507$ (respectively $2, 4, 1, 8$ ideals). One explicitly specified common smooth four-bump profile is zero at every other integer norm. Both the 2262 literal rows and the separately labelled 2148 guaranteed nonprincipal subset are evaluated; the latter is not asserted to be the whole nonprincipal family. There are 95 offdiagonal (h, q) blocks, 65 with $h = 1$, and 210 ordered learned coordinates. Every row case retains 4530 raw column zeros. Real and complex profiles give four requested-10^{-12} certificates, using $53, 55, 53, 55$ metric updates respectively: 216 executed updates, with every full rational state saved. A separate 200 executed ordinary coordinate updates are budget-limited and carry no requested-error claim. The standalone replay is linked from the book and project page. No analytic saving, positive determinant, zero-location assertion, upstream quantum training, new Lean declaration or GRH conclusion follows.

33 Finite phase features, gradients and exact verification

For each block let $G_q = (\mathbb{Z}/6\mathbb{Z})^{\omega(q)}$. A literal admissible row determines local phases e_p at every prime of q; rows with a q-collision or a common-prime collision with h are deleted before forming the exact sparse histogram $n(e)$. Orientation r has

$$\theta_r(e) = \zeta_6^{\sum_{p|q} s_p e_p}, \qquad s_p = 1 \ (p \mid r), \quad s_p = -1 \ (p \mid t), \qquad \Gamma_r = \sum_e n(e) \theta_r(e).$$

Different active orientations have distinct exponent vectors modulo 6, so their characters are orthonormal for the uniform measure on G_q. The model and target projection are

$$F_g(e) = \sum_{r \in J_q} g_r \overline{\theta_r(e)}, \qquad F_\Gamma(e) = \sum_{r \in J_q} \Gamma_r \overline{\theta_r(e)}.$$

Indeed F_Γ is the orthogonal projection of $|G_q| n$ onto this character subspace, and $\mathbb{E}_{G_q} F_g \theta_r = g_r$. The positive spectral metric assigns weight λ_r to this orthonormal basis. Its squared error is exactly

$\sum_r \lambda_r |g_r - \Gamma_r|^2$. Thus the run evaluates the full finite feature-projection loss by Parseval; it does not claim to have executed a sampled or stochastic training loss. This projection need not be a nonnegative density. A Hermitian readout matrix need not be positive semidefinite.

Gradient and contraction. Use ζ_6 coordinates $g = a + b\zeta_6$, so $|g|^2 = a^2 + ab + b^2$ and $M = \left(\begin{smallmatrix} 1 & 1/2 \\ 1/2 & 1 \end{smallmatrix}\right)$. For the objective $L/2$, the ordinary real-coordinate gradient in block i is $\lambda_i M(x_i - x_i^*)$. In the block metric $\lambda_i M$, the gradient is $x_i - x_i^*$. A metric step of size $1/2$ therefore halves every complex error and quarters L. Conjugate initial coordinates and this step preserve the Hermitian reverse orientations. The ordinary trajectory treats all ordered coordinates independently; asymmetric weights need not preserve that symmetry. For ordinary steps $\eta = 2/(3\lambda_{\max})$, the Hessian eigenvalues lie between $\lambda_{\min}/2$ and $3\lambda_{\max}/2$, whence

$$L_{t+1} \le \left(1 - \frac{\lambda_{\min}}{3\lambda_{\max}}\right)^2 L_t.$$

This last bound is conservative; the saved 50 ordinary steps per case are not asserted to attain the requested tolerance.

Proof of Proposition 14. Subtract (38) with the same exact diagonal. Weighted Cauchy–Schwarz, applied only to the error, gives

$$\left|\sum_i \mu(q_i) d_i^X (g_i - \Gamma_i)\right|^2 \le \left(\sum_i \frac{|d_i^X|^2}{\lambda_i}\right)\left(\sum_i \lambda_i |g_i - \Gamma_i|^2\right).$$

Here $|\mu(q_i)| = 1$ because q_i is squarefree; the predicted sum itself keeps its signs and phases. Each summand in the second factor is nonnegative, which proves the coordinate bound. The exact affine update gives $g^{(t)} - \Gamma = 2^{-t}(g^{(0)} - \Gamma)$, proving contraction. For finite rational positive λ_i, K_X, L_0 and $\tau > 0$, some integer t makes all three bounds at most τ^2. The zero-error case stops immediately. No varying-family constants enter this finite assertion. □

The activated original family. Let

$$\phi(t) = \begin{cases} \exp(1 - 1/(1-t^2)), & |t| < 1, \\ 0, & |t| \ge 1, \end{cases} \qquad W(y) = \sum_{N \in \{7507, 7519, 8281, 10507\}} w_N \phi(4(6000y - N)).$$

These disjoint bumps are supported in $(1, 2)$ and $\phi(0) = 1$. For the real case $w_N = 1$; for the complex case $w_N = 1 + \zeta_6(4N - 30000)/6000$. At integer norms $W(n/6000) = w_n$ at the four activated norms and is zero elsewhere. Hence all active original squarefree columns are present, rather than an undeclared column submatrix. No exponential evaluation is needed at these exact nodes. The profile is a fixed finite demonstration and does not establish a profile-uniform analytic estimate. The counts r_U are $0, 2, 4, 2$ at the four norms. The guaranteed nonprincipal subset removes precisely those literal rows whose outside-S norm is 1; no classification of the omitted 114 rows as a whole is asserted. Every nonunit outside-S factor has a prime exponent from 1 to 3, since its prime norm is at least 7 and $7^4 > 625$; its nonzero residue modulo 6 gives a nontrivial local character. All six units and all original S-power rows belong to the separately saved full literal case.

The seed 20261105 selects deterministic rational conjugate initial coordinates by SHA256 of the ordered column identifiers. There is no random sampling. Exact $\mathbb{Q}(\zeta_6)$ arithmetic stores every metric and ordinary state. Training accepts only when both observable certificates and the

uniform-coordinate certificate are at most τ^2, preventing a scalar cancellation from masquerading as accurate correlation learning. The real full-row targets are $E_{CC} = 2393/125$ and $E_{C\mu} = 431/50$; the guaranteed-subset targets are $4293/250$ and $771/100$. Complex-profile targets and the complete phase histograms are saved alongside. Each case has 64 nonzero negative-$\mu(q)$ C–C target orientations. Replacing all signs by plus is an explicit failing control, not an estimator.

Reproduction and scope. The complete run is reproducible.

The source is `examples/arithmetic_learning/learn.py`; its README states the exact smooth profile, CLI, resource guard and loss convention. The accepted immutable output is `outputs/experiments/2026-11-05-arithmetic-learning-v2/`. All source/helper hashes, exact input bytes and artifact hashes are bound before and after execution and checked on replay. Existing outputs, nonpositive or binary-float tolerances, and stale source/input/case bytes are refused. The preserved v1 logging failure is not the accepted run. The focused verifier independently reconstructs 31668 local symbols and 450 Gram targets, then checks the full saved metric trajectory. Independent mathematical review and exact code checks are recorded separately; finite arithmetic checks are not a proof of an asymptotic signed estimate. Full exact training labels are supervision, so this is a certified finite reference architecture, without an efficiency, generalization, PSD, zero-location or GRH claim. Analytic and categorical Lean remain unfinished.

34 Why the common profile must stay fixed

The unsigned envelope of Section 30 and the unresolved signed estimate have different profile quantifiers. We give an obstruction to an H-sized bound uniform through the sup norm alone. It uses the *actual* coefficients and nonprincipal rows, rather than replacing the columns by arbitrary weights. It does not contradict an estimate for one fixed smooth profile or one with adequate derivative dependence.

Fix $K = \mathbb{Q}(\omega)$, $\omega^2 + \omega + 1 = 0$, principal $\nu = 1$, and a finite good set S containing the primes above $2, 3$. Write $\zeta = 1 + \omega$, $\delta = 1 + 2\omega$, so ζ has order six and $\delta^2 = -3$. Retain the literal raw symbols and normalization

$$I_u(W) = D^{-1/2} \sum_{m \text{ squarefree } (m,S)=1} A_U(m)\chi_m(u)W(Nm/D), \qquad E_u = I_u - P_u.$$

Here and below squarefree means *ideal* squarefree. For $U^2 = D$ the accepted high-window identity gives $A_U(m) = \mu(m)(r_U(m) - 1)$ on $D < Nm < 2D$, with $r_U(m)$ even and $|A_U(m)| \geq 1$. The multiset of prime-ideal norms of a squarefree ideal is determined by its rational norm. Thus $A_U(m) = A_N$ depends only on $N = Nm$, including divisor-norm multiplicities. Moreover

$$\eta(m) := \chi_m(\zeta) = \zeta^{(Nm-1)/6}, \qquad (m, S) = 1. \tag{39}$$

The primitive character η has conductor $(2)^2(\delta)^3$, of norm 432; its powers $1, 5$ have conductor norm 432, powers $2, 4$ norm 27, and power 3 norm 16. All five powers are nonprincipal. These row conductors do not remove the raw zero at any prime dividing a sixth factor.

Put $\phi(t) = \exp(1 - 1/(1 - t^2))$ for $|t| < 1$ and zero otherwise. For integer D set $w_N = \operatorname{sgn}(A_N)\overline{\zeta^{(N-1)/6}}$ when a good squarefree ideal of norm N exists and $D < N < 2D$, and zero otherwise. Define one common, complex profile

$$W_D(y) = \sum_{D<N<2D} w_N \phi(4(Dy - N)). \tag{40}$$

Its disjoint supports lie strictly inside $(1,2)$, its sup norm is one for large D, and $\|W_D^{(j)}\|_\infty \ll_j D^j$. In particular the endpoints have value zero. This is a profile varying with scale, independent of the row, not a separate choice for each ideal or each row.

Proposition 15 (A moving-profile lower bound)**.** *For every fixed $0<\theta<1$, take $H=D^\theta$, $U=\sqrt{D}$, and $X=H^{1/6}$. As integer $D\to\infty$,*

$$\sum_{0<Nu\le H} |E_u(W_D)|^2 \ge (c_S+o(1))DH^{1/6}, \qquad c_S>0. \tag{41}$$

This lower bound already comes from literal rows $u_v=\zeta\beta_v^6$, where v is an arbitrary good ideal of norm at most X and β_v is a generator. These rows have nonprincipal primitive character η, so $P_{u_v}=0$. Their exact values are

$$E_{u_v}(W_D) = D^{-1/2} \sum_{\substack{m \text{ squarefree, } (m,Sv)=1 \\ D<Nm<2D}} |A_U(m)|. \tag{42}$$

Every collision with v remains deleted, even when the primitive conductor drops to 432. No squarefree condition is imposed on v.

At $H=D^{3/4}$ the lower exponent is $9/8$, whereas the current upper exponent is $7/6$. Thus this does not establish sharpness of $7/6$ or a new upper saving. For $0<\theta<2/3$, the term $DH^{1/6}$ dominates the current unsigned envelope; it is necessary up to subpower losses in this *varying-profile* class, even on nonprincipal rows.

For $k\ge 0$ put $M_k(W)=\sum_{j=0}^k \|(y\partial_y)^j W\|_\infty$. If, for every $\varepsilon>0$, a bound uniform over all common profiles with constants independent of W, D were to assert

$$\sum_{0<Nu\le H} |E_u(W)|^2 \ll_\varepsilon M_k(W)^\alpha (HD)^\varepsilon H, \qquad \alpha\ge 0,$$

then necessarily $k\alpha \ge 1-5\theta/6$. At $\theta=3/4$ this is $k\alpha\ge 3/8$. For a single fixed ε the necessary condition is only $k\alpha+(1+\theta)\varepsilon \ge 1-5\theta/6$. These are necessary cost comparisons, not sufficient estimates. A fixed-W result, a class with bounded seminorms, or the existing $M_4(W)^2$ dependence is not refuted.

The same rows expose a concrete centering error in the joint completion: removing the orientation mean of the actual correlation can delete a nonzero signed contribution on *nonprincipal* rows. The extended proof retains the literal unit phase, gives an exact four-orientation example, and explains why orientation projection alone supplies no saving. The standalone profile and coherent-mean inspector shows the complete finite norm groups, arbitrary sixth-factor masks and exact residual quantities. Its finite tables are diagnostics of this construction, not asymptotic estimates, training or GRH evidence.

35 Proof and joint-mean boundary of the moving-profile construction

35.1 Norm groups and primitive conductors

For a split rational prime p, a squarefree ideal contains zero, one or both conjugate primes. Its rational norm exponent is respectively $0,1,2$; at exponent one there are two choices, and at exponent two its divisor norms include two copies of p. For an inert prime, the exponent is zero or two, and

the nontrivial choice is the single prime ideal of norm p^2. Thus the rational norm determines μ, the divisor-norm multiset and r_U, even though it need not determine the ideal. For example norm 49 has squarefree ideal divisor norms $1, 7, 7, 49$, whereas the inert prime of norm 25 has divisor norms $1, 25$. Additional fixed S deletions remove some ideals without changing this constancy among those remaining. The high-window formula and its parity are the already proved facts of Section 20, not a new coefficient estimate.

At a good prime ideal $\mathfrak{p}$ of norm q, Euler's definition gives $\chi_{\mathfrak{p}}(\zeta) = \zeta^{(q-1)/6}$. All such norms are 1 modulo 6; for $a = 1 + 6r, b = 1 + 6s$ one has $(ab-1)/6 = r + s + 6rs$. Multiplicativity proves (39). Write $\chi_4(N) = (-1)^{(N-1)/2}$ for odd N. On good ideals,

$$\eta = \chi_4(N)\omega^{(N-1)/3},$$
$$\eta^2 = \omega^{2(N-1)/3} = \omega^{-(N-1)/3}, \qquad \eta^3 = \chi_4(N).$$

The inverse orientation in the middle expression is essential.

For completeness the local conductors are exact, not merely upper bounds. The quadratic norm character factors modulo $(2)^2 = (4)$. Since $\mathrm{Tr}(\delta\gamma) \in 3\mathbb{Z}$ for every $\gamma \in \mathcal{O}$, expanding the norm gives $N(\alpha + \delta^3\beta) \equiv N\alpha \pmod 9$. The cubic norm character therefore factors modulo δ^3. The coprime local witnesses

$$\alpha_2 = 13 + 6\omega, \quad N\alpha_2 = 127, \quad \alpha_2 - 1 = 2\omega\delta^3,$$
$$\alpha_3 = -11, \quad N\alpha_3 = 121, \quad \alpha_3 - 1 = 4\delta^2.$$

show minimality. The first is 1 modulo 2, is $1 + 2\omega$ modulo 4, and is 1 modulo δ^3; its η value is -1. The second is 1 modulo 4 and $-2 = 1 + \delta^2$ modulo δ^3, but not 1 modulo δ^3; its η value is ω and its square is ω^2. Thus the quadratic local exponent is two and the cubic one is three. The two moduli are coprime, giving the stated norms $432, 27, 16$ for the nontrivial powers. Global units have norm one, so these residue characters are unit trivial and descend to ideal Hecke characters in this class-number-one field. A norm-seven prime has value ζ, proving order six. A larger finite S does not principalize the character: CRT gives representatives of the same nontrivial ray classes avoiding the extra excluded primes. Primitive conductor reduction never restores raw primes excluded by S or the sixth factor.

35.2 Exact alignment, with arbitrary overlapping sixth factors

The bumps in (40) have radius $1/(4D)$ and center N/D. Distinct supports are disjoint. At an active integer node their value is exactly w_N, and all endpoints are zero. Differentiation gives $\|W_D^{(j)}\|_\infty \le (4D)^j \|\phi^{(j)}\|_\infty$. Expanding $(y\partial_y)^j$ into ordinary derivatives and using $1 < y < 2$ gives $M_k(W_D) \ll_k D^k$.

Every nonzero ideal v has a generator β_v. The literal element $u_v = \zeta\beta_v^6$ is independent of that choice, since all units have sixth power one; distinct ideals yield distinct rows. Its norm is $(Nv)^6$. For squarefree good m, exact raw multiplicativity says

$$\chi_m(u_v) = \eta(m)\mathbf{1}_{(m,v)=1}.$$

Repeated primes of v are allowed and its collisions with m remain zero. The primitive row character is always η, so the actual principal indicator in (15) is zero when $\nu = 1$. Substituting the nodal profile proves (42) with no factor six, unit averaging or residue replacement.

35.3 The coprime-triple density and the selected $h = 1$ block

Put $\kappa = \pi/(3\sqrt{3})$. Lattice counting, with six unit generators per ideal, gives

$$A_K(Y) := \#\{a : Na \le Y\} = \kappa Y + O(\sqrt{Y} + 1), \qquad A_K(Y) \le CY \quad (Y \ge 1).$$

This is the existing elementary ideal count. Let $T(D, X)$ count ordered triples (m, n, v) with m, n squarefree good, $D < Nm, Nn < 2D$, $Nv \le X$, v good and arbitrary, and the three ideals pairwise coprime. For $D, X \to \infty$, $X \le D$, we claim

$$\begin{aligned} T(D, X) &= (\kappa^3 \mathcal{C}_S + o(1))D^2 X, \\ \mathcal{C}_S &= \prod_{\mathfrak{p} \in S} (1 - t_{\mathfrak{p}})^3 \prod_{\mathfrak{p} \notin S} (1 - 5t_{\mathfrak{p}}^2 + 6t_{\mathfrak{p}}^3 - 2t_{\mathfrak{p}}^4) > 0, \\ t_{\mathfrak{p}} &= (N\mathfrak{p})^{-1}. \end{aligned} \tag{43}$$

Here a product over S means its prime ideals. To prove this, first impose the local restrictions at a fixed finite set of primes containing S. They are finite Boolean combinations of divisibility by $\mathfrak{p}$ and $\mathfrak{p}^2$ in m, n, and by $\mathfrak{p}$ in v. Finite inclusion-exclusion and the count for fixed ideal multiples give the product of the local densities, with relative error tending to zero. At a good prime the allowed cases are no positive valuation, a single valuation one in m or n, or an arbitrary positive valuation in v. The factor is therefore

$$(1 - t)^3 \left(1 + 2t + \frac{t}{1 - t}\right) = (1 - t)^2(1 + 2t - 2t^2) = 1 - 5t^2 + 6t^3 - 2t^4.$$

At S all three valuations must vanish, giving $(1 - t)^3$. The geometric series for v is essential; v is not squarefree.

An omitted failure belongs to one of five events: $\mathfrak{p}^2 \mid m$, $\mathfrak{p}^2 \mid n$, $\mathfrak{p} \mid (m, n)$, $\mathfrak{p} \mid (m, v)$ or $\mathfrak{p} \mid (n, v)$. Divisibility counts and $A_K(Y) \le CY$ bound each by $O(D^2X/(N\mathfrak{p})^2)$, uniformly; if a quotient scale is below one its count is zero. Thus the tail after norm P is

$$O\left(D^2 X \sum_{N\mathfrak{p} > P} (N\mathfrak{p})^{-2}\right) = o_P(D^2 X).$$

The series converges by the ideal count (or $\zeta_K(2) < \infty$). Take the large-scale limit first at fixed P, then $P \to \infty$. Each local factor is positive and differs from one by $O(t^2)$, proving (43). Strict annulus endpoints have no effect on the main term by the same lattice count. The secondary range $X/2 \le Nv < X$ has half this limiting volume, though it may be empty at small X.

On the selected literal unit class every expanded pair term in (42) is nonnegative. Keeping only $(m, n) = 1$ gives its exact $h = 1$ A–A contribution

$$\mathcal{E}^{AA}_{\text{selected},h=1} = D^{-1} \sum_{v \text{ good } Nv \le X} \sum_{\substack{m,n \text{ squarefree,good } D < Nm, Nn < 2D \\ (m,n)=(m,v)=(n,v)=1}} |A_U(m) A_U(n)| \ge D^{-1} T(D, X).$$

Consequently, with $X = H^{1/6}$,

$$\sum_{0 < Nu \le H} |E_u|^2 \ge \sum_{Nv \le X,\ v \text{ good}} |E_{u_v}|^2 \ge \mathcal{E}^{AA}_{\text{selected},h=1} \ge (\kappa^3 \mathcal{C}_S + o(1))DX.$$

The first inequality is positivity of the *outer squared row sum*. The selected $h = 1$ block does not lower-bound the signed $h = 1$ sum over other literal rows or a unit-complete family. No positive-semidefinite claim about a gcd-truncated Gram matrix is needed. This is A–A, not a positive-density assertion for C–C long-long weights, which may vanish. This proves Proposition 15 with $c_S = \kappa^3 \mathcal{C}_S$. The derivative tradeoff follows by comparing powers in $DH^{1/6} \ll D^{k\alpha + (1+\theta)\varepsilon} H$ and then letting $\varepsilon \downarrow 0$ only when the proposed bound holds for every positive ε.

35.4 Projection preserves the coupled problem, not a saving

Use the exact joint notation of Section 29. For fixed h, q with $q \neq 1$ (including $h = 1$), let $n = \varphi(q)$, $\ell = N(hq)$, $Q = Nq$, and

$$c = \frac{Q^2 \xi_q(-\alpha)\tau_\beta(\xi_q)}{\ell^6},$$
$$\Gamma_r^F = c \sum_K B_F(K)\overline{\theta_r(K)}, \qquad A(K) = \sum_{r \in J} a_r \overline{\theta_r(K)}.$$

The active orientations J include all original restrictions, and the literal unit phase belongs to a_r. The distinct θ_r are orthonormal for $n^{-1}\sum_K$. If P projects onto their span, then

$$PB_F = (cn)^{-1} \sum_{r \in J} \Gamma_r^F \theta_r, \qquad \sum_K |PB_F(K)|^2 = (|c|^2 n)^{-1} \sum_r |\Gamma_r^F|^2.$$

Replacing B_F by PB_F in the full bilinear pairing is exact. Cauchy–Schwarz is precisely $|c\sum_K B_F A| \leq \|a\|_2 \|\Gamma^F\|_2$. The conductor and group-size factors cancel. Projection rank alone gives no arithmetic saving for this particular coupled vector.

For $s = \omega(q)$ let $C = \{K : x_{\mathfrak{p}}(K)^2 = 1 \text{ for all } \mathfrak{p} \mid q\}$. All orientations agree on C. If M_C denotes multiplication by its indicator, the matrix in the orthonormal orientation basis is

$$PM_C P = 3^{-s}\mathbf{1}\mathbf{1}^*, \qquad \lambda = |J|/3^s, \qquad \|(1-P)M_C P\|^2 = \lambda(1-\lambda).$$

Thus coherent restriction does not commute with projection. For example

$$\left| \sum_C (B_F - PB_F)A \right| \leq \sqrt{\lambda(1-\lambda)} \|(1-P)B_F\|_2 \|A\|_2,$$

where all norms in this inequality are counting ℓ^2 norms. Normalized norms would require an additional factor n on the right. Counting a coherent subset alone is not a lower bound for its actual signed pairing.

35.5 A nonzero orientation mean on actual nonprincipal rows

For the selected rows let V be any finite set of good ideals, put the first five row factors equal to one, and retain the sixth-factor mask. For each h, q define $Z_{h,q} = \#\{v \in V : (v, hq) = 1\}$. After removing the separately displayed literal unit phase, the exact correlation is

$$\Gamma_r^F = Z_{h,q} = c \sum_K B_V(K)\overline{\theta_r(K)} \quad \text{for every active orientation } r.$$

To pass from the published C–C identity to A–A replace each r_U by $r_U - 1$. At $h = 1$, $m = r, n = q/r$, the coefficient is

$$a_r^{AA} = (r_U(m) - 1)(r_U(n) - 1) W_m \overline{W_n} \theta_{m,n}(\zeta), \qquad \mu(q) a_r^{AA} = |A_U(m) A_U(n)|.$$

This retains $\mu(q)$, the common complex profile and its conjugation, all deleted common primes, and the literal ζ phase. If $Z_{1,q} > 0$ and at least one active interior A–A coefficient is nonzero, the actual pairing is positive, whereas $\sum_r |\Gamma_r^F - \text{mean}\,\Gamma^F|^2 = 0$. An orientation-variance-only replacement with no mean term is therefore false. The mean is not a principal residue: these rows have character η, with nonprincipal powers. This does not refute a properly formulated fixed-profile signed estimate.

A fresh exact example uses $D = 4000$, $U = 63$ and chosen prime ideals of norms $7, 13, 79, 67, 97$. Their product has four active orientations: the pairs 6097/7663, 7189/6499 and both reversals. Here $r_U = 2$, and A_U is -1 on the three-prime columns and $+1$ on the two-prime columns. The four good ideals with $10 \le Nv \le 20$ have norms $13, 13, 19, 19$. Exactly one collides with the chosen 13-prime, so $Z = 3$ for all four orientations. The exact selected $h = 1$ contribution is $12/4000 = 3/1000$ despite zero orientation variance. Dropping the literal ζ phase incorrectly changes the coefficient sum from 4 to -1, giving $-3/4000$. The coherent fraction is $4/243$ and the squared leakage is $956/59049$. The residue group and infinite Fourier tail are not estimated by this finite check. The fixed row conductor 432 must be distinguished from the pair-column conductor $q = mn$, whose norm here and in the asymptotic annulus is comparable to D^2.

The exact implementation at `examples/profile_uniformity/` keeps complete norm groups and arbitrary sixth factors, including distinct ideals of norm 49 with different masks. The saved six-scale examples, exact joint fixture and their independent checks certify finite identities only. Full-union masks, coherent means, fixed-profile uniformity and the remaining mixed/pure upper-envelope neighborhoods still require a new signed estimate. No determinant comparison, zero location, analytic or categorical Lean theorem, quantum result or GRH proof is obtained.

36 A bound for the signed coprime slice and an exact character twist

The coprime-column contribution can be negative, so it needs its own bound. Fix the original ν, S, one common possibly complex smooth W supported in $[1, 2]$, $D > 1$, $U \ge 1$ and $U^2 \le D$. Consider an actual block

$$u = u_0 v_1 v_2^2, \quad A_j \le N v_j < 2A_j \ (j = 1, 2), \quad Nu \le H, \qquad R = A_1 A_2, \quad M = \max(A_1, A_2).$$

The good ideals v_1, v_2 are squarefree and coprime, with their literal primary/ray representatives; all other good factors are the unit ideal. Let F be the actual indicator, including capacity and every row mask. In the remaining neighborhoods the first factor is nonunit, so the primitive row character is nonprincipal and $E_u = I_u$ there. For $X, Y \in \{A_U, C_U, \mu\}$ define the actual signed slice

$$S_1^{XY} = \frac{1}{D} \sum_u F(u) \sum_{\substack{m,n \text{ squarefree,good} \\ (m,n)=1}} X(m) Y(n) \nu(m) \overline{\nu(n)} \chi_m(u) \overline{\chi_n(u)} W(Nm/D) \overline{W(Nn/D)}. \quad (44)$$

All columns in this formula are the original annular columns, including raw zero sentinels. No positive-semidefinite claim about the coprime truncation is made.

Proposition 16 (A bound for the actual signed $h = 1$ slice)**.** *Uniformly in the original cutoff and this actual block,*

$$|S_1^{XY}| \ll_{\nu,S,\epsilon} \|W\|_\infty^2 (HD)^\epsilon \min\left\{ R + \frac{DR}{M} + \frac{D^{2/3} R}{M^{1/3}},\ D + R^2 \right\}. \quad (45)$$

The exact identity underlying the estimate is

$$\begin{aligned} D S_1^{XY} &= \sum_{\substack{d \text{ squarefree} \\ (d,S)=1}} \mu(d) \sum_{u \text{ in the actual block} \atop (d,u)=1} T_d^X(u) \overline{T_d^Y(u)}, \\ T_d^X(u) &= \sum_{\substack{a \text{ squarefree} \\ (a,dS)=1}} X(da) \nu(a) \chi_a(u) W(Nd\,Na/D). \end{aligned} \quad (46)$$

In particular $X(da)$ is not replaced by $X(d)X(a)$, and the mask $(d,u)=1$ remains.

The proof applies the accepted block sieves inside each positive row norm in (46), then explicitly loses the d signs in an absolute triangle. At $H=D^{3/4}$ the two remaining power-scale neighborhoods $A_1 \asymp D^{5/12}, A_2 \asymp D^{1/6}$ and $A_1 \asymp D^{3/4}, A_2 \asymp 1$, with feasible capacity constants, both retain power $7/6$. The desired $3/4$ still needs a signed saving of $5/12$. This proves a separate signed-slice bound; it does not improve the published global envelope or supply the missing cancellation.

There is also an exact arithmetic description of the two-pivot phase. Put $\delta = 1+2\omega$ and, at each good prime ideal p, set $j(p) = J_p(\chi_p, \chi_p^2)$. Extend j multiplicatively on squarefree columns. Let $\Psi(a)$ be the unique generator of an ideal prime to 2δ satisfying $\Psi(a) \equiv 1 \pmod{2\delta}$.

Proposition 17 (The actual Jacobi column twist)**.** *The map Ψ is an algebraic Hecke character with exact finite conductor (2δ), of norm 12, and principal convention $\Psi((\alpha)) = \alpha$ for $\alpha \equiv 1 \pmod{2\delta}$ (infinity type $(1,0)$). For good squarefree columns,*

$$(j(p)) = p, \quad -j(p) = \Psi(p), \qquad \mu(m)j(m) = \Psi(m). \tag{47}$$

Define $\vartheta(m) = \Psi(m)/\sqrt{Nm}$ and $\rho(m) = j(m)/\sqrt{Nm} = \mu(m)\vartheta(m)$ on those columns. In the high window the actual gauged coefficients are $b_A = (r_U - 1)\vartheta\nu W$, $b_C = r_U\vartheta\nu W$ and $b_\mu = \vartheta\nu W$.

This does not replace the original row kernel. Its compensating inverse gauge remains, so all ordinary coefficient/operator norms are unchanged. The new character has nonzero infinity type and is not silently covered by a theorem stated only for finite-order twists. The identities in (47) involving μ must not be extended to arbitrary nonsquarefree sixth factors.

A fresh fixed-profile example takes $\nu = 1$, S the primes above $2, 3$, literal unit one and primary row generators 1 modulo 3. It lies in an actual $A_6 = 1$ block at $D = 60000$, $U = 244$, $H = 3833$, $A_1 = 37$, $A_2 = 7$. All 20 rows of the stated literal class are nonprincipal. A selected q has exactly two legal column orientations, of norms $97051/96091$ and its reversal. Their correlations in the ω basis are $(-4,-2)$ and $(-2,2)$, with mean -3. Their selected contribution is $S_{1,q}^{AA} = -1/10000$; deleting the mean gives zero. The centered squared norm is 6, not zero, and the proposed variance-only Cauchy bound would incorrectly require $36 \le 12$. This refutes that centering rule in a fixed-profile row block, not the asymptotic target. Other q contributions remain. The coprime-kernel inspector shows the actual rows, raw zeros, two-column Gram matrix and compensated gauge; no training or arithmetic spectrum is inferred from that finite view.

37 Proof of the signed-slice budget and its surviving analytic burden

37.1 Gcd inversion before an absolute bound

Insert $\mathbf{1}_{(m,n)=1} = \sum_{d \mid m,\ d \mid n} \mu(d)$ in (44), and write $m = da, n = db$. Squarefreeness forces $(a,d) = (b,d) = 1$. The common factors give $\nu(d)\overline{\nu(d)} = 1$ and $\chi_d(u)\overline{\chi_d(u)} = \mathbf{1}_{(d,u)=1}$. This proves (46), retaining the original coefficient $X(da)$ and the full row-zero mask. For $X = Y$ each inner squared row sum is nonnegative; its $\mu(d)$-weighted sum need not be.

Put $L = D/Nd$. Whenever a column survives, $L \ge 1/2$. The actual divisor-assignment bound and ideal counting give

$$\sum_{\substack{a \text{ squarefree} \\ (a,dS)=1}} |X(da)\nu(a)W(Nd\,Na/D)|^2 \ll_\eta D^\eta \|W\|_\infty^2 L.$$

Indeed $N(da) \le 2D$, $|X(da)| \ll_\eta D^\eta$ uniformly in d, U, and there are $O(L)$ ideals of norm at most $2L$ for $L \ge 1/2$. This includes C_U by $C_U = A_U + \mu$ on the active annulus. There is no multiplicativity assertion for $X(da)$. For $L \ge 1$, the powered block bound and the ordinary-conductor comparison of Section 30, followed by Cauchy for the cross term, imply

$$|G_d^{XY}| \ll (HD)^\eta \|W\|_\infty^2 L \min\left\{R + \frac{LR}{M} + \frac{L^{2/3}R}{M^{1/3}},\ L + R^2\right\}, \tag{48}$$

where G_d^{XY} denotes the inner row sum in (46). The raw (d, u) restriction is only enlarged inside these positive norms; u_0 is a fixed column multiplier of modulus at most one. For $1/2 \le L < 1$, only the unit column can survive. The direct row count is $O(R)$, and the right side of (48) is at least a constant times R since $R \ge 1$ and $L \ge 1/2$. Thus this endpoint does not require a sieve with an unstated subunit length.

Now take the absolute triangle in d explicitly. The two estimates may be summed separately and their minimum taken. The elementary ideal sums

$$\sum_{Nd \le 2D} \frac{1}{Nd} \ll \log(2D), \qquad \sum_d (Nd)^{-2} \ll 1, \qquad \sum_d (Nd)^{-5/3} \ll 1$$

turn the first estimate into $R\log(2D) + DR/M + D^{2/3}R/M^{1/3}$ after division by D; the second gives $D + R^2\log(2D)$. Absorb the logarithms with smaller internal exponents to obtain (45). All raw zeros, literal/ray/squarefree/capacity restrictions were retained before the stated positive enlargements. No bound for an indefinite submatrix was inferred from positivity of a larger Gram matrix.

At the mixed power-scale point the powered exponents are $7/12, 7/6, 10/9$, and the ordinary exponents are $1, 7/6$. At the pure point they are $3/4, 1, 7/6$ versus $1, 3/2$. These are upper-bound exponents with feasible dyadic constants, not actual energy lower bounds. A sufficient $h = 1$ target for the A–A combination is

$$\left|\sum_d \mu(d) G_d^{AA}\right| \ll_{\nu,S,W,\epsilon} DH(HD)^\epsilon. \tag{49}$$

At $H = D^{3/4}$ this asks for power $7/4$ before division by D, while (45) gives $13/6$, leaving $5/12$. This is sufficient for this slice, not a necessary bound for every separate C–C/C–μ term if their signed assembly cancels further. The quantities G_d depend on prime supports, $X(da)$ and row deletions; they are not an independent smooth function of Nd to which a scalar Möbius partial-sum estimate can simply be applied.

37.2 The correctly oriented Jacobi generator

At a good prime ideal p, let $n = Np$ and $\chi = \chi_p$ be the literal positive sextic residue character. In $\mathcal{O}/p$, its value reduces to $x^{(n-1)/6}$. For $j(p) = J(\chi, \chi^2)$, the elementary Gauss–Jacobi identity and the three nontrivial Gauss norms give $|j(p)|^2 = n$. Reduction in that *same* residue field gives

$$j(p) \equiv \sum_x x^{(n-1)/6}(1-x)^{(n-1)/3} = 0 \pmod p.$$

Every exponent after expansion lies strictly between 0 and $n-1$, so the finite-field power sums vanish. Hence $(j(p)) = p$. This includes an inert prime of norm p_0^2, and inverting the literal character conjugates the ideal at a split prime.

This raw sum equals the old $J(\chi^2,\chi^3)$: the quotient of the two Gauss–Jacobi expressions is $G(\chi)G(\chi^5)/G(\chi^3)^2=1$. It differs from the theta-curve Frobenius normalization by its explicitly stated factor $-\chi^3(-1)$; that sign is not imported silently. No new Gauss or Jacobi norm theorem is claimed.

Let $f=(2\delta)$. The six units map bijectively onto $(\mathcal{O}/f)^\times$: the components at 2 and δ have orders three and two, with complementary kernels. Since $\mathcal{O}$ is a PID, every ideal prime to f has a unique generator $\Psi(a)\equiv 1 \pmod f$. Uniqueness gives multiplicativity. If $u_f(\alpha)$ is the unit having the residue of α modulo f, then

$$\Psi((\alpha))=\alpha/u_f(\alpha), \qquad |\Psi(a)|^2=Na.$$

This explicitly defines the finite part and the principal infinity-type convention. The finite conductor divides f. A proper divisor omits a component: the unit -1 is 1 modulo 2, and ω is 1 modulo δ. A smaller defining modulus would require the unit ideal to have value -1 or ω. Thus the finite conductor is exactly f.

To identify the actual $j(p)$, now reduce in the *character-value* ring, a different operation from reduction modulo p above. Modulo δ, a sixth root v satisfies $v=v^3$ and $v^2=1$. The zero branches at $x=0,1$ are kept, so $j(p)\equiv\sum_{x\neq 0,1}\chi^3(x)=-1 \pmod\delta$. Modulo 2, $v=v^4$, and $J(\chi,\chi^2)\equiv J(\chi^4,\chi^2)=-\chi^4(-1)=-1$. The inverse-character Jacobi identity here follows by the bijection $x/(1-x)$ onto the nonzero field elements other than -1. Consequently $-j(p)\equiv 1 \pmod f$, proving $-j(p)=\Psi(p)$. Multiplication proves (47) on squarefree columns. For nonsquarefree ideals μ vanishes, so that last identity must not be extended to arbitrary sixth factors. This construction establishes the fixed character directly; it assumes no elliptic modularity, spectral comparison, or distribution of prime arguments.

37.3 The actual two-factor completion, not a new row kernel

Freeze the literal u_0,v_3,v_4,v_5,v_6 and put $Z=u_0v_3^3v_4^4v_5^5v_6^6$. The last factor is arbitrary, with its allowed overlaps. For $h=1$, write $q=mn\neq 1$, $Q=Nq$, $\theta=\chi_m\overline{\chi_n}$ and $\xi_q=\theta^3$. For $q=(\beta)$ use the original inverse-different pairing $e_\beta(z)=\exp(2\pi i\,\mathrm{Tr}(z/(\delta\beta)))$. Let $F(v_1,v_2)$ be the exact common finite weight, retaining the first-five coprimality, primary/ray and capacity restrictions. It must be independent of the column orientation. Define

$$\widehat{F}_q(k_1,k_2)=\sum_{v_1,v_2}F(v_1,v_2)e_\beta(-k_1v_1-k_2v_2),$$
$$T_{q,F}(\theta)=\sum_{k_1,k_2\in(\mathcal{O}/q)^\times}\widehat{F}_q(k_1,k_2)\overline{\theta(k_1k_2^2)}.$$

Periodization sums every actual representative; no injectivity assumption is needed. The primitive powers θ,θ^2 have zero Gauss sums at nonunit frequencies, including at inert primes. Exact inversion gives

$$\sum F\theta(v_1)\theta(v_2)^2=Q^{-2}\tau_\beta(\theta)\tau_\beta(\theta^2)T_{q,F}(\theta)=Q^{-2}J_q(\theta,\theta^2)\tau_\beta(\xi_q)T_{q,F}(\theta).$$

Jacobi sums CRT without additive-generator multipliers, so $J_q=j(m)\overline{j(n)}=\sqrt{Q}\,\rho(m)\overline{\rho(n)}$. The quadratic central phase is the already recorded period-four law

$$\gamma(\beta)=\tau_\beta(\xi_q)/\sqrt{Q}=\tfrac12\{1+i^{-b}+i^a+i^{b-a}\}, \qquad \beta=a+b\omega.$$

For composite squarefree q, identifying $\tau_\beta(\xi_q)$ with $\sum_x e_\beta(x^2)$ requires CRT expansion of the local root counts; $1+\xi_q$ is not the root-count function on a composite ring. The known quadratic Gaussian evaluation is reused, not retested as a new phase theorem. Its pairing is exactly the displayed e_β.

With $b_X = \rho X \nu W$, the actual two-pivot identity is

$$DS_1^{XY} = \sum_q \frac{\gamma(\beta_q)}{Q} \sum_{mn=q} b_X(m)\overline{b_Y(n)}\theta_{m,n}(Z)T_{q,F}(\theta_{m,n}). \tag{50}$$

The two Gauss magnitudes, equivalently the central Gauss and Jacobi magnitudes, give exactly $1/Q$. There is no extra $\xi_q(-1)$ in this two-factor identity; the full six-factor formula of Section 28 retains that sign. If β changes to $u\beta$, both γ and T acquire the same quadratic sign $\xi_q(u)$, which cancels in their product. If Z shares any prime with q, $\theta(Z) = 0$, including at a sixth-power collision. No dropped primitive-conductor prime is restored.

Multiplying columns by ρ changes the kernel by its compensating inverse unitary gauge. It preserves operator norms and eigenvalues even for an indefinite coprime-truncated matrix. A fixed-character coefficient estimate cannot replace a bound for the actual q-dependent $T_{q,F}$. At the pure endpoint $v_2 = c$ with $(c, q) = 1$, exact summation gives

$$T_{q,F}(\theta) = \theta(c)^2\tau_\beta(\overline{\theta^2}) \sum_{k_1 \in (\mathcal{O}/q)^\times} \widehat{F}_1(k_1)\overline{\theta(k_1)}.$$

Since $\theta^2(-1) = 1$, the extra Gauss pair multiplies to Q, returning the original one-pivot scalar $\tau_\beta(\theta)/Q$. A nonunit c gives the original zero. Thus a putative central-phase gain cannot be retained after discarding a literal length-one Fourier factor.

37.4 A central reciprocal remains after the gauge

On squarefree columns put $B_U = \mu A_U$ and $\vartheta = \Psi/\sqrt{N}$. The original ternary assignment gives exactly

$$B_U(m) = \sum_{\substack{abc=m \\ Na,Nb \le U}} \mu(c), \qquad \rho(m)A_U(m) = \vartheta(m)B_U(m).$$

All factors are squarefree and pairwise coprime. Fix a nonzero numerator z and keep the full deleted set $T = S \cup \operatorname{supp}(z)$. Write Ξ for the unitary primitive character of $\vartheta\nu\kappa_z$, with these raw deletions retained. Its nonzero infinity type prevents a silent reuse of a finite-order-only moment theorem. For $\Re s > 1$,

$$J_U^{\sharp,T}(s) = \sum_{m \text{ squarefree},\ (m,T)=1} B_U(m)\Xi(m)(Nm)^{-s} = N_U^T(s)/L^T(s,\Xi),$$

$$N_U^T(s) = \sum_{\substack{a,b \text{ squarefree} \\ (a,b)=1,\ (ab,T)=1 \\ Na,Nb \le U}} \frac{\Xi(ab)}{(Nab)^s} \prod_{p \mid ab} (1 - \Xi(p)(Np)^{-s})^{-1}.$$

Fixing a, b leaves the Möbius Euler product in c; this proves the identity by absolute convergence. The numerator is finite and analytic for $\Re s > 0$. At $U = 1$ it is one, so the series is exactly $1/L^T(s,\Xi)$. That boundary identity is not a high-window counterexample. It shows why cancellation of central reciprocal poles cannot simply be inferred from the column-character identification.

For $x_p = \Xi(p)(Np)^{-s}$, $\alpha_p = (Np)^{-v}$ and $\beta_p = (Np)^{-w}$, the gauged three-placement Euler factor is $1 + x_p(\alpha_p + \beta_p - 1)$, with the opposite signs from the ungauged factor in Section 20. In the initial absolute domain it gives

$$\prod_{p \notin T} \{1 + x_p(\alpha_p + \beta_p - 1)\} = \frac{L^T(s+v,\Xi)L^T(s+w,\Xi)}{L^T(s,\Xi)} \prod_{p \notin T} H_p,$$

where

$$H_p = \frac{\{1 + x_p(\alpha_p + \beta_p - 1)\}(1 - x_p\alpha_p)(1 - x_p\beta_p)}{1 - x_p}.$$

Expansion gives

$$H_p - 1 = \frac{x_p^2(\alpha_p + \beta_p - \alpha_p^2 - \beta_p^2 - \alpha_p\beta_p)}{1 - x_p} + \frac{x_p^3\alpha_p\beta_p(\alpha_p + \beta_p - 1)}{1 - x_p}.$$

Thus the correction converges normally in the open region $\Re s > 0$ where the real parts of all of

$$2s + v,\ 2s + w,\ 2s + 2v,\ 2s + 2w,\ 2s + v + w,$$
$$3s + 2v + w,\ 3s + v + 2w,\ 3s + v + w$$

exceed one. In particular, for $\Re s \geq 1/2 + \eta$ and $\Re v, \Re w \geq 0$, its terms differ from one by $O_\eta((Np)^{-1-2\eta})$. Deleted primes only remove factors. Finite correction factors may vanish; no zero-free product is asserted. The displayed factorization is first an absolute-domain identity and is used meromorphically only where the functions have been continued. At $v = w = 0$ it simplifies to $L^T(s, \Xi)/L^T(2s, \Xi^2)$, a special slice cancellation which does not remove the general central reciprocal. No finite-height inversion is declared exact without truncation errors; no new contour, reciprocal family moment or analytic saving is proved. Most importantly, this fixed-z coefficient series is not the actual q-dependent joint weight in (50).

37.5 The fixed-profile remaining-sector diagnostic

Use the fixed smoothstep $s(t) = e^{-1/t}/(e^{-1/t} + e^{-1/(1-t)})$ for $0 < t < 1$, extended by zero and one at the two sides, and set $W(y) = s(8(y-1))s(8(2-y))$. It is smooth, supported in $[1, 2]$, and identically one on $[9/8, 15/8]$, independently of D. Take $\nu = 1$ and S the primes above $2, 3$. At $D = 60000, U = 244, H = 3833$, select the actual block $A_1 = 37, A_2 = 7$, literal unit one and primary row generators 1 modulo 3. Complete enumeration gives 44 candidate pairs, with two coprimality and 22 capacity rejections, leaving 20 distinct nonprincipal rows. The first outside-S valuation is one in each row.

Choose $\mathfrak{p}_p = (p, \omega - \lambda_p)$ for

$$(p, \lambda_p) \in \{(37, 10), (43, 6), (61, 13), (307, 17), (313, 98)\}.$$

Their product q has exactly two legal annular column orientations: 97051/96091 and its reversal. Direct small- and long-divisor assignments give $r_U = (0, 2)$, $\mu = (-1, 1)$, $C_U = (0, 2)$ and $A_U = (1, 1)$. Both profile values are exactly one. Twelve orientation/row evaluations are raw zeros. Exact residue-field exponentiation yields

$$\Gamma_1 = -4 - 2\omega, \qquad \Gamma_2 = -2 + 2\omega, \qquad \Gamma_{\mathrm{av}} = -3.$$

Here Γ_{av} is the orientation average. With $a^{AA} = (-1, -1)$ and $\mu(q) = -1$,

$$S_{1,q}^{AA} = -6/60000 = -1/10000, \qquad \sum_r a_r(\Gamma_r - \Gamma_{\mathrm{av}}) = 0.$$

The centered norm is $\sum_r |\Gamma_r - \Gamma_{\mathrm{av}}|^2 = 6$ and $\sum_r |a_r|^2 = 2$, whereas $|\sum_r a_r\Gamma_r|^2 = 36 > 12$. A variance-only Cauchy bound therefore fails even with nonzero variance. It is not a negative full row energy or an asymptotic counterexample.

The corresponding *selected two-column* Gram matrix has diagonal $14, 20$ and the displayed conjugate off-diagonal entries. Its eigenvalues are $17 \pm \sqrt{21}$ and its A quadratic form is 28 before division by D. Removing the diagonal for coprimality gives eigenvalues $\pm 2\sqrt{3}$ and the A form -6. Other original columns and q values have not been suppressed in a full-family claim: this display is a selected-q contribution with its complete stated row block. All finite data and scalar identities have independent exact checks; no finite spectrum or learned label implies a determinant or GRH theorem. The next estimate must retain the signed mean, reciprocal burden and actual coupled q weights, closing the still-open 5/12 gap.

38 The signed divisor mean and a complete-column laboratory

The preceding signed-slice estimate loses the sign of the common ideal divisor. We now retain that variable in an exact analytic reduction and in a complete finite computation. Neither supplies the missing 5/12 saving at $H = D^{3/4}$. The computations below use a new fixed profile; they do not fill in the uncomputed columns of the earlier broad-plateau profile.

Let $\mathcal{C}$ be original squarefree good ideal columns in $[D, 2D]$, and let $\mathcal{R}$ be an actual finite literal row block, with nonnegative multiplicities w_u. All shared-prime zeros of the original $\chi_m(u)$ remain. For arbitrary complex coefficient vectors define the linear-first form

$$DQ_1(x, y) = \sum_{u \in \mathcal{R}} w_u \sum_{\substack{m,n \in \mathcal{C} \\ (m,n)=1}} x_m \overline{y_n} \chi_m(u) \overline{\chi_n(u)}. \tag{51}$$

The original component is obtained with $x_m = X_U(m)\nu(m)W(Nm/D)$ and the corresponding y_m. On the nonprincipal row blocks used below the original residue term vanishes. Equation (51) is a raw coprime form on a generic block; principal rows would still require their separate residue assembly.

For every ideal divisor of an active column put

$$F_d(x; u) = \sum_{m \in \mathcal{C},\ d|m} x_m \chi_m(u), \qquad Z_d(x, y) = \sum_u w_u F_d(x; u) \overline{F_d(y; u)}.$$

The exact identity is

$$DQ_1(x, y) = \sum_d \mu(d) Z_d(x, y). \tag{52}$$

There is no extra $\mu(q)$: the vectors already contain the original coefficients. Full divisors $d = m$ are essential to cancel self-pairs. Divisors and squarefreeness refer to ideals, including distinct ideals with the same rational norm.

Proposition 18 (A weighted divisor-incidence lift)**.** *Suppose $\lambda_d \geq 0$ is an arbitrary-coefficient bound*

$$\sum_u w_u |F_d(x; u)|^2 \leq \lambda_d \sum_{m \in \mathcal{C},\ d|m} |x_m|^2.$$

With $B(m) = \sum_{d|m} \lambda_d$ and $\|x\|_B^2 = \sum_m B(m)|x_m|^2$, one has

$$|DQ_1(x, y)| \leq \|x\|_B \|y\|_B. \tag{53}$$

For a finite literal phase table one may take the exactly computed

$$\lambda_d = E_d := \sum_u w_u \#\{m \in \mathcal{C} : d \mid m,\ \chi_m(u) \neq 0\}.$$

This last choice is a Frobenius certificate, not a new analytic estimate. For $x = t + e$ it gives the finite perturbation bound

$$|Q_1(x,x) - Q_1(t,t)| \le D^{-1}(2\|t\|_B\|e\|_B + \|e\|_B^2). \tag{54}$$

Proof. Insert ideal-divisor inversion in (51). Apply row Cauchy to each Z_d and weighted Cauchy to the resulting sum over d. Then $\sum_d \lambda_d\|x_d\|_2^2 = \|x\|_B^2$. The squared Frobenius norm of the actual F_d matrix is exactly E_d. Expand $Q_1(t+e,t+e) - Q_1(t,t)$ to obtain the last assertion. Positivity of the coprime-truncated matrix was never assumed. □

Using an accepted positive-row sieve bound for λ_d only recovers the old exponent up to ideal-divisor factors. It does not estimate the signed sum before the triangle in (52).

There is also a genuinely signed scalar expansion. Fix squarefree residual ideals a, b, which may overlap, and an actual row u. Keep the full deleted set $\mathcal{T} = S \cup \operatorname{supp}(abu)$ and set

$$\mathcal{Z}_{U;a,b,\mathcal{T}}(s) = \sum_{\substack{d \text{ squarefree} \\ (d,\mathcal{T})=1}} \frac{\mu(d)A_U(da)A_U(db)}{(Nd)^s}, \qquad \Re s > 1.$$

For four positive cutoff variables put $P_p = 1-(Np)^{-v_1}-(Np)^{-v_2}$ and $Q_p = 1-(Np)^{-w_1}-(Np)^{-w_2}$. The common-divisor Euler factor is $1-(Np)^{-s}P_pQ_p$. Its first-order factorization has four numerator and five denominator zeta factors:

$$\prod_{p\notin\mathcal{T}}(1-(Np)^{-s}P_pQ_p) = \frac{\prod_{i=1}^2 \zeta_K^{\mathcal{T}}(s+v_i)\zeta_K^{\mathcal{T}}(s+w_i)}{\zeta_K^{\mathcal{T}}(s)\prod_{i,j=1}^2 \zeta_K^{\mathcal{T}}(s+v_i+w_j)}\,\mathcal{H}_{\mathcal{T}}(s;v,w). \tag{55}$$

The correction is normally convergent at $\Re s \ge \frac12+\eta$, $\Re v_i, \Re w_j \ge 0$, for each fixed $\eta > 0$. This says nothing about zeros or residues of the zeta ratio. Finite-height Perron kernels, with their explicit endpoint/error costs, recover the cutoffs by successive limits.

A direct finite cutoff expansion gives more explicit information:

$$\mathcal{Z}_{U;a,b,\mathcal{T}}(s) = \frac{\mathcal{N}_{U;a,b,\mathcal{T}}(s)}{\zeta_K^{\mathcal{T}}(s)}. \tag{56}$$

The numerator is a finite sum over eight nonzero joint short-divisor states, holomorphic on $\Re s > 0$. Uniformly for $\sigma \ge \sigma_0 > 0$,

$$|\mathcal{N}_{U;a,b,\mathcal{T}}(\sigma+i\tau)| \ll_{K,\sigma_0,\epsilon} (Na\,Nb)^{\epsilon}U^{4\max(1-\sigma,0)+\epsilon}.$$

The costs near $\sigma = 1/2$ persist. The remaining long state has the Möbius sign, but the eight finite short states cannot be omitted. A fixed-residual reciprocal zeta factor also does not prove an uncancelled pole in the complete coherent sum.

Exact complete-column reference. Choose once $\delta = 1/240000$ and $W_\delta(y) = s((y-1)/\delta)s((2-y)/\delta)$, where s is the smooth step from zero to one. For the finite grid $D \le 60000$ every strict integer norm node has $W_\delta = 1$; endpoints vanish. Its derivatives cost $O_k(\delta^{-k})$. The three cases have different declared row blocks, not a common asymptotic scaling experiment. They use principal ν, S above $2, 3$, literal class one and nonprincipal actual rows.

D	columns	divisors	rows	S_1^{AA}
144	41	53	5	$-19/36$
576	164	212	11	$-221/288$
60000	16953	21750	20	$-5119/5000$

In the last case the exact even/odd Möbius masses are 543416 and 604844, giving numerator -61428 before division by D. All 339060 raw column–row values, including 47908 zeros, and all nine cross components are retained. In particular 1361 columns have nonsquarefree rational norm while remaining squarefree as ideals. The signed-divisor inspector exposes the actual states and dimensions. These finite negative slices are neither whole-family asymptotic counterexamples nor a positive arithmetic determinant or GRH result.

39 Proofs and limits of the signed-divisor development

39.1 Actual divisor incidence and a finite error certificate

The identity $\mathbf{1}_{(m,n)=1} = \sum_{d|m,\ d|n} \mu(d)$ is an identity of ideal divisors. Expanding the finite sums proves (52). When $m = da$ is squarefree, $(d, a) = 1$. If $(d, u) = 1$, multiplication by $\chi_d(u)$ is a rowwise unit phase; if $(d, u) > 1$ it is the original zero. For original $x_m = X_U(m)\nu(m)W(Nm/D)$, precisely $F_d(x; u) = \nu(d)\chi_d(u)T_d^X(u)$ on $(d, u) = 1$, and it is zero otherwise. Thus their row norms agree on the actual restricted support. For general vectors the coefficient x_{da} remains unfactored. No identity $X_U(da) = X_U(d)X_U(a)$ is used. Residual a, b may overlap.

Write $\|x_d\|_2^2 = \sum_{d|m} |x_m|^2$. The hypotheses of Proposition 18 give

$$\begin{aligned} |DQ_1(x, y)| &\le \sum_d |Z_d(x, y)| \le \sum_d \lambda_d \|x_d\|_2 \|y_d\|_2 \\ &\le \left(\sum_d \lambda_d \|x_d\|_2^2 \right)^{1/2} \left(\sum_d \lambda_d \|y_d\|_2^2 \right)^{1/2} = \|x\|_B \|y\|_B. \end{aligned}$$

This is a divisor-incidence lift of positive-row bounds to an indefinite coprime form. It is not an inference from positivity of a larger Gram matrix. In particular $B(m) \le \lambda_{\max} \tau_{\text{ideal}}(m)$, with $\tau_{\text{ideal}}(m) = 2^{\omega(m)}$ on squarefree ideals. The subpolynomial divisor bound recovers the existing exponent; this triangle does not supply the desired signed cancellation.

For the original unweighted two-pivot block ($w_u = 1$), the accepted arbitrary-coefficient inputs at $L = D/Nd \ge 1$ allow

$$\lambda_d \ll_\epsilon (HD)^\epsilon \min \left\{ R + \frac{LR}{M} + \frac{L^{2/3}R}{M^{1/3}},\ L + R^2 \right\}, \qquad R = A_1 A_2, \quad M = \max(A_1, A_2).$$

Their constants are classical-input constants, not computed numbers. Any removal of a d-row mask occurs only inside the positive norm. For $1/2 \le L < 1$, $Na \in [L, 2L]$ permits at most the unit ideal; its operator norm is bounded directly by the actual row measure. For an empty subsequence take $\lambda_d = 0$.

For the fully computable choice $\lambda_d = E_d$, the matrix entries are $\sqrt{w_u}\chi_m(u)$ on d-divisible columns. The squared Frobenius norm is precisely the stated raw nonzero-entry count. It is valid for arbitrary coefficient vectors, not just the saved trajectory. No unknown sieve constant enters that finite certificate. For approximations t_x, t_y with errors $e_x = x - t_x, e_y = y - t_y$,

$$|Q_1(x, y) - Q_1(t_x, t_y)| \le D^{-1}(\|e_x\|_B \|y\|_B + \|t_x\|_B \|e_y\|_B).$$

The AA specialization follows by expanding three terms. To verify it without floating square roots, put $T = \|t\|_B^2, E = \|e\|_B^2$ and $L = D|Q_1(x, x) - Q_1(t, t)| - E$. If $L > 0$, the certificate is the equivalent real inequality $L^2 \le 4TE$; if $L \le 0$ it is immediate. In the saved exact $\mathbf{Z}[\omega]$ computations with rational weights, these are integer/rational checks rather than floating-root tests. These are finite coefficient-error statements. They do not prove a varying-family remainder estimate, spectral positivity or a determinant comparison.

39.2 A signed common-divisor Euler product

From here restore ALL original annular squarefree good columns, not an arbitrary subset from the vector lemma, and use the actual unweighted raw row block of Section 36. On the stated nonprincipal blocks it is the corresponding component of the residue-subtracted moment. On a block containing principal rows the raw identity remains valid, but the original principal correction is additional and is not discarded. Expand the exact d identity while keeping its sign. For fixed a, b, u the inner sum is

$$\sum_{\substack{d \text{ squarefree} \\ (d,Sabu)=1}} \mu(d) A_U(da) A_U(db) W(Nd\,Na/D)\overline{W(Nd\,Nb/D)}. \tag{57}$$

The exact outer coefficient is $\nu(a)\overline{\nu(b)}\chi_a(u)\overline{\chi_b(u)}$. The residual a, b are individually squarefree but need not be coprime. If their quotient character loses a common prime, that prime remains in $\mathcal{T} = S \cup \operatorname{supp}(abu)$ and in every original zero. The common d character cancels only to $\mathbf{1}_{(d,u)=1}$.

For $\Re s > 1$ the series $\mathcal{Z}$ is absolutely convergent, with coefficient majorant $3^{\omega(a)+\omega(b)}9^{\omega(d)}$, uniformly in U. Use four positive Perron variables. Write $\mathcal{E}_{\mathcal{T}}(s; v, w) = \prod_{p \notin \mathcal{T}}(1-(Np)^{-s}P_pQ_p)$. The fixed residual prefactor is $P_a(v)Q_b(w) = \prod_{p|a} P_p(v) \prod_{p|b} Q_p(w)$; a common prime of a, b contributes both factors. The common d prime has nine joint cutoff positions and the sign $\mu(p) = -1$, giving local factor $1 - zPQ$ with $z = (Np)^{-s}$, $\alpha_i = (Np)^{-v_i}$ and $\beta_j = (Np)^{-w_j}$. Thus the exact correction in (55) is

$$H_p = \frac{(1 - zPQ)\prod_{i=1}^{2}(1 - z\alpha_i)\prod_{j=1}^{2}(1 - z\beta_j)}{(1 - z)\prod_{i,j=1}^{2}(1 - z\alpha_i\beta_j)}. \tag{58}$$

The linear coefficient of numerator and denominator is $-1 - \sum_{i,j} \alpha_i\beta_j$, fixing all nine zeta signs. At all four auxiliary variables zero the product is $1/\zeta_K^{\mathcal{T}}(s)$ and $H_p = 1$. At $v_2 = w_2 = 0$ it is $1/\zeta_K^{\mathcal{T}}(s + v_1 + w_1)$, again with $H_p = 1$. These are generating-function checks, not finite-cutoff identifications.

For $\Re s \geq 1/2 + \eta$, $\Re v_i, \Re w_j \geq 0$, the degree-five numerator and denominator of (58) have equal constant and linear coefficients. The sums of absolute coefficient majorants in degrees two through five are at most 146 and 26. Since $|PQ| \leq 9$ and $|\alpha_i|, |\beta_j| \leq 1$,

$$|H_p - 1| \leq C_\eta (Np)^{-1-2\eta}, \qquad C_\eta = \frac{172}{(1 - 7^{-1/2-\eta})^5}. \tag{59}$$

Indeed the five denominator factors have modulus at least $1 - |z|$; good prime ideal norms are at least seven. Ideal counting proves normal convergence, uniform in deleted primes and vertical phases, and

$$|\mathcal{H}_{\mathcal{T}}| \leq \exp\left(C_\eta \sum_p (Np)^{-1-2\eta}\right).$$

Local correction factors can vanish; no inverse-correction bound is asserted. Normal convergence controls neither the four numerator nor five reciprocal zeta factors in the critical region. Their poles, zeros and crossed residues cannot be suppressed.

39.3 Finite Perron heights and the exact eight-state numerator

Set $V = \lfloor U \rfloor + 1/2$ and $\Delta_V = \min_{n \geq 1} |\log(V/n)| \geq 1/(2V+1)$. For four finite vertical heights T_j on positive real lines c_j, let $I_T(s)$ be the fourfold integral of $P_aQ_b\mathcal{E}_{\mathcal{T}}(s; v, w)$ against $\prod_j V^{z_j}\, dz_j/(2\pi i z_j)$.

In the initial half-plane $\sigma = \Re s > 1$, finite-height Fubini is justified by the absolute placement majorant. Put

$$A_{\mathrm{abs}}(\sigma, c) = \prod_{p|a}(1 + q^{-c_1} + q^{-c_2}) \prod_{p|b}(1 + q^{-c_3} + q^{-c_4}) \\ \cdot \prod_{p \notin \mathcal{T}} \{1 + q^{-\sigma}(1 + q^{-c_1} + q^{-c_2})(1 + q^{-c_3} + q^{-c_4})\}, \quad q = Np.$$

Then

$$|I_T(s) - \mathcal{Z}_{U;a,b,\mathcal{T}}(s)| \le V^{\sum c_j} A_{\mathrm{abs}}(\sigma, c) \left\{ \prod_{j=1}^{4} \left(1 + \frac{2}{\pi T_j \Delta_V}\right) - 1 \right\}, \tag{60}$$

and $A_{\mathrm{abs}} \le 3^{\omega(a)+\omega(b)} \zeta_K(\sigma)^9$. For one short integer norm n, the indicator is at most $(V/n)^c$, and the truncated Perron error is at most $2(V/n)^c/(\pi T \Delta_V)$. Expand the difference of four products and sum every nonempty error subset to prove (60). Consequently successive Perron limits recover the exact cutoffs. The half-integer gap also gives a height-uniform $C_U(V/n)^c$ bound, allowing dominated convergence in the initial absolutely convergent series. Four infinite Perron integrals are not declared absolutely convergent or interchanged silently with Mellin inversion.

There is an entirely finite alternative. Partition $a = a_0 a_1 a_2$ and $b = b_0 b_1 b_2$, internally coprime in each column. Cross-column overlaps remain allowed. A common d prime occupies one state $(i, j) \in \{0, 1, 2\}^2$. Separate its $(0, 0)$ part e from eight pairwise coprime squarefree atomic ideals h_{ij} avoiding $\mathcal{T}$, and let h be their product. Define

$$h_i^L = \prod_{j=0}^{2} h_{ij} \ (i = 1, 2), \qquad h_j^R = \prod_{i=0}^{2} h_{ij} \ (j = 1, 2).$$

An atomic h_{ij} with both indices nonzero belongs to one short factor in each column. These four short factors must not be forced globally disjoint. Let $\mathcal{A}_U$ comprise precisely these tuples satisfying $N(a_i h_i^L) \le U$ and $N(b_j h_j^R) \le U$ for $i, j = 1, 2$. The numerator in (56) is exactly

$$\mathcal{N}_{U;a,b,\mathcal{T}}(s) = \sum_{\mathcal{A}_U} \frac{\mu(a_1 a_2)\mu(b_1 b_2)\mu(h_{11}h_{12}h_{21}h_{22})}{(Nh)^s} \prod_{p|h}(1 - (Np)^{-s})^{-1}. \tag{61}$$

At a common prime the sign is $(-1)^{1+\mathbf{1}_{i>0}+\mathbf{1}_{j>0}}$. It is negative for $(0, 0)$, positive for the four single-short states, and negative for the four double-short states. Summing the remaining squarefree e, coprime to $\mathcal{T}h$, gives $1/\zeta_K^{\mathcal{T}}(s)$ times the local product in (61). This proves the exact identity, including its sign.

Multiplying the four cutoff inequalities yields $Nh \le U^4$; double states are counted twice in that product. Hence the numerator is finite and holomorphic on $\Re s > 0$. At $U = 1$ it equals one: this boundary is not a high-window counterexample or a global pole obstruction. For $\sigma \ge \sigma_0 > 0$ and every $\epsilon > 0$,

$$|\mathcal{N}_{U;a,b,\mathcal{T}}(\sigma + i\tau)| \ll_{K,\sigma_0,\epsilon} (Na\, Nb)^{\epsilon} U^{4\max(1-\sigma,0)+\epsilon}. \tag{62}$$

There are at most $3^{\omega(a)+\omega(b)}$ residual partitions and $8^{\omega(h)}$ atomic placements at fixed h. The finite local product is $\ll_{\sigma_0,\epsilon} (Nh)^{\epsilon}$. Apply the ideal-divisor bound and ideal counting up to U^4, using smaller internal exponents; a logarithm at the transition is absorbed in U^{ϵ}. Uniformity as $\sigma_0 \to 0$ is not claimed. For $\sigma = 1/2 + \eta$ with $0 < \eta < 1/2$, the old-type cutoff cost is $U^{2-4\eta+\epsilon}$, not a new saving.

39.4 A uniform common-profile Mellin embedding of the actual mean

For residual a, b put $M = \max(Na, Nb)$ and $\varrho = \min(Na, Nb)/M$. At $\varrho \le 1/2$ the product of the two profile values vanishes, including equality because of endpoint vanishing. Otherwise put

$$\Phi_\varrho(t) = W(t)\overline{W(\varrho t)}$$

when $Na \ge Nb$, and use its conjugate for the reversed orientation. It is supported in $[1/\varrho, 2] \subset [1, 2]$. With $M_k(W) = \max_{j \le k} \|(t\, d/dt)^j W\|_\infty$,

$$|\widehat{\Phi}_\varrho(\sigma + i\tau)| \le C_{I,k} M_k(W)^2 (1 + |\tau|)^{-k} \tag{63}$$

uniformly for σ in any fixed compact strip I and all $\varrho \in [1/2, 1]$. Leibniz and log-coordinate integration by parts prove this bound; no ratio-dependent change of W is made. The transform is entire. For $c > 1$, ordinary Mellin inversion after the finite numerator identity gives the exact scalar (57) as

$$\frac{1}{2\pi i} \int_{\Re s = c} \widehat{\Phi}_\varrho(s) (D/M)^s \frac{\mathcal{N}_{U;a,b,\mathcal{T}}(s)}{\zeta_K^{\mathcal{T}}(s)}\, ds. \tag{64}$$

This integral is vertically absolutely convergent in its initial half-plane. The original profile forces $M \le 2D$, so the residual a, b and actual row sums are finite.

Keep the full local deletion factor

$$\frac{1}{\zeta_K^{\mathcal{T}}(s)} = \frac{1}{\zeta_K(s)} \prod_{p \in \mathcal{T}} (1 - (Np)^{-s})^{-1}.$$

At fixed $\sigma_0 > 0$ it costs $\ll_{S,\sigma_0,\epsilon} (N \operatorname{rad}(abu))^\epsilon$, hence at most $(HD^2)^\epsilon$ on actual support. This estimates the finite mask dependence, not the central reciprocal. Multiplying (64) by every original outer coefficient and summing the actual rows and residuals gives

$$DS_1^{AA} = \frac{1}{2\pi i} \int_{\Re s = c} \frac{D^s \mathcal{A}_{D,H,U,W}(s)}{\zeta_K(s)}\, ds. \tag{65}$$

Explicitly, with each profile orientation retained,

$$\mathcal{A}_{D,H,U,W}(s) = \sum_{u \in \mathcal{R}} \sum_{\substack{a,b \text{ squarefree,good} \\ Na, Nb \le 2D,\ \varrho > 1/2}} \nu(a)\overline{\nu(b)}\chi_a(u)\overline{\chi_b(u)}\, M^{-s}\, \widehat{\Phi}_\varrho(s) \cdot \mathcal{N}_{U;a,b,\mathcal{T}}(s) \prod_{p \in \mathcal{T}} (1 - (Np)^{-s})^{-1}.$$

This is a *finite coherent sum*, not an independent marginal moment or a positive sum. All raw zeros remain in its outer coefficients. For fixed finite D, H, U it is holomorphic on $\Re s > 0$ and rapidly decreasing on fixed vertical strips there, with constants depending on those finite parameters. No uniform target estimate for it is supplied. This is an exact signed embedding, not a zero-location theorem.

39.5 A nonzero short-state correction and the unproved step

An actual annular control uses good prime ideals of norms $7, 13, 19, 31, 43$ with primary generators $(1, 3), (4, 3), (-2, 3), (1, 6), (7, 6)$ in the $1, \omega$ basis. Take $D = 1500, U = 25$, $d = p_7$, $a = p_{13}p_{19}$, $b = p_{13}p_{31}$ and the nonprincipal row $u = p_{43}$ in the pure block $A_1 = 37, A_2 = 1, H = 241$, literal

class one. The full raw supports are disjoint from u, and $N(da) = 1729, N(db) = 2821$ lie in the annulus. A single fixed smooth profile with plateau $[1.1, 1.9]$ equals one at both nodes. Direct assignments give $A_U(a) = A_U(b) = -1$, $A_U(da) = 1$, $A_U(db) = -1$. The actual signed d coefficient is $+1$. The $(0,0)$ state alone would give -1; states $(1,0)$ and $(2,0)$ each contribute $+1$. Dropping all short states therefore misses a nonzero $e = 1$ correction. This illustrates the new eight-state identity, not a new discovery of coefficient nonmultiplicativity.

The original pair has gcd p_7p_{13}. Its contributions across $d = 1, p_7, p_{13}, p_7p_{13}$ cancel exactly. Thus a fixed-residual reciprocal does not prove an uncancelled pole or an obstruction in the complete coherent mean. The missing step is to estimate that full mean by $DH(HD)^\epsilon$, retaining fixed-profile/seminorm dependence, raw masks and conductor/ray/capacity restrictions. Scalar Mertens cancellation in the long state, a normal correction product, or an absolute cutoff numerator bound does not prove it. Principal and short assembly, independent positive determinant/zero-location bridges and separate Dirichlet, Hecke, Dedekind, elliptic and Artin obligations remain open. No holomorphy or desired zero location is assumed to close them.

39.6 Exact arithmetic and source-bound computational scope

The evaluator streams each actual row and accumulates original $x_m\chi_m(u)$ into every ideal-divisor bucket $d \mid m$. It then forms all nine cross products for A_U, C_U, μ and sums with $\mu(d)$. Write $n_{\mathrm{col}} = |\mathcal{C}|$, $I = \sum_m \tau_{\mathrm{ideal}}(m)$, K for the distinct divisors and r for the listed rows. This bucket stage costs $O(r(I + K))$ arithmetic operations with a fixed three channels. Its streaming storage is $O(K)$ plus incidences and either a streamed row or separately retained phase table. Character construction, factorization and integer bit lengths are additional costs, not declared free. For coefficients bounded by C, bucket coordinates are $O(n_{\mathrm{col}}C)$, products $O(n_{\mathrm{col}}^2C^2)$ and accumulated values $O(rn_{\mathrm{col}}^2C^2)$ for unit row weights. For compressed integer multiplicities replace r in this bit-size bound by $\sum_u w_u$ and account for weight multiplication.

For $D = 60000$, the actual phase table has shape 20×16953; the incidence has 73420 nonzeros over 21750 divisor buckets, and F_d has $20 \times 21750 \times 3$ complex states. There are 271688516 ordered coprime column pairs. The naive all-pair row loop would visit 5748084180 pairs and was *not executed* at this size. This is an operation count, not an unexecuted timing benchmark. Independent direct pair evaluation covers the two complete small families; the large family is checked linearly from literal field arithmetic and divisor-poset inversion.

The exact large results also include

$$S_1^{CC} = \frac{661}{7500}, \qquad S_1^{C\mu} = \frac{13931}{30000}, \qquad S_1^{\mu\mu} = -\frac{229}{1250}.$$

For the ordered vectors (A_U, C_U, μ) the full pairing table is

$$\frac{1}{30000}\begin{pmatrix} -30714 & -11287 & 19427 \\ -11287 & 2644 & 13931 \\ 19427 & 13931 & -5496 \end{pmatrix}.$$

This is the table of three original coefficient vectors, not a spectral carrier for GRH. The conjugate μC entry and the AA identity retain their actual linear-first convention. All three saved computations use source v1. A separate v2 helper fixes its unused $N = U = 1$ high-window assertion; every one of the 17158 executed columns has $N > U$, so no accepted row/table or source snapshot was silently replaced. The fixed-profile flat-node implementation refuses requests above its stated $D \le 60000$ certificate. Neither these three different finite blocks nor their measured runtimes imply an asymptotic saving.

40 A target-sized budget for rough-prime masks

The actual signed divisor mean admits a useful quantitative simplification: the rough-prime coprimality mask can be restored with an explicit correction bound. The remaining main term still retains the Möbius sign and the coherent row mean; it is not estimated by this result.

Work over $\mathbb{Z}[\omega]$. Fix a finite S containing the primes above $2, 3$, a unitary multiplicative ν on good ideals, and one common complex profile W supported on $[1, 2]$, smooth and zero at both endpoints. Put $B = \|W\|_\infty$ and assume

$$1 \le U, \qquad U^2 \le D, \qquad 2D \le 7U^2, \qquad H \ge 1. \tag{66}$$

Let $\mathcal{R}$ be the actual finite family of nonzero literal rows, with original weights $w_u \ge 0$, $1 \le Nu \le H$, and $R_w = \sum_u w_u$. All original internal squarefree/coprime conditions, unit and S-ray classes, capacity and sixth-factor masks remain in this family. Write S_1^{AA} for the raw form (51) with $x_m = A_U(m)\nu(m)W(Nm/D)$, where

$$A_U(g) = \sum_{\substack{xy \mid g,\ (x,y)=1 \\ Nx \le U,\ Ny \le U}} \mu(x)\mu(y). \tag{67}$$

Here the columns in (51) are the COMPLETE original good squarefree annulus $[D, 2D]$, with the common W as its norm cutoff and the original full-good-ideal/ν, S formulation. No selected-column indicator or additional column-ray filter is allowed in this reduction; all actual row unit/ray/capacity filters remain. Principal-row residues, if present, require their separate assembly. For the intended $A_1 > 1$ blocks the defining conductor of ν is supported in S, and the rows are nonprincipal.

Let f range over good squarefree U-smooth ideals, put $F = Nf$, and set

$$T_f(t; u) = \sum_{\substack{a \text{ good squarefree} \\ (a,f)=1}} A_U(fa)\nu(a)\chi_a(u)W(FtNa/D). \tag{68}$$

Here $\chi_a(u) = 0$ on $(a, u) > 1$. For a prime ideal p, let $T_f^{(p)}(t; u)$ impose the additional condition $(a, p) = 1$. The smooth factor f is distinct from the original column-pair gcd sector $h = 1$. Define the following *unnormalized* expressions:

$$K_0 = \sum_u w_u \sum_{\substack{f \text{ good SF, } U\text{-smooth} \\ (f,u)=1}} \mu(f)|T_f(1; u)|^2, \tag{69}$$

$$K_T = \sum_u w_u \sum_{\substack{f \text{ good SF, } U\text{-smooth} \\ (f,u)=1}} \mu(f) \sum_{\substack{p \notin S,\ Np > U \\ p \nmid u}} |T_f^{(p)}(Np; u)|^2, \tag{70}$$

$$K_S = \sum_u w_u \sum_{\substack{f \text{ good SF, } U\text{-smooth} \\ (f,u)=1}} \mu(f) \sum_{p \notin S,\ Np > U} |T_f(Np; u)|^2. \tag{71}$$

Throughout, p denotes a prime ideal. None of these expressions is declared positive, since the $\mu(f)$ signs remain.

The exact decomposition of the common ideal divisor $d = fe$ has at most one rough prime when the residual pair is nonunit. The unit residual pair also permits two distinct rough primes, necessarily with $f = 1$. Its exact correction is

$$C_2 = \sum_{\substack{\{p,q\} \text{ unordered, } p \ne q \\ p,q \notin S,\ Np, Nq > U}} |W(NpNq/D)|^2 \sum_{\substack{u \in \mathcal{R} \\ (u,pq)=1}} w_u. \tag{72}$$

Conjugate primes of equal rational norm are included. One has

$$DS_1^{AA} = K_0 - K_T + C_2. \tag{73}$$

In K_S, terms with $p \mid a$ or $p \mid u$ are explicitly auxiliary: fpa can fail to be squarefree. They use $A_U(fa)$ as written, not a new original column or a substitution for $A_U(fpa)$.

Theorem 19 (Explicit rough-prime mask budget)**.** *Put* $\kappa_U(u) = \#\{p \notin S : p \mid u,\ Np > U\}$ *and* $\Delta = (K_S - K_T)/D$. *Under* (66)*,*

$$|\Delta| \le B^2 \left(229376 \sum_u w_u \kappa_U(u) + 2048 R_w \right), \tag{74}$$

$$0 \le C_2/D \le 1792 B^2 R_w. \tag{75}$$

Consequently

$$S_1^{AA} = \frac{K_0 - K_S}{D} + E, \tag{76}$$

$$|E| \le B^2 \left(229376 \sum_u w_u \kappa_U(u) + 3840 R_w \right). \tag{77}$$

If $H \le D^{3/4}$*, then every row has* $\kappa_U(u) \le 1$*, and in particular*

$$|E| \le 233216 B^2 R_w \le 2^{18} B^2 R_w. \tag{78}$$

Only the rough-prime deletion has been restored; the smooth-divisor mask, literal row family and original coefficients remain. The explicit bound uses the exact restricted identity

$$A_U(g) = \mu(g)(2\mathbf{1}_{Ng \le U} - 1) \quad (Ng \le 7U,\ g \text{ good squarefree}).$$

Every active term of $T_f(Np; u)$ has $N(fa) < 7U$, hence $|A_U(fa)| = 1$ there. This is not a shrinkage or proxy for the whole original coefficient family. Only absolute correction bounds use it; the signed coefficients in the clean main remain unchanged.

For general H, one has $\kappa_U(u) \le 2\log(2H)/\log(2U)$, also at $U = 1$. For an actual two-pivot block with $R_w \ll A_1 A_2 \le H$ and $H \le D^{3/4}$, (78) gives an $O(B^2 H)$ error with no divisor-power or profile-derivative loss. The remaining target problem for this raw A–A slice is the clean signed main $(K_0 - K_S)/D$. It has not been bounded here; the 7/6 envelope and missing 5/12 improvement for the complete original remainder are unchanged. The argument neither deletes a coherent mean nor treats the coprime-truncated form as positive.

Three exact finite controls. These use principal ν, S above $2, 3$, literal unit class one and complete original good squarefree annuli. The first two row variables are squarefree, good and coprime and may include the unit ideal; all other factors equal one.

D	U	H	A_1	A_2	columns	rows
200	10	53	13	1	52	5
512	16	107	19	1	139	7
900	30	384	7	7	250	2

Put $\eta_0(t) = \exp(-1/t)$ for $t > 0$ and $\eta_0(t) = 0$ for $t \le 0$, and define

$$s(t) = \frac{\eta_0(t)}{\eta_0(t) + \eta_0(1-t)}.$$

Choose once $\delta = 1/4096$ and $W_\delta(t) = s((t-1)/\delta)s((2-t)/\delta)$. For integer $D \le 1024$, strict integer norm nodes have weight one and endpoints have weight zero; its derivatives cost $O_k(4096^k)$. The following numbers are RAW, before division by D:

D	K_0	K_S	K_T	E_{row}	E_{res}	C_2	DS_1^{AA}
200	339	235	209	20	6	21	151
512	366	781	711	48	22	89	−256
900	436	410	366	0	44	28	98

In every case $K_S - K_T = E_{\mathrm{row}} + E_{\mathrm{res}}$ and $DS_1^{AA} = K_0 - K_T + C_2$. The displayed nonnegative values of $E_{\mathrm{row}}, E_{\mathrm{res}}$ are finite observations; no general nonnegativity is assumed for either signed correction. Only C_2 is nonnegative by definition. The first two capacities are $\lfloor D^{3/4} \rfloor$. The third $H = 384$ is a declared noncritical fixture capacity, not $\lfloor 900^{3/4} \rfloor$; both of its rows have $\kappa_U = 0$. All ideal and raw-zero masks, unit residual states and explicitly auxiliary non-squarefree p^2 terms are retained. The exact control summary records the scopes and exact fractions after one division by D. These are finite identities, not a proof or numerical fit of the analytic bounds, and do not estimate the signed clean main.

41 Proof of the rough-prime mask budget

41.1 The common ideal divisor and the exact rough part

Here (51) uses the complete original good squarefree annulus and common W, without an additional selected-column or column-ray indicator. The actual row filters are unchanged. The identity $\mathbf{1}_{(m,n)=1} = \sum_{d \mid m,\ d \mid n} \mu(d)$ in (51) is an identity of ideal divisors. Write $m = da, n = db$. The residual ideals a, b are squarefree but may overlap; squarefreeness of the original columns requires $(d, ab) = 1$. On $(d, u) = 1$ the original common phases satisfy

$$|\nu(d)|^2 \chi_d(u)\overline{\chi_d(u)} = 1.$$

On $(d, u) > 1$ the product is the original raw zero. Hence the full deleted set is $\mathcal{T} = S \cup \mathrm{supp}(abu)$, including residual primes where $\chi_a\overline{\chi_b}$ might become imprimitive. With

$$F_{a,b}(x) = W(xNa/D)\overline{W(xNb/D)}, \qquad M = \max(Na, Nb), \quad \alpha = \min(Na, Nb),$$

the inner signed sum is

$$\sum_{\substack{d \ \mathrm{SF} \\ (d,\mathcal{T})=1}} \mu(d) A_U(da) A_U(db) F_{a,b}(Nd). \tag{79}$$

If its profile is nonzero, then

$$D/\alpha < Nd < 2D/M. \tag{80}$$

All endpoints have zero profile value. In particular the profile is identically zero if $\alpha \le M/2$. No independent marginal row family or changed profile has been substituted.

Factor $d = fe$ uniquely, with f squarefree U-smooth and e squarefree with every prime norm greater than U. Whenever ez is squarefree,

$$A_U(ez) = A_U(z). \tag{81}$$

Indeed neither short divisor in (67) can contain a prime of e; deleting e preserves the two cutoffs and both Möbius signs in a bijection of the original assignments. There is no factor $\mu(e)$ in (81). This invariance is specific to A_U; it does not hold unchanged for the other coefficient channels. The values $A_U(fa), A_U(fb)$ are not factorized.

If $M > 1$, goodness gives $M \geq 7$. From (66) and (80), $Nd \leq 2D/M \leq U^2$. Thus e is either 1 or one prime ideal, and (79) is exactly

$$\sum_{\substack{f \text{ good SF, } U\text{-smooth} \\ (f,\mathcal{T})=1}} \mu(f) A_U(fa) A_U(fb) \left(F_{a,b}(Nf) - \sum_{\substack{p \notin \mathcal{T} \\ Np > U}} F_{a,b}(NfNp) \right). \tag{82}$$

The rough part here belongs to the *common divisor* d; the full columns can have additional rough primes in a or b.

If $M = 1$, then $a = b = 1$. Three rough prime ideals would have norm product strictly greater than $7U^2 \geq 2D$: two contribute more than U^2, and the third good prime has norm at least 7. Two rough primes are possible, but then a nonunit f would similarly give $Nf\,Ne > 7U^2 \geq 2D$. Hence that case has $f = 1$ only. Its sign is $\mu(pq) = 1$ and its coefficient is $A_U(pq) = A_U(1) = 1$. Summing this case over the unchanged literal rows gives (72); summing all other cases of (82) over a, b gives the completed squares in (69)–(70). This proves (73). The unit ideal among residuals and conjugate split primes must both be retained. The condition $f = 1$ is not a termwise assertion that the original pair has gcd 1: the Möbius expansion cancels noncoprime original pairs across its divisor sectors.

41.2 An exact disjoint error identity

We restore only $(p, abuS) = 1$ to $p \notin S$. In particular the conditions $(f, abuS) = 1$ and every raw $\chi_a(u)$ remain in the two residual sums. The full deleted union has the disjoint partition

$$\mathbf{1}_{p|abu} = \mathbf{1}_{p|u} + \mathbf{1}_{p\nmid u}\mathbf{1}_{p|a \text{ or } p|b}. \tag{83}$$

Therefore $K_S - K_T = E_{\text{row}} + E_{\text{res}}$, where the errors are *unnormalized*:

$$E_{\text{row}} = \sum_u w_u \sum_{\substack{f \text{ good SF, } U\text{-smooth} \\ (f,u)=1}} \mu(f) \sum_{\substack{p \notin S,\ Np > U \\ p|u}} |T_f(Np; u)|^2, \tag{84}$$

$$E_{\text{res}} = \sum_u w_u \sum_{\substack{f \text{ good SF, } U\text{-smooth} \\ (f,u)=1}} \mu(f) \sum_{\substack{p \notin S,\ Np > U \\ p\nmid u}} \left(|T_f(Np; u)|^2 - |T_f^{(p)}(Np; u)|^2 \right). \tag{85}$$

The two errors are signed; their bounds below use absolute values, and neither is assumed nonnegative. At $p \mid u$, all residuals divisible by p already have raw character zero. This is consistent with the full square in (84); it does not make that square zero. These formulas keep the complete row/residual union without double counting.

Suppose a nonzero omitted residual has $p \mid a$, and write $a = pc$. Its profile requires $Nf\,(Np)^2 Nc < 2D$. If fc is nonunit, then

$$Nf\,(Np)^2 Nc > 7U^2 \geq 2D,$$

a contradiction. Thus necessarily $f = c = 1$ and $a = p$. This is a support statement about the explicitly auxiliary extension; the formal column $fpa = p^2$ is not an original squarefree column. With $P = Np$, set

$$z_p(u) = \nu(p)\chi_p(u) W(P^2/D), \qquad T_p(u) = T_1(P; u). \tag{86}$$

Here $A_U(p) = 1$, since $P > U$. The difference $T_f - T_f^{(p)}$ vanishes unless $f = 1$, when it equals z_p. Consequently

$$E_{\mathrm{res}} = \sum_u w_u \sum_{\substack{p \nmid S,\ Np > U \\ p \nmid u}} \left(2\,\mathrm{Re}(T_p(u)\overline{z_p(u)}) - |z_p(u)|^2 \right). \tag{87}$$

The subtraction $-|z_p|^2$ is the exact intersection $p \mid a$ and $p \mid b$; omitting it would count that intersection twice. Complex profiles and all remaining literal phases are allowed. If $z_p \neq 0$, then

$$\sqrt{D} < P < \sqrt{2D}. \tag{88}$$

The condition $P > U$ is automatic here from $U^2 \leq D$.

41.3 Elementary ideal and coefficient estimates

We give all counting inputs; no prime distribution theorem or Mertens estimate is needed. The Eisenstein norm is Euclidean. After rounding both real coordinates of a complex quotient in the basis $1, \omega$, the error $s + t\omega$ has $|s|, |t| \leq 1/2$ and norm at most $3/4 < 1$. A least-norm nonzero element of any ideal therefore generates it by Euclidean division. Multiplication by $a + b\omega$ has determinant $a^2 - ab + b^2$, giving the principal ideal norm. Since $a^2 - ab + b^2 \geq (a^2 + b^2)/2$, counting all possible generators, with harmless overcounting by units, proves for $X \geq 1$ that

$$\#\{a : Na \leq X\} \leq (2\sqrt{2X} + 1)^2 < 16X. \tag{89}$$

An annulus $L < Na < 2L$ therefore contains at most $32L$ ideals when $2L > 1$. If $2L \leq 1$ its profile contribution is zero, including the unit endpoint when $2L = 1$. This proves the needed uniform bound even when $1/2 < L < 1$; the unit ideal must not be dropped in that range.

There is an exact coefficient identity on the restricted support:

$$A_U(g) = \mu(g)\big(2\mathbf{1}_{Ng \leq U} - 1\big) \quad (g \text{ good squarefree},\ Ng \leq 7U). \tag{90}$$

Every proper divisor of a nonunit good g has norm at most $Ng/7 \leq U$. If $Ng \leq U$, all ternary assignments are allowed and their sum is $\prod_{p \mid g}(1 - 1 - 1) = \mu(g)$. If $Ng > U$, only the assignments $(x, y) = (g, 1), (1, g)$ are omitted; each has sign $\mu(g)$, leaving $-\mu(g)$. The unit ideal gives $A_U(1) = 1$ and is included. No approximation or multiplicativity of the cutoff coefficient is used.

In an active term of (68), $P = Np > U$ gives

$$N(fa) < 2D/P < 2D/U \leq 7U.$$

Therefore $|A_U(fa)| = 1$ exactly there. This statement concerns only the rough-prime support, not the entire original coefficient family. With $L = D/(FP)$, the ideal count and $|\chi_a(u)| \leq 1$ now give

$$|T_f(P; u)| \leq 32B\,\frac{D}{FP}. \tag{91}$$

The same statement holds if the sum is empty, and covers $L < 1$ by the preceding endpoint discussion. A dyadic use of (89) gives

$$\sum_f \frac{1}{(Nf)^2} \leq \sum_{j \geq 0} 16\,2^{j+1} 2^{-2j} = 64, \tag{92}$$

where the left side may even include all ideals. The smooth squarefree restriction only decreases this bound.

41.4 The normalized row-mask budget

Apply triangle only to (84), use (91), and then (92). The raw square contributes D^2/P^2; dividing by D once gives

$$\frac{|E_{\mathrm{row}}|}{D} \le 65536B^2 \sum_u w_u \sum_{\substack{p\mid u,\ p\notin S\\ Np>U}} \frac{D}{(Np)^2}. \tag{93}$$

By (66), $D/(Np)^2 \le D/U^2 \le 7/2$. Let $\kappa_U(u)$ count the distinct good rough prime ideals dividing u. Each $P = Np$ is at least 7 and greater than U, so $P^2 > 7U \ge 2U$. The complete ideal norm of the literal row therefore gives

$$\kappa_U(u) \le \frac{2\log(Nu)}{\log(2U)} \le \frac{2\log(2H)}{\log(2U)}. \tag{94}$$

The denominator is positive also at $U = 1$; unit rows have $\kappa_U = 0$. All sixth-factor divisors remain in this bound, and no rational-norm gcd is substituted for an ideal mask. Using $D/(Np)^2 \le 7/2$ in (93) gives the sharper form

$$\frac{|E_{\mathrm{row}}|}{D} \le 229376B^2 \sum_u w_u \kappa_U(u). \tag{95}$$

The logarithmic bound (94) is available for arbitrary H.

41.5 The residual intersection budget

Use (87), $|z_p| \le B$, and (91) at $f = 1$. After the single division by D,

$$\frac{|E_{\mathrm{res}}|}{D} \le B^2 R_w \sum_{\sqrt{D}<Np<\sqrt{2D}} \left(\frac{64}{Np} + \frac{1}{D}\right). \tag{96}$$

Here absolute values bound only the correction; the exact identity (87) retains its original phases. Since $D \ge 1$ and $Np \le \sqrt{2D} \le 2D$, one has $1/D \le 2/(Np)$. Also

$$\sum_{\sqrt{D}<Np<\sqrt{2D}} \frac{1}{Np} \le 16\sqrt{2} \tag{97}$$

by (89). Thus

$$|E_{\mathrm{res}}|/D \le 1056\sqrt{2}\, B^2 R_w \le 2048B^2 R_w.$$

Together with (95), this proves (74).

41.6 The exact unit-residual correction and the clean main

In (72), each contributing pair has $P, Q > U$ and $PQ \le 2D$. Hence both prime norms are less than $2D/U$. There are at most $32D/U$ such prime ideals by (89), so the number of unordered distinct pairs is at most $512D^2/U^2$. Only inside this *nonnegative* expression enlarge the masked row count to R_w and the profile to B^2. Thus

$$0 \le \frac{C_2}{D} \le 512\frac{D}{U^2} B^2 R_w \le 1792B^2 R_w,$$

proving (75). Conjugate primes of equal rational norm remain included; no rational-squarefree restriction or new diagonal exclusion is imposed. This is a bookkeeping bound for the exact

correction, not an oscillatory saving or a lower bound for the full signed slice. Other residual sectors can cancel it.

Now (73) gives (76) with $E = \Delta + C_2/D$, and the explicit bounds give (77). Suppose $H \le D^{3/4}$. If two distinct good rough primes divided a row, then $H \ge Nu \ge (Np)(Nq) \ge 49$, but also

$$H \ge Nu \ge (Np)(Nq) > U^2 \ge 2D/7.$$

Combining with $H \le D^{3/4}$ gives $D^{1/4} < 7/2$ and hence $H < (7/2)^3 = 343/8 < 49$, a contradiction. Thus $\kappa_U(u) \le 1$ at every scale in this capacity range. Substitution in (77) gives $|E| \le 233216B^2R_w \le 2^{18}B^2R_w$, proving (78). This is an exact endpoint argument, not an asymptotic approximation.

For a two-pivot block, counting each of its two dyadic ideal variables by (89) gives $R_w \ll A_1A_2$ when the original fixed unit classes and multiplicities are bounded. Internal coprimality and capacity reduce that count; grouping literal fibers preserves total multiplicity. Under $A_1A_2 \le H \le D^{3/4}$ the correction is therefore $O(B^2H)$. The exact local coefficient identity and these correction bounds do not estimate the remaining signed main $(K_0 - K_S)/D$. Its original $A_U(fa)$, coherent rows, common fixed profile and Möbius mean remain the analytic obligation; no new global exponent follows.

41.7 Finite control scope

The tables in the shared core retain complete original squarefree ideal columns and raw character zeros. Their residual-prime extension is a separate auxiliary sum, including the non-squarefree p^2 terms needed for (87); those are not original columns. The exact intersection subtraction and unit residual states with $1/2 < L < 1$ are checked separately. The common profile has fixed $\delta = 1/4096$ and large fixed derivative constants. The capacities and row blocks differ between cases, and the $H = 384$ case is outside the critical capacity range. These finite equalities neither determine the analytic constants nor imply an asymptotic saving.

41.8 A coherent prime-substitution budget, and its remaining cost

This subsection requires $U \ge 7$. It does not extend the prime-density argument to $U = 1$, although the coefficient and mask lemmas above hold for all $U \ge 1$. For real $t \ge U$, put

$$G_f(t) = \sum_{\substack{u \in \mathcal{R} \\ (f,u)=1}} w_u |T_f(t;u)|^2, \qquad M_1(W) = \max\{B, \|xW'(x)\|_\infty\}.$$

The active residuals satisfy $Na \le 2D/(FU)$, so their union is finite; G_f is smooth and vanishes for $t \ge 2D/F$. Since the rough prime in K_S is unmasked, it enters only this radial common profile. Let

$$E_S(t) = \pi_K^S(t) - \mathrm{Li}(t), \qquad \delta_S(U) = \sup_{t \ge U} \frac{|E_S(t)|}{t}, \qquad \mathcal{J} = \sum_f \mu(f) \int_U^\infty \frac{G_f(t)}{\log t}\, dt,$$

where π_K^S counts prime *ideals* outside the fixed finite S with multiplicity, including inert prime norms and distinct conjugate ideals. Stieltjes partial summation gives the exact coherent identity

$$\begin{aligned} K_S &= \mathcal{J} + \mathcal{E}_{\mathrm{PIT}}, \\ \mathcal{E}_{\mathrm{PIT}} &= \sum_f \mu(f) \left(-G_f(U)E_S(U) - \int_U^\infty E_S(t) G_f'(t)\, dt \right). \end{aligned} \tag{98}$$

The lower term uses $\pi_K^S(U)$, so the condition $Np > U$ is exact even when a prime norm equals U. A conjugate pair contributes a jump of two. This equality precedes any absolute sum over f.

Let $\lambda(L)$ be a TRUE arbitrary-coefficient squared norm bound for the specified positive row operator on good squarefree ideals of norm in $[L, 2L]$, including the actual weights. For $1/2 \le L < 1$, set $\lambda_*(L) = R_w$ by the direct unit-column bound; for $L \ge 1$, use $\lambda_*(L) = \lambda(L)$. No columns are active for $L < 1/2$. With $L = D/(Ft)$, differentiate only the common W: multiplication by $t\,d/dt$ replaces $W(x)$ by $xW'(x)$. The exact restricted identity (90), ideal counting and weighted row Cauchy give

$$G_f(t) + |tG_f'(t)| \ll M_1(W)^2 \lambda_*(L) L.$$

Any enlargement of a common f-row mask here occurs only inside these positive norm estimates. Substitution in (98), with $L_f = D/(FU)$, proves

$$\frac{|\mathcal{E}_{\mathrm{PIT}}|}{D} \ll M_1(W)^2 \delta_S(U) \sum_{\substack{f \text{ good SF, } U\text{-smooth} \\ F < 2D/U}} \frac{1}{F} \left(\lambda_*(L_f) + \int_{1/2}^{L_f} \lambda_*(L)\, \frac{dL}{L} \right). \tag{99}$$

The original Möbius sign remains in K_S and $\mathcal{J}$ themselves; the absolute f sum appears only in this error majorant.

Specialize now ONLY to the established positive input for the actual $A_6 = 1$ unweighted two-pivot block, or bounded natural literal weights for which the same input has been established. Put

$$R = A_1 A_2, \qquad M_{\mathrm{row}} = \max(A_1, A_2), \qquad A_1 A_2^2 \le H.$$

At $L \ge 1$ that input is

$$\lambda(L) \ll_\epsilon (HD)^\epsilon \min \left\{ R + \frac{LR}{M_{\mathrm{row}}} + \frac{L^{2/3} R}{M_{\mathrm{row}}^{1/3}},\ L + R^2 \right\}. \tag{100}$$

It is not asserted for arbitrary nonnegative weights, enlarged primitive conductor families or arbitrary sixth-factor enlargements. Using its first branch in (99) gives

$$\frac{|\mathcal{E}_{\mathrm{PIT}}|}{D} \ll_\epsilon (HD)^\epsilon M_1(W)^2 \delta_S(U) \mathcal{B}_{\mathrm{op}}, \tag{101}$$

$$\mathcal{B}_{\mathrm{op}} = R \log^2(2D) + \frac{D}{U} \frac{R}{M_{\mathrm{row}}} + \left(\frac{D}{U} \right)^{2/3} \frac{R}{M_{\mathrm{row}}^{1/3}}. \tag{102}$$

The logarithmic term uses $\sum_{Nf \le X} (Nf)^{-1} \ll \log(2X)$; the other terms use convergent ideal sums of powers 2 and 5/3. The same positive bounds give $|\mathcal{J}|/D \ll_\epsilon (HD)^\epsilon B^2 \mathcal{B}_{\mathrm{op}} / \log U$. These are inherited majorants, not a new signed gain or an estimate of $K_0 - \mathcal{J}$.

Capacity implies

$$\mathcal{B}_{\mathrm{op}} \ll H \log^2(2D) + (D/U) H^{1/3} + (D/U)^{2/3} H^{2/3}.$$

At $H = D^{3/4}$ and $U \asymp \sqrt{D}$, these terms have powers $3/4, 3/4, 5/6$. The last power occurs in this majorant already on pure scales $A_1 \asymp H, A_2 = 1$ with fixed capacity slack. It is not an actual-error lower bound.

For the fixed field, Thorner–Zaman, *Algebra & Number Theory* 13 (2019), 1039–1068, Theorem 1.1, applied to the trivial extension K/K with absolute degree $[K : \mathbb{Q}] = 2$ and absolute discriminant 3, supplies for sufficiently large U a bound

$$\delta_S(U) \ll_{K,S} \exp(-c_K \sqrt{\log U}).$$

The possible exceptional real zero is fixed: retain its secondary $\mathrm{Li}(x^\beta)$ term, or absorb it only into fixed-field constants and the threshold. Finite S deletions are included. No prime theorem uniform in the growing deleted set or row conductor is used. The DOI, publisher URL and pinned reading-cache hashes are in the public citation record.

A fixed-field subpower prime error does not make the available $D^{5/6+\epsilon}\delta_S(U)$ majorant certify the target $D^{3/4+\epsilon}$ for every $\epsilon > 0$. Actual signed errors might cancel. Even an exact prime density would leave the coherent signed $K_0 - \mathcal{J}$ unestimated. The elementary correction theorem uses only B, whereas the prime-substitution error requires $M_1(W)^2$. The complete original remainder's $7/6$ envelope remains unchanged.

42 The coherent signed main and its remaining inputs

The physical sixth-power-free comparison in Section 44 provides a new positive upper bound. We now collect the exact coherent reductions that it can enter. They preserve the original mean and all raw masks; none proves the target by a scalar cancellation assertion.

Use the complete original good squarefree annulus and common complex profile of Section 40. Write $\mathcal{U}$ for the actual nonzero literal rows, $w_u \ge 0$, $R_w = \sum_u w_u$ and $B = \|W\|_\infty$. Retain fixed S, ν, all internal row masks and (66). Every identity below is a raw A–A identity; principal residues still require their separate assembly. Classical operator specializations are only for the original unweighted $A_6 = 1$ rows, or their already justified bounded natural multiplicities.

First there is a useful distinction between two auxiliary kernels. Let K_{all} be K_S of (71) with the U-smooth restriction on its squarefree f removed, retaining $(f,u) = 1$ and exactly $A_U(fa)$ in every T_f. Set

$$\begin{aligned} P_2 &= \sum_u w_u \sum_{\substack{p\notin S\\ Np>U}} \sum_{\substack{q\notin S,\ Nq>U\\ q\nmid u}} |W(NpNq/D)|^2, \\ K_S &= K_{\mathrm{all}} + P_2, \qquad 0 \le P_2/D \le 3584 B^2 R_w. \end{aligned} \tag{103}$$

Here (p,q) is ordered, $p = q$ is allowed, and only q has the stated row mask. This is distinct from the original unordered, distinct-prime C_2 in (72). In particular

$$S_1^{AA} = \frac{K_0 - K_{\mathrm{all}}}{D} + \Delta + \frac{C_2 - P_2}{D}. \tag{104}$$

At $H \le D^{3/4}$ the error in this *new* auxiliary main is at most $236800 B^2 R_w$. The old $233216 B^2 R_w$ bound for $(K_0 - K_S)/D$ is unchanged. The physical comparison in (126) estimates the new combined main, in its stated $A_6 = 1$ scope, without estimating its two summands separately.

The prime-discrepancy clauses below additionally require $U \ge 7$. Define

$$E_S(t) = \pi_K^S(t) - \mathrm{Li}_2(t), \qquad \mathrm{Li}_2(t) = \int_2^t \frac{dv}{\log v}, \qquad M_1(W) = \max\{B, \|xW'(x)\|_\infty\}.$$

The count is of prime ideals outside fixed S, including distinct conjugate ideals and inert prime norms. Put

$$Q(t) = \sum_{\substack{f \text{ good SF}\\ f\ U\text{-smooth}}} \mu(f) \sum_{\substack{u\in\mathcal{U}\\ (f,u)=1}} w_u |T_f(t;u)|^2, \qquad \mathcal{J} = \int_U^\infty \frac{Q(t)}{\log t}\, dt.$$

Proposition 20 (A joint form retaining its mean)**.** *On the finite space of good squarefree g with $Ng \le 2D/U$, let*

$$b_g(t) = A_U(g)\nu(g)W(Ng\,t/D), \quad (\mathsf{R}c)(u) = \sqrt{w_u}\sum_g c_g\chi_g(u), \quad (\mathsf{P}_f c)_g = \mathbf{1}_{f|g}c_g.$$

With ordinary counting ℓ^2 norms,

$$Q(t) = b(t)^*\mathcal{M}_U b(t), \qquad \mathcal{M}_U = \sum_{f\ U\text{-smooth}} \mu(f)\mathsf{P}_f\mathsf{R}^*\mathsf{R}\mathsf{P}_f, \quad t \ge U. \tag{105}$$

The matrix is self-adjoint, with no positivity assertion. The exact prime substitution is

$$K_S = \mathcal{J} + \mathcal{E}_{\mathrm{PIT}}, \qquad \mathcal{E}_{\mathrm{PIT}} = -E_S(U)Q(U) - \int_U^\infty E_S(t)Q'(t)\,dt. \tag{106}$$

The endpoint excludes $Np = U$ exactly. It must remain.

For any true arbitrary-coefficient row norm $\lambda_*(L)$, use the direct unit-column value R_w on $1/2 \le L < 1$. On the active range $1 \le L \le D/U$, the current original two-pivot bound is the minimum in Proposition 26; it is not a theorem for arbitrary row weights. Let $\lambda^{\#}(D/T)$ dominate this norm on the two adjacent active length dyads. Then

$$\frac{|\mathcal{E}_{\mathrm{PIT}}|}{D} \ll_\epsilon (HD)^\epsilon M_1(W)^2\left\{\frac{|E_S(U)|}{U}\lambda_*(D/U) + \sum_T \mathfrak{d}_2(T)\lambda^{\#}(D/T)\right\}, \tag{107}$$

where $T = 2^jU < 2D$ and $\mathfrak{d}_2(T)^2 = \int_T^{2T}|E_S(t)|^2t^{-3}\,dt$. This is a sufficient-input budget, not an asserted power estimate for $\mathfrak{d}_2$ or for the endpoint. The pointwise version retains $\lambda_*(D/U) + \int_{1/2}^{D/U}\lambda_*(L)\,dL/L$, with the current minimum inside that integral.

On good rough-prime columns the exact quadratic-form compression is $\Pi_{\mathcal{P}}\mathcal{M}_U\Pi_{\mathcal{P}} = \Pi_{\mathcal{P}}\mathsf{R}^*\mathsf{R}\Pi_{\mathcal{P}}$. This is not an invariant-subspace assertion: off-block couplings remain and can cancel its contribution. Radial Mellin Parseval likewise retains equal-norm ideal correlations, including their derivatives; it neither orthogonalizes orientations nor permits deleting their mean. Even exact prime density leaves $K_0 - \mathcal{J}$ to estimate.

Proposition 21 (A sufficient smooth-divisor tail cutoff)**.** *For the original unweighted $A_6 = 1$ two-pivot rows put $R = A_1A_2$, $M_{\mathrm{row}} = \max(A_1, A_2)$. Restrict the smooth f sums in K_0, K_S to $Nf > Y \ge 1$. The inherited positive input gives*

$$\begin{aligned}\frac{|K_0^{>Y}|}{D} &\ll_\epsilon (HD)^\epsilon B^2\left\{R + \frac{DR}{M_{\mathrm{row}}Y} + \frac{D^{2/3}R}{M_{\mathrm{row}}^{1/3}Y^{2/3}}\right\},\\ K_S^{>Y} &= 0 \qquad (Y \ge 2D/U).\end{aligned} \tag{108}$$

If $R \ll H$, a sufficient target-size tail cutoff is

$$Y \ge \max\left\{1, \frac{2D}{U}, \frac{DR}{M_{\mathrm{row}}H}, \frac{DR^{3/2}}{M_{\mathrm{row}}^{1/2}H^{3/2}}\right\}.$$

This is a coarse inherited sufficient bound, not signed cancellation or an optimal localization. The smooth incidence divisor f is not the original column-pair gcd h.

Theorem 22 (A conditional input diagnostic)**.** *Fix S and $0 < \alpha < 1/2$. If for every $\epsilon > 0$,*

$$\mathfrak{d}_2(T) \ll_{\epsilon,K,S} T^{-\alpha+\epsilon}$$

holds for all sufficiently large real T, then

$$\zeta_K(s) \neq 0 \qquad (\operatorname{Re} s > \max\{1-\alpha, 1/2\},\ s \neq 1).$$

This is a strict-region implication from an explicitly hypothetical input, not a new zero-free theorem or a necessary condition for the actual signed target.

With the current physical input the pure witness has the sufficient PIT majorant 19/24, and the mixed witness retains the smaller 7/9 from the minimum. Independent scalar powers $U^{-1/12}$ and $U^{-1/18}$ would suffice for those majorants, with compatible endpoint control; their all-scale mean-square versions would entail the strict strips 11/12 and 17/18 above. None is asserted available. The old pure 5/6 budget and $U^{-1/6}$ input, with conditional strip 5/6, are historical values for the former operator bound. Proposition 24 gives the current global 9/8 comparison, and Proposition 25 the scoped $A_6 = 1$ comparison 13/12. The target, principal/short assembly beyond the stated scopes, independent determinant/zero-location bridges and full analytic/category Lean formalization remain open.

43 Proofs for the coherent main reductions

Throughout, squarefree and prime mean ideal squarefree and prime. The original symbol is extended by zero, the complete good annulus is unchanged, and the same complex W is used for every row. The common smooth divisor f is an incidence variable after ideal Möbius inversion; it is not the original column-pair gcd sector $h = 1$.

43.1 Removing only the auxiliary smooth restriction

The high-window support of a rough-prime term satisfies

$$N(fa), N(fb) < 2D/U \leq 7U.$$

The accepted exact coefficient identity (90) therefore applies to these squarefree products, including the unit:

$$A_U(g) = \mu(g)(2\mathbf{1}_{Ng\leq U} - 1), \qquad Ng \leq 7U. \tag{109}$$

For completeness, every proper divisor of a nonunit good g has norm at most $Ng/7 \leq U$. The unrestricted ordered disjoint-divisor sum is $\prod_{p|g}(1-1-1) = \mu(g)$. If $Ng > U$, precisely the two assignments $(g,1),(1,g)$ are omitted, each with weight $\mu(g)$. If $Ng \leq U$ none is omitted; $g = 1$ is direct. This does not simplify the general annular K_0 coefficients, whose norms can reach $2D$.

For a nonunit residual pair put $M = \max(Na, Nb) \geq 7$. On active rough-prime support, with $F = Nf$ and $P = Np > U$, $FMP < 2D \leq 7U^2$, hence $F < U$. Such f is automatically U-smooth. Thus dropping its smoothness changes only $a = b = 1$. There $F < 2D/P < 7U$. If a squarefree good f contains a prime q with $Nq > U$, any additional nonunit factor would give $F > 7U$. The omitted nonsmooth f is consequently exactly q, with $\mu(q) = -1$ and $|A_U(q)| = 1$. This proves the sign and all masks in (103). The outer p was already deliberately unmasked in K_S; only $q \nmid u$ comes from the original f-row mask. In particular $p = q$ gives a permitted auxiliary p^2 term, not an original squarefree column.

A nonzero pair has both prime norms below $2D/U$. Using (89), the number of ordered pairs is at most $(32D/U)^2$. Therefore

$$P_2/D \le 1024B^2R_wD/U^2 \le 3584B^2R_w.$$

Substitute $K_S = K_{\mathrm{all}} + P_2$ into (76). The old critical error $233216B^2R_w$ and the new charged $3584B^2R_w$ give $236800B^2R_w$ in (104). The original C_2 remains unordered and distinct, with both row masks; it is not this ordered correction. No lower bound for the full indefinite raw slice is inferred from either nonnegative correction.

43.2 An exact two-jump masked Mertens functional

Expand the nonunit residual part of K_S before any absolute sum in f, a conductor or a row. Retain

$$\mathcal{T} = S \cup \operatorname{supp}(abu), \qquad \eta_{a,b}(u) = \nu(a)\overline{\nu(b)}\chi_a(u)\overline{\chi_b(u)}.$$

The ideals a, b are individually squarefree but may overlap. Every common support prime remains in $\mathcal{T}$, even if their quotient character loses that prime. The rough p in K_S may belong to $\mathcal{T}$; this is its defined auxiliary enlargement. Only the common f phases cancel, leaving the full $(f, \mathcal{T}) = 1$ mask. By (109), the remaining coefficient is

$$\mu(f)A_U(fa)A_U(fb) = \mu(f)\mu(a)\mu(b)(2\mathbf{1}_{FNa\le U} - 1)(2\mathbf{1}_{FNb\le U} - 1).$$

Put $\alpha = \min(Na, Nb)$, $M = \max(Na, Nb)$ and

$$X = \frac{D}{PM}, \quad r = \frac{\alpha}{M}, \quad c_0 = \frac{U}{M}, \quad d_0 = \frac{U}{\alpha}.$$

At $r \le 1/2$ the profile product vanishes, including equality. Otherwise let $\Phi_r(v) = W(v)\overline{W(rv)}$ when $Na = M$, and use its conjugate in the reversed orientation. Set $h(t) = \Phi_r(t/X)$, supported on $[X/r, 2X]$ with zero endpoints. The exact cutoff product is

$$(2\mathbf{1}_{FNa\le U} - 1)(2\mathbf{1}_{FNb\le U} - 1) = 1 - 2\mathbf{1}_{c_0<F\le d_0}.$$

No half-weight occurs at a norm cutoff. With the closed-prefix sum

$$M_{\mathcal{T}}(y) = \sum_{\substack{Nf\le y\\(f,\mathcal{T})=1}} \mu(f), \qquad M_{\mathcal{T}}(y) = 0 \quad (y < 1),$$

the scalar factor, multiplied outside by $\mu(a)\mu(b)\eta_{a,b}(u)$, is

$$\mathcal{I}_{a,b,u}(P) = \sum_{(f,\mathcal{T})=1} \mu(f)h(Nf)(1 - 2\mathbf{1}_{c_0<Nf\le d_0}). \tag{110}$$

Ordinary finite Abel summation gives, for every real y,

$$\sum_{\substack{Nf\le y\\(f,\mathcal{T})=1}} \mu(f)h(Nf) = h(y)M_{\mathcal{T}}(y) - \int_0^y M_{\mathcal{T}}(t)h'(t)\,dt.$$

Subtracting the two closed prefixes proves the endpoint-sensitive identity

$$\begin{aligned}\mathcal{I}_{a,b,u}(P) = &-\int_0^\infty M_{\mathcal{T}}(t)h'(t)\,dt + 2\int_{c_0}^{d_0} M_{\mathcal{T}}(t)h'(t)\,dt\\ &+ 2h(c_0)M_{\mathcal{T}}(c_0) - 2h(d_0)M_{\mathcal{T}}(d_0).\end{aligned} \tag{111}$$

When $c_0 = d_0$ all extra terms cancel. A boundary outside the profile has $h = 0$. Combining the two integrals gives a multiplier of absolute value one almost everywhere, rather than three separate derivative costs. Log-coordinate differentiation yields

$$\int |h'(t)|\,dt \le 2\log 2\, M_1(W)^2, \qquad 2|h(c_0)| + 2|h(d_0)| \le 4B^2.$$

Consequently

$$|\mathcal{I}_{a,b,u}(P)| \le 6M_1(W)^2 \sup_{X\le t\le 2X} |M_{\mathcal{T}}(t)|. \tag{112}$$

The full outer residual and row sum is still coherent.

A useful explicit uniformity quantity is

$$m_Q(Z) = \sup_{\substack{E \text{ finite good prime ideals} \\ N\prod_{p\in E} p\le Q}} \sup_{t\ge Z} \frac{|M_{S\cup E}(t)|}{t}.$$

It always satisfies $m_Q(Z) \le 16$ by (89). For the original residuals,

$$N\,\mathrm{rad}(abu)_{S^c} \le Na\,Nb\,Nu \le (2D/P)^2 H \le 14DH.$$

Restrict the expanded nonunit K_S to $X \ge Z \ge 1/2$. This cuts the residual f-length X, not $F = Nf$ in the tail below. Equation (112) bounds each scalar by $12M_1(W)^2 X m_{14DH}(Z)$. For fixed P put $V = D/P$. Ideal counting gives

$$\#\{(a,b) : \max(Na, Nb) \le t\} \le 256t^2, \qquad \sum_{\max(Na,Nb)\le T} \frac{1}{\max(Na,Nb)} \le 512T.$$

The second inequality follows by partial summation and is vacuous for $T < 1$. Since $M \le V/Z$, the absolute residual cost is at most $6144M_1(W)^2 m_{14DH}(Z)V^2/Z$. Also $\sum_{Np>U}(Np)^{-2} \le 32/U$, again by ideal counting. Only now taking the residual, prime and row triangle proves

$$\frac{|K_S[\text{nonunit}; X \ge Z]|}{D} \le 196608M_1(W)^2 R_w \frac{D}{UZ}\, m_{14DH}(Z). \tag{113}$$

No arbitrary-weight sieve is hidden here: the only row input is $|\eta_{a,b}(u)| \le 1$ and R_w. A *hypothetical* uniform estimate $|M_{S\cup E}(t)| \ll_{\theta,\epsilon,S} (N\prod_{p\in E} p)^{\epsilon} t^{\theta}$, $0 < \theta \le 1$, would give, for $Z \ge 1$, the restricted bound

$$\frac{|K_S[\text{nonunit}; X \ge Z]|}{D} \ll M_1(W)^2 (DH)^{\epsilon} R_w (D/U) Z^{\theta-2}.$$

No $\theta < 1$ hypothesis is assumed. On the region $2X \le 7$, every nonunit good f has zero profile (the possible norm-7 endpoint is zero). Hence only $f = 1$ survives, and (110) is exactly $h(1)(1-2\mathbf{1}_{c_0<1\le d_0})$. For instance $a = b$ gives $|W(1/X)|^2$. Thus even a strong long-sum Mertens estimate cannot supply cancellation in this one-term region after the outer residual triangle. This is a limitation of that scalar route, not a lower bound for the complete signed sum; other residual and row sectors may cancel it.

43.3 Deleted-prime transfer and its fixed-field subpower input

The deletion has an exact fixed-zeta interpretation. Put $M_K(x) = \sum_{Na\le x} \mu(a)$, and let $\mathcal{D}_{\mathcal{T}}$ be *all* ideals supported on $\mathcal{T}$, including prime powers. For $\sigma > 0$ set

$$E_{\mathcal{T}}(\sigma) = \prod_{p\in\mathcal{T}} (1 - (Np)^{-\sigma})^{-1}.$$

Absolute Euler multiplication initially on $\operatorname{Re} s > 1$ gives the coefficient identity and hence the finite convolution

$$M_{\mathcal{T}}(x) = \sum_{\substack{d \in \mathcal{D}_{\mathcal{T}} \\ Nd \le x}} M_K(x/Nd). \tag{114}$$

Restricting d to squarefree ideals would be false. This is an artificial finite Euler deletion, not a new primitive Hecke conductor. Its finite factors have possible poles only on $\operatorname{Re} s = 0$.

Choose C_K with $A_K(x) \le C_K x$, and define $\rho_K(y) = \sup_{v \ge y} |M_K(v)|/(C_K v) \le 1$. For $x \ge 1$ and $0 < \delta < 1$, split (114) at $Nd = \sqrt{x}$. On the first part use $\rho_K(\sqrt{x})$ and $\sum (Nd)^{-1} \le E_{\mathcal{T}}(1-\delta)$. Rankin's inequality bounds the second positive sum by $x^{-\delta/2} E_{\mathcal{T}}(1-\delta)$. Thus

$$|M_{\mathcal{T}}(x)| \le C_K x E_{\mathcal{T}}(1-\delta)\{\rho_K(\sqrt{x}) + x^{-\delta/2}\}. \tag{115}$$

Its minimum with the trivial $C_K x$ bound is always valid.

For clarity, the fixed-field analytic input used to estimate ρ_K is separate from this elementary transfer. The already pinned Thorner–Zaman, Theorem 3.1 gives the classical Dedekind zero-free region, with a possible single fixed exceptional real zero. Specialize only to the trivial extension K/K, of absolute degree 2 and discriminant 3. Shrink the fixed constant so any fixed exceptional zero is outside the chosen region, retaining its dependence in the fixed threshold. For $|t| \ge 2$, the Euler bound just right of 1, fixed-degree functional equation/convexity in a strip of width $O(1/\log(|t|+3))$, and Borel–Carathéodory for an analytic logarithm on smaller zero-free disks give

$$|1/\zeta_K(\sigma + it)| \ll_K \log^A(|t|+3), \qquad \sigma \ge 1 - c_K/\log(|t|+3).$$

Indeed the upper bound for $\operatorname{Re} \log \zeta_K$ and its value at a center to the right of 1 are $O(\log\log(|t|+3))$; a smaller disk has the same bound for its absolute logarithm. Bounded t is handled in a fixed zero-free neighborhood of 1, where the reciprocal is holomorphic also at the pole of ζ_K. No logarithm is continued through that pole.

Here is the Perron error needed for the deduction. The planar fundamental-cell count in this class-number-one field gives $A_K(x+h) - A_K(x) \ll_K h + \sqrt{x} + 1$ for $h \ge 0$. Apply truncated Perron to $1/\zeta_K$ at $x' = \lfloor x \rfloor + 1/2$, whose summatory value equals $M_K(x)$. The coefficients are bounded in absolute value by the ideal counts. Dyadic annuli around x' control the near-integer terms by $x \log^A(2x)/T + \sqrt{x}\log^A(2x)$; the remote terms have the first cost. Shift to $1 - c_K/\log(T+3)$, which lies inside the shrunken zero-free region for every $|t| \le T$. There is no residue at 1, since the reciprocal has a zero there. The vertical and horizontal bounds give

$$|M_K(x)| \ll_K x \log^A(2x) \left\{ \exp\left(-\frac{c_K \log x}{\log(T+3)} \right) + T^{-1} \right\} + \sqrt{x}\log^A(2x).$$

Choosing $T = \exp(c'_K \sqrt{\log(2x)})$ and absorbing logarithms proves

$$|M_K(x)| \ll_K x \exp(-c''_K \sqrt{\log(2x)}). \tag{116}$$

Small x is absorbed in the fixed constant. This is fixed-field subpower cancellation; it asserts no fixed power or uniformity over varying primitive conductors or fields.

The remaining deletion factor can be quantified elementarily. Put $R_{\mathcal{T}} = N \prod_{p \in \mathcal{T} \setminus S} p$ and $Z_0 = \log(e^3 R_{\mathcal{T}}) \ge 3$. For $0 < \delta \le 1/2$, split the prime sum for $\log E_{\mathcal{T}}(1-\delta)$ at Z_0^2. Ideal counting bounds the small part by $O_K(Z_0^{2\delta} \log Z_0)$. For the large part use $\sum_{p \in \mathcal{T} \setminus S} \log Np \le \log R_{\mathcal{T}}$ and the monotonicity of $v^{1-\delta}\log v$; its bound is $O(\log R_{\mathcal{T}}/(Z_0^{2-2\delta}\log Z_0)) = O(1)$. Fixed S factors cost a fixed constant. Hence

$$\log E_{\mathcal{T}}(1-\delta) \le C_{K,S}\{1 + Z_0^{2\delta}\log Z_0\}.$$

If $\delta \log Z_0 \le 1$, this is $E_{\mathcal{T}}(1-\delta) \ll Z_0^A$. Taking $\delta = \min\{c/\sqrt{\log(2x)}, 1/\log Z_0, 1/2\}$ in (115) proves, for $x \ge 3$,

$$|M_{\mathcal{T}}(x)| \ll_{K,S} x Z_0^A \exp\left(-c \min\left\{\sqrt{\log(2x)}, \frac{\log(2x)}{\log Z_0}\right\}\right). \tag{117}$$

For $R_{\mathcal{T}} \le D^C$ and $x \ge D^v$, fixed $C, v > 0$, this gives a uniform $x \log^A(2D) \exp(-c\sqrt{\log(2x)})$ bound. It is useful on long scalar ranges, but gives no power saving and no cancellation in the unit-f range above. The exact product bound (115) retains every raw deletion before any convenient larger estimate.

43.4 The joint matrix, raw zeros and its compression

All spaces here carry ordinary counting ℓ^2 norms. On the fixed finite short-column space of Proposition 20, $g = fa$ is squarefree precisely when $(f, a) = 1$. For $(f, u) = 1$, multiplicativity gives

$$(\mathsf{R}\mathsf{P}_f b)(u) = \sqrt{w_u}\nu(f)\chi_f(u)T_f(t;u).$$

If $(f, u) > 1$, every g in that projected column support has the original zero $\chi_g(u) = 0$. Thus its row is exactly zero, rather than being restored through a primitive character. Squaring proves $G_f(t) = \|\mathsf{R}\mathsf{P}_f b(t)\|^2$ and (105). The Gram entries remain

$$(\mathsf{R}^*\mathsf{R})_{g,h} = \sum_u w_u \overline{\chi_g(u)}\chi_h(u).$$

No coherent orientation mean is removed.

The divisor sum in each matrix entry is the exact Schur factor

$$C_U(g,h) = \sum_{\substack{f \mid (g,h) \\ f \ U\text{-smooth}}} \mu(f) = \mathbf{1}_{(g,h) \text{ has no prime norm } \le U}.$$

On active $t \ge U$ support, $Ng, Nh < 7U$. Any such g containing a rough prime must be that prime alone. It follows that

$$C_U(g,h) = \mathbf{1}_{(g,h)=1} + \mathbf{1}_{g=h=p,\ p\notin S,\ Np>U}. \tag{118}$$

For the projection $\Pi_{\mathcal{P}}$ onto rough-prime columns, $\mathsf{P}_f\Pi_{\mathcal{P}} = 0$ for every nonunit smooth f and $\mathsf{P}_1\Pi_{\mathcal{P}} = \Pi_{\mathcal{P}}$. Therefore

$$\Pi_{\mathcal{P}}\mathcal{M}_U\Pi_{\mathcal{P}} = \Pi_{\mathcal{P}}\mathsf{R}^*\mathsf{R}\Pi_{\mathcal{P}}. \tag{119}$$

This is a quadratic-form *compression*, not an invariant restriction: in general $\mathcal{M}_U\Pi_{\mathcal{P}}$ and $\Pi_{\mathcal{P}}\mathcal{M}_U$ differ. Off-block couplings and other columns can cancel its contribution. It is no lower bound for full Q, the prime error or the original raw slice.

On any common active length dyad $[L, 2L]$, a true squared row norm $\lambda_*(L)$ gives, for arbitrary coefficients c there,

$$|c^*\mathcal{M}_U c| \le \sum_f \|\mathsf{R}\mathsf{P}_f c\|^2 \le \lambda_*(L)\sum_g d_U(g)|c_g|^2, \qquad d_U(g) = \#\{f \mid g : f \ U\text{-smooth}\}.$$

This is a positive incidence-frame majorant, not Möbius cancellation. Since g is squarefree, $d_U(g) \le 2^{\omega(g)} \ll_\eta (Ng)^\eta$. Self-adjointness bounds the compressed operator norm by the same $\lambda_*(L)\max d_U$ and permits polarization for cross terms. For $L = D/t$, the exact short coefficient and ideal counting give

$$\|b(t)\|^2 \ll B^2 L, \qquad \|tb'(t)\|^2 \ll M_1(W)^2 L.$$

The direct unit endpoint covers $1/2 \le L < 1$; no column survives below $1/2$. Differentiating the fixed matrix, not a moving support matrix, gives $Q'(t) = 2\,\mathrm{Re}(b'(t)^* \mathcal{M}_U b(t))$ and

$$|Q(t)| + |tQ'(t)| \ll_\eta (HD)^\eta M_1(W)^2 \lambda_*(D/t) D/t. \tag{120}$$

Only here may the classical specialization be inserted, in exactly its original natural-row scope. For arbitrary weights a true weighted norm is needed; the elementary Frobenius alternative is $\lambda_*(L) \ll R_w L$.

43.5 Exact Stieltjes boundary and squarefunction criterion

The finite Q vanishes for $t \ge 2D$, including its endpoint. Stieltjes integration over (U, ∞) counts precisely $Np > U$; comparing its prime measure with $dt/\log t$ proves (106). Changing the normalization of Li by a constant changes the two boundary/integral terms by opposite amounts and changes no identity. We fix Li_2 here. The original raw slice is consequently

$$S_1^{AA} = \frac{K_0 - \mathcal{J}}{D} + \Delta + \frac{C_2}{D} - \frac{\mathcal{E}_{\mathrm{PIT}}}{D}.$$

It keeps the original masks and the full mean. The prime measure contains only fixed S, never a growing deleted set $S \cup \mathrm{supp}(abu)$.

Let $\delta_S(U) = \sup_{t \ge U} |E_S(t)|/t$. Use (120) and substitute $L = D/t$ to obtain

$$\frac{|\mathcal{E}_{\mathrm{PIT}}|}{D} \ll_\epsilon (HD)^\epsilon M_1(W)^2 \delta_S(U) \left\{ \lambda_*(D/U) + \int_{1/2}^{D/U} \lambda_*(L)\, \frac{dL}{L} \right\}. \tag{121}$$

One may keep $\max d_U$ explicitly instead of absorbing it into a subpower; no uniform logarithmic improvement is thereby claimed. With the current physical input of Proposition 26, the minimum of its three branches remains inside this integral. All constants and internal small exponents are chosen uniformly on the neighboring capacity-feasible blocks.

For $I_T = [T, 2T] \cap [U, 2D]$, let $\lambda^\#(D/T)$ dominate $\lambda_*(L)$ on $L \in [D/(2T), D/T]$, using the direct endpoint when necessary. Equation (120) implies

$$\|Q'\|_{L^2(I_T)} \ll (HD)^\eta M_1(W)^2 \lambda^\#(D/T) D/T^{3/2}.$$

Our weighted squarefunction convention gives $\|E_S\|_{L^2(I_T)} \le (2T)^{3/2} \mathfrak{d}_2(T)$. Cauchy on each dyad and the retained endpoint prove (107). The weighted definition is comparable by absolute constants to $T^{-3/2}\|E_S\|_{L^2([T,2T])}$. Global Cauchy across unrelated dyads need not improve this bound. The sufficient target condition is that its braces be $\ll H(HD)^\epsilon$, including the pointwise term. It is an explicit input to investigate, not an established estimate or an independence assumption between E_S and the signed derivative.

43.6 Mellin Parseval retains equal-norm coherence

Extend only the same finite short-column vector to $t > 0$ for this identity; below U it need not equal the unrestricted original T_f with additional longer columns. No signed equality for K_S is enlarged below U. For fixed f, u let $c_g = A_U(g)\nu(g)\chi_g(u)\mathbf{1}_{f|g}$, and use

$$\widehat{W}(i\tau) = \int_0^\infty W(x) x^{i\tau} \frac{dx}{x}.$$

Ordinary Mellin Plancherel for the finite amplitude gives exactly

$$\begin{aligned}&\int_0^\infty \left|\sum_g c_g W(Ng\,t/D)\right|^2 \frac{dt}{t}\\ &= \frac{1}{2\pi}\int_{\mathbb{R}} |\widehat{W}(i\tau)|^2 \left|\sum_g c_g (D/Ng)^{i\tau}\right|^2 d\tau.\end{aligned} \tag{122}$$

For the derivative of this *amplitude*, replace $|\widehat{W}|^2$ by $\tau^2|\widehat{W}|^2$. This does not assert a diagonal Mellin formula for the square of the quadratic derivative Q'. Finite signed summation in f and the actual rows preserves the matrix form at the norm-twisted coefficient vector. No absolute-value or centering operation is inserted into that equality.

If distinct ideals g, h have $Ng = Nh$, their radial profiles and all their t derivatives are identical. Their radial cross correlation is the same nonzero profile norm as the diagonal when $W \neq 0$. The Mellin phases $(Ng)^{-i\tau}$ also coincide; distinct conjugate split ideals provide actual such orientations. Thus neither radial Parseval nor the derivative removes their coherent mean. The safe decomposition is

$$\sum_r a_r \Gamma_r = \Gamma_{\mathrm{av}} \sum_r a_r + \sum_r a_r(\Gamma_r - \Gamma_{\mathrm{av}}).$$

Both terms remain. The already recorded fixed-profile example in Section 36 has centered variance sum 6, coefficient norm squared 2, but full scalar squared $36 > 12$; it is not rerun or identified with this prime error. It illustrates why a variance-only replacement is invalid. Applying an existing arbitrary-coefficient norm at every τ in (122) leaves its operator power unchanged; the twist has modulus one.

43.7 Proof of the inherited smooth-divisor localization

Now specialize only to the original unweighted $A_6 = 1$ two-pivot rows. The exact K_0 positive summand is

$$G_f(1) = \sum_{\substack{u\in\mathcal{U}\\(f,u)=1}} |T_f(1;u)|^2.$$

Here $A_U(fa)$ is in the general annular domain and stays unfactored. The divisor-assignment bound and ideal counting give its coefficient squared norm $\ll_\eta B^2 D^\eta L$, $L = D/F$. The accepted powered subsequence norm (48) gives

$$G_f(1) \ll (HD)^\eta B^2 \left\{ LR + \frac{L^2 R}{M_{\mathrm{row}}} + \frac{L^{5/3} R}{M_{\mathrm{row}}^{1/3}} \right\}.$$

For $1/2 < L < 1$ count the unit residual directly. Its $O(RD^\eta B^2)$ cost is covered by the same expression up to a constant; for $L = 1/2$ the endpoint is zero. No subunit sieve is used. Now take the triangle only in the tail $F > Y$, and enlarge the nonnegative squarefree smooth-ideal upper sums to all ideals. Partial summation from ideal counting yields

$$\sum_{Y<F<2D} F^{-1} \ll \log(2D), \quad \sum_{F>Y} F^{-2} \ll Y^{-1}, \quad \sum_{F>Y} F^{-5/3} \ll Y^{-2/3}.$$

Division by D and absorption of logarithms and internal assignment powers prove (108). Every rough-prime K_S term has $FPNa > FU$, so $F > Y \geq 2D/U$ forces zero profile even in its auxiliary nonsquarefree cases. The C_2 and P_2 unit corrections are not silently removed by this tail split.

When $R \ll H$, solving the last two monomial inequalities for a sufficient Y gives the displayed cutoff in Proposition 21. At the mixed witness it is enough to take $Y = D^{13/24}$ with fixed capacity slack, giving tail exponents $7/12, 5/8, 3/4$. At the pure witness $Y = D^{5/8}$ gives $3/4, 3/8, 3/4$. Both exceed the rough support exponent $1/2$ for large D. These are inherited sufficient upper cutoffs, not energy lower bounds or assertions of current optimality after the new physical input. All small smooth divisors, including $f = 1$, remain; no signed main estimate follows from the positive tail. The current whole-form comparison in Proposition 25 is a different consequence of the new external norm input.

43.8 Proof of the conditional zero-free input diagnostic

Assume exactly the all-scale, all-ϵ hypothesis of Theorem 22. For $\operatorname{Re} s > 1$ put

$$F_S(s) = \int_2^\infty E_S(t) t^{-s-1}\, dt.$$

On a dyad, Cauchy gives

$$\int_T^{2T} |E_S(t)| t^{-\sigma-1}\, dt \le \mathfrak{d}_2(T) \left(\int_T^{2T} t^{1-2\sigma}\, dt \right)^{1/2} \le C_\sigma \mathfrak{d}_2(T) T^{1-\sigma}.$$

At $\sigma = 1$ use the logarithmic limit of the finite integral. On a compact subset of $\sigma > 1 - \alpha$, choose ϵ smaller than its distance to that boundary. The dyadic majorant is a summable geometric series, also after any fixed power of $\log t$. The same integral therefore converges normally, with all holomorphic derivatives, on $\sigma > 1 - \alpha$. Its bounded initial interval is entire. No contour has crossed an unknown zero.

For $\sigma > 1$, the prime-ideal series satisfies

$$P_S(s) = \sum_{p \notin S} (Np)^{-s} = L_0(s) + sF_S(s), \qquad L_0(s) = \int_2^\infty \frac{t^{-s}}{\log t}\, dt.$$

In this field $\pi_K^S(2) = 0$ and $\operatorname{Li}_2(2) = 0$, so no lower boundary atom is missing. Direct differentiation gives $L_0'(s) = -2^{1-s}/(s-1)$. Only this meromorphic derivative is continued; no single-valued logarithmic primitive is asserted through $s = 1$. Consequently

$$-P_S'(s) = \frac{2^{1-s}}{s-1} - (sF_S(s))'.$$

For $\zeta_K^S(s) = \zeta_K(s) \prod_{p \in S} (1 - (Np)^{-s})$, its initial logarithmic derivative is

$$-\frac{(\zeta_K^S)'(s)}{\zeta_K^S(s)} = -P_S'(s) + H_S(s), \quad H_S(s) = \sum_{p \notin S} \sum_{k \ge 2} \log(Np)(Np)^{-ks}.$$

The higher-prime-power series is normally holomorphic for $\sigma > 1/2$: its geometric tail is bounded on compact subsets by $C \sum_p \log(Np)(Np)^{-2\sigma}$, an absolutely convergent ideal series. Inert prime ideals are included, not mistaken for absent rational primes. The right side has no possible pole in $\sigma > \max\{1 - \alpha, 1/2\}$ except $s = 1$. Standard continuation of the fixed Dedekind zeta, or $\zeta_K(s) = \zeta(s) L(s, \chi_{-3})$, identifies it with the actual logarithmic derivative. A zero elsewhere would give a pole with nonzero integer residue, impossible there. Finite deleted Euler factors vanish only on $\sigma = 0$, so they change no zero in this region. This proves Theorem 22. A fixed exceptional real

zero, if inside the strict region, would be excluded by the hypothetical input itself; it is not assumed absent.

For orientation, the old and current sufficient scalar power accounting is recorded below. Every input and ensuing strip in the table is conditional, and endpoint control is additional to a mean-square norm. Here $\sigma = \operatorname{Re} s$ and the input column means $U^{-\alpha}$.

certificate	budget	input α	strict strip
old pure (historical)	5/6	1/6	$\sigma > 5/6$
current pure	19/24	1/12	$\sigma > 11/12$
current mixed	7/9	1/18	$\sigma > 17/18$

The current values use Proposition 26, not the historical pure 5/6 operator budget. Since $U \asymp D^{1/2}$, subtracting the target 3/4 from the respective PIT powers and doubling gives 1/12 and 1/18. A pointwise all-scale power for δ_S would imply the corresponding squarefunction power; the converse and its pointwise endpoint are not supplied. The accepted fixed-field subpower PIT is not one of these independent powers.

The new physical input does prove its stated global 9/8 and $A_6 = 1$ signed/combined 13/12 comparisons. The coherent reductions above do not improve them to the target, nor establish desired positivity, general Artin holomorphy, determinant/zero-location bridges, a varying-profile fixed-seminorm theorem or new Lean results. Finite formal controls and source-bound reviews check bookkeeping; they do not replace any infinite analytic hypothesis. Full primary reading caches remain local: only public citation/version/hash metadata accompanies these proofs.

44 A sieve comparison indexed by the physical radicand

This section applies a new external preprint input to the original rows. It is an application of that theorem, not a new proof of a large sieve, a signed target estimate, or an analytic Lean result. Throughout, $K = \mathbb{Q}(\omega)$, S is fixed and contains the primes above $2, 3$ and the original conductor exclusions, and ν is the original fixed finite-order Hecke column character. Let $D, H \geq 1$, $U \geq 1$, $U^2 \leq D$, and let W be one common possibly complex smooth profile supported on $[1, 2]$, zero at its endpoints. Write $B_W = \|W\|_\infty$. The estimates below concern the original unweighted literal rows, or already justified bounded natural multiplicities. They do not insert arbitrary positive row weights into an unweighted theorem.

Pinned external input. De Faveri's arXiv:2610.04045v1, Theorem 1.1, bounds the squared row norm $\Theta_n(A, B)$ for the Section 2.3 family of order-dividing-n Hecke characters, with both indices nth-power-free ideals prime to fixed S:

$$\Theta_n(A, B) \ll_\epsilon (AB)^\epsilon (A + B + A^{1-1/n} B^{2/n} + A^{2/n} B^{1-1/n}). \tag{123}$$

Here $n \geq 3$, K contains μ_n, and the chosen fixed S-integer ring has class number one. The constants may depend on the fixed family construction. We use $n = 6$, not a conjectural $A + B$ bound; the additional μ_{2n} condition of that paper's Theorem 1.4 is not a hypothesis of Theorem 1.1. Set

$$\Phi_6(A, B) = A + B + A^{5/6} B^{1/3} + A^{1/3} B^{5/6},$$
$$\Phi_{\mathrm{all}}(H, D) = H + DH^{1/6} + H^{5/6} D^{1/3} + H^{1/3} D^{5/6}.$$

Proposition 23 (Exact raw-family identification)**.** *There is an allowed fixed choice of the data in* (123) *for which, for every good ideal* $\mathfrak{a}$*, with unique generator* $\alpha_{\mathfrak{a}} \equiv 1 \pmod 3$*,*

$$\chi_{\mathfrak{a}}(\mathfrak{m}) = \left(\frac{\alpha_{\mathfrak{a}}}{\mathfrak{m}}\right)_6$$

including the original zero at each common prime. In particular, the row index is the physical sixth-power-free ideal, including valuations $2, 3, 4, 5$*, rather than its radical conductor.*

The extended proof specifies the data and removes the possible unit ambiguity. No extra column-ray selection is made. For the unique outside-S decomposition of an actual literal row,

$$\begin{aligned} u &= u_0 \alpha_{\mathfrak{a}} \alpha_{\mathfrak{r}_6}^6, & \mathfrak{a} &= \mathfrak{r}_1 \mathfrak{r}_2^2 \mathfrak{r}_3^3 \mathfrak{r}_4^4 \mathfrak{r}_5^5, \\ \chi_{\mathfrak{m}}(u) &= \chi_{\mathfrak{m}}(u_0) \chi_{\mathfrak{a}}(\mathfrak{m}) \mathbf{1}_{(\mathfrak{m},\mathfrak{r}_6)=1}, & Nu &= Nu_0 \, N\mathfrak{a} \, (N\mathfrak{r}_6)^6. \end{aligned} \tag{124}$$

The first five factors are squarefree and pairwise coprime; $\mathfrak{r}_6$ is arbitrary and may overlap them. All original union, common-prime, literal, ray and capacity restrictions remain.

Proposition 24 (Physical-row comparison)**.** *For the original raw sums* I_u *of* (14)*, and for the original residue-subtracted* $E_u = I_u - P_u$ *with fixed* ν, S*,*

$$\sum_{0<Nu\le H} |I_u|^2 + \sum_{0<Nu\le H} |E_u|^2 \ll_{\nu,S,\epsilon} B_W^2 (HD)^{\epsilon} \Phi_{\mathrm{all}}(H, D).$$

For $1 \le H \le D$ *one has*

$$\Phi_{\mathrm{all}}(H, D) \le 4DH^{1/6}. \tag{125}$$

Thus the critical positive envelope at $H = D^{3/4}$ *is* $D^{9/8+\epsilon}$*, improving the previous* $7/6$ *comparison by* $1/24$ *while still exceeding the* $3/4$ *target by* $3/8$*.*

The agreement with the varying-profile obstruction of Section 34 concerns only a supremum-norm-controlled class allowing W_D to vary. It is not a lower bound or a sharpness assertion for one fixed W, nor a GRH counterexample.

Proposition 25 (The actual $A_6 = 1$ signed slice)**.** *For an original unweighted two-pivot block, retaining all its actual filters and the complete annular columns, the raw form* S_1^{AA} *of* (44) *satisfies*

$$|S_1^{AA}| \ll_{\nu,S,\epsilon} B_W^2 (HD)^{\epsilon} \Phi_6(H, D).$$

The same bound holds after summing the fixed-S parts of such rows: the D term costs a fixed polylogarithmic count, absorbed in $(HD)^{\epsilon}$*, not a convergent exponent-zero Euler sum. At* $H = D^{3/4}$ *the available power is* $13/12$*, an improvement of* $1/12$ *from the old* $7/6$ *majorant, but still* $1/3$ *above the target.*

The proof inserts (123) into the exact original common-divisor identity (46); $A_U(dn)$ is not factored and is not put into the rough short-coefficient domain. Under the additional high-window and capacity hypotheses of Section 42, its exact ordered rough correction gives

$$\left|\frac{K_0 - K_{\mathrm{all}}}{D}\right| \le |S_1^{AA}| + 236800 B_W^2 R_w. \tag{126}$$

This estimates the actual combined form through its exact identity; it does not take separate main-term triangles in K_0, K_{all}, and does not prove signed cancellation at the target scale.

Proposition 26 (Updated short positive certificate)**.** *On an original unweighted block, put* $R = A_1A_2$, $M = \max(A_1, A_2)$ *and* $H_b = \min\{H/Nu_0, 8A_1A_2^2\}$ *for fixed literal/S factor* u_0*. Empty blocks contribute zero; a nonempty block has* $H_b \geq 1$*. For* $1 \leq L \leq D$ *the actual squared row-operator norm has the bound*

$$\lambda_{\mathrm{now}}(L) \ll (HD)^{\epsilon} \min \begin{Bmatrix} R + \dfrac{LR}{M} + \dfrac{L^{2/3}R}{M^{1/3}}, & L + R^2, \\ \Phi_6(H_b, L) \end{Bmatrix}.$$

For $1/2 \leq L < 1$ *use the direct unit-column norm* R_w*. This new norm may be substituted into the mean-retaining joint Abel criterion of Proposition 20, with its lower endpoint intact. At* $H = D^{3/4}$, $U \asymp D^{1/2}$*, its pure witness costs* 19/24*; at the displayed mixed witness the minimum retains the old* 7/9 *bound, which beats the generic* 19/24*.*

In this positive certificate sufficient independent scalar discrepancy powers are respectively

$$U^{-1/12+\epsilon} \quad \text{and} \quad U^{-1/18+\epsilon}. \tag{127}$$

If imposed as all-scale mean-square inputs, these would entail the conditional zero-free strips $\Re s > 11/12$ and $\Re s > 17/18$ of Theorem 22. Neither input nor strip is proved here; they are not necessary conditions for the actual signed target. The endpoint and coherent $K_0 - J$ remain. Smooth-tail localization (Proposition 21) does not remove $f = 1$ or supply this estimate.

45 Proof of the physical-row sieve application

45.1 The precise literature input and the primary generator

The external input is de Faveri, *Optimal large sieve for fixed order characters*, arXiv:2610.04045v1, 2 October 2026, Theorem 1.1, printed page 2. Input (123) is invoked with the family in its Sections 2.2–2.5; its deep proof is not reproduced or claimed as a new sieve or Lean theorem here. The primary PDF SHA-256 is the concatenation of these two lines:

`97720b3cacccbedb9c2014520ad4d128`
`ae23afd0a781bdeff27dab4e86b7c726`

The public citation receipt records the version and text hash. Full PDF/text reading snapshots are local cache and excluded from public closure.

The ring $\mathcal{O} = \mathbb{Z}[\omega]$ is a PID. Its six units map bijectively onto $(\mathcal{O}/(3))^{\times}$, which has six elements. Consequently every fractional ideal prime to 3 has a unique generator $\alpha_{\mathfrak{a}} \equiv 1 \pmod 3$, interpreted locally at 3. This choice is multiplicative, since a product of two such generators is again 1 modulo 3 and generates the product ideal.

Choose the source's fixed modulus $\mathfrak{c}_0$ sufficiently divisible at S that its stipulated local sixth-power condition holds, and choose it divisible by (3). In the source's Section 2.3 choose ordinary ideal representatives E_0 prime to S, and set $m_{E_0} = \alpha_{E_0}$. Global principality makes these allowed generators after S-localization too. The prescribed products then give $m_E = \alpha_E$ for each representative E. For coprime good ideals $\mathfrak{a}, \mathfrak{m}$, the construction gives

$$\mathfrak{a} = (x)E\mathfrak{g}^6, \qquad x \equiv 1 \pmod{\mathfrak{c}_0}, \qquad (\mathfrak{g}, \mathfrak{m}) = 1, \qquad \chi_{\mathfrak{a}}(\mathfrak{m}) = \left(\frac{xm_E}{\mathfrak{m}}\right)_6.$$

Choose an ordinary fractional generator $\gamma_{\mathfrak{g}}$ of $\mathfrak{g}$. Ordinary ideal equality implies

$$\alpha_{\mathfrak{a}} = \eta x m_E \gamma_{\mathfrak{g}}^6, \qquad \eta \in \mu_6. \tag{128}$$

All factors are units at 3, and $\gamma_{\mathfrak{g}}^6 \equiv 1 \pmod 3$ because the residue unit group has order six. The other three displayed generators are 1 modulo 3. Hence $\eta \equiv 1 \pmod 3$ and injectivity on μ_6 forces $\eta = 1$. Taking the sixth-power residue symbol proves the exact identification in Proposition 23. For a common prime the source's induced character and the original raw symbol are both zero. Thus no primitive-conductor transfer can silently restore a nonunit value.

The field contains μ_6; it need not contain μ_{12} for the Theorem 1.1 input. Its characters have order dividing six and trivial infinite type, so physical valuations $2, 3, 4, 5$ are included. They are not replaced by a squarefree radical or indexed by its norm. The fixed choice of $S, \mathfrak{c}_0, E, m_E$ is independent of D, H, U and of all coefficient masks below. We fix the complex embedding to agree with the original sextic symbol. This is a deliberate allowed construction, not an assertion that every noncanonical choice in the source is identical.

45.2 The complete raw row family and the sixth factor

At each prime outside S, divide its valuation in (u) by six. The remainder defines the sixth-power-free ideal $\mathfrak{a}$, and the quotient defines $\mathfrak{r}_6$. The quotient may exceed one and may be positive at a prime of $\mathfrak{a}$. Thus $\mathfrak{r}_6$ is neither assumed squarefree nor made coprime to the first five factors. With the multiplicative primary generators, the quotient $u_0 = u/(\alpha_{\mathfrak{a}}\alpha_{\mathfrak{r}_6}^6)$ is an integral S-supported element, including its literal unit. This decomposition is unique and has the norm relation in (124).

For a squarefree good column $\mathfrak{m}$, the sixth-power factor has symbol one when $(\mathfrak{m}, \mathfrak{r}_6) = 1$ and zero otherwise. The exact family identification gives (124), even when $\mathfrak{a}, \mathfrak{r}_6$ overlap. In particular a prime common to both factors is still an original raw zero, not two independent masks. For fixed $u_0, \mathfrak{r}_6$ set

$$c_{\mathfrak{m}} = A_U(\mathfrak{m})\nu(\mathfrak{m})\chi_{\mathfrak{m}}(u_0)W(N\mathfrak{m}/D)\mathbf{1}_{(\mathfrak{m},\mathfrak{r}_6)=1}.$$

These coefficients are independent of the remaining row $\mathfrak{a}$. They vanish outside the COMPLETE original good squarefree annulus; putting zero on the extra sixth-power-free indices is permitted. The actual row restrictions can only decrease this positive norm, and

$$\sum_{\mathfrak{m}} |c_{\mathfrak{m}}|^2 \ll_\epsilon B_W^2 D^{1+\epsilon}. \tag{129}$$

Indeed $|A_U(\mathfrak{m})| \le 3^{\omega(\mathfrak{m})}$ uniformly in U, and there are $O_K(D)$ ideal columns. This is the original assignment bound; it does not assert $A_U(dn) = A_U(d)A_U(n)$ or use the rough short-coefficient formula on K_0.

The fixed ν and literal/S phase have modulus one on good columns. A row-independent unitary Jacobi factor of nonzero infinity type could also be included in these arbitrary coefficients. This does not turn it into an order-six row character, supply a theorem for varying infinite types, or improve the operator norm by a gauge. Arbitrary positive row weights are not part of this application; they require a true weighted operator bound in the exact identities of Proposition 20.

45.3 The raw global sum and the principal assembly

For fixed $u_0, \mathfrak{r}_6$, use (123) with

$$A = \frac{H}{Nu_0(N\mathfrak{r}_6)^6}, \qquad B = 2D.$$

If $A < 1$ there is no remaining ideal row. If $A = 1$ the unit row remains; it is not dropped as a subunit endpoint. Equations (124) and (129), and the factor D^{-1} in $|I_u|^2$, give a bound by

$B_W^2(HD)^\epsilon\Phi_6(A,D)$. Now sum all $\mathfrak{r}_6$ of norm at most $(H/Nu_0)^{1/6}$. The $A, A^{5/6}D^{1/3}, A^{1/3}D^{5/6}$ terms use respectively

$$\sum_{\mathfrak{r}_6}(N\mathfrak{r}_6)^{-6}, \qquad \sum_{\mathfrak{r}_6}(N\mathfrak{r}_6)^{-5}, \qquad \sum_{\mathfrak{r}_6}(N\mathfrak{r}_6)^{-2},$$

which converge. The D term instead uses the ideal count $O_K((H/Nu_0)^{1/6})$. Thus the four resulting terms are

$$\begin{aligned} &H(Nu_0)^{-1}, \qquad DH^{1/6}(Nu_0)^{-1/6}, \\ &H^{5/6}D^{1/3}(Nu_0)^{-5/6}, \qquad H^{1/3}D^{5/6}(Nu_0)^{-1/3}. \end{aligned} \tag{130}$$

Every exponent on the fixed finite S-supported semigroup is positive, so these sums converge geometrically, with the six literal units included. This proves the raw part of Proposition 24, with every arbitrary sixth-factor overlap still represented by its coefficient mask.

For the residue-subtracted family, use the original formula (15). Its deleted Euler residue is bounded by a fixed constant and

$$\sum_d \frac{|\beta_U(d)|}{Nd} \le \left(\sum_{Na\le U}\frac{1}{Na}\right)^2 \ll_K \log^2(2U).$$

Therefore $|P_u| \ll_{\nu,S} B_W\sqrt{D}\log^2(2U)$ on a principal row. Outside S, such a row has all valuations divisible by six, since the original ν has its defining conductor in S. The already established count is $O_{\nu,S,\epsilon}(H^{1/6}(HD)^\epsilon)$. Consequently

$$\sum_u |P_u|^2 \ll_{\nu,S,\epsilon} B_W^2(HD)^\epsilon DH^{1/6}. \tag{131}$$

Using $|I-P|^2 \le 2|I|^2+2|P|^2$ proves the stated E_u bound. No Artin holomorphy, spectral positivity or zero-location assertion is used. For $1 \le H \le D$, each of the four terms in Φ_{all} is at most $DH^{1/6}$, proving (125). The varying-W_D construction of Section 34 gives a compatible supremum-norm-class obstruction; it does not make a fixed-W lower bound.

45.4 The signed $A_6 = 1$ common-divisor form

For a fixed u_0 the first-five physical ideal is sixth-power-free. For the two-pivot block it is exactly $v_1v_2^2$, of norm at most H/Nu_0, not Nv_1Nv_2. Apply the positive norm (123) inside the exact common-divisor identity (46). For each squarefree good d, put $L = D/Nd$ and keep the original coefficient $A_U(dn)$, the column mask $(n, dS) = 1$, and the raw row mask $(d, u) = 1$. The row mask may be enlarged only inside the positive bound. The assignment bound and ideal count give coefficient squared norm $\ll_\epsilon B_W^2D^\epsilon L$ for $L \ge 1$. For $1/2 \le L < 1$ the only possible good residual is the unit; its direct row count is bounded by the same polynomial up to an absolute factor. There are no active columns for $L < 1/2$. Thus, only now taking the triangle over d,

$$|S_1^{AA}| \ll B_W^2(HD)^\epsilon \sum_{Nd\le 2D}\frac{1}{Nd}\Phi_6(H/Nu_0, D/Nd).$$

The four ideal sums have exponents $1, 2, 4/3, 11/6$. The first is logarithmic, and the others converge, which proves the fixed-u_0 assertion of Proposition 25. If all S-supported u_0 are included, the terms with positive Nu_0 exponent sum geometrically. The D term has exponent zero and costs only the

finite count $O_S((\log(2H))^{|S|})$, absorbed into $(HD)^\epsilon$. It is not summed as a convergent exponent-zero Euler product.

The accepted identity of Section 42 gives

$$S_1^{AA} = \frac{K_0 - K_{\mathrm{all}}}{D} + \mathcal{E}_{\mathrm{corr}}, \qquad |\mathcal{E}_{\mathrm{corr}}| \le 236800 B_W^2 R_w$$

in its stated high window and at $H \le D^{3/4}$. This retains its ordered (p, q) correction, including $p = q$ and the asymmetric restored prime mask; it is distinct from the unordered unit correction C_2. Taking this already bounded correction gives (126). Neither K_0 nor K_{all} is separately bounded by a substituted marginal mean. The resulting 13/12 bound is an improved positive comparison, not the target signed estimate.

45.5 The updated short certificate and its unresolved input

For a nonempty block $H_b \ge 1$. At $1 \le L \le D$, the two already established block bounds and the new physical-row bound are all valid positive estimates of the SAME actual unweighted operator. Their minimum is Proposition 26. The factor 8 in H_b is the literal dyadic bound $Nv_1(Nv_2)^2 < 8A_1A_2^2$, with the original capacity retained. No radical conductor is substituted for that physical norm. At $1/2 \le L < 1$ retain the actual unit-column norm. For the exact mean-retaining Abel identity of Proposition 20, write $X_0 = D/U$ and $M_1(W) = \max(\|W\|_\infty, \|xW'\|_\infty)$. Its sufficient positive discrepancy certificate is

$$\frac{|\mathcal{E}_{\mathrm{PIT}}|}{D} \ll (HD)^\epsilon M_1(W)^2 \delta_S(U) \left\{ \lambda_{\mathrm{now},*}(X_0) + \int_{1/2}^{X_0} \lambda_{\mathrm{now},*}(L) \frac{dL}{L} \right\}.$$

Here $U \ge 7$ is retained for the Li/Stieltjes definition, and the lower endpoint $E_S(U)Q(U)$ is still part of the exact identity. The minimum must stay inside the integral for general neighboring capacity-feasible blocks. It is not a variance-only substitution. At the critical line, the four new short monomial powers are $3/4, 1/2, 19/24, 2/3$. The pure minimum is 19/24; in the mixed $A_1 \asymp D^{5/12}, A_2 \asymp D^{1/6}$ neighborhood the old 7/9 entry remains smaller. Equating those majorant powers with the 3/4 target gives the sufficient scalar exponents in (127). These are not actual-error lower bounds. The conditional strip implication in Theorem 22 therefore has boundaries 11/12 and 17/18 for these proposed inputs; no such input or strip has been established here. In particular, pointwise endpoint control and the coherent signed $K_0 - J$ are not supplied by an independent mean-square assumption alone.

The new primary input improves a positive norm certificate. Equal-norm ideal orientations, coherent means, original common profiles, full-union raw zeros and all fixed-conductor qualifications remain. The actual signed target, independent determinant bridge and analytic/category Lean work remain open.

46 Coherent dual bands, actual cutoffs and arithmetic localization

The physical-row comparison in Section 44 leaves a signed problem. We now retain its exact primitive roots and sharp row capacity, prove quantitative finite-prefix and divisor-band comparisons, and localize a coherent sector that can already be controlled. The remaining unequal-norm signed moment is stated explicitly. These results do not improve the whole signed power 13/12 or the global unsigned power 9/8. The deep sieve input remains the pinned external theorem (123); its cubic specialization is used below.

46.1 The original form and the scope of the dual representation

Let $K = \mathbb{Q}(\omega)$, $\mathcal{O} = \mathbb{Z}[\omega]$ and $\delta = 1 + 2\omega = i\sqrt{3}$. The fixed set S contains the primes above $2, 3$ and the original exclusions. All columns are the complete good squarefree *ideal* annulus $D \le N\mathfrak{m} \le 2D$, $D > 1$. Distinct ideals of the same norm remain distinct. Put

$$\begin{aligned} A_U(\mathfrak{m}) &= \sum_{\substack{\mathfrak{rs}\mid\mathfrak{m}\\(\mathfrak{r},\mathfrak{s})=1\\ N\mathfrak{r},N\mathfrak{s}\le U}} \mu(\mathfrak{r})\mu(\mathfrak{s}), \\ c_{\mathfrak{m}} &= A_U(\mathfrak{m})\nu(\mathfrak{m})R_{\mathfrak{m}}(u_0)W(N\mathfrak{m}/D). \end{aligned} \tag{132}$$

Here $1 \le U \le \sqrt{D}$, ν is the original fixed unitary finite-order column character, and W is one common possibly complex smooth profile supported on $[1, 2]$, zero at its endpoints. Write $B_W = \|W\|_\infty$. For a good ideal $\mathfrak{v}$, let $\alpha_{\mathfrak{v}} \equiv 1 \pmod 3$ be its unique primary generator and set $R_{\mathfrak{m}}(\mathfrak{v}) = (\alpha_{\mathfrak{v}}/\mathfrak{m})_6$. The literal symbol is zero at a common prime. A primitive extension at an S-argument below is a different object, not restoration of a deleted raw row value.

Let $A_1, A_2 \ge 1$. Only the original unweighted two-pivot rows are used:

$$u = u_0\alpha_{\mathfrak{v}_1}\alpha_{\mathfrak{v}_2}^2, \qquad A_i \le N\mathfrak{v}_i < 2A_i, \qquad Nu_0N\mathfrak{v}_1(N\mathfrak{v}_2)^2 \le H.$$

The ideals $\mathfrak{v}_1, \mathfrak{v}_2$ are good, squarefree and coprime; all actual literal and ray restrictions remain. The other four factors are the unit ideal. The fixed u_0 is literal unit/S-supported, without good-prime support. Extra good frozen factors, a sum over varying u_0, and arbitrary row weights are not included. Write $R = A_1A_2$ and let R_w be the actual number of rows. The normalization is

$$\begin{aligned} \mathcal{P}_D &= DS_1^{AA} \\ &= \sum_{(\mathfrak{v}_1,\mathfrak{v}_2)\ \text{actual}} \sum_{(\mathfrak{m},\mathfrak{n})=1} c_{\mathfrak{m}}\overline{c_{\mathfrak{n}}}R_{\mathfrak{m}}(\mathfrak{v}_1\mathfrak{v}_2^2)\overline{R_{\mathfrak{n}}(\mathfrak{v}_1\mathfrak{v}_2^2)}. \end{aligned} \tag{133}$$

The original Möbius signs are already in the two A_U coefficients; no additional $\mu(\mathfrak{mn})$ is inserted. This raw coprime form is not a positive energy or principal-term assembly.

46.2 Primitive normalization and an exact finite-cutoff kernel

For each source class A, use the fixed decomposition $\widetilde{\chi}_{\mathfrak{m}} = \xi_{\mathfrak{m}}\psi_A$, $\mathrm{cond}(\xi_{\mathfrak{m}}) = \mathfrak{m}$, and $\beta_{\mathfrak{m}} = \epsilon(\xi_{\mathfrak{m}})\psi_A(\mathfrak{m})$. Individual ξ and ψ may have nonzero fixed angular and radial unitary types; they are not declared sixth-order. The finite source reciprocity factors and the *whole* original row-ray projector are expanded in their complete finite character basis, with coefficients $h_{AB\rho}(\mathfrak{v}_2)$. On coprime class-$A, B$ columns the resulting primitive character is

$$\begin{aligned} &\sigma = \mathrm{prim}(\rho\psi_A\overline{\psi_B}), \quad \mathfrak{f}_S = \mathrm{cond}(\sigma), \\ &\eta = \xi_{\mathfrak{m}}\overline{\xi_{\mathfrak{n}}}\sigma, \qquad\qquad F = N\mathfrak{f}_S\, N\mathfrak{m}\, N\mathfrak{n}. \end{aligned} \tag{134}$$

Its total infinity type is zero. The actual S conductor, including any drop, is retained.

Proposition 27 (Primitive roots with their types retained)**.** *For a primitive unitary Hecke character $\chi((\alpha)) = \theta(\alpha)\,\mathrm{phase}(\alpha)^k N(\alpha)^{it}$, with conductor (γ), let*

$$\tau_\gamma(\theta) = \sum_{x \bmod \gamma} \theta(x)\exp\bigl(2\pi i\,\mathrm{Tr}(x/(\delta\gamma))\bigr).$$

Then

$$\epsilon(\chi) = (3N\mathfrak{f})^{it} i^{-|k|} \operatorname{phase}(\delta\gamma)^k \frac{\tau_\gamma(\theta)}{\sqrt{N\mathfrak{f}}}. \tag{135}$$

The unit-conductor/principal formula has its usual meromorphic terms. For pairwise coprime primitive conductors,

$$\epsilon\Big(\prod_j \chi_j\Big) = i^{\sum_j |k_j| - |\sum_j k_j|} \prod_j \epsilon(\chi_j) \prod_{i\neq j} \chi_i(\mathfrak{f}_j). \tag{136}$$

Consequently the actual character in (134) *satisfies*

$$\begin{aligned} \epsilon(\eta) &= \kappa_{AB\rho} u_{\mathfrak{m}} \overline{v_{\mathfrak{n}}}, \\ u_{\mathfrak{m}} &= \beta_{\mathfrak{m}} \xi_{\mathfrak{m}}(\mathfrak{f}_S) \rho(\mathfrak{m}), \\ v_{\mathfrak{n}} &= \beta_{\mathfrak{n}} \xi_{\mathfrak{n}}(\mathfrak{f}_S) \overline{\rho(\mathfrak{n})}. \end{aligned} \tag{137}$$

The scalar $\kappa_{AB\rho}$ is fixed. These are two distinct unitary column gauges. The primitive values $\sigma(\mathfrak{e})$ and $\sigma(\mathfrak{h})$ at S, including zero, remain separately in the kernel; they cannot be replaced by ρ there.

This closes the normalization hypothesis for these stated conventions. A bounded, invertible change of phases gives no power saving by itself.

For a fixed actual $\mathfrak{v}_2$, define the exact integer endpoints

$$a = \lceil A_1 \rceil, \qquad b = \min\{\lceil 2A_1 \rceil - 1, \lfloor H/(Nu_0(N\mathfrak{v}_2)^2) \rfloor\}.$$

An empty interval contributes zero. Choose one fixed nondecreasing smooth step s_0, equal to zero on $(-\infty, 0]$ and one on $[1, \infty)$, and put

$$G_{\mathfrak{v}_2}(x) = s_0(4(A_1 x - a + 1/2)) s_0(4(b + 1/2 - A_1 x)). \tag{138}$$

It is the exact row indicator at every ideal norm node, has amplitude at most one, support in $[1/2, 5/2]$ and total variation at most two. Its derivatives cost powers of A_1; it is not the fixed column W. Set $C_0 = 3/(4\pi^2)$ and

$$\dot{G}(X) = C_0^{-1/2} \int_0^\infty G(t) J_0(2\sqrt{Xt/C_0})\, dt. \tag{139}$$

Proposition 28 (The actual signed kernel before absolute norms)**.** *For every nonzero class/ray component, the good conductor of η is $\mathfrak{mn}$, so η is nonprincipal. Put $P_S = \prod_{p\in S} p$ and*

$$L_{\mathfrak{e}} = \frac{A_1}{(N\mathfrak{k})^2 N\mathfrak{l} N\mathfrak{e}}, \qquad K_{\mathfrak{l}} = \sqrt{b/N\mathfrak{l}}.$$

Squarefree-row inversion, finite S deletion and primitive Poisson give (133) *by summing the following kernel:*

$$\begin{aligned} &c_{\mathfrak{m}} \overline{c_{\mathfrak{n}}} R_{\mathfrak{m}}(\mathfrak{v}_2^2) \overline{R_{\mathfrak{n}}(\mathfrak{v}_2^2)} h_{AB\rho}(\mathfrak{v}_2) \\ &\quad \cdot \mu(\mathfrak{k})\mu(\mathfrak{l})\mu(\mathfrak{e}) \eta(\mathfrak{k}^2\mathfrak{l}\mathfrak{e}) \epsilon(\eta) \frac{L_{\mathfrak{e}}}{\sqrt{F}} \sum_{\mathfrak{y}} \dot{G}_{\mathfrak{v}_2}(L_{\mathfrak{e}} N\mathfrak{y}/F) \overline{\eta(\mathfrak{y})}. \end{aligned} \tag{140}$$

The outer sums are over every actual $\mathfrak{v}_2$, every class/ray component and original coprime column pair, all $\mathfrak{l} \mid \mathfrak{v}_2$ with $N\mathfrak{l} \le b$, all good squarefree $\mathfrak{k}$ with $(\mathfrak{k}, \mathfrak{v}_2) = 1$ and $N\mathfrak{k} \le K_{\mathfrak{l}}$, and every $\mathfrak{e} \mid P_S$. The dual sum is over all *nonzero integral ideals, including S-supported ones. It converges at each fixed parameter set with the actual derivative constants of G. The original pre-S cutoff is never extended in this identity.*

All signs, means, product norms and raw zeros occur before a triangle. A higher induced modulus does not have a unit Gauss root automatically: its nonprimitive additive frequencies may vanish.

46.3 What finite prefixes and actual divisor bands now satisfy

Write

$$\begin{aligned}\Phi(Z,L) &= Z + LZ^{1/6} + Z^{5/6}L^{1/3} + Z^{1/3}L^{5/6},\\ I(t) &= \sum_{j\ge 0}\frac{t^j}{(j!)^2}.\end{aligned} \tag{141}$$

Let $\mathcal{P}_{\le X}$ retain all terms of (140) with $N\mathfrak{y} \le X$, including every source/ray/S component.

Proposition 29 (All-component exact finite prefixes)**.** *For every finite $X \ge 1$,*

$$\frac{|\mathcal{P}_{\le X}|}{D} \ll B_W^2(DXA_2)^{\epsilon}\frac{A_1A_2}{D}\Phi(X,D). \tag{142}$$

A separate exact all-orders estimate is

$$|\mathcal{P}_{\le X}| \ll B_W^2(DXA_2)^{\epsilon}A_1A_2\Phi(X,D)I\left(\frac{(5/2)A_1X}{C_0D^2}\right). \tag{143}$$

No completed Euler model, low-frequency approximation or $X \le \sqrt{H}$ hypothesis is used in (142)*. On a fixed same-source-class/$\rho = 1$ component, the additional critical-half band estimate is*

$$\frac{|\mathcal{P}_{C,1}(Y,2Y]|}{D} \ll B_W^2(DYA_2)^{\epsilon}A_2\sqrt{A_1/Y}\,\Phi(Y,D)(1+\log(2A_1))^2. \tag{144}$$

Constants depend on the fixed S, ray/source data, ν, step and ϵ, not on any hidden uniform smoothness of the sharp G.

At $H \asymp D^{3/4}$, the mixed witness $(A_1,A_2) \asymp (D^{5/12},D^{1/6})$ has a target-sized full finite prefix $X \asymp D$. The pure prefix at $X = D$ has upper power $11/12$. At $X = D^{3/8}$ the pure/mixed powers are $13/16$ and $31/48$. These growing bounds cannot be sent to $X = \infty$.

Proposition 30 (Finite transfer of the actual cutoff and profile)**.** *Let $X \ge 2$, $1 \le U \le X$, and let $K = (K_{\mathfrak{d}\mathfrak{a}})$ be a finite matrix indexed by good squarefree ideals of norms at most X, zero unless $(\mathfrak{d},\mathfrak{a}) = 1$ and $N\mathfrak{d}N\mathfrak{a} \le X$. All additional masks may remain in K. Set*

$$\begin{aligned}G_{\mathcal{I}\mathcal{J}} &= \max_{\mathfrak{d}\in\mathcal{I}} 3^{\omega(\mathfrak{d})}\max_{\mathfrak{a}\in\mathcal{J}} 3^{\omega(\mathfrak{a})},\\ c_P &= 1/\log(4X), \qquad L_T = \operatorname{asinh}(T/c_P)/\pi.\end{aligned}$$

Empty matrices are zero. For every finite $T \ge 8X$,

$$\|[A_U(\mathfrak{d}\mathfrak{a})K_{\mathfrak{d}\mathfrak{a}}]\|_{\mathrm{op}} \le G_{\mathcal{I}\mathcal{J}}\left(e^2L_T^2 + \frac{32X^{3/2}}{T}\right)\|K\|_{\mathrm{op}}. \tag{145}$$

For $T = (4X)^4$ this is $O_\epsilon(X^\epsilon\log^2(4X))\|K\|$. The common complex multiplier $W(N\mathfrak{d}N\mathfrak{a}/D)$ costs

$$\begin{aligned}M(W) &= \frac{1}{2\pi}\int_{\mathbb{R}}|\widehat{W}(it)|\,dt, \qquad \widehat{W}(s) = \int_0^\infty W(x)x^s\,\frac{dx}{x},\\ M(W) &\le \frac{\log 2}{\pi}(B_W + \|xW'(x) + x^2W''(x)\|_\infty).\end{aligned} \tag{146}$$

Two occurrences in a bilinear/Gram estimate cost the square of the transfer norm. This is not $A_U(\mathfrak{d}\mathfrak{a}) = A_U(\mathfrak{d})A_U(\mathfrak{a})$ or coefficient shrinkage.

Proposition 31 (A cubic comparison for an actual auxiliary band)**.** *Fix one same-source-class/$\rho = 1$ component of the auxiliary individual-column extension of* (140)*. Insert the exact coprime identity $\sum_{\mathfrak{d}|(\mathfrak{m},\mathfrak{n})} \mu(\mathfrak{d})$, and retain the* identity-source-class *divisors $E < N\mathfrak{d} \le 2E$, $1 \le E \le D/2$. Denote the resulting exact signed dual band by $\mathcal{P}^{\mathrm{id}}_{E,C}(Y, 2Y]$. Put $L = D/E$ and $\Theta_3(E, L) = E + L + (EL)^{2/3}$. For fixed $0 < \lambda < 1/2$,*

$$\begin{aligned}\frac{|\mathcal{P}^{\mathrm{id}}_{E,C}(Y, 2Y]|}{D} &\ll M(W)^2 (DYA_2)^{\epsilon} A_2 A_1^{1-\lambda} D^{2\lambda-1} Y^{-\lambda} \\ &\quad \cdot \min\{\Phi(Y, L),\ (Y/E)\Theta_3(E, L)\}.\end{aligned} \tag{147}$$

The arithmetic identity behind the new branch is

$$\beta_{\mathfrak{d}\mathfrak{a}} = \epsilon(\widetilde{\chi}_{\mathfrak{d}})\beta_{\mathfrak{a}}\alpha(1, C)\widetilde{\chi}_{\mathfrak{a}}(\mathfrak{d})^2, \qquad [\mathfrak{d}] = 1, \quad [\mathfrak{a}] = C. \tag{148}$$

It exposes the physical cubic residue matrix in $\mathfrak{d}, \mathfrak{a}$. The external input is de Faveri, Theorem 1.1 and Example 2.2, now at $n = 3$, not a conductor-parametric sieve.

This improves a *scoped norm comparator*: at $E \asymp D^{2/3}$ the second branch is $O(Y)$. It still uses $|\mu(\mathfrak{d})| \le 1$ at its final Cauchy step. Such $\mathfrak{d}$ are auxiliary divisors, not the original column-pair gcd h. All divisor classes, small $\mathfrak{d}$ and the full high-dual sum are not bounded here.

46.4 Coherent means that can be localized, and those that remain

Proposition 32 (The full equal-norm sector)**.** *Original coprime squarefree columns with $N\mathfrak{m} = N\mathfrak{n} = N > 1$ satisfy $\mathfrak{n} = \overline{\mathfrak{m}}$, with rational-squarefree N supported on split rational primes. The original good masks require* both *primes above each such rational prime to avoid S; S need not be enlarged or conjugation-stable. There are $O(D)$ ordered annular pairs. Define S_{eq} as the normalized contribution to S_1^{AA}, namely D^{-1} times their full primal contribution to $\mathcal{P}_D$ in* (133)*. Then*

$$|S_{\mathrm{eq}}| \ll B_W^2 R_w D^{\epsilon} \ll B_W^2 R (DH)^{\epsilon}. \tag{149}$$

Before the factor D^{-1}, their ordered orientation sum at norm N is

$$\begin{aligned}a_U(N)^2 |W(N/D)|^2 \prod_{p|N} \{a_p \chi_p(t) + \overline{a_p \chi_p(t)}\}, \\ t = N u_0 N\mathfrak{v}_1 (N\mathfrak{v}_2)^2,\end{aligned} \tag{150}$$

where $a_p = \nu(\pi_p)\overline{\nu(\bar{\pi}_p)}$, $\chi_p(t) = (t/\pi_p)_6$, and $a_U(N) = \sum_{rs|N,(r,s)=1,r,s\le U} \mu_{\mathbb{Z}}(r)\mu_{\mathbb{Z}}(s)$. This includes all ordered orientations and every raw collision zero. For the all-component primitive-ideal prefix/band,

$$\frac{|\mathcal{P}_{\mathrm{eq}}(\le Y)| + |\mathcal{P}_{\mathrm{eq}}(Y, 2Y]|}{D} \ll B_W^2 (DYA_2)^{\epsilon} RY/D. \tag{151}$$

The signed complement is bounded by the whole primal term minus that prefix, not by an absolute high-mode summation.

Since $R \ll H$ on nonempty blocks, the desired full fixed-W target is equivalent, up to (149), to that for unequal norms. Its desired constants may depend on fixed W or appropriate specified seminorms. Formula (150) is neither identically zero nor positive in general.

Proposition 33 (Literal-ray recombination in the pure sector)**.** *In the pure case $\mathfrak{v}_2 = \mathcal{O}$, choose an S-supported ray modulus $\mathfrak{c} = (\gamma_c)$ divisible by* (6)*. For an original coprime pair put $\mathfrak{q} = \mathfrak{mn}$, $Q = N\mathfrak{q}$, $\gamma_q = \alpha_{\mathfrak{m}}\alpha_{\mathfrak{n}}$, and let $\vartheta_{\mathfrak{q}}(x) = (x/\mathfrak{m})_6\overline{(x/\mathfrak{n})_6}$ be the primitive finite character modulo $\mathfrak{q}$. Let $P_{\mathfrak{k}^2}$ be the full literal primary/good-S/ray indicator on $\mathcal{O}/\mathfrak{c}$ after $v_1 = k^2 z$, with finite additive transform $\widehat{P}$. The exact common-modulus Fourier coefficient is*

$$\widehat{F}_{\mathfrak{k}^2}(\beta) = \tau_{\mathfrak{q}}\vartheta_{\mathfrak{q}}(\gamma_c)\overline{\vartheta_{\mathfrak{q}}(\beta)}\widehat{P}_{\mathfrak{k}^2}(\beta\gamma_q^{-1} \bmod \mathfrak{c}). \tag{152}$$

Its literal-lattice radial Poisson prefactor is $L_{\mathfrak{k}}/(N\mathfrak{c}Q)$, $L_{\mathfrak{k}} = A_1/(N\mathfrak{k})^2$. The S-local Fourier-mean sublattice $\beta = \gamma_c\zeta$ has coefficient $r_{\mathfrak{k}^2}\tau_{\mathfrak{q}}\overline{\vartheta_{\mathfrak{q}}(\zeta)}$, where $r = \sum P$. The six-unit projection annihilates every nontrivial unit restriction and retains the unit-trivial branch, whose condition is

$$N\mathfrak{m} \equiv N\mathfrak{n} \pmod{36}. \tag{153}$$

For the basic S above $2, 3$ with no extra ray condition, $\mathfrak{c} = (6)$ has $r = 3$ and ideal coefficient $6r/N\mathfrak{c} = 1/2$. This is not the global $\beta = 0$ mode, which vanishes. All S-supported dual ideals remain, with values defined by any generator once the character is unit-trivial.

Literal additive bands and the earlier primitive-ideal bands have different fixed-S frequency support/rescalings; no unmatched windows are identified. Unequal norms congruent modulo 36 and all other literal residue channels remain. Finite ray recombination supplies no automatic zero mean or signed target estimate.

Remaining arithmetic obligation. A finite-prefix comparison, the sparse equal-norm sector and the cubic auxiliary-band comparison do not bound the complete signed unequal-norm kernel. The required additional estimate must retain all small divisors/source classes and the literal S mean together with complementary residues before the final conductor/divisor/dual triangle. No desired positivity, energy lower bound, new zero-free strip, determinant comparison, analytic Lean or GRH claim follows.

47 Proofs for the coherent-band comparisons

These proofs use de Faveri's pinned arXiv:2610.04045v1, Theorem 1.1 and its stated character family as external input. The primitive Poisson convention agrees with its Lemma 3.1; the formulas below also derive the normalization. This is not a proof of that deep large-sieve theorem. The source version and independent proof reviews are bound by the accompanying source map; full third-party reading caches are not public source dependencies.

47.1 Finite Fourier transforms and the primitive Gaussian

For $\mathfrak{f} = (\gamma)$ use the primitive additive pairing

$$e_\gamma(x) = \exp\bigl(2\pi i \operatorname{Tr}(x/(\delta\gamma))\bigr), \qquad \delta = 1 + 2\omega.$$

Let θ be primitive modulo $\mathfrak{f}$, extended by zero at nonunits of this *actual* conductor. For every $\beta \in \mathcal{O}$,

$$\sum_{x \bmod \gamma} \theta(x)e_\gamma(\beta x) = \overline{\theta(\beta)}\tau_\gamma(\theta). \tag{154}$$

At a unit this is change of variable. At a nonunit frequency the additive character factors through a proper conductor quotient. Primitivity supplies a unit in the kernel of unit reduction on which

θ is nontrivial. Multiplication by that unit leaves the additive factor unchanged and forces the sum to vanish. Unit reduction is surjective prime by prime and hence by CRT. Parseval for this nondegenerate finite additive pairing gives

$$\sum_{\beta \bmod \mathfrak{f}} \left| \sum_x \theta(x) e_\gamma(\beta x) \right|^2 = N\mathfrak{f}\,\varphi(\mathfrak{f}), \qquad |\tau_\gamma(\theta)|^2 = N\mathfrak{f}.$$

For the unit conductor use its one-element convention $\tau = 1$. For coprime γ_1, γ_2, the two CRT frequencies are γ_2^{-1} and γ_1^{-1}, so

$$\tau_{\gamma_1\gamma_2}(\theta_1\theta_2) = \theta_1(\gamma_2)\theta_2(\gamma_1)\tau_{\gamma_1}(\theta_1)\tau_{\gamma_2}(\theta_2). \tag{155}$$

The product is primitive because each ramified local conductor occurs in only one factor.

Here are the archimedean factors that a raw Gauss sum alone would miss. Put $d = |k|$ and

$$P_k(z) = \begin{cases} z^k, & k \geq 0, \\ \bar{z}^{-k}, & k < 0. \end{cases}$$

With Euclidean Fourier transform $\widehat{g}(w) = \int_{\mathbb{C}} g(z) e^{-2\pi i \operatorname{Re}(z\bar{w})}\, dz$, differentiating the two-dimensional Gaussian gives

$$\widehat{P_k(z)e^{-\pi b|z|^2}}(w) = (-i)^d b^{-d-1} P_k(w) e^{-\pi|w|^2/b}, \qquad b > 0. \tag{156}$$

Holomorphic or antiholomorphic harmonicity removes the lower contraction terms. The sign is $(-i)^{|k|}$ for either sign of k. Let $\chi_0 = \chi N^{-it}$ and $a_\gamma = 2/(\sqrt{3}|\gamma|)$. Its theta series is

$$\Theta_{\chi_0}(t) = \sum_{\alpha \in \mathcal{O}} \theta(\alpha) P_k(\alpha) e^{-\pi a_\gamma t |\alpha|^2}.$$

The lattice $\gamma\mathcal{O}$ has covolume $(\sqrt{3}/2)|\gamma|^2$ and dual lattice $2/(\bar{\gamma}\bar{\delta})\mathcal{O}$. Write a dual point as $w = 2\bar{\beta}/(\bar{\gamma}\bar{\delta}) = c\bar{\beta}$. Then $|c| = a_\gamma$ and $\operatorname{phase}(c) = \operatorname{phase}(\delta\gamma)$. Equations (154) and (156), applied in each residue class, yield

$$\Theta_{\chi_0}(t) = i^{-d} \operatorname{phase}(\delta\gamma)^k \frac{\tau_\gamma(\theta)}{\sqrt{N\mathfrak{f}}} t^{-d-1} \Theta_{\bar{\chi}_0}(1/t). \tag{157}$$

The inverse different and the conductor generator have not been replaced by a raw character value at δ.

Use $\Gamma_{\mathbb{C}}(s) = 2(2\pi)^{-s}\Gamma(s)$ and

$$\Lambda(s, \chi) = (3N\mathfrak{f})^{s/2} \Gamma_{\mathbb{C}}(s - it + d/2) L(s, \chi).$$

Every nonzero ideal has six generators; their contributions agree because $\theta(u)u^k = 1$. The Mellin transform of the theta series, with exponent $s + d/2$, is therefore exactly $3(3N\mathfrak{f})^{d/4}\Lambda(s, \chi_0)$. The same constant occurs on the conjugate side. Mellin integration of (157) proves the root formula for $t = 0$. At nonunit conductor the zero frequency vanishes by primitivity. At unit conductor and $k \neq 0$, $P_k(0) = 0$. At unit conductor and $k = 0$, retain the Gaussian zero term when splitting its Mellin integral at one; this gives the usual meromorphic polar terms and root one. They are not dropped from the functional equation.

Changing γ to $u\gamma$ multiplies its Gauss sum by $\theta(u)$ and its angular factor by u^k, whose product is one. Thus the formula is independent of the conductor generator. Complex conjugation gives

$\tau_\gamma(\bar\theta)=\theta(-1)\overline{\tau_\gamma(\theta)}$, where $\theta(-1)=(-1)^k$; together with the angular factor this proves $\epsilon(\bar\chi)=\overline{\epsilon(\chi)}$. Finally

$$\Lambda(s,\chi_0N^{it})=(3N\mathfrak{f})^{it/2}\Lambda(s-it,\chi_0),\qquad \epsilon(\chi_0N^{it})=(3N\mathfrak{f})^{it}\epsilon(\chi_0).$$

This proves (135) also for radial unitary twists.

In (155), the angular cross factors $\mathrm{phase}(\gamma_2)^{k_1}\,\mathrm{phase}(\gamma_1)^{k_2}$ convert the finite residue values into ideal values $\chi_1(\mathfrak{f}_2)\chi_2(\mathfrak{f}_1)$. The remaining i factor is $i^{|k_1|+|k_2|-|k_1+k_2|}$. Radial cross factors satisfy

$$\prod_j(3N\mathfrak{f}_j)^{it_j}\prod_{i\neq j}(N\mathfrak{f}_j)^{it_i}=\left(3N\prod_j\mathfrak{f}_j\right)^{i\sum_j t_j}.$$

Iteration proves (136). Coprimality is a hypothesis on the *actual primitive ideal conductors*, not on rational norms or higher induced moduli.

47.2 Transport to every original class and ray component

Let the types of $\xi_\mathfrak{m},\xi_\mathfrak{n}$ be $(k_A,t_A),(k_B,t_B)$. The character σ in (134) has types (k_B-k_A,t_B-t_A). The three primitive conductors $\mathfrak{m},\mathfrak{n},\mathfrak{f}_S$ are pairwise coprime; no good local ramification can cancel against the other two. Applying (136) gives

$$\begin{aligned}\epsilon(\eta)&=\gamma_{AB\rho}\epsilon(\xi_\mathfrak{m})\overline{\epsilon(\xi_\mathfrak{n})}\xi_\mathfrak{m}(\mathfrak{n}\mathfrak{f}_S)\overline{\xi_\mathfrak{n}(\mathfrak{m}\mathfrak{f}_S)}\sigma(\mathfrak{m}\mathfrak{n}),\\ \gamma_{AB\rho}&=i^{|k_A|+|k_B|+|k_B-k_A|}\epsilon(\sigma).\end{aligned}\tag{158}$$

Source reciprocity is $\widetilde\chi_\mathfrak{m}(\mathfrak{n})\overline{\widetilde\chi_\mathfrak{n}(\mathfrak{m})}=\alpha(A,B)$. Substituting $\widetilde\chi=\xi\psi$ in (158) cancels all crossed ψ values and proves (137), with $\kappa_{AB\rho}=\alpha(A,B)\gamma_{AB\rho}$. In particular the remaining factor is $\rho(\mathfrak{m})\rho(\mathfrak{n})$, not an assumed class constant; ρ need not factor through the source class group.

For good dual $\mathfrak{z}$, write $\widetilde\chi_\mathfrak{m}(\mathfrak{z})=\gamma_A(\mathfrak{z})R_\mathfrak{m}(\mathfrak{z})$, with the fixed finite reciprocity phase in that class. Then

$$\begin{aligned}\overline{\sigma(\mathfrak{z})\xi_\mathfrak{m}(\mathfrak{z})}\xi_\mathfrak{n}(\mathfrak{z})&=\Omega_{AB\rho}(\mathfrak{z})\overline{R_\mathfrak{m}(\mathfrak{z})}R_\mathfrak{n}(\mathfrak{z}),\\ \Omega_{AB\rho}&=\bar\sigma\,\psi_A\bar\psi_B\,\bar\gamma_A\gamma_B,\qquad |\Omega_{AB\rho}(\mathfrak{z})|=1.\end{aligned}\tag{159}$$

The equality includes raw collision zeros. This is a fixed unitary *row* phase for each component, not equality of the two column classes. At S the separate scalar $\sigma(\mathfrak{k}^2\mathfrak{l}\mathfrak{e})\overline{\sigma(\mathfrak{h})}$ is kept, including zero. There is no division by such a scalar, and no substitution of $\rho(\mathfrak{e}/\mathfrak{h})$ at a conductor drop. Off the original coprime column support, products of individual column gauges below are only an auxiliary extension. They are not roots or primitive conductors of overlapping pairs.

47.3 The sharp capacity, square-divisor signs and exact Poisson

The endpoints in (138) are integers. The two smooth transitions lie between ideal norm nodes, so its value at $x=N\mathfrak{v}_1/A_1$ is exactly one or zero according to the original interval and capacity. For nonempty fibres these transitions are disjoint. Their monotonicity gives total variation two; their jth derivatives cost powers of A_1. The elementary identity

$$\sum_{\mathfrak{v}\ \text{good}}\mu(\mathfrak{v})^2\mathbf{1}_{(\mathfrak{v},\mathfrak{v}_2)=1}B(\mathfrak{v})=\sum_{\substack{\mathfrak{k}\ \text{good,SF}\\(\mathfrak{k},\mathfrak{v}_2)=1}}\mu(\mathfrak{k})\sum_{\mathfrak{l}\mid\mathfrak{v}_2}\mu(\mathfrak{l})\sum_{\mathfrak{z}\ \text{good}}B(\mathfrak{k}^2\mathfrak{l}\mathfrak{z})\tag{160}$$

follows by expanding the squarefree indicator and then the remaining coprimality indicator. Auxiliary $\mathfrak{z}$ need not be squarefree, and may overlap $\mathfrak{k}$ or $\mathfrak{l}$. The original norm support fixes $N\mathfrak{l}\le b$ and

$N\mathfrak{k} \le \sqrt{b/N\mathfrak{l}}$ *before* Poisson. Including more $\mathfrak{k}$ would add zero primal terms whose individual dual windows need not be zero.

On good $\mathfrak{v}_1$, expand exactly

$$\mathbf{1}_{\mathrm{ray}} R_{\mathfrak{m}}(\mathfrak{v}_1)\overline{R_{\mathfrak{n}}(\mathfrak{v}_1)} = \sum_{\rho} h_{AB\rho}(\mathfrak{v}_2)\eta_{\mathfrak{m},\mathfrak{n},\rho}(\mathfrak{v}_1).$$

This includes every reciprocity/ray component, not a selected same class term. With $L = A_1/((N\mathfrak{k})^2 N\mathfrak{l})$, good-$S$ deletion is the *finite* identity

$$\sum_{\mathfrak{z}\ \mathrm{good}} G(N\mathfrak{z}/L)\eta(\mathfrak{z}) = \sum_{\mathfrak{e}\mid P_S} \mu(\mathfrak{e})\eta(\mathfrak{e}) \sum_{\mathfrak{y}} G(N\mathfrak{y}/(L/N\mathfrak{e}))\eta(\mathfrak{y}). \tag{161}$$

Primitive zeros at S suppress the corresponding terms automatically. They do not assign forbidden raw values to an original row.

For a nonprincipal primitive type-zero η of conductor norm F, its normalized functional equation and ordinary Mellin inversion give

$$\sum_{\mathfrak{y}} G(N\mathfrak{y}/L)\eta(\mathfrak{y}) = \epsilon(\eta)\frac{L}{\sqrt{F}} \sum_{\mathfrak{y}} \dot{G}(LN\mathfrak{y}/F)\overline{\eta(\mathfrak{y})}. \tag{162}$$

To see the field constant explicitly, put

$$\gamma_K(s) = C_0^{s-1/2}\frac{\Gamma(s)}{\Gamma(1-s)}, \qquad \widehat{G}(s) = \int_0^\infty G(t)t^s\,\frac{dt}{t}.$$

The Mellin transform of (139) is $\gamma_K(s)\widehat{G}(1-s)$. The Bessel identity may first be justified absolutely in $0 < \operatorname{Re} s < 1/4$; smooth compact support and repeated Hankel integration by parts continue it to $\operatorname{Re} s > 0$ at each fixed G. Mellin inversion of the primal ideal sum, its primitive functional equation, and the change $s \mapsto 1-s$ now give (162) with exactly $L/\sqrt{F}$. There is no principal residue since η has nonempty good ramification. The intermediate contour moves involve the entire nonprincipal completed function and the fixed smooth transform, not an inverse L function or an assumed zero-free strip. This also recovers the pinned source Lemma 3.1 convention.

At a good prime ideal of norm P, a nontrivial primitive character modulo p lifted to p^j, $j \ge 2$, has raw Gauss sum zero for a primitive additive frequency. More generally it vanishes unless the frequency has valuation $j-1$, when its value is $P^{j-1}\bar{\chi}(\beta/\pi^{j-1})\tau_p(\chi)$ with compatible additive normalizations. This follows by summing the lifts first. Thus neither lower character order nor a higher physical modulus licenses a unit root. Applying (162) only at the actual F in (134), and inserting (161) into (160), proves the exact signed kernel. Splitting every dual ideal uniquely as $\mathfrak{y} = \mathfrak{h}\mathfrak{z}$, with $\mathfrak{h}$ supported on S and $\mathfrak{z}$ good, retains all modes.

47.4 The joint bilinear estimate behind finite prefixes

The external sixth-order norm gives the following arbitrary-column- coefficient comparison, for every $Z, L \ge 1$:

$$\sum_{\substack{\mathfrak{z}\ \mathrm{good}\\ N\mathfrak{z}\le Z}} \left| \sum_{\substack{\mathfrak{a}\ \mathrm{SF,good}\\ N\mathfrak{a}\asymp L}} b_{\mathfrak{a}} R_{\mathfrak{a}}(\mathfrak{z}) \right|^2 \ll (ZL)^{\epsilon}\Phi(Z,L)\sum_{\mathfrak{a}} |b_{\mathfrak{a}}|^2. \tag{163}$$

For completeness, extract $\mathfrak{z} = \mathfrak{r}\mathfrak{s}^6$ with sixth-power-free $\mathfrak{r}$. Freeze $\mathfrak{s}$ as the original column collision mask $(\mathfrak{a},\mathfrak{s}) = 1$ and apply (123) to $\mathfrak{r}$. The four sums have factors $\sum (N\mathfrak{s})^{-6}$, the count $N\mathfrak{s} \le Z^{1/6}$,

$\sum(N\mathfrak{s})^{-5}$ and $\sum(N\mathfrak{s})^{-2}$, respectively. Ideal counting proves (163). This is not an arbitrary *row*-weight theorem or an extension of a residue-subtracted energy estimate. For $1/2 \le L < 1$ use the direct unit residual column; below $1/2$ there is no annular residual column. Primitive ξ is not treated as sixth-order: convert to raw R inside each fixed residual-class positive norm, where its common ψ row phase cancels.

At $s = \lambda + i\tau$, use (137) to form the two *distinct* vectors

$$\begin{aligned} b_{\mathfrak{m}}(s) &= c_{\mathfrak{m}} u_{\mathfrak{m}} R_{\mathfrak{m}}(\mathfrak{v}_2^2)\xi_{\mathfrak{m}}(\mathfrak{k}^2\mathfrak{l}\mathfrak{e})\overline{\xi_{\mathfrak{m}}(\mathfrak{h})}(N\mathfrak{m})^{s-1/2}, \\ d_{\mathfrak{n}}(s) &= c_{\mathfrak{n}} v_{\mathfrak{n}} R_{\mathfrak{n}}(\mathfrak{v}_2^2)\xi_{\mathfrak{n}}(\mathfrak{k}^2\mathfrak{l}\mathfrak{e})\overline{\xi_{\mathfrak{n}}(\mathfrak{h})}(N\mathfrak{n})^{s-1/2}. \end{aligned} \tag{164}$$

The pairing is $b_{\mathfrak{m}}(s)\overline{d_{\mathfrak{n}}(\bar{s})}$, not a PSD form at s. The exact remaining primitive-conductor factor is

$$(N\mathfrak{f}_S)^{s-1/2}\left(\frac{A_1}{N\mathfrak{l}N\mathfrak{e}}\right)^{1-s}, \tag{165}$$

with scalar $\sigma(\mathfrak{k}^2\mathfrak{l}\mathfrak{e})\bar{\sigma}(\mathfrak{h})$ and finite $\mathfrak{k}$ weight $(N\mathfrak{k})^{-2+2s}$. Its S exponent is $s - 1/2$, not its inverse. On $0 < \lambda < 1/2$ its modulus is at most one.

Insert exact gcd inversion before comparison and put $\mathfrak{m} = \mathfrak{d}\mathfrak{a}$, $\mathfrak{n} = \mathfrak{d}\mathfrak{b}$. Residual $\mathfrak{a}, \mathfrak{b}$ may overlap each other but are individually coprime to $\mathfrak{d}$. Raw multiplicativity gives

$$\overline{R_{\mathfrak{d}\mathfrak{a}}(\mathfrak{z})}R_{\mathfrak{d}\mathfrak{b}}(\mathfrak{z}) = \mathbf{1}_{(\mathfrak{d},\mathfrak{z})=1}\overline{R_{\mathfrak{a}}(\mathfrak{z})}R_{\mathfrak{b}}(\mathfrak{z}).$$

Their source classes are $A[\mathfrak{d}]^{-1}$ and $B[\mathfrak{d}]^{-1}$, generally different. Bilinear Cauchy in the row variable, including (159), gives two separate positive norms. Only *inside* them may the raw $(\mathfrak{d}, \mathfrak{z})$ mask be dropped. Uniformly in the fixed column phases and τ,

$$\sum_{\mathfrak{a}} |b_{\mathfrak{d}\mathfrak{a}}(s)|^2, \quad \sum_{\mathfrak{b}} |d_{\mathfrak{d}\mathfrak{b}}(\bar{s})|^2 \ll B_W^2 D^{2\lambda+\epsilon}/N\mathfrak{d}. \tag{166}$$

Indeed $|A_U(\mathfrak{m})| \le 3^{\omega(\mathfrak{m})}$ and an annular subsequence has $O(D/N\mathfrak{d})$ columns. Neither $A_U(\mathfrak{d}\mathfrak{a})$ nor $\beta_{\mathfrak{d}\mathfrak{a}}$ is factored. For the small residual interval the direct unit count has the same bound.

Apply (163) to the two separate vectors. The divisor sums have exponents $1, 2, 4/3, 11/6$:

$$\sum_{N\mathfrak{d}\le 2D} \frac{1}{N\mathfrak{d}}\Phi(Z, D/N\mathfrak{d}) \ll \log(2D)\Phi(Z, D), \tag{167}$$

with the direct endpoint understood. For $Y < N\mathfrak{h}N\mathfrak{z} \le 2Y$, the weight modulus is at most $Y^{-\lambda}$ and the legal positive row prefix is $N\mathfrak{z} \le 2Y/N\mathfrak{h} \ge 1$. The four positive row exponents in Φ give

$$\sum_{\mathfrak{h}\ S\text{-supported},\, N\mathfrak{h}\le 2Y} \Phi(2Y/N\mathfrak{h}, D) \ll_S \Phi(Y, D).$$

This is a summation over *all* S-supported modes, not an exponent-zero convergent sum or a missing raw S value.

The sharp G requires an honest transform bound. Integration by parts against its finite variation gives

$$\widehat{G}(1-s) = -\frac{1}{1-s}\int t^{1-s}\, dG(t).$$

For fixed $0 < \lambda < 1/2$, Stirling's ratio estimate is $|\gamma_K(\lambda + i\tau)| \ll_\lambda (1 + |\tau|)^{2\lambda-1}$. Thus

$$\int_{\mathbb{R}} |\widehat{G}(1-\lambda-i\tau)\gamma_K(\lambda+i\tau)|\, d\tau \ll_\lambda 1, \tag{168}$$

using the integrable tail $(1+|\tau|)^{2\lambda-2}$ and $\mathrm{TV}(G)\le 2$. No derivative seminorm growing with A_1 is hidden here. Only after the joint estimate may the original finite $\mathfrak{k}$ sum be bounded by $\zeta_K^S(2-2\lambda)$; the identity has not been Euler-completed. Finite $\mathfrak{e}$, the original $\mathfrak{l}\mid\mathfrak{v}_2$ divisor sum, and unweighted actual $\mathfrak{v}_2$ counting cost a fixed constant and $A_2(DA_2)^{\epsilon}$. Equations (166)– (168) give

$$\frac{|\mathcal{P}_D(Y,2Y]|}{D}\ll_\lambda B_W^2(DYA_2)^{\epsilon}A_2A_1^{1-\lambda}D^{2\lambda-1}Y^{-\lambda}\Phi(Y,D). \tag{169}$$

Every finite source/ray component is reconstructed using the distinct gauges and bilinear argument, rather than a same-class inference.

For $0<\lambda<1/6$, each of the four powers of Y after multiplication by $Y^{-\lambda}$ is positive. Finite dyadic summation, a last partial band and the dual unit separately therefore give the same bound with $Y=X$ for every finite prefix $X\ge 1$. For a requested final loss $\epsilon_0>0$, fix $\lambda=\min\{1/12,\epsilon_0/4\}$, independently of all scales, and use the positive-norm loss $\epsilon_0/2$. Since $A_1,X\ge 1$,

$$D^{2\lambda}(A_1X)^{-\lambda}\le D^{\epsilon_0/2},\qquad (DXA_2)^{\epsilon_0/2}D^{\epsilon_0/2}\le (DXA_2)^{\epsilon_0}.$$

This proves (142); neither endpoint $\lambda=0$ nor $\lambda=1/2$ is invoked.

For the independent all-orders bound, expand J_0 exactly at every finite dual prefix. The jth moment satisfies $|M_j|=|\int t^jG(t)dt|\le 2(5/2)^j$. The resulting original-column factors $(D/N\mathfrak{m})^j$ and $(D/N\mathfrak{n})^j$ are contractions, including on $\mathfrak{m}=\mathfrak{d}\mathfrak{a}$; they are not replaced by $(D/N\mathfrak{a})^j$. The exact S factor is $(N\mathfrak{f}_S)^{-j-1/2}$, and the finite $\mathfrak{k}$ norm exponent is $2+2j$ with the *same* character phase $\eta(\mathfrak{k})^2$ as before. After gcd inversion, weighted bilinear Cauchy may bound $(N\mathfrak{h}N_{\mathfrak{z}})^j\le X^j$ inside each positive norm; the product of the square roots costs X^j, not X^{2j}. All remaining divisor/S/subsequence constants are uniform in j. The jth joint bound is the prefactor of (143) times $((5/2)A_1X/(C_0D^2))^j/(j!)^2$. Summing this absolutely convergent series proves (143) without an infinite-dual interchange.

For completeness, the scoped critical-half estimate uses $H(v)=e^{v/2}G(e^v)$. Its Fourier transform is $\widehat{G}(1/2-i\tau)$. The fixed smooth step and its two transitions give $\|H\|_1+\|H'\|_1\ll 1$ and $\|H''\|_1\ll 1+A_1$. The minimum of the three Fourier bounds $O(1)$, $O(1/|\tau|)$ and $O(A_1/\tau^2)$ has integral $O(1+\log(2A_1))$. On the half line $|\gamma_K(1/2+i\tau)|=1$ exactly. The original finite square-divisor sum is harmonic, $\sum_{N\mathfrak{k}\le K_{\mathfrak{l}}}(N\mathfrak{k})^{-1}\ll 1+\log(2A_1)$; it is not an infinite $\zeta_K(1)$. The same joint argument in the stated same-class/ $\rho=1$ component proves (144). This estimate is not summed over all dual bands to infer a whole target.

47.5 Finite-height Schur transfer, including the operator error

For $x\ne 1$ define

$$I_T(x)=\frac{1}{2\pi}\int_{-T}^{T}\frac{x^{c_P+it}}{c_P+it}\,dt.$$

Fourier inversion of $e^{-c_Pv}\mathbf{1}_{v\ge 0}$ gives its full vertical limit $\mathbf{1}_{x>1}$. Integration by parts in each oscillatory tail gives

$$|I_T(x)-\mathbf{1}_{x>1}|\le\frac{2x^{c_P}}{\pi T|\log x|}. \tag{170}$$

There is no pole-free inverse-L contour assertion. Set $U^{\#}=\lfloor U\rfloor+1/2$. Every divisor norm $n\le X$ is a positive integer, so

$$|\log(U^{\#}/n)|\ge\frac{1}{2X+1}\ge\frac{1}{3X},\qquad (U^{\#}/n)^{c_P}\le e.$$

The indicator is exactly $n \le U$, even when U is an integer. Since $6e/\pi < 6$, (170) is at most $e_T = 6X/T \le 3/4$. A product of the two indicators therefore has error at most $2e_T + e_T^2 \le 18X/T$.

For squarefree $\mathfrak{m}$ put

$$H_{\mathfrak{m}}(v,w) = \prod_{p \mid \mathfrak{m}} (1 - (Np)^{-v} - (Np)^{-w}).$$

There are exactly $3^{\omega(\mathfrak{m})}$ ordered disjoint divisor assignments. On the supported coprime products only, $H_{\mathfrak{d}\mathfrak{a}} = H_{\mathfrak{d}} H_{\mathfrak{a}}$, so

$$\begin{aligned} A_U(\mathfrak{d}\mathfrak{a}) &= \frac{1}{(2\pi i)^2} \int_{c_P - iT}^{c_P + iT} \int_{c_P - iT}^{c_P + iT} \frac{(U^{\#})^{v+w}}{vw} H_{\mathfrak{d}}(v,w) H_{\mathfrak{a}}(v,w)\, dv\, dw + R_T(\mathfrak{d},\mathfrak{a}), \\ |R_T(\mathfrak{d},\mathfrak{a})| &\le (18X/T) 3^{\omega(\mathfrak{d})} 3^{\omega(\mathfrak{a})}. \end{aligned} \tag{171}$$

Forbidden entries of K remain zero. The integral is the actual operator $\operatorname{diag}(H_{\mathfrak{d}}) K \operatorname{diag}(H_{\mathfrak{a}})$, with an ordinary right diagonal, not a Hermitian assumption. The two diagonal norms cost $G_{\mathcal{I}\mathcal{J}}$, $(U^{\#})^{2c_P} \le e^2$, and each normalized absolute vertical integral is L_T.

A uniform entry error alone does not multiply operator norms. Instead, with counting ℓ^2 norms,

$$\|R_T \circ K\|_{\mathrm{op}} \le \|R_T \circ K\|_F \le \max |R_T|\, \|K\|_F \le \max |R_T| \sqrt{\min(\#\mathcal{I}, \#\mathcal{J})}\, \|K\|_{\mathrm{op}}.$$

The accepted fixed-field ideal bound $\#\{N\mathfrak{a} \le X\} \le 2X$ gives an error at most $18\sqrt{2} X^{3/2} G_{\mathcal{I}\mathcal{J}}/T$, less than the displayed 32 constant. This proves (145). Splitting the finite set of prime ideals $Np < 3^{1/\eta}$ from the rest proves $3^{\omega(\mathfrak{m})} \ll_\eta (N\mathfrak{m})^\eta$. Use $\eta = \epsilon/2$ and $T = (4X)^4$ to obtain the subpower/log bound. This is a mathematical finite-matrix transfer, not a claim of fast numerical Perron evaluation at that height.

Let $F_0(v) = W(e^v)$. Two integrations by parts give $|\widehat{F_0}(t)| \le \min\{\|F_0\|_1, \|F_0''\|_1/t^2\}$. Split at $|t| = 1$ to get

$$M(W) \le (\|F_0\|_1 + \|F_0''\|_1)/\pi.$$

Here $F_0''(v) = xW'(x) + x^2 W''(x)$ and its support has length $\log 2$, proving (146). Absolute Fourier inversion expresses $W(N\mathfrak{d}N\mathfrak{a}/D)$ as an integral of two unitary norm diagonals and a scalar. Its Schur multiplier norm is at most $M(W)$. Compose with (171); two coefficient occurrences cost the square. All original zeros remain throughout, and no factorization of a primitive $\beta_{\mathfrak{d}\mathfrak{a}}$ is supplied by this lemma.

47.6 The cubic transposition in an identity divisor class

Fix the scope of Proposition 31. On good arguments the primary-generator source construction is index-multiplicative. For $[\mathfrak{d}] = 1$, comparison with the unit source class removes the S conductor of $\widetilde{\chi}_{\mathfrak{d}}$; its type is zero. For $[\mathfrak{a}] = C$ and $(\mathfrak{d},\mathfrak{a}) = 1$,

$$\xi_{\mathfrak{d}\mathfrak{a}} = \widetilde{\chi}_{\mathfrak{d}} \xi_{\mathfrak{a}}.$$

The primitive conductors are coprime and their angular correction in (136) is one. Multiplication by the actual $\psi_C(\mathfrak{d}\mathfrak{a})$ yields

$$\beta_{\mathfrak{d}\mathfrak{a}} = \epsilon(\widetilde{\chi}_{\mathfrak{d}}) \beta_{\mathfrak{a}} \widetilde{\chi}_{\mathfrak{d}}(\mathfrak{a}) \widetilde{\chi}_{\mathfrak{a}}(\mathfrak{d}).$$

Source reciprocity is $\widetilde{\chi}_{\mathfrak{d}}(\mathfrak{a}) \overline{\widetilde{\chi}_{\mathfrak{a}}(\mathfrak{d})} = \alpha(1, C)$, proving (148). In the product for residuals $\mathfrak{a}, \mathfrak{b}$, their common ϵ and α cancel:

$$\beta_{\mathfrak{d}\mathfrak{a}} \overline{\beta_{\mathfrak{d}\mathfrak{b}}} = \beta_{\mathfrak{a}} \overline{\beta_{\mathfrak{b}}} (\widetilde{\chi}_{\mathfrak{a}} \overline{\widetilde{\chi}_{\mathfrak{b}}})^2(\mathfrak{d}). \tag{172}$$

If $\mathfrak{g} = (\mathfrak{a}, \mathfrak{b})$, the primitive cubic conductor is $(\mathfrak{a}/\mathfrak{g})(\mathfrak{b}/\mathfrak{g})$, nonprincipal when $\mathfrak{a} \neq \mathfrak{b}$. The same source class cancels its S component. Common good primes have dropped only from that primitive conductor; the full $\mathfrak{ab}$ union still remains excluded from $\mathfrak{d}$.

The other common $\mathfrak{d}$ phases from ν, u_0 and the $\mathfrak{k}, \mathfrak{l}, \mathfrak{v}_2, \mathfrak{e}, \mathfrak{h}$ modulations cancel in the pair. Their raw zeros, together with the $\mathfrak{z}$ factor, retain

$$(\mathfrak{d}, S\mathfrak{a}\mathfrak{b}\mathfrak{k}\mathfrak{v}_2\mathfrak{z}) = 1.$$

This notation means coprimality with the union of the displayed prime supports. Primitive values at $\mathfrak{e}, \mathfrak{h}$ are legitimate unitary phases of a good-conductor character. The radial factor $(N\mathfrak{d})^{2s-1}$ is kept, not called a cancelled phase.

At fixed $\mathfrak{z}, s$ the core matrix is

$$K_{\mathfrak{d}\mathfrak{a}}^{(\mathfrak{z})} = \mathbf{1}_{[\mathfrak{d}]=1,[\mathfrak{a}]=C,(\mathfrak{d},\mathfrak{k}\mathfrak{v}_2\mathfrak{z})=1}\widetilde{\chi}_{\mathfrak{a}}(\mathfrak{d})^2\overline{R_{\mathfrak{a}}(\mathfrak{z})}. \tag{173}$$

Its cubic character already has the $(\mathfrak{d}, \mathfrak{a})$ zero. All additional fixed residual phases are column contractions. On good arguments the squared sextic symbol is the actual primary cubic symbol $(\alpha_{\mathfrak{a}}/\mathfrak{d})_3$. The pinned external Example 2.2 constructs precisely this family with primary generators congruent to one modulo 3. Both variables are squarefree and hence cube-free; source-sixth-class selectors and deleting the prime above 2 only restrict the positive matrix. Theorem 1.1 at $n = 3$, or its transpose, therefore gives

$$\|K^{(\mathfrak{z})}\|_{\mathrm{op}}^2 \ll_\eta (EL)^\eta \Theta_3(E, L). \tag{174}$$

The lengths are physical ideal norms E, L, not varying conductor labels.

The original annular condition implies $L/2 \leq N\mathfrak{a} \leq 2L$, a bounded number of dyads with lower endpoint at least one. On the encompassing rectangle $N\mathfrak{d}N\mathfrak{a} \leq 4D$, apply Proposition 30 with $X = 4D$. This is the necessary repair: $A_U(\mathfrak{d}\mathfrak{a})$ depends on *both* matrix indices and cannot be put into a fixed coefficient vector by fiat. Composing the common W costs $M(W)$. With $0 < \lambda < 1/2$, the squared row radial factor is $O(E^{2\lambda-1})$, and the column squared norm is

$$\sum_{L/2\leq N\mathfrak{a}\leq 2L} (N\mathfrak{a})^{2\lambda-1} \ll_\lambda L^{2\lambda}.$$

The product is $D^{2\lambda}/E$. Thus the two Mellin vectors have squared norms at most

$$M(W)^2 D^{2\lambda+\eta}\Theta_3(E, L)/E.$$

Cauchy in $\mathfrak{d}$ now uses $|\mu(\mathfrak{d})| \leq 1$, explicitly losing rather than exploiting its sign. A final triangle over $N\mathfrak{z} \leq Z$ gives the new inner branch

$$|T_E(Z; s)| \ll M(W)^2 D^{2\lambda+\eta}(Z/E)\Theta_3(E, L). \tag{175}$$

The accepted separate positive-divisor argument gives $|T_E(Z; s)| \ll M(W)^2 D^{2\lambda+\eta}\Phi(Z, L)$: there are $O(E)$ divisors and each actual subsequence norm is $O(M(W)^2 D^{2\lambda+\eta}/E)$. Both estimates apply to the same quantity, so their minimum is legal.

For $Y < N\mathfrak{h}N\mathfrak{z} \leq 2Y$, put $Z_{\mathfrak{h}} = 2Y/N\mathfrak{h} \geq 1$. The new branch uses the convergent positive exponent-one sum $\sum_{\mathfrak{h}} Z_{\mathfrak{h}} \ll_S Y$; the old branch uses its four positive Φ exponents. Sum each branch to obtain the minimum after summation. The same finite $\mathfrak{k}$ majorant, $\mathfrak{e}/\mathfrak{l}$ bounds, actual A_2 count and fixed-line BV integral as above now prove (147), including the once-only $1/D$ normalization. No all-divisor or whole-dual conclusion follows.

On rough prime divisors $\mathfrak{d} = p$, $Np > U$, no short divisor can contain p, so $A_U(p\mathfrak{a}) = A_U(\mathfrak{a})$ exactly. The remaining fixed-residual arithmetic sum has the shape

$$-A_U(\mathfrak{a})A_U(\mathfrak{b}) \sum_{\substack{E<Np\le 2E,\ Np>U,\ [p]=1\\ p\notin S\cup\mathrm{supp}(\mathfrak{a}\mathfrak{b}\mathfrak{k}\mathfrak{v}_2\mathfrak{z})}} \lambda_{\mathfrak{a}\mathfrak{b}}(p)(Np)^{2s-1} \cdot W(NpN\mathfrak{a}/D)\overline{W(NpN\mathfrak{b}/D)}, \tag{176}$$

up to the already fixed residual phases, with $\lambda_{\mathfrak{a}\mathfrak{b}} = (\widetilde{\chi}_{\mathfrak{a}}\overline{\widetilde{\chi}_{\mathfrak{b}}})^2$. Here $\mu(p) = -1$ is constant. A fixed-field scalar Mertens estimate is not a uniform bound for these varying cubic prime twists. The auxiliary principal case $\mathfrak{a} = \mathfrak{b}$ may cancel against other divisor sectors; it is no physical lower obstruction.

47.7 Classification and proof of the full equal-norm sector

Factor the common rational norm prime by prime. At an inert prime there is only one ideal of norm p^2; equal norms force the same squarefree valuation in both columns and coprimality forces it to zero. The same argument excludes a ramified prime. At a split prime write the two pairs of valuations as $(x, y), (z, w) \in \{0, 1\}^2$. Then $x + y = z + w$ and $xz = yw = 0$. Sum two is forbidden; sum one forces opposite choices. Hence $\mathfrak{n} = \bar{\mathfrak{m}}$ and the common norm is rational-squarefree supported on split primes. Both primes above each such prime must already be good. This does not assume that the original S is conjugation-stable or enlarge it. Each eligible $\mathfrak{m}$ has exactly one partner, so the ordered pair count is at most $\#\{N\mathfrak{m} \le 2D\} \le 4D$. All other rational-nonsquarefree norms remain in the *full* annulus; only this sector has been classified.

The ternary bound is uniform in the original cutoff: $|A_U(\mathfrak{m})| \le 3^{\omega(\mathfrak{m})} \ll_\eta D^\eta$. For each actual row the absolute triangle in this sparse sector costs $O(B_W^2 D^{1+\epsilon})$ before normalization. Divide once by D and sum the actual rows to obtain (149). Original row masks give $R_w \ll R$, and a nonempty capacity block has $Nu_0A_1A_2^2 \le H$, hence $R \le H/(Nu_0A_2) \le H$. For each fixed W, $S_1^{AA} = S_{\mathrm{eq}} + S_{\mathrm{neq}}$ proves the target equivalence stated in the core, with constants allowed to depend on that fixed profile.

Conjugating the residue-field definition gives, including zeros,

$$R_{\bar{\mathfrak{m}}}(x) = \overline{R_{\mathfrak{m}}(\bar{x})}, \qquad R_{\mathfrak{m}}(x)\overline{R_{\bar{\mathfrak{m}}}(x)} = R_{\mathfrak{m}}(Nx). \tag{177}$$

For a split rational p, $p \mid Nx$ exactly when x meets at least one of its two conjugate primes, so the right side retains the original *union* zero. For an orientation of rational- squarefree norm N, restriction to rational integers gives a primitive Dirichlet character $\psi_{\mathfrak{m}} \pmod N$: $\mathbb{Z}/N\mathbb{Z} \longrightarrow \mathcal{O}/\mathfrak{m}$ is an isomorphism, and each local character has exact order six. Its ideal norm pullback has conductor $(N) = \mathfrak{m}\bar{\mathfrak{m}}$ and infinity type zero. There is no smaller radical or common truncated-frequency assertion.

Choose one prime π_p above each admissible split rational p. Every divisor of an orientation $\mathfrak{m}$ has a unique rational divisor norm and the same Möbius parity. Thus $A_U(\mathfrak{m}) = a_U(N)$ with the *two original* cutoffs, and the partner has the same coefficient. At $t = Nu_0N\mathfrak{v}_1(N\mathfrak{v}_2)^2$, (177) makes the local orientation factors $a_p\chi_p(t)$ and its conjugate. Expanding their product gives all $2^{\omega(N)}$ ordered orientations exactly, proving (150). Both reversals are included, not divided by two. The common complex profile gives $|W(N/D)|^2$. For $\nu = 1$ and no raw collision the character factors are $2\,\mathrm{Re}\,\mu_6 \in \{2, 1, -1, -2\}$, so they need not vanish or be positive; a_U or W can still vanish. This is not a row-energy lower bound.

For (151), use the exact primitive representation *restricted first* to these $O(D)$ pairs. Its conductor norm is $N\mathfrak{f}_S N^2 \ge D^2$, and $|\dot{G}| \le 2C_0^{-1/2}$ by $|J_0(t)| \le 1$ on real t. The Poisson amplitude

is at most $A_1/((N\mathfrak{k})^2 N\mathfrak{l} N\mathfrak{c}D)$. There are $O(Y)$ integral dual ideals in a prefix or band, counting every S-supported ideal once. The finite square-divisor sum is bounded by its convergent positive Nk^{-2} majorant; finite S/ray and $\mathfrak{l}$ divisor costs give $A_2(DA_2)^\epsilon$. With the sparse pair count and coefficient product $B_W^2 D^\epsilon$, the numerator is $O(B_W^2(DYA_2)^\epsilon RY)$. Division by D proves the bound. This uses no general row-weight norm and no root cancellation.

At fixed parameters, exact Poisson gives $DS_{\mathrm{eq}} = \mathcal{P}_{\mathrm{eq}}(\le Y) + \mathcal{P}_{\mathrm{eq}}(> Y)$, with its actual smooth derivative constants. Therefore

$$\frac{|\mathcal{P}_{\mathrm{eq}}(> Y)|}{D} \ll B_W^2\{D^\epsilon R_w + (DYA_2)^\epsilon RY/D\}. \tag{178}$$

This is a *signed* whole-minus-prefix identity and triangle, not absolute summation of high modes. The condition $RY/D \le H$ gives pure $Y \asymp D$ and mixed $Y \asymp D^{7/6}$ target-sized prefix/complement sectors. Unequal norms and arbitrary larger individual windows remain unestimated.

47.8 A common literal modulus and the surviving ray mean

This last construction concerns the pure block only. The primary section on good ideals is multiplicative because $\mu_6 \to (\mathcal{O}/3)^\times$ is bijective. Let $\mathfrak{c} = (\gamma_c)$ be fixed, S-supported and divisible by (6), large enough for all actual row-ray congruences. For $\mathfrak{q} = \mathfrak{m}\mathfrak{n}$ coprime to it, put

$$\vartheta_{\mathfrak{q}}(x) = (x/\mathfrak{m})_6\overline{(x/\mathfrak{n})_6}, \qquad \tau_{\mathfrak{q}} = \sum_{x \bmod \mathfrak{q}} \vartheta_{\mathfrak{q}}(x) e_{\gamma_q}(x).$$

Every good local factor is nontrivial sextic, so its finite conductor is exactly $\mathfrak{q}$, including nonunit frequency zeros. It need not yet be unit-trivial or an ideal character. The finite indicator $P_{\mathfrak{k}^2}$ imposes residual generator $x \equiv 1 \pmod 3$, original S-goodness and the actual ray condition on $\alpha_{\mathfrak{k}}^2 x$. It is independent of $\mathfrak{m}, \mathfrak{n}$. Literal square-divisor inversion is

$$\mathcal{P}_D = \sum_{(\mathfrak{m},\mathfrak{n})=1} c_{\mathfrak{m}}\bar{c}_{\mathfrak{n}} \sum_{\substack{\mathfrak{k}\ \text{good,SF} \\ N\mathfrak{k} \le \sqrt{b}}} \mu(\mathfrak{k})\vartheta_{\mathfrak{q}}(\alpha_{\mathfrak{k}}^2) \sum_{x \in \mathcal{O}} P_{\mathfrak{k}^2}(x \bmod \mathfrak{c})\vartheta_{\mathfrak{q}}(x) G(Nx/L_{\mathfrak{k}}). \tag{179}$$

The original $(\mathfrak{k}, \mathfrak{q})$ zero is kept. Auxiliary lattice x need not generate a squarefree ideal.

Use the same positive additive-pairing sign as above. CRT separates the common-modulus Fourier transform into the $\mathfrak{q}$ sum at frequency $\beta\gamma_c^{-1}$ and the $\mathfrak{c}$ sum at $\beta\gamma_q^{-1}$. Primitive finite Gauss Fourier proves (152) for *every* β, including its zero at $(\beta, \mathfrak{q}) > 1$. This is the finite recombination of all ray components; it is not a unit root at the larger induced modulus $\mathfrak{c}\mathfrak{q}$. For $g_L(x) = G(|x|^2/L)$ the Euclidean radial transform is

$$\widehat{g}_L(w) = \pi L \int_0^\infty G(t) J_0(2\pi\sqrt{Lt}|w|)\, dt.$$

At the common dual vector $w = 2\bar{\beta}/(\bar{\delta}\bar{\gamma}_c\bar{\gamma}_q)$ this equals $(\sqrt{3}/2)L\dot{G}(LN\beta/(N\mathfrak{c}Q))$. The lattice covolume is $(\sqrt{3}/2)N\mathfrak{c}Q$. Consequently, writing $F_{\mathfrak{k}^2} = P_{\mathfrak{k}^2}\vartheta_{\mathfrak{q}}$,

$$\sum_{x \in \mathcal{O}} F_{\mathfrak{k}^2}(x) G(Nx/L) = \frac{L}{N\mathfrak{c}Q} \sum_{\beta \in \mathcal{O}} \widehat{F}_{\mathfrak{k}^2}(\beta)\dot{G}(LN\beta/(N\mathfrak{c}Q)). \tag{180}$$

No class/root phase is discarded: the original finite function itself has been transformed. Primitive-conductor descriptions induced to the common modulus must reproduce precisely this frequency support, rather than a falsely flat lifted Gauss sum.

Combining the formulas gives the explicit signed literal additive band

$$\mathcal{B}_Y = \sum_{(\mathfrak{m},\mathfrak{n})=1} c_{\mathfrak{m}}\bar{c}_{\mathfrak{n}} \sum_{\substack{\mathfrak{k}\ \text{good,SF}\\ N\mathfrak{k}\le\sqrt{b}}} \frac{\mu(\mathfrak{k})L_{\mathfrak{k}}\tau_{\mathfrak{q}}\vartheta_{\mathfrak{q}}(\gamma_c)}{N\mathfrak{c}Q}\vartheta_{\mathfrak{q}}(\alpha_{\mathfrak{k}}^2) \cdot \sum_{Y<N\beta\le 2Y} \overline{\vartheta_{\mathfrak{q}}(\beta)}\widehat{P}_{\mathfrak{k}^2}(\beta\gamma_q^{-1})\dot{G}(L_{\mathfrak{k}}N\beta/(N\mathfrak{c}Q)). \tag{181}$$

All finite $\mu(\mathfrak{k})$, the original unfactored coefficients and the full pair/ray/raw masks precede any absolute norm. There is no gcd triangle or positive completion in this identity. The frequency band is literal and unit-invariant; it is not simply the old primitive-ideal band under a new name.

On $\beta = \gamma_c\zeta$, the fixed-residue Fourier frequency is zero. The factor $\vartheta_{\mathfrak{q}}(\gamma_c)$ cancels and

$$\widehat{F}_{\mathfrak{k}^2}(\gamma_c\zeta) = r_{\mathfrak{k}^2}\tau_{\mathfrak{q}}\overline{\vartheta_{\mathfrak{q}}(\zeta)}.$$

This S-local mean is not the global $\beta = 0$, which is zero by primitive nontriviality. A six-unit orbit sum in this radial band vanishes if $\vartheta_{\mathfrak{q}}|_{\mu_6}$ is nontrivial. Otherwise it is six times the sum over *all* integral ideals; values there are defined with any generator, not a nonexistent primary generator for an S-supported ideal. At $\varepsilon = 1+\omega$, local residue powers give

$$\vartheta_{\mathfrak{q}}(\varepsilon) = \varepsilon^{(N\mathfrak{m}-N\mathfrak{n})/6}.$$

Every good prime-ideal norm is one modulo six, and $(\prod Np-1)/6 \equiv \sum(Np-1)/6 \pmod 6$, proving (153). In the surviving case $\vartheta_{\mathfrak{q}}$ is a primitive type-zero ideal character and $\epsilon_{\mathfrak{q}} = \tau_{\mathfrak{q}}/\sqrt{Q}$. The exact retained mean channel is

$$\mathcal{B}_Y^{S\text{-mean}} = \sum_{\substack{(\mathfrak{m},\mathfrak{n})=1\\ N\mathfrak{m}\equiv N\mathfrak{n}\ (36)}} c_{\mathfrak{m}}\bar{c}_{\mathfrak{n}} \sum_{\substack{\mathfrak{k}\ \text{good,SF}\\ N\mathfrak{k}\le\sqrt{b}}} \mu(\mathfrak{k})\vartheta_{\mathfrak{q}}(\alpha_{\mathfrak{k}}^2)\frac{6r_{\mathfrak{k}^2}}{N\mathfrak{c}}\frac{L_{\mathfrak{k}}}{\sqrt{Q}}\epsilon_{\mathfrak{q}} \cdot \sum_{Y/N\mathfrak{c}<N\mathfrak{z}\le 2Y/N\mathfrak{c}} \overline{\vartheta_{\mathfrak{q}}(\mathfrak{z})}\dot{G}(L_{\mathfrak{k}}N\mathfrak{z}/Q). \tag{182}$$

This includes nonunit and S-supported dual modes. The $\mathfrak{k}$ phase is cubic, $\vartheta_{\mathfrak{q}}(\mathfrak{k})^2$, not one, and its original cutoff has not been Euler-completed.

For the basic S above $2,3$ and no additional ray condition, modulo (6) the allowed coordinates $a+b\omega$ are $(1,0),(1,3),(4,3)$. The fourth primary residue $(4,0)$ is excluded by the original 2-adic coprimality condition. Thus $r=3$, $N\mathfrak{c}=36$, and $6r/N\mathfrak{c} = 1/2$ exactly. Dropping that condition would give the incorrect $2/3$. The new exact local Fourier check has $\sum|\widehat{P}|^2 = 36\cdot 3 = 108$; it is a finite S-ring consistency check, not an arithmetic-family or asymptotic band test.

Proposition 32 controls the whole primal equal-norm sector. That alone does not control this sector in an individual literal band or S-mean channel. To formulate a literal unequal-norm problem, impose $N\mathfrak{m} \ne N\mathfrak{n}$ on the original pair sum before applying (179) and (180). The projected formula retains unequal norms congruent modulo 36 as well as all complementary literal residues. Neither the full $\mathcal{B}_Y$ nor its separately projected unequal-norm version is proved small here. At a pure critical block a sufficient band estimate would be

$$|\mathcal{B}_Y| \ll_{S,\nu,W,\epsilon} B_W^2 DA_1(DA_1Y)^{\epsilon},$$

with the fixed-profile dependence retained. This is not necessary for the whole kernel, where bands may cancel. Finite ray/root recombination alone does not provide it; nor does a nonzero local mean give a lower bound for the whole band. This identifies the remaining coherent unequal-norm arithmetic problem without deleting its means or changing any original conductor, capacity or profile.

48 An arithmetic Gauss tail with the original finite cutoff

The coherent kernel of Section 46 contains a difficult root-weighted mean. A large part of its *short-divisor expansion* can now be bounded by retaining a cubic interaction before taking a triangle. This gives a signed cutoff comparison, not a smaller norm for the original coefficients. A direct exact-discrete-frequency version is included as well. The remaining small-short-divisor moment and the separate discrete-frequency error are stated below. The whole signed power 13/12 and global unsigned power 9/8 do not change.

48.1 Original objects and a fixed common profile

Keep $K = \mathbb{Q}(\omega)$, $\mathcal{O} = \mathbb{Z}[\omega]$, $\delta = 1 + 2\omega$, and the fixed exclusions S of Section 46.1. A good ideal is prime to S. All columns are the complete good squarefree ideal annulus $D \le N\mathfrak{m} \le 2D$, $D > 1$. Equal norms do not identify ideals. For a real cutoff $1 \le T \le \sqrt{D}$ put

$$\begin{aligned} A_T(\mathfrak{m}) &= \sum_{\substack{\mathfrak{a}\mathfrak{b}\mid\mathfrak{m}\\(\mathfrak{a},\mathfrak{b})=1\\ N\mathfrak{a},N\mathfrak{b}\le T}} \mu(\mathfrak{a})\mu(\mathfrak{b}),\\ c_f(\mathfrak{m}) &= f(\mathfrak{m})\nu(\mathfrak{m})R_{\mathfrak{m}}(u_0)W(N\mathfrak{m}/D). \end{aligned} \tag{183}$$

Here ν is the original fixed unitary finite-order character; u_0 is the fixed literal unit/S-supported factor with no good support. The actual coefficient is c_{A_U}, $1 \le U \le \sqrt{D}$. The function W is one fixed, possibly complex, smooth common profile supported on $[1, 2]$, zero at the endpoints. Define

$$\begin{aligned} B_W &= \|W\|_\infty, \qquad \widehat{W}(t) = \int_0^\infty W(x)x^{-it}\,\frac{dx}{x},\\ \mathcal{M}(W) &= \frac{1}{2\pi}\int_{\mathbb{R}} |\widehat{W}(t)|\,dt. \end{aligned} \tag{184}$$

This explicit seminorm is finite and controls B_W. It is a seminorm of the column profile, not of the sharp nodal row function.

Only the original unweighted *pure* two-pivot block is used: $A_2 = 1$, $\mathfrak{v}_2 = \mathcal{O}$, and the other four factors are units. The first row is good and squarefree, with the original literal/ray conditions and capacity. Its integer endpoints are

$$a = \lceil A_1 \rceil, \qquad b = \min\{\lceil 2A_1 \rceil - 1, \lfloor H/Nu_0 \rfloor\}, \qquad K = \sqrt{b}. \tag{185}$$

An empty interval contributes zero throughout. Otherwise $b \ge a \ge 1$. Use the exact G of (138), its transform $\dot{G}$ in (139), and $C_0 = 3/(4\pi^2)$. In particular $0 \le G \le 1$, $\operatorname{supp} G \subset [1/2, 5/2]$, and $\int G \le 2$. The original *pre-S* cutoff K is never recomputed after another deletion.

Let $\mathcal{I}$ be the original ordered coprime column pairs. One may also impose the *original* unequal-norm projection; the same choice is then used in every expression below. For each source pair A, B and original ray character ρ use the primitive η, $\mathfrak{f}_S$ and F of (134). Thus $F = N\mathfrak{f}_S N\mathfrak{m}N\mathfrak{n}$, and η has total infinity type zero. Put

$$\begin{aligned} T_q &= S \cup \operatorname{supp}(\mathfrak{m}\mathfrak{n}),\\ Q_s^{T_q}(K;\eta^2) &= \sum_{\substack{\mathfrak{k}\ \text{good,SF}\\ N\mathfrak{k}\le K,\ (\mathfrak{k},T_q)=1}} \frac{\mu(\mathfrak{k})\eta(\mathfrak{k})^2}{(N\mathfrak{k})^s},\\ P_S(\eta) &= \sum_{\mathfrak{e}\mid P_S} \frac{\mu(\mathfrak{e})\eta(\mathfrak{e})}{N\mathfrak{e}}, \qquad P_S = \prod_{p\in S} p. \end{aligned} \tag{186}$$

The symbol $(\mathfrak{k}, T_q) = 1$ means avoidance of every prime in that set. Squaring η never restores a deleted prime. Values at $\mathfrak{e}$ and at any later S-supported frequency are separate primitive values, including zero.

48.2 The actual retained mean and its finite leading model

In the exact primitive kernel (140), write each dual ideal uniquely as $\mathfrak{y} = \mathfrak{h}\mathfrak{r}\mathfrak{s}^6$: $\mathfrak{h}$ is S-supported, $\mathfrak{r}$ is good and sixth-free, and $\mathfrak{s}$ is good. Overlap of $\mathfrak{r}, \mathfrak{s}$ is allowed. For $X \geq 64$ set

$$V = X^{1/6}/2, \qquad b_{\mathfrak{w}} = (X/N\mathfrak{w})^{1/6}, \qquad t_{\mathfrak{w}} = \big((N\mathfrak{w})^{-1/6} - 1/2\big)_+. \tag{187}$$

The top band is $V < N\mathfrak{s} \leq 2V$, $N\mathfrak{y} \leq X$. Replacing its $\mathfrak{s}$ count by its principal fixed-ray density defines the retained mean M_{top}. This operation keeps, rather than deletes, its coherent main term. The selected branches are $\rho^6 = 1$. At this top window $N\mathfrak{h}N\mathfrak{r} < 64 < 7^6$, so $\mathfrak{w} = \mathfrak{h}\mathfrak{r}$ runs over *all* integral ideals of norm below 64, including S-supported ideals.

Set

$$\kappa_S = \frac{\pi}{3\sqrt{3}} \prod_{p \in S} (1 - (Np)^{-1}), \qquad g(\mathfrak{m}) = \prod_{p \mid \mathfrak{m}} (1 - (Np)^{-1}). \tag{188}$$

On original coprime support the density is $\kappa_S g(\mathfrak{m}) g(\mathfrak{n})$. For $\theta = (A, B, \rho, \mathfrak{m}, \mathfrak{n}, \mathfrak{e}, \mathfrak{w})$ in the preceding sets, with $\rho^6 = 1$, define

$$\begin{aligned} d_\theta[f_1, f_2] = \frac{\kappa_S A_1}{D} h_{AB\rho}(\mathcal{O}) c_{f_1}(\mathfrak{m}) \overline{c_{f_2}(\mathfrak{n})} \frac{g(\mathfrak{m}) g(\mathfrak{n}) \epsilon(\eta)}{\sqrt{F}} \\ \cdot \frac{\mu(\mathfrak{e})\eta(\mathfrak{e})}{N\mathfrak{e}} \overline{\eta(\mathfrak{w})}, \qquad \lambda_\theta = \frac{A_1 N\mathfrak{w}}{N\mathfrak{e} F}. \end{aligned} \tag{189}$$

The exact definition of the normalized retained mean is

$$\begin{aligned} \mathcal{M}_{\mathrm{top}} := \frac{M_{\mathrm{top}}}{D} = \sum_\theta d_\theta[A_U, A_U] \sum_{\substack{\mathfrak{k} \text{ good,SF} \\ N\mathfrak{k} \leq K,\ (\mathfrak{k}, T_q) = 1}} \frac{\mu(\mathfrak{k})\eta(\mathfrak{k})^2}{(N\mathfrak{k})^2} \\ \cdot \int_V^{b_{\mathfrak{w}}} \dot{G}\left(\frac{\lambda_\theta t^6}{(N\mathfrak{k})^2} \right) dt. \end{aligned} \tag{190}$$

All sums here are finite. In particular there is no $\mathfrak{k} \mid \mathfrak{w}$ or $(\mathfrak{k}, \mathfrak{w}) = 1$ condition. The finite original source/ray coefficients $h_{AB\rho}(\mathcal{O})$ have not been centered. An argument at $\mathfrak{e}$ is not divided by a value at $\mathfrak{w}$.

Put $g_0 = C_0^{-1/2} \int G$ and define

$$\mathcal{L}_0[f_1, f_2] = g_0 X^{1/6} \sum_\theta d_\theta[f_1, f_2]\, t_{\mathfrak{w}} Q_2^{T_q}(K; \eta^2). \tag{191}$$

Replacing both column functions here is an explicitly defined auxiliary model, not a change to the original primal data.

Proposition 34 (Matched leading model with its original cutoff)**.** *For the actual coefficients and every finite $X \geq 64$,*

$$|\mathcal{M}_{\mathrm{top}} - \mathcal{L}_0[A_U, A_U]| \ll B_W^2 (DX)^\epsilon \frac{A_1^2 X^{7/6}}{D^2}. \tag{192}$$

The constant uses no derivative bound for the sharp G. Both interval endpoints and the original finite K occur in $\mathcal{L}_0$. This estimates the model error, not the leading signed model.

48.3 The root is an incomplete cubic Gauss weight

For each good ideal use its primary generator $\alpha_{\mathfrak{a}} \equiv 1 \pmod 3$ and write $\chi_{j,\mathfrak{a}}$ for its positive residue character of order $j = 3, 6$. For good squarefree $\mathfrak{m}$ put

$$\begin{aligned} G_{\mathfrak{m}} &= \frac{1}{\sqrt{N\mathfrak{m}}} \sum_{x \bmod \alpha_{\mathfrak{m}}} \chi_{3,\mathfrak{m}}(x) \exp\Big(2\pi i \operatorname{Tr} \frac{x}{\delta \alpha_{\mathfrak{m}}}\Big), \\ g_{\mathfrak{m}}(z) &= \overline{\chi_{3,\mathfrak{m}}(z)} G_{\mathfrak{m}}. \end{aligned} \tag{193}$$

The second expression is the normalized finite Gauss sum at frequency z, including zero when $(\mathfrak{m}, z) > 1$. For $\mathfrak{q} = \mathfrak{m}\mathfrak{n}$ put $\theta_q = \chi_{6,\mathfrak{m}} \overline{\chi_{6,\mathfrak{n}}}$. Let θ_S be the *actual finite S part* of the total type-zero η, at its own primitive conductor $\mathfrak{f}_S = (\gamma_S)$. Then

$$\mathcal{S}_{\mathfrak{m},\mathfrak{n},\rho} = \theta_S(\alpha_{\mathfrak{m}}\alpha_{\mathfrak{n}}) \theta_q(\gamma_S) \frac{\tau_{\gamma_S}(\theta_S)}{\sqrt{N\mathfrak{f}_S}}. \tag{194}$$

The unit-conductor convention is one. This scalar has modulus one; it is not obtained by discarding the infinity type of a constituent. The accepted Jacobi character in (47) gives $\vartheta = \Psi/\sqrt{N}$, with conductor (2δ) of norm 12 and nonzero infinity type. The exact root formula is

$$\begin{aligned} \epsilon(\eta) &= \mathcal{S}_{\mathfrak{m},\mathfrak{n},\rho} \Gamma(\alpha_{\mathfrak{m}}\alpha_{\mathfrak{n}}) \mu(\mathfrak{m}) \mu(\mathfrak{n}) \vartheta(\mathfrak{m}) \overline{\vartheta(\mathfrak{n})}\, \overline{G_{\mathfrak{m}}} G_{\mathfrak{n}}, \\ \Gamma(a + b\omega) &= \tfrac{1}{2}\{1 + i^{-b} + i^{a} + i^{b-a}\}. \end{aligned} \tag{195}$$

No extra quadratic sign is inserted. In the original kernel, $\eta(\mathfrak{k}^2)$ and the second-row phase, when present, produce the raw cubic frequency $\overline{g_{\mathfrak{m}}(\alpha_{\mathfrak{k}}\alpha_{\mathfrak{v}_2})} g_{\mathfrak{n}}(\alpha_{\mathfrak{k}}\alpha_{\mathfrak{v}_2})$; here $\mathfrak{v}_2 = \mathcal{O}$. This incomplete weighted frequency set is not a complete residue average.

48.4 A signed short-divisor tail from a physical cubic matrix

For squarefree $\mathfrak{m}$ there is an exact sign-preserving assignment identity:

$$\mu(\mathfrak{m}) A_T(\mathfrak{m}) = \sum_{\substack{\mathfrak{a}\mathfrak{b}\mathfrak{c} = \mathfrak{m} \\ N\mathfrak{a}, N\mathfrak{b} \le T}} \mu(\mathfrak{c}). \tag{196}$$

The factors are pairwise coprime and squarefree. For $Y \le U$ the difference of cutoff boxes has two disjoint pieces. After interchanging $\mathfrak{a}, \mathfrak{b}$ in the second and setting $\mathfrak{z} = \mathfrak{b}\mathfrak{c}$, the exact coefficient is

$$T_U(\mathfrak{z}) + T_Y(\mathfrak{z}), \qquad T_L(\mathfrak{z}) = \sum_{\substack{\mathfrak{b} \mid \mathfrak{z} \\ N\mathfrak{b} \le L}} \mu(\mathfrak{z}/\mathfrak{b}). \tag{197}$$

It is independent of $\mathfrak{a}$. The conjugate Gauss CRT identity retains the cross interaction

$$\overline{G_{\mathfrak{a}\mathfrak{z}}} = \overline{G_{\mathfrak{a}} G_{\mathfrak{z}}} \chi_{3,\mathfrak{a}}(\alpha_{\mathfrak{z}}), \qquad (\mathfrak{a}, \mathfrak{z}) = 1. \tag{198}$$

The raw cubic entry is zero at a common prime.

The deep input is explicitly external: de Faveri, arXiv:2610.04045v1, Theorem 1.1 and Example 2.2, at $n = 3$. For the physical primary cubic matrix on good squarefree ideals of norms comparable to A, Z it gives

$$\big\| [\chi_{3,\mathfrak{a}}(\alpha_{\mathfrak{z}})] \big\|_{\mathrm{op}}^2 \ll_{\epsilon} (AZ)^{\epsilon} \{A + Z + (AZ)^{2/3}\}. \tag{199}$$

Both indices are actual ideals and hence cube-free, not arbitrary primitive-conductor labels. This is the same pinned input as the cubic comparison in Proposition 31.

Theorem 35 (Actual Gauss-weighted cutoff tail)**.** *For any bounded multiplicative ideal function λ on good ideals, with $|\lambda| \le 1$, and $1 \le Y \le U \le \sqrt{D}$,*

$$\begin{aligned}\left|\sum_{\mathfrak{m}\ \mathrm{good,SF}} \frac{\mu(\mathfrak{m})(A_U - A_Y)(\mathfrak{m})\overline{G_{\mathfrak{m}}}\lambda(\mathfrak{m})}{\sqrt{N\mathfrak{m}}} W(N\mathfrak{m}/D)\right| \\ \ll_\epsilon \mathcal{M}(W) D^\epsilon \{D^{1/3} + \sqrt{D/Y}\}.\end{aligned} \tag{200}$$

The constant is uniform in the conductor or raw deletion set in λ. For the actual pure leading model, with every original finite K, class/ray/root/S component and coprime pair retained,

$$|\mathcal{L}_0[A_U, A_U] - \mathcal{L}_0[A_Y, A_Y]| \ll B_W \mathcal{M}(W) D^\epsilon A_1 X^{1/6}\{Y^{-1/2} + D^{-1/6}\}. \tag{201}$$

The original unequal-norm projection is also allowed.

This uses the cubic interaction before the short-divisor triangle. It does not assert pointwise or coefficient-norm shrinkage of A_U, nor a separate cancellation theorem for $\mu(\mathfrak{c})$.

48.5 What has been reduced, and what remains

Combining Proposition 34 and Theorem 35 gives

$$\begin{aligned}|\mathcal{M}_{\mathrm{top}} - \mathcal{L}_0[A_Y, A_Y]| \ll (DX)^\epsilon \Big\{ B_W^2 \frac{A_1^2 X^{7/6}}{D^2} \\ + B_W \mathcal{M}(W) A_1 X^{1/6}(Y^{-1/2} + D^{-1/6})\Big\}.\end{aligned} \tag{202}$$

At $A_1 \asymp H = D^{3/4}$, $X = D^{3/8}$ and $Y = \min\{U, D^{1/8}\}$, the errors are target-sized. If $Y = U$ the coefficient tail is exactly zero; otherwise the three displayed powers are $-1/16$, $3/4$ and $31/48$. The remaining exact coefficients satisfy

$$\mu(\mathfrak{m})A_Y(\mathfrak{m}) = \sum_{\substack{\mathfrak{abc}=\mathfrak{m}\\ N\mathfrak{a}, N\mathfrak{b} \le Y}} \mu(\mathfrak{c}), \qquad N\mathfrak{c} \ge D/Y^2 \ge D^{3/4}. \tag{203}$$

The small-short-divisor, long residual-Möbius cubic Gauss bilinear moment $\mathcal{L}_0[A_Y, A_Y]$ is still unestimated at the target.

48.6 The exact discrete top prefix, with every ray retained

There is also a direct reduction that does not introduce the principal-density split. Its exact support is

$$\mathcal{T}_X = \{(\mathfrak{w}, \mathfrak{s}) : N\mathfrak{w} < 64,\ \mathfrak{s} \text{ good},\ V < N\mathfrak{s} \le 2V,\ N\mathfrak{w}(N\mathfrak{s})^6 \le X\}. \tag{204}$$

There is no squarefree restriction on $\mathfrak{s}$. Overlap of $\mathfrak{w}, \mathfrak{s}$ and of $\mathfrak{k}, \mathfrak{s}$ remains allowed. Let $\widetilde{\Theta}$ be the same outer set as θ above, but with *every* original ρ, and put

$$\begin{aligned}\widetilde{d}_\theta[f_1, f_2] = \frac{A_1}{D} h_{AB\rho}(\mathcal{O}) c_{f_1}(\mathfrak{m})\overline{c_{f_2}(\mathfrak{n})}\frac{\epsilon(\eta)}{\sqrt{F}} \\ \cdot \frac{\mu(\mathfrak{e})\eta(\mathfrak{e})}{N\mathfrak{e}}\overline{\eta(\mathfrak{w})}.\end{aligned} \tag{205}$$

This is an explicit definition, not division by a density or a primitive value. With the original $\mathfrak{k}$ support as in (186), define the normalized exact piece

$$\begin{aligned}\mathcal{P}_X[f_1,f_2]=\sum_{\theta\in\widetilde{\Theta}}\widetilde{d}_\theta[f_1,f_2]\sum_{\mathfrak{k}}\frac{\mu(\mathfrak{k})\eta(\mathfrak{k})^2}{(N\mathfrak{k})^2}\\ \cdot\sum_{\mathfrak{s}:\ (\mathfrak{w},\mathfrak{s})\in\mathcal{T}_X}\bar{\rho}(\mathfrak{s})^6\mathbf{1}_{(\mathfrak{s},\mathfrak{mn})=1}\dot{G}\left(\frac{\lambda_\theta(N\mathfrak{s})^6}{(N\mathfrak{k})^2}\right).\end{aligned}\tag{206}$$

Then $\mathcal{P}_X[A_U,A_U]$ is the actual original top primitive-prefix contribution divided by D. It is a signed representation-fixed piece, not a positive energy or the whole infinite dual sum.

Let $\mathcal{L}_X^K[f_1,f_2]$ be exactly (206) with $\dot{G}$ replaced by g_0. Only in that leading model can the finite $\mathfrak{k}$ sum be factored as $Q_2^{T_q}(K;\eta^2)$. The discrete $\mathfrak{s}$ sum and its endpoints remain. Unlike $\mathcal{L}_0$, this model has no $\kappa_S g(\mathfrak{m})g(\mathfrak{n})$ density, no $\rho^6=1$ projection and no integral replacing the $\mathfrak{s}$ family.

Proposition 36 (Direct original-prefix cutoff reduction)**.** *Put $\Lambda=A_1X/D^2$. With precisely the original support,*

$$|\mathcal{P}_X[A_U,A_U]-\mathcal{L}_X^K[A_U,A_U]|\ll B_W^2(DX)^\epsilon A_1V\Lambda,\tag{207}$$

and

$$\begin{aligned}|\mathcal{L}_X^K[A_U,A_U]-\mathcal{L}_X^K[A_Y,A_Y]|\ll B_W\mathcal{M}(W)(DX)^\epsilon A_1V\\ \cdot\{Y^{-1/2}+D^{-1/6}\}.\end{aligned}\tag{208}$$

In particular the actual original piece satisfies

$$\begin{aligned}|\mathcal{P}_X[A_U,A_U]-\mathcal{L}_X^K[A_Y,A_Y]|\ll(DX)^\epsilon A_1V\{B_W^2\Lambda\\ +B_W\mathcal{M}(W)(Y^{-1/2}+D^{-1/6})\}.\end{aligned}\tag{209}$$

The original unequal-norm projection is allowed.

At the critical pure parameters used above the three powers are again $-1/16$, $3/4$ and $31/48$, with exact zero coefficient tail when $Y=U$. This is a reduction of the *actual discrete* piece to a different, still unestimated, small-short-divisor model. That model retains every ρ, every nonunit S mode, the exact $\mathfrak{w},\mathfrak{s}$ support, root/conductor data, original finite K and both long residual-Möbius column coefficients.

Separate open errors and scope. The accepted fixed-ray disk count gives the distinct top-band certificate

$$\frac{|\mathcal{P}_{\mathrm{top}}-M_{\mathrm{top}}|}{D}\ll B_W^2(DX)^\epsilon A_1X^{1/18}\{1+\sqrt{A_1X/D^2}\}.\tag{210}$$

Its pure critical power is $37/48>3/4$; this is an upper certificate, not a lower bound for the actual error. The direct reduction in Proposition 36 does not introduce that split: its discrete model retains those effects. No improved bound for the separate discrepancy, for the remaining small-cutoff models, or for the other original sectors is proved here. The result is an application of a named external cubic sieve to an actual arithmetic tail, not a new sieve, a complete analytic Lean formalization, a determinant/zero-location bridge or a proof of GRH. Fixed-profile constants and the explicit $\mathcal{M}(W)$ cannot be replaced by a uniform supnorm-only claim for varying profiles.

49 Proofs for the finite-cutoff Gauss tail

The arithmetic norm input is the explicitly pinned de Faveri, arXiv:2610.04045v1, Theorem 1.1 and Example 2.2, at order three. It is used on its actual primary cube-free ideal family. We do not reprove that deep theorem. The primitive root conventions are those already established in Proposition 27, and the Jacobi orientation is that of (47). The accompanying source map binds the accepted research proofs and their independent reviews. Full third-party reading caches are not public proof inputs.

49.1 The coherent main retains all sixth-part and ray data

For a good ideal $\mathfrak{s}$, the original fixed-source transport in (134) gives

$$\overline{\eta(\mathfrak{s}^6)} = \overline{\rho(\mathfrak{s})}^6 \mathbf{1}_{(\mathfrak{s},\mathfrak{mn})=1}. \tag{211}$$

This is an assertion about the total character on good arguments, not about the order of the individual possibly angular/radial $\xi_{\mathfrak{m}}$. Write $\psi = \bar{\rho}^6$, and let δ_ψ be one if it is principal on good ideals and zero otherwise. The conductor of ψ belongs to fixed S data.

Here is the fixed-ray counting input, with its full deletion. For $t \geq 0$, the classical disk estimate gives

$$\begin{aligned} N_{\psi,q}(t) &:= \sum_{\substack{\mathfrak{s} \text{ good} \\ N\mathfrak{s}\leq t,\ (\mathfrak{s},\mathfrak{q})=1}} \psi(\mathfrak{s}) \\ &= \delta_\psi \kappa_S g(\mathfrak{q}) t + E_{\psi,q}(t), \\ |E_{\psi,q}(t)| &\ll_{S,\psi} t^{1/3} \sum_{\mathfrak{d}|\mathfrak{q}} (N\mathfrak{d})^{-1/3} \ll_\epsilon D^\epsilon t^{1/3}, \qquad N\mathfrak{q} \leq 4D^2. \end{aligned} \tag{212}$$

We recall a proof to specify the uniformity being used. For a fixed ideal modulus, count elements in each translated lattice coset inside a disk of radius R. Convolve the disk indicator with a fixed compact smooth nonnegative mollifier of radius Δ. The smoothed disks of radii $R - \Delta$ and $R + \Delta$ sandwich the original disk. Their zero modes differ from its area by $O(R\Delta + \Delta^2)$.

For a nonzero dual vector ξ, polar integration gives

$$\widehat{\mathbf{1}_{|z|\leq R}}(\xi) = \frac{R J_1(2\pi R|\xi|)}{|\xi|}, \qquad |J_1(u)| \ll \min\{1, u^{-1/2}\}. \tag{213}$$

The decay follows directly from the angular integral for J_1: for $u \geq 1$, remove arcs of length $u^{-1/2}$ around the two stationary points. Their total contribution is $O(u^{-1/2})$. On the complement integrate once by parts in the phase; the boundary terms and the integral of the derivative of the reciprocal phase derivative are also $O(u^{-1/2})$. For $u \leq 1$ the angular formula gives the bounded estimate. Thus no unproved cancellation of lattice points is inserted. The mollifier supplies arbitrarily rapid decay beyond $|\xi| \asymp \Delta^{-1}$. Counting points of the fixed dual lattice in unit-width radial shells yields

$$\sum_{\xi \neq 0} |\widehat{\mathbf{1}_{|z|\leq R}}(\xi) \widehat{\varphi}(\Delta\xi)| \ll R^{1/2}\Delta^{-1/2}. \tag{214}$$

The translation only changes unit phases. Choose Δ comparable to $R^{-1/3}$ for large R. The boundary and dual errors are both $O(R^{2/3})$. Bounded R is absorbed into the fixed constant.

A type-zero finite ray character is a fixed periodic generator function, trivial on the six units. Apply the preceding count to its finitely many residue cosets, and only then form the possibly

complex weighted sum. Division by six and the covolume $\sqrt{3}/2$ of $\mathcal{O}$ give the coefficient $\pi/(3\sqrt{3})$. The principal coset average is the good-S density; a nonprincipal character has mean zero. For $t \ge 1$ this proves the undeleted error $O(t^{1/3})$. For $0 < t < 1$ the ideal count is zero and $t \le t^{1/3}$, so the same ideal estimate remains valid; $t = 0$ is exact.

Now use the *entire* finite inclusion–exclusion:

$$N_{\psi,q}(t) = \sum_{\mathfrak{d}\mid\mathfrak{q}} \mu(\mathfrak{d})\psi(\mathfrak{d}) N_{\psi,1}(t/N\mathfrak{d}). \tag{215}$$

Divisors with $N\mathfrak{d} > t$ are still present, with their subunit arguments. For principal ψ the main sum is $\kappa_S t \prod_{p\mid q}(1-(Np)^{-1})$; otherwise it is zero. The error in (212) follows. Its final divisor product is $O_\epsilon(D^\epsilon)$ by separating finitely many small primes and using $1 + P^{-1/3} \le P^\eta$ at all sufficiently large P. This is ideal-squarefreeness, not rational-squarefreeness.

For fixed $\mathfrak{h}, \mathfrak{r}$, the original sum over $V < N\mathfrak{s} \le b_{\mathfrak{h}\mathfrak{r}}$ has weight $\dot{G}(c\,t^6)$. In the pure block the scale is

$$c = \frac{A_1 N\mathfrak{h} N\mathfrak{r}}{(N\mathfrak{k})^2 N\mathfrak{e} F}. \tag{216}$$

Closed prefixes at both endpoints give the exact error formula

$$w(b)E_{\psi,q}(b) - w(V)E_{\psi,q}(V) - \int_V^b E_{\psi,q}(t)w'(t)\,dt, \qquad w(t) = \dot{G}(ct^6). \tag{217}$$

There are no half-weights at ideal norm ties. The angular representations of J_0, J_1 imply, with $\Lambda = A_1 X/D^2$,

$$\|\dot{G}\|_\infty \ll 1, \qquad |x\dot{G}'(x)| \ll \sqrt{x}, \qquad \int_V^b |w'(t)|\,dt \ll \sqrt{\Lambda}. \tag{218}$$

Indeed $ct^6 \le \Lambda$ and $b/V \le 2$; differentiating the Bessel factor, not G, proves the last assertion. Consequently the error at each fixed outer set is $O(D^\epsilon V^{1/3}(1+\sqrt{\Lambda}))$.

At the top window $N\mathfrak{h}N\mathfrak{r} < 64$. Every good ideal of norm below 64 is sixth-free, because a nontrivial good prime has norm at least seven. Therefore the unique decomposition into its S part and good part runs exactly over all $\mathfrak{w}$ in (190). The principal condition is precisely $\rho^6 = 1$. On the original coprime support $g(\mathfrak{m}\mathfrak{n}) = g(\mathfrak{m})g(\mathfrak{n})$. This proves the stated finite formula for M_{top}/D.

For the separate error, $|A_U| \le 3^{\omega(\mathfrak{m})} \ll_\epsilon D^\epsilon$, and there are $O(D)$ original annular ideals. After normalizing once by D, $A_1/(D\sqrt{F}) \le A_1/D^2$ cancels the two column counts. The S/ray/$\mathfrak{e}$ sets are fixed, the top $\mathfrak{w}$ set is finite, and the *original* finite $\mathfrak{k}$ sum is majorized by $\sum (N\mathfrak{k})^{-2}$. Since $V^{1/3} \asymp X^{1/18}$, this proves (210). Only its error uses these triangles; the mean in (190) stays signed.

49.2 Exact good-modulus Gauss factors and separate primitive S values

The finite cubic sum at frequency z is

$$\frac{1}{\sqrt{N\mathfrak{m}}} \sum_{x \bmod \alpha_{\mathfrak{m}}} \chi_{3,\mathfrak{m}}(x) \exp\!\left(2\pi i \operatorname{Tr} \frac{zx}{\delta\alpha_{\mathfrak{m}}}\right) = \overline{\chi_{3,\mathfrak{m}}(z)} G_{\mathfrak{m}}. \tag{219}$$

For a unit frequency this is change of variable. At a collision, one local additive frequency is zero and its nonprincipal local cubic character has zero sum. CRT proves the assertion for every good squarefree ideal, including inert primes and nonsquarefree rational norms. The normalized sum has modulus one at unit frequencies.

The total η has type zero, so the previously proved root formula is its finite Gauss sum divided by $\sqrt{F}$. CRT at the coprime primitive conductors $\mathfrak{f}_S, \mathfrak{q}$ gives exactly (194) times $\tau_q(\theta_q)/\sqrt{Nq}$.

Although their product is unit-trivial, the finite S and good-q characters need not be unit-trivial separately. No extra unit average is taken. The actual finite S character, including its conductor drop, is used at its own primitive modulus. At unit conductor its Gauss sum is one.

On good squarefree q, all local powers $\theta_q, \theta_q^2, \theta_q^3$ are nonprincipal and primitive. The Gauss–Jacobi identity therefore gives

$$\tau_q(\theta_q)\tau_q(\theta_q^2) = J_q(\theta_q, \theta_q^2)\tau_q(\theta_q^3). \tag{220}$$

Over a residue field this follows by setting $x + y = t$ in the product of Gauss sums; the zero-t term vanishes because the product character is nonprincipal. Multiplication by CRT proves the composite identity: its additive cross multipliers on the left combine to those of θ_q^3 on the right. The Jacobi sum itself has no additive-generator multiplier.

The correctly oriented result already proved in Section 36 is

$$\begin{aligned} J_q(\theta_q, \theta_q^2) &= j(\mathfrak{m})\overline{j(\mathfrak{n})}, \\ j(\mathfrak{m})/\sqrt{N\mathfrak{m}} &= \mu(\mathfrak{m})\vartheta(\mathfrak{m}), \qquad \vartheta = \Psi/\sqrt{N}, \quad \operatorname{cond}\Psi = (2\delta). \end{aligned} \tag{221}$$

Its sign is fixed by $-j(p) \equiv 1 \pmod{2\delta}$ and $(j(p)) = p$ for the positive sextic character. The reduction argument includes inert primes. Unique generators congruent to one modulo 2δ then identify $-j(p)$ with $\Psi(p)$. Multiplication yields the displayed identity only on squarefree columns. In particular Ψ has its actual infinity type $(1, 0)$; it is not a finite-order gauge.

Cubic reciprocity for the same primary section and the fact that $\chi_3(-1) = 1$ yield

$$\begin{aligned} \chi_{3,\mathfrak{m}}(\alpha_{\mathfrak{n}}) &= \chi_{3,\mathfrak{n}}(\alpha_{\mathfrak{m}}), \\ \tau_q(\theta_q^2)/\sqrt{Nq} &= G_{\mathfrak{m}}\overline{G_{\mathfrak{n}}}. \end{aligned} \tag{222}$$

The cubic cross factors cancel, while $\tau(\bar{\chi}_3) = \overline{\tau(\chi_3)}$. The quadratic Gaussian evaluation in the same two-factor completion is $\tau_q(\theta_q^3)/\sqrt{Nq} = \Gamma(\alpha_{\mathfrak{m}}\alpha_{\mathfrak{n}})$, with the period-four formula in (195). For composite q, this uses that Gauss evaluation with its exact pairing; $1 + \theta_q^3$ is not the root-count function of a composite residue ring. Combining these facts with (220) proves (195). There is no additional $\theta_q^3(-1)$ factor.

For the original good $\mathfrak{k}$,

$$\begin{aligned} \eta(\mathfrak{k}^2) &= \theta_S(\alpha_{\mathfrak{k}})^2\chi_{3,\mathfrak{m}}(\alpha_{\mathfrak{k}})\overline{\chi_{3,\mathfrak{n}}(\alpha_{\mathfrak{k}})}, \\ \overline{G_{\mathfrak{m}}}G_{\mathfrak{n}}\chi_{3,\mathfrak{m}}(z)\overline{\chi_{3,\mathfrak{n}}(z)} &= \overline{g_{\mathfrak{m}}(z)}g_{\mathfrak{n}}(z). \end{aligned} \tag{223}$$

Before the pure specialization the row factor $R_{\mathfrak{m}}(\mathfrak{v}_2^2)\overline{R_{\mathfrak{n}}(\mathfrak{v}_2^2)}$ makes $z = \alpha_{\mathfrak{k}}\alpha_{\mathfrak{v}_2}$. All raw collisions are zeros. The separate original $(\mathfrak{k}, \mathfrak{v}_2) = 1$ condition is not replaced by a frequency assertion. The primitive factors at $\mathfrak{e}$ and $\mathfrak{w}$ are $\eta(\mathfrak{e})$ and $\overline{\eta(\mathfrak{w})}$. When expressed with generators, their S parts are the separate factors $\theta_S(\gamma_{\mathfrak{e}})$ and $\overline{\theta_S(\gamma_{\mathfrak{w}})}$. If either is zero the component vanishes. No quotient of these arguments, division by a primitive zero, or primary section for an S-supported ideal is used. Changes of these generators cancel in the total unit-trivial character.

49.3 Matched finite endpoint models and a real-axis remainder

For $z \geq 0$ and an integer $J \geq 0$, the angular cosine formula and its real Taylor remainder give

$$\left| J_0(2\sqrt{z}) - \sum_{j=0}^{J} \frac{(-z)^j}{(j!)^2} \right| \leq \frac{z^{J+1}}{((J+1)!)^2}. \tag{224}$$

Indeed use the cosine Taylor polynomial through degree $2J+1$ inside its angular integral. The remainder is at most $|2\sqrt{z}\cos\theta|^{2J+2}/(2J+2)!$, and the normalized angular moment of $\cos^{2j}\theta$ is $(2j)!/(4^j(j!)^2)$. This proves both the coefficients and the bound for every nonnegative z, without an exponential majorant or a sign assertion about the error.

Set

$$m_j=\int_0^\infty v^jG(v)\,dv,\qquad I_j(\mathfrak w)=\frac{b_{\mathfrak w}^{6j+1}-V^{6j+1}}{6j+1}=\int_V^{b_{\mathfrak w}}t^{6j}\,dt. \tag{225}$$

The exact order-J signed model of (190) is

$$\mathcal L_J[f_1,f_2]=C_0^{-1/2}\sum_\theta d_\theta[f_1,f_2]\sum_{j=0}^{J}\frac{(-1)^jm_j\lambda_\theta^jI_j(\mathfrak w)}{C_0^j(j!)^2}\,Q_{2+2j}^{T_q}(K;\eta^2). \tag{226}$$

The character in every Q_s is still η^2. Both finite endpoints, every primitive root and every original finite $\mathfrak k$ cutoff remain. In particular $I_0=X^{1/6}t_{\mathfrak w}$ and (226) at $J=0$ is exactly (191).

Applying (224) before the original $\mathfrak k$ sum and the t integral gives a per-θ bound

$$\frac{C_0^{-1/2}m_{J+1}(\lambda_\theta/C_0)^{J+1}I_{J+1}(\mathfrak w)}{((J+1)!)^2}\sum_{\mathfrak k\ \text{original}}(N\mathfrak k)^{-4-2J}. \tag{227}$$

Only this certified remainder takes absolute values. On the full support $\lambda_\theta b_{\mathfrak w}^6=A_1X/(N\mathfrak eF)\le\Lambda$. Since $m_{J+1}\le 2(5/2)^{J+1}$ and $\sum(N\mathfrak k)^{-4-2J}$ is bounded by one constant independent of J and all deletions, put $\Omega=(5/2)\Lambda/C_0$. Counting the two actual annular column sets only in the error yields

$$|\mathcal M_{\mathrm{top}}-\mathcal L_J[A_U,A_U]|\ll B_W^2(DX)^\epsilon A_1X^{1/6}\frac{\Omega^{J+1}}{((J+1)!)^2}. \tag{228}$$

Here the two $O(D)$ column counts cancel $1/(D\sqrt F)\le D^{-2}$, the S/source/ray sets are fixed, the $\mathfrak w$ set has bounded size, and each interval has length at most V. The constant in this inequality is independent of J and of derivatives of G. At $J=0$ it proves (192). The infinite Bessel hierarchy was already available; the statement used here is the finite real-axis certificate. It bounds the remainder, not the signed polynomial.

49.4 All source and ray phases give bounded multiplicative columns

Freeze $\mathfrak n,\mathfrak e,\mathfrak w,\mathfrak k$ and one source/ray component in the pure model. After (195), every $\mathfrak m$ dependence apart from $\mu(\mathfrak m)f(\mathfrak m)\overline{G_{\mathfrak m}}W(N\mathfrak m/D)/\sqrt{N\mathfrak m}$ is a product of bounded multiplicative functions, except for two fixed finite selectors. The multiplicative factors include ϑ,ν,g, fixed-numerator residue symbols at $u_0,\gamma_S,\mathfrak k,\mathfrak e,\mathfrak w$, the value $\theta_S(\alpha_{\mathfrak m})$, and the original $(\mathfrak m,\mathfrak n)=1$ deletion. They include zeros: a numerator valuation divisible by six does not restore a collision. Primary generators multiply on the required coprime products.

Let $\mathcal C$ be the fixed source group and $\mathcal H=(\mathcal O/(4))^\times$, of order 12. Finite character expansions give

$$\begin{aligned}\mathbf 1_{[\mathfrak m]=A}&=\frac{1}{|\mathcal C|}\sum_{\chi\in\widehat{\mathcal C}}\overline{\chi(A)}\chi([\mathfrak m]),\\ \Gamma(x\alpha_{\mathfrak n})&=\sum_{\chi\in\widehat{\mathcal H}}a_{\chi,\mathfrak n}\chi(x),\qquad\sum_\chi|a_{\chi,\mathfrak n}|\le\sqrt{12}.\end{aligned} \tag{229}$$

The first coefficient mass is one. The second follows by normalized Parseval and Cauchy since the function has modulus one on $\mathcal H$. These bounds are independent of $\mathfrak n,\mathfrak k,D$. Extending a fixed formula

off its source class only extends a coefficient function; the finite selector sum restores the original class. No primitive root of a wrong-class or overlapping-conductor pair is asserted.

Consequently every actual fixed-other-column expression is a finite linear combination, of fixed total coefficient mass, of the linear forms in (200). The associated λ may have growing conductor and other-column deletion, but $|\lambda| \le 1$ and multiplicativity on coprime products suffice. This does not flatten the original roots or make the coprime matrix positive. It is not a silent extension of a same-class, $\rho = 1$ formula.

49.5 Two exact cutoff boxes expose a physical cubic matrix

To prove (196), in each term of (183) set $\mathfrak{c} = \mathfrak{m}/(\mathfrak{a}\mathfrak{b})$. Squarefreeness makes the factors pairwise coprime and $\mu(\mathfrak{m})\mu(\mathfrak{a})\mu(\mathfrak{b}) = \mu(\mathfrak{c})$. The cutoff difference is the exact disjoint assignment

$$\begin{aligned}&\mathbf{1}_{N\mathfrak{a},N\mathfrak{b}\le U} - \mathbf{1}_{N\mathfrak{a},N\mathfrak{b}\le Y}\\ &\quad = \mathbf{1}_{Y<N\mathfrak{a}\le U,\, N\mathfrak{b}\le U} + \mathbf{1}_{N\mathfrak{a}\le Y,\, Y<N\mathfrak{b}\le U}.\end{aligned} \tag{230}$$

Swap $\mathfrak{a}, \mathfrak{b}$ in the second piece and put $\mathfrak{z} = \mathfrak{b}\mathfrak{c}$ in both. On coprime primary $\mathfrak{a}, \mathfrak{z}$, Gauss CRT has the cross $\chi_{3,\mathfrak{a}}(\alpha_{\mathfrak{z}})^2$. Conjugation and cubic reciprocity give exactly (198). Thus the left side of (200), before taking its absolute value, is

$$\begin{aligned}\sum_{\substack{\mathfrak{a},\mathfrak{z}\ \text{good,SF}\\ Y<N\mathfrak{a}\le U}} &\frac{\overline{G_{\mathfrak{a}}}\lambda(\mathfrak{a})}{\sqrt{N\mathfrak{a}}}\,\frac{\overline{G_{\mathfrak{z}}}\lambda(\mathfrak{z})[T_U(\mathfrak{z}) + T_Y(\mathfrak{z})]}{\sqrt{N\mathfrak{z}}}\\ &\cdot \chi_{3,\mathfrak{a}}(\alpha_{\mathfrak{z}})W(N\mathfrak{a}N\mathfrak{z}/D).\end{aligned} \tag{231}$$

At a common prime the cubic entry is zero. Hence this encompassing SF rectangle encodes the original squarefreeness of $\mathfrak{m} = \mathfrak{a}\mathfrak{z}$, without inventing a primitive overlapping root. Also $\mathbf{1}_{(\mathfrak{a}\mathfrak{z},\mathfrak{n})=1} = \mathbf{1}_{(\mathfrak{a},\mathfrak{n})=1}\mathbf{1}_{(\mathfrak{z},\mathfrak{n})=1}$, and the other fixed-numerator collision masks split in the same way. In particular the full other-column mask survives. There is no assumption that A_U or T_L is multiplicative. The elementary estimate $|T_U(\mathfrak{z}) + T_Y(\mathfrak{z})| \le 2\tau(\mathfrak{z})$ is uniform in both cutoffs.

49.6 Cubic Cauchy before the distinguished-divisor triangle

Use dyads $N\mathfrak{a} \asymp A$ in $(Y, U]$. Only $N\mathfrak{z} \asymp Z$, $Z \asymp D/A$, can meet the original profile. Both indices in (231) are good squarefree ideals, hence cube-free. Primary cubic reciprocity identifies the entry with the physical family of (199), or its transpose. Additional row/column masks are coefficient contractions inside that arbitrary-coefficient norm. No character indexed only by a growing conductor is substituted.

Ideal counting and the divisor estimate show, with arbitrarily small exponent losses,

$$\begin{aligned}\sum_{N\mathfrak{a}\asymp A}\left|\frac{\overline{G_{\mathfrak{a}}}\lambda(\mathfrak{a})}{\sqrt{N\mathfrak{a}}}\right|^2 &\ll 1,\\ \sum_{N\mathfrak{z}\asymp Z}\left|\frac{\overline{G_{\mathfrak{z}}}\lambda(\mathfrak{z})[T_U(\mathfrak{z}) + T_Y(\mathfrak{z})]}{\sqrt{N\mathfrak{z}}}\right|^2 &\ll_\epsilon D^\epsilon.\end{aligned} \tag{232}$$

These are coefficient bounds after the exact signed grouping, not a new bound for the norm of A_U.

The exact Mellin separation of the one common occurrence of W is

$$W(N\mathfrak{a}N\mathfrak{z}/D) = \frac{1}{2\pi}\int_{\mathbb{R}} \widehat{W}(t)D^{-it}(N\mathfrak{a})^{it}(N\mathfrak{z})^{it}\,dt. \tag{233}$$

Its powers are unitary on this line, so it costs $\mathcal{M}(W)$ once. The arbitrary-coefficient cubic operator norm and (232) therefore give

$$|\text{one dyadic form}| \ll_\epsilon \mathcal{M}(W) D^\epsilon \sqrt{A + D/A + D^{2/3}}. \tag{234}$$

This is a bilinear matrix inequality, not positivity of the original coprime form. Summing the finite dyads with $A \geq Y/2$ and $A \leq U \leq \sqrt{D}$ gives

$$\mathcal{M}(W) D^\epsilon \{\sqrt{U} + \sqrt{D/Y} + D^{1/3}\} \ll \mathcal{M}(W) D^\epsilon \{\sqrt{D/Y} + D^{1/3}\}. \tag{235}$$

Dyadic logarithms and divisor losses are absorbed by distributing a prescribed ϵ among the preceding bounds. Partial dyads and real cutoff ties are literal coefficient indicators. Empty tails, including $Y = U$, are exactly zero. This proves (200).

49.7 Full pair support and the original finite k reassembly

Fix the other column $\mathfrak{n}$. Its raw coprimality and numerator deletions are already in λ by Section 49.4. An optional condition $N\mathfrak{m} \neq N\mathfrak{n}$ is not multiplicative. The number of ideals of a fixed rational norm N is at most the integer divisor function: at a split prime an exponent e gives at most $e+1$ choices, at an inert prime at most one, and at a ramified prime one. Consequently there are $O_\epsilon(D^\epsilon)$ ideals in the forbidden norm fibre, with all orientations separate. Since $|A_U| + |A_Y| \ll_\epsilon D^\epsilon$, subtracting it costs

$$O_\epsilon(B_W D^{-1/2+\epsilon}) \quad \text{in the fixed-other-column sum.} \tag{236}$$

This is below (200); $\mathcal{M}(W) \geq B_W$ by Fourier inversion. No whole equal-norm energy is subtracted from a different frequency sector.

The untouched original column has the elementary bound

$$\sum_{\mathfrak{n} \text{ original}} \frac{|A_T(\mathfrak{n}) W(N\mathfrak{n}/D)|}{\sqrt{N\mathfrak{n}}} \ll_\epsilon B_W D^{1/2+\epsilon}, \qquad T \in \{U, Y\}. \tag{237}$$

Take this triangle only after (200). In the leading model all S, source/ray and $N\mathfrak{w} < 64$ coefficient masses are fixed; g_0 is bounded independently of capacity derivatives, and $F^{-1/2}$ supplies the two square-root norm weights already used. The normalization is once by $1/D$. For each actual fixed $\mathfrak{k}$, the column estimate is uniform in its bounded multiplicative numerator phase. Thus its *original finite* sum is bounded by $\sum (N\mathfrak{k})^{-2} < \infty$, with no $(N\mathfrak{k})^\epsilon$ or altered conductor cost. One does not need to complete that sum.

Equivalently, define $\mathcal{G}_D^K$ by $\mathcal{L}_0[f_1, f_2] = A_1 X^{1/6} \mathcal{G}_D^K[f_1, f_2]$. The preceding estimates give

$$|\mathcal{G}_D^K[A_U - A_Y, A_T]| \ll B_W \mathcal{M}(W) D^\epsilon \{D^{-1/6} + Y^{-1/2}\}. \tag{238}$$

For a tail on the second column, conjugate the scalar expression and reverse the roles. This sends $G_\mathfrak{n}$ to $\overline{G_\mathfrak{n}}$, W to $\bar{W}$ and the finite coefficients to their conjugates. The conjugate cubic matrix has the same norm, and $\mathcal{M}(\bar{W}) = \mathcal{M}(W)$. No Hermitian assertion about an individual source block is required. Finally use the exact bilinear identity

$$\begin{aligned} \mathcal{G}_D^K[A_U, A_U] - \mathcal{G}_D^K[A_Y, A_Y] = {} & \mathcal{G}_D^K[A_U - A_Y, A_U] \\ & + \mathcal{G}_D^K[A_Y, A_U - A_Y]. \end{aligned} \tag{239}$$

This proves (201) with all original finite cutoffs, primitive S zeros and pair masks.

49.8 The critical reduction is not a bound for its residual mean

Equations (192) and (201) prove (202). When $U \ge D^{1/8}$, substituting $A_1 \asymp D^{3/4}$, $X = D^{3/8}$, $Y = D^{1/8}$ gives the exact exponent calculations

$$\begin{aligned} 2(3/4) + (7/6)(3/8) - 2 &= -1/16, \\ 3/4 + (1/6)(3/8) - (1/2)(1/8) &= 3/4, \\ 3/4 + (1/6)(3/8) - 1/6 &= 31/48. \end{aligned} \tag{240}$$

If $U < D^{1/8}$ choose $Y = U$; the coefficient difference then vanishes identically, rather than replacing $Y^{-1/2}$ by a smaller number. The small-cutoff moment remains exactly the auxiliary expression with that Y. Since both short divisors have norms at most Y, their residual in (203) has norm at least D/Y^2. All other-column and finite-frequency couplings remain in its two-column sum.

The current separate discrepancy has critical power

$$3/4 + (1/18)(3/8) = 37/48 > 3/4. \tag{241}$$

This upper certificate is not a noncancellation theorem. A bound for the small-cutoff leading model alone does not cancel or improve it. Both it and the remaining original sectors require additional signed arithmetic.

49.9 Direct discrete frequencies require no principal-density approximation

For each original dual ideal, its good valuations determine uniquely $\mathfrak{y} = \mathfrak{w}\mathfrak{s}^6$ with good $\mathfrak{s}$ and good part of $\mathfrak{w}$ sixth-free. At $V = X^{1/6}/2$, the lower-open condition $N\mathfrak{s} > V$ implies $N\mathfrak{w} < 64$. Conversely every ideal of norm below 64 has sixth-free good part, since the least good prime norm is seven. This proves the exact support (204) with all its nonunit S modes and possible overlaps. No squarefreeness of $\mathfrak{s}$ or new coprimality with $\mathfrak{k}$ is introduced.

Equation (211) holds for *every* original ρ. Retaining that factor and the actual discrete $\mathfrak{s}$ family in (140) gives (206). The prefactor is exactly $A_1/(D\sqrt{F})$, including the once-only normalization. All primitive factors at $\mathfrak{e}, \mathfrak{w}$ are still separate, and a good collision remains zero even when its exponent is a sixth power.

The pointwise real-axis remainder (224) at $J = 0$ implies

$$|\dot{G}(x) - g_0| \le C_0^{-3/2} x \int_0^\infty vG(v)\,dv \le 3C_0^{-3/2} x, \qquad x \ge 0. \tag{242}$$

Every argument in (206) is at most $\Lambda = A_1 X/D^2$. Count only in this error: there are $O(D^2)$ annular column pairs, $A_1/(D\sqrt{F}) \le A_1/D^2$, the original coefficient products are $O(B_W^2 D^\epsilon)$, and the original finite $\mathfrak{k}$ absolute budget is bounded by $\sum (N\mathfrak{k})^{-2}$. The $\mathfrak{e}$/source/ray and $\mathfrak{w}$ sets have fixed finite mass. Since $V \ge 1$, the exact good $\mathfrak{s}$ annulus contains $O(V)$ ideals. This proves (207), without a ray count, interpolated endpoint, density or arbitrary row-weight norm.

It remains to justify applying the column-tail theorem to every original ray. The good-q powers in the proof of (195) are local powers $1, 2, 3$ of θ_q; that proof never uses $\rho^6 = 1$. For every original component the total η is type zero and its finite θ_S is fixed at its actual primitive S conductor. Hence the finite selector expansion in (229) and its coefficient mass are valid on all these components. No new assertion of sixth order for an individual ξ or replacement of a primitive S Gauss sum is made.

Now freeze $\mathfrak{s}$ as well as $\mathfrak{n}, \mathfrak{k}, \mathfrak{e}, \mathfrak{w}$. Relative to Section 49.4, the only new $\mathfrak{m}$ factor is $\mathbf{1}_{(\mathfrak{m},\mathfrak{s})=1}$. It is a bounded multiplicative ideal function and splits exactly in the grouped column as

$$\mathbf{1}_{(\mathfrak{a}\mathfrak{z},\mathfrak{s})=1} = \mathbf{1}_{(\mathfrak{a},\mathfrak{s})=1} \mathbf{1}_{(\mathfrak{z},\mathfrak{s})=1}. \tag{243}$$

The $\mathfrak{n}/\mathfrak{s}$ mask is a scalar which may be zero. The fixed $\bar{\rho}(\mathfrak{s})^6$ phase is another bounded scalar. All earlier full-union raw deletions and separate primitive S values remain. Thus (200) is uniform in $\mathfrak{s}$ as well as in $\mathfrak{n}, \mathfrak{k}$, with no $(N\mathfrak{s})^{\epsilon}$ or $(N\mathfrak{k})^{\epsilon}$ cost. Their growing deleted primes are coefficient diagonals in the same physical cubic matrix.

There is no density in the new exact discrete model. Use the underlying arbitrary-λ theorem with the factor $g(\mathfrak{m})$ omitted, rather than dividing the earlier density-weighted moment by $g(q)$. If a $\mathfrak{k}$ domain has been written with q deletion, first express it on the common original good-SF set $N\mathfrak{k} \le K$ with its exact raw $\eta(\mathfrak{k})^2$ zero. This permits a fixed $\mathfrak{k}$ column argument, restores no prime and adds no $(\mathfrak{k}, \mathfrak{s})$ condition.

Apply the signed column estimate before the other column triangle. The latter is (237); after once-only $1/D$ it yields $B_W \mathcal{M}(W) D^{\epsilon}(Y^{-1/2} + D^{-1/6})$ for each fixed $\mathfrak{s}, \mathfrak{w}, \mathfrak{k}, \mathfrak{e}$ and finite component, with the original $A_1 g_0$ factor. Sum the original finite $\mathfrak{k}$ norm weights and the fixed $\mathfrak{e}, \mathfrak{w}$ components. Only now use the $O(V)$ count of the exact $\mathfrak{s}$ family. Reversed tails use conjugation, as in Section 49.7, and the original unequal-norm selector uses the same separate fixed-norm fibre bound (236). The exact identity

$$\begin{aligned} \mathcal{L}_X^K[A_U, A_U] - \mathcal{L}_X^K[A_Y, A_Y] = \mathcal{L}_X^K[A_U - A_Y, A_U] \\ + \mathcal{L}_X^K[A_Y, A_U - A_Y] \end{aligned} \tag{244}$$

then proves (208). Combining it with the one Bessel error proves (209).

This direct proof does not assert a smaller value for the separate $37/48$ certificate in the continuous mean route; it never introduces that split. All of its nonprincipal-ray and exact-discrete effects remain in $\mathcal{L}_X^K[A_Y, A_Y]$. At $Y \le D^{1/8}$ the residual column norms are at least $D/Y^2 \ge D^{3/4}$, just as in (203). The new model is still unbounded at target size. It is different from the principal-density $\mathcal{L}_0[A_Y, A_Y]$.

A capacity-empty integer interval, or empty exact $\mathcal{T}_X$, gives zero by definition. Absence of actual good squarefree rows in a *nonempty* integer interval alone does not imply that each dual prefix piece is zero; cancellation with other frequencies can be required. No full dual complement, other block, whole signed assembly or determinant/zero-location obligation is resolved by this direct small-cutoff reduction.

Why the completed shortcuts do not supply that arithmetic. The inverse-different cubic Gauss frequency is incomplete: expanding it leaves the actual weighted primary $\mathfrak{k}$ set, not a complete residue system modulo $\mathfrak{m}\mathfrak{n}$. The bounded number of $\mathfrak{w}$ modes supplies no such completion. Likewise, expanding both short divisors and taking their absolute triangles before the cubic cross does not give the proved estimate. Even an explicitly hypothetical uniform inner exponent $\alpha < 1$ would produce the diagnostic

$$D^{\alpha - 1/2} \left(\sum_{N\mathfrak{a} \le U} (N\mathfrak{a})^{-\alpha} \right)^2 \ll_{\alpha} D^{\alpha - 1/2} U^{2(1-\alpha)}. \tag{245}$$

At $U \asymp \sqrt{D}$ this is $D^{1/2}$. It is a conditional failure of that method, not an available smooth-coefficient Gauss theorem or an impossibility result. The actual proof above retains the physical cubic cross until its joint matrix norm. It bounds one specified signed tail and leaves the long residual-Möbius, finite-$\mathfrak{k}$ small-short moment as the next arithmetic problem. No whole signed, global unsigned, determinant, zero-location, analytic Lean or GRH conclusion is implied.

50 Ramified theta transforms and the original pure dual prefix

The exact discrete prefix introduced in Section 48 can be estimated with its original coefficient A_U. The argument constructs the previously missing finite-character transform at the ramified prime and groups all original short and Heath–Brown factors before taking a short-factor triangle. The resulting power is 101/144 for this specified pure finite dual prefix. The whole signed comparison 13/12, the global unsigned comparison 9/8, and the remaining infinite-frequency assembly retain their earlier status.

The retained principal-density mean in the earlier discussion is a different expression with its own comparison problem. The result here uses every original ray component and the actual discrete sixth-part set. It supplies a new proof for the actual finite discrete prefix rather than modifying the historical cutoff stages.

50.1 The normalized object and its original support

Put $\omega = e^{2\pi i/3}$, $K = \mathbb{Q}(\omega)$, $\mathcal{O} = \mathbb{Z}[\omega]$ and $\lambda = 1 + 2\omega = i\sqrt{3}$. A good ideal avoids the original fixed set S, containing the primes above $2, 3$. All ideal columns are good and squarefree; different ideals of equal norm remain different columns. Use the unique multiplicative primary generator $\alpha_{\mathfrak{a}} \equiv 1 \pmod 3$. The original character ν has fixed finite order, with its conductor exclusions in S; no order-dividing-six hypothesis is imposed on ν.

Assume

$$D \geq 2^{64}, \quad 1 \leq U \leq \sqrt{D}, \quad A_1 \asymp H = D^{3/4}, \quad X = D^{3/8}, \quad V = X^{1/6}/2. \tag{246}$$

Only the original unweighted pure two-pivot block is asserted: $A_2 = 1$, its second good factor and the other four good factors are units, and u_0 is the fixed literal unit/S-supported factor with no good prime support. Keep its original capacity endpoints, nodal function G, and pre-S cutoff $K_{\mathrm{cut}} = \sqrt{b}$. An empty capacity interval gives zero. None of these data is recomputed during a coefficient decomposition.

Let W be one fixed possibly complex smooth function supported on $[1, 2]$, with $W(1) = W(2) = 0$, and write

$$\begin{aligned} B_W &= \|W\|_\infty, \qquad \widehat{W}(t) = \int_0^\infty W(x) x^{-it} \frac{dx}{x}, \\ \mathcal{M}(W) &= \frac{1}{2\pi} \int_{\mathbb{R}} |\widehat{W}(t)|\, dt, \qquad\qquad V_1(W) = B_W + \int_1^2 |W'(x)|\, dx. \end{aligned} \tag{247}$$

These are seminorms of the common column profile. No uniform derivative hypothesis on the sharp capacity interpolant is added. The exact frequency set is

$$\mathcal{T}_X = \left\{ (\mathfrak{w}, \mathfrak{s}) : \begin{array}{l} \mathfrak{w} \text{ integral, } N\mathfrak{w} < 64, \quad \mathfrak{s} \text{ good integral,} \\ V < N\mathfrak{s} \leq 2V, \quad N\mathfrak{w}(N\mathfrak{s})^6 \leq X \end{array} \right\}. \tag{248}$$

There is no squarefreeness requirement on $\mathfrak{s}$. The permitted overlaps among $\mathfrak{k}, \mathfrak{s}, \mathfrak{w}$ remain permitted, including nonunit S-supported $\mathfrak{w}$.

Here is a full definition of the piece being bounded. Let $\Omega = (A, B, \rho, \mathfrak{m}, \mathfrak{n})$ range over the original finite source/ray components and their ordered coprime good squarefree columns in $[D, 2D]$, with the original source classes. An original unequal-norm projection may be imposed consistently. Use the actual primitive character η_Ω of (134), its actual S-conductor $\mathfrak{f}_S$, and

$$F_\Omega = N\mathfrak{f}_S\, N\mathfrak{m}\, N\mathfrak{n}, \qquad \mathfrak{q} = \mathfrak{m}\mathfrak{n}, \qquad c_U(\mathfrak{m}) = A_U(\mathfrak{m}) \nu(\mathfrak{m}) R_{\mathfrak{m}}(u_0) W(N\mathfrak{m}/D).$$

The numerator value $R_{\mathfrak{m}}(u_0)$ includes its original raw zeros. Put $P_S = \prod_{\mathfrak{p}\in S}\mathfrak{p}$. With $C_0 = 3/(4\pi^2)$, the normalized prefix is

$$\begin{aligned}\mathcal{P}_X[A_U, A_U] = \frac{A_1}{D}\sum_{\Omega} h_{AB\rho}(\mathcal{O})c_U(\mathfrak{m})\overline{c_U(\mathfrak{n})}\frac{\varepsilon(\eta_\Omega)}{\sqrt{F_\Omega}} \\ \cdot \sum_{\mathfrak{e}\mid P_S}\frac{\mu(\mathfrak{e})\eta_\Omega(\mathfrak{e})}{N\mathfrak{e}}\sum_{\substack{\mathfrak{k}\ \text{good,SF}\\ N\mathfrak{k}\le K_{\mathrm{cut}}}}\frac{\mu(\mathfrak{k})\eta_\Omega(\mathfrak{k})^2}{(N\mathfrak{k})^2} \\ \cdot \sum_{(\mathfrak{w},\mathfrak{s})\in\mathcal{T}_X}\overline{\eta_\Omega(\mathfrak{w})}\,\overline{\rho(\mathfrak{s})}^{\,6}\mathbf{1}_{(\mathfrak{s},\mathfrak{q})=1}\dot{G}\left(\frac{A_1 N\mathfrak{w}(N\mathfrak{s})^6}{(N\mathfrak{k})^2 N\mathfrak{e}F_\Omega}\right).\end{aligned} \tag{249}$$

All sums are finite. The ray label ρ here is distinguished from the Jacobi column weight below. Values at $\mathfrak{e}$ and $\mathfrak{w}$ are separate primitive values, including zero; a conductor drop does not restore a deleted prime. The original $1/D$ occurs once. The raw numerator is $D\mathcal{P}_X$.

Define $\mathcal{L}_X^{K_{\mathrm{cut}}}$ by replacing $\dot{G}$ in (249) by $g_0 = C_0^{-1/2}\int_0^\infty G(v)\,dv$, with every other datum fixed. For the original coefficients the established real-axis estimate is

$$|\mathcal{P}_X - \mathcal{L}_X^{K_{\mathrm{cut}}}| \ll B_W^2(DX)^{\epsilon}A_1V\Lambda, \qquad \Lambda = A_1X/D^2. \tag{250}$$

It uses $0 \le G \le 1$, its fixed compact support and a pointwise Bessel remainder. The value g_0 may vanish and is never divided out.

50.2 Actual finite characters and a sharp ramified row

The actual root calculation leaves one angular Jacobi gauge and a cubic Gauss root in each varying column. On good squarefree ideals the Jacobi function is $\rho_J = \mu\vartheta$, where $\vartheta = \Psi/\sqrt{N}$ has infinity type $(1,0)$; here $\Psi(\mathfrak{a})$ is the chosen generator 1 mod 2λ. The arithmetic μ belongs in the coefficient opening. There is exactly one varying ϑ, or its conjugate; an additional finite-order assertion about ρ_J would be false.

Write $\chi_{3,\mathfrak{a}}$ for the positive cubic residue character, with its zero at every shared prime. The source Gauss root is

$$G_{\mathfrak{a}}^{\mathrm{src}} = \frac{1}{\sqrt{N\mathfrak{a}}}\sum_{x \bmod \alpha_{\mathfrak{a}}}\chi_{3,\mathfrak{a}}(x)e^{2\pi i\,\mathrm{Tr}(x/\alpha_{\mathfrak{a}})}. \tag{251}$$

The actual root uses $\lambda\alpha_{\mathfrak{a}}$ in this additive denominator. The normalization of the angular factor is explicit:

$$\begin{aligned}\Theta_\ell(\mathfrak{a}) &= (\alpha_{\mathfrak{a}}/|\alpha_{\mathfrak{a}}|)^\ell, \\ \vartheta(\mathfrak{a}) &= \Theta_1(\mathfrak{a})\overline{\chi_{3,\mathfrak{a}}(2)}, \\ G_{\mathfrak{a}}^{\mathrm{act}} &= \chi_{3,\mathfrak{a}}(\lambda)G_{\mathfrak{a}}^{\mathrm{src}}.\end{aligned} \tag{252}$$

Thus the two fixed corrections may be written as source frequency $12t$. The factor 12 belongs to the fixed finite character data; $12t$ is not declared a good squarefree moving level.

After exact product-character expansion of the joint finite S/source/Γ selectors, every remaining own-variable phase has the form $\Theta_{\pm1}\kappa$, where κ belongs to a fixed finite residue family, with its original modulus-zero extension. It includes higher 3-primary conductors of ν. The finite source quotient is a fixed ray quotient; fixed numerator symbols are finite Kummer characters. The full Fourier sums retain their actual coefficient masses and reconstruct the primitive $\mathfrak{e}/\mathfrak{w}$ zeros. Growing other-column deletions remain literal masks.

For a physical good squarefree frequency $\mathfrak{t}$, an exact interval $I \subset [T, 2T]$, and a finite good deletion ideal Q, write the actual row in either required conjugate orientation as

$$\mathcal{F}_\psi(T; \mathfrak{t}; Q, I) = \sum_{\substack{\mathfrak{a} \text{ good,SF} \\ N\mathfrak{a} \in I,\ (\mathfrak{a}, Q)=1}} \psi(\mathfrak{a}) \overline{G_{\mathfrak{a}}^{\mathrm{act}}} \chi_{3,\mathfrak{a}}(\alpha_{\mathfrak{t}}) \frac{W(N\mathfrak{a}/T)}{\sqrt{N\mathfrak{a}}}. \tag{253}$$

Here ψ is from this actual fixed family, including its angular factor. It is not an arbitrary bounded multiplicative function.

Theorem 37 (The actual fixed-family sharp row)**.** *For every such member, both conjugate orientations and every physical good squarefree* $\mathfrak{t}$*, including the unit,*

$$|\mathcal{F}_\psi(T; \mathfrak{t}; 1, I)| \ll_{\mathrm{fixed},\epsilon} V_1(W)(2 + TN\mathfrak{t})^\epsilon \min\{\sqrt{T},\ T^{1/4}(N\mathfrak{t})^{3/8}\}, \qquad T \geq 1. \tag{254}$$

The constant may depend on the fixed conductor/character family. A literal unit calculation extends the estimate to $1/2 \leq T < 1$*; for* $T \leq 1/2$ *the row is zero with the stated endpoint convention.*

Proof. Use the unconditional theta modularity and known 27 integral cusps in Dunn–Radziwiłł, arXiv:2109.07463v3, Section 5 and Appendix A. Its weight-zero cubic theta θ is automorphic with Kubota multiplier χ on $\Gamma_2 = \mathrm{SL}_2(\mathbb{Z})\Gamma_1(3)$, where $\Gamma_1(3)$ is the principal congruence subgroup of $\mathrm{SL}_2(\mathcal{O})$. Fix its known representatives γ_j and $F_j = \theta \circ \gamma_j$. The correct ramified coefficient projector acts on $\lambda^3\mu$. Its primary selector is $\mathbf{1}_{\lambda^3\mu \equiv 1 \bmod 3}$. A finite Fourier expansion writes it, and the away-3 periodic twist, as actual rational translates of theta.

For $a = A/C$ reduced, choose $\gamma = \left(\begin{smallmatrix} A & B \\ C & D_0 \end{smallmatrix}\right)$ of determinant one, and choose j so that $\delta = \gamma\gamma_j^{-1} \in \Gamma_2$. Use $U(x) = \left(\begin{smallmatrix} 1 & x \\ 0 & 1 \end{smallmatrix}\right)$ and $W_C = \left(\begin{smallmatrix} 0 & -1/C \\ C & 0 \end{smallmatrix}\right)$. Then

$$\gamma^{-1}U(a) = U(-D_0/C)W_{-C}, \qquad \theta(U(a)w) = \chi(\delta)F_j(U(-D_0/C)W_{-C}w).$$

The multiplier is its literal Kubota value. Pure nonzero angular derivatives kill the pulled-back radial constants. Their Mellin identity has factor $(NC)^{1-2s}$, including the Jacobian. The explicit cusp coefficients have absolute Dirichlet norm at most

$$3^8 \zeta_K^{(\lambda)}(\sigma)\zeta_K^{(\lambda)}(3\sigma - \tfrac{1}{2}), \qquad 1 < \sigma < 13/12.$$

Normalized finite Fourier mass is $O_{\mathrm{fixed}}(\sqrt{N\mathfrak{t}})$ and each reduced denominator has norm $O_{\mathrm{fixed}}(N\mathfrak{t})$. This gives the conservative dual cost $(N\mathfrak{t})^{3/2+\epsilon}$.

Coefficientwise cube completion retains the raw $L(3s - \frac{1}{2}, \kappa^3\Theta_{3\ell})$, with every stipulated Euler omission. Apply convexity to this entire completed numerator first. Its reciprocal is used only where $\Re(3s - \frac{1}{2}) > 1$, by absolutely convergent Euler products. Exact half-integer Perron thresholds and partial summation give (254). The optimized height is $(T/(N\mathfrak{t})^{3/2})^{1/4}$ when this is at least one; the elementary count gives the other branch. The extended edition retains the constants and the zeroth-angular residues explicitly. □

50.3 The full original coefficient and its finite prefix

On the actual contributing annulus $N\mathfrak{m} > U$, put

$$B_U(\mathfrak{m}) := \mu(\mathfrak{m})A_U(\mathfrak{m}) = \sum_{\substack{\mathfrak{abc}=\mathfrak{m} \\ N\mathfrak{a}, N\mathfrak{b} \leq U}} \mu(\mathfrak{c}). \tag{255}$$

This B_U is not the project's original C_U. Equation (255) is not asserted for arbitrary small columns or the unit. Unit factors inside an annular factorization remain present.

Theorem 38 (Original pure finite dual-prefix bound)**.** *Under* (246)*, with the complete original support and fixed data of* (249)*,*

$$|\mathcal{P}_X[A_U, A_U]| \ll_{\text{fixed},\epsilon} (DXA_1)^{\epsilon} A_1 V\{B_W[\mathcal{M}(W) + V_1(W)]D^{-1/9} + B_W^2[\Lambda + D^{-1}]\}. \tag{256}$$

In particular its normalized leading power is

$$13/16 - 1/9 = 101/144 < 3/4, \qquad 3/4 - 101/144 = 7/144. \tag{257}$$

The raw numerator has one additional factor D.

Proof. Open (255) and the exact $J = 5$ finite Heath–Brown identity at $z_0 = (2D)^{1/5}$. There are at most eleven factors, including the original $\mathfrak{a}, \mathfrak{b}$. Choose their grouping from fixed independent dyadic scales. A balanced whole-factor partition has a physical cubic matrix bound $\mathcal{M}(W)D^{1/3+\epsilon}$, using the pinned de Faveri, arXiv:2610.04045v1, Theorem 1.1 and Example 2.2, at order three.

An unbalanced box has one long unrestricted coefficient-one factor $\mathfrak{v}$. At $D \geq 2^{64}$ neither an original short factor nor a short Möbius factor can occupy that scale. The product $\mathfrak{p}$ of all other factors has norm $O(D^{1/3})$ and its intrinsic signed assignment coefficient is divisor-bounded. This retains the $j = 1$ boxes, which are all balanced.

For $N\mathfrak{p} > Z$ the cubic matrix gives $\mathcal{M}(W)D^{\epsilon}[D^{1/3} + \sqrt{D/Z}]$. For $N\mathfrak{p} \leq Z$, freeze the original opposite column $\mathfrak{n}$ first. Keep $\mathbf{1}_{(\mathfrak{p},\mathfrak{n})=1}$ and all prefix $\mathfrak{k}/\mathfrak{s}/\mathfrak{w}$ zeros in its actual coefficient. The remaining long-variable deletion is $Q = \mathrm{rad}_{\text{good}}(\mathfrak{p}\mathfrak{n}\mathfrak{k}\mathfrak{s}\mathfrak{w})$. Its exact divisor $\mathfrak{h}$ changes frequency to $\mathfrak{p}\mathfrak{k}\mathfrak{h}$ and length to $D/(N\mathfrak{p}N\mathfrak{h})$. Including the original prefix weight, Theorem 37 gives

$$\frac{1}{\sqrt{N\mathfrak{p}N\mathfrak{h}}}\left(\frac{D}{N\mathfrak{p}N\mathfrak{h}}\right)^{1/4}(N(\mathfrak{p}\mathfrak{k}\mathfrak{h}))^{3/8} = D^{1/4}(N\mathfrak{k})^{3/8}(N\mathfrak{p})^{-3/8}(N\mathfrak{h})^{-3/8}. \tag{258}$$

Thus the small-prefix column costs $V_1(W)D^{\epsilon}D^{1/4}(N\mathfrak{k})^{3/8+\epsilon}Z^{5/8}$. Choose $Z = D^{2/9}$; both column costs have power $7/18$.

Only after that bound take the opposite-column elementary $B_W D^{1/2+\epsilon}$ count. Use the original A_1/D once, sum the original finite $\mathfrak{k}$ with the convergent positive majorant $\sum(N\mathfrak{k})^{-13/8+\epsilon}$, and count the exact $\mathfrak{s}$ family by $O(V)$. Optional unequal norms cost the same fixed-norm fibre bound $O(B_W^2 A_1 V D^{-1+\epsilon})$. Finally (250) applies directly to A_U, A_U. No A_U-to-A_Y replacement or common-long-divisor opening is used. □

The estimate has fixed-conductor and explicit profile-seminorm dependence. It does not supply a supremum-norm estimate for varying W_D, an arbitrary weighted-row operator, the infinite dual complement, mixed or neighboring blocks, short/principal assembly, an independent positive determinant, or a zero-location theorem. The Dirichlet, Hecke, Dedekind, elliptic and general Artin obligations remain separate. No analytic/category Lean formalization or GRH proof is claimed.

51 Proofs of the ramified transform and full-coefficient prefix bound

The external theta input is the unconditional Section 5 and Appendix A of Dunn–Radziwiłł, *Bias in cubic Gauss sums: Patterson's conjecture*, arXiv:2109.07463v3, 14 May 2024. Its GRH-conditional prime theorems are not used. The matrix input is the already admitted physical ideal family in de Faveri, *Optimal large sieve for fixed order characters*, arXiv:2610.04045v1, Theorem 1.1 and Example 2.2, at order three. The theta modularity and deep cubic norm remain external theorems. The constructions and estimates below specify how they apply to the actual finite-character family.

51.1 Fixed residue data, angular order and raw evaluations

The full varying-column phase list, after the arithmetic μA_U opening, consists of one ϑ or its conjugate, the cubic Gauss root, the original finite-order ν, finite source/ray factors, the mod-4 Fourier factors of Γ, the actual primitive S-type factors, and residue symbols at the fixed numerators $u_0, \gamma_S, \mathfrak{e}, \mathfrak{w}$. The set $N\mathfrak{w} < 64$, the S-types and the $\mathfrak{e}$ set are fixed finite sets. The growing $\mathfrak{k}$ numerator belongs to the cubic frequency instead of this list.

There is a common fixed residue modulus for this list. A finite-order Hecke character has infinity type zero and factors through its actual conductor. The source group is explicitly a ray quotient: Section 2.3.2 of de Faveri defines $\mathcal{C} = H_{\mathfrak{c}} \otimes \mathbb{Z}/6\mathbb{Z}$ for one sufficiently deep S-supported ray modulus $\mathfrak{c}$. The chosen primary generators realize these ray characters as finite residue characters. The Γ factors are characters of $(\mathcal{O}/(4))^{\times}$.

The local sixth-power tests also have a fixed finite modulus. For a prime $\mathfrak{p}$, Hensel's criterion at 1 gives $1 + \mathfrak{p}^e \subset (K_{\mathfrak{p}}^{\times})^6$ when $e > 2v_{\mathfrak{p}}(6)$. One may use $e_{(2)} = 3$ and $e_{(\lambda)} = 5$. These are moduli for testing local classes, not replacements for primitive character evaluation moduli. For a fixed nonzero integral numerator ζ, its n-th residue function, $n \mid 6$, is a finite Kummer character for $K(\zeta^{1/n})/K$. A safe sextic conductor bound is supplied by the polynomial discriminant ideal $(6)^6(\zeta)^5$. A field/order discriminant divides this separable polynomial discriminant, including when a proper factor supplies the field. A fractional numerator can be cleared by a fixed sixth power, retaining its original numerator and denominator exclusions. This uses finite one-dimensional Kummer characters, without a general Artin-holomorphy hypothesis.

A joint finite S-type or primitive-value selector is not assumed multiplicative. If $\mathcal{H}$ is its fixed finite residue group, write its actual function as

$$\begin{aligned} f(x,y) &= \sum_{\chi,\psi \in \widehat{\mathcal{H}}} b_{\chi,\psi}\chi(x)\psi(y), \\ b_{\chi,\psi} &= |\mathcal{H}|^{-2} \sum_{x,y \in \mathcal{H}} f(x,y)\overline{\chi(x)\psi(y)}, \\ \sum_{\chi,\psi} |b_{\chi,\psi}| &\le |\mathcal{H}|\,\|f\|_{\infty}. \end{aligned} \tag{259}$$

Normalized Parseval and Cauchy prove the last line. This expansion precedes freezing a primitive type or separating column factors. Forbidden primitive $\mathfrak{e}/\mathfrak{w}$ values are zero in f; individual Fourier terms can be nonzero there, and only their full sum reconstructs the zero. No term is individually renamed an original nonzero row.

Thus one may use a common fixed raw modulus F_{raw} and functions

$$\kappa_{\mathrm{raw}}(\mathfrak{a}) = \kappa_{\mathrm{prim}}(\alpha_{\mathfrak{a}})\mathbf{1}_{(\mathfrak{a},F_{\mathrm{raw}})=1}. \tag{260}$$

They are imprimitive residue characters on that modulus. All stipulated prime zeros survive a primitive conductor drop.

The Jacobi function $\rho_J = \mu\vartheta$ is not one of these finite characters. Its μ has already entered the arithmetic coefficient B_U, and its unique remaining angular factor is ϑ, with order $+1$ or -1 in the two conjugate orientations. Common or prefix Jacobi factors stay in their actual outer coefficients when those variables are frozen. An arbitrary bounded multiplicative function, a growing opposite column, a Möbius coefficient or the nonperiodic density $\prod_{\mathfrak{p}\mid\mathfrak{a}}(1-(N\mathfrak{p})^{-1})$ is not made into a member of this fixed family. That density does not occur in $\mathcal{P}_X$.

51.2 Known integral cusps and an explicit absolute coefficient norm

Let $\Gamma = \mathrm{SL}_2(\mathcal{O})$, $\Gamma_1(3) = \ker(\Gamma \to \mathrm{SL}_2(\mathcal{O}/(3)))$ and $\Gamma_2 = \mathrm{SL}_2(\mathbb{Z})\Gamma_1(3)$. The source theta function is a weight-zero function on hyperbolic three-space with $\theta(\delta w) = \chi(\delta)\theta(w)$ for $\delta \in \Gamma_2$. The multiplier χ is the source Kubota character. The quotient $\Gamma_2\backslash\Gamma$ has 27 cosets. One can see the finite count from $\mathcal{O}/(3) \simeq \mathbb{F}_3[\varepsilon]/(\varepsilon^2)$: the reduction kernel in SL_2 has 3^3 elements, and $\mathrm{SL}_2(\mathbb{F}_3)$ has 24.

Fix the published integral representatives γ_j and put $F_j(w) = \theta(\gamma_j w)$. For the exact chosen function define its actual Fourier coefficient $a_j(\mu)$ by

$$F_j(z,v) = \sigma_j v^{2/3} + \sum_{\mu\neq 0} a_j(\mu) v K_{1/3}(4\pi|\mu|v) e(\mathrm{Tr}(\mu z)), \qquad e(x) = e^{2\pi i x}. \tag{261}$$

For those representatives $\sigma_j = 3^{5/2}/2$ for $j \le 9$ and zero for $j \ge 10$. The support lies in $\lambda^{-4}\mathcal{O}$. The symbol a_j always means this actual coefficient; a source table using $d_j(\mu)$, $d_j(-\mu)$ or a conjugated function is converted before using an exact phase formula.

The source formulas (5.7), (5.13), (5.14) and Appendix A express these coefficients through τ, τ_1, τ_2, unit rotations, fixed unit factors and additive phases. The first nine families have the τ absolute norm and the other eighteen have the τ_1, τ_2 norms. Rotations, conjugation and index reversal preserve absolute norm.

Proposition 39 (Known-cusp absolute norm)**.** *For $1 < \sigma < 13/12$, every one of these actual cusp families satisfies*

$$\mathcal{A}_j(\sigma) := \sum_{\mu\neq 0} \frac{|a_j(\mu)|}{(N\mu)^\sigma} \le 3^8 \zeta_K^{(\lambda)}(\sigma)\zeta_K^{(\lambda)}(3\sigma - \tfrac{1}{2}). \tag{262}$$

The right side is independent of a moving good level or a translation. Here $\zeta_K^{(\lambda)}$ omits the prime above 3.

Proof. In the τ formula write $c, d \equiv 1 \pmod 3$, with c squarefree and d arbitrary, possibly sharing primes with c. For the first family $\mu = \pm\epsilon\lambda^{3n-4}cd^3$, $n \ge 1$, there are at most six sign/unit branches and

$$|\tau(\mu)| \le 3^{n/2+2}\sqrt{Nd}.$$

For the second family $\mu = \pm\lambda^{3n-3}cd^3$, $n \ge 0$, there are at most two branches and $|\tau(\mu)| \le 3^{n/2+5/2}\sqrt{Nd}$. These follow from the literal unnormalized primitive Gauss magnitude $|g(a,c)| = \sqrt{Nc}$ and the source factor $|d/c|$; the fixed Gauss numerators are units modulo good c. Retain the separate c and cube variables. A generic $(N\mu)^{1/6}$ coefficient bound would lose convergence near 1.

Put $z = 3^{1/2-3\sigma}$ and $Z_{\mathrm{SF}}(\sigma) = \sum_{c\ \mathrm{primary,SF}} (Nc)^{-\sigma}$, with the prime λ excluded. Then

$$\mathcal{A}_\tau(\sigma) \le \left\{ \frac{6\,3^{2+4\sigma}z + 2\,3^{5/2+3\sigma}}{1-z} \right\} Z_{\mathrm{SF}}(\sigma)\zeta_K^{(\lambda)}(3\sigma - \tfrac{1}{2}). \tag{263}$$

The τ_1, τ_2 families each have three unit branches at $\lambda^{-4}cd^3$, with magnitude at most $9\sqrt{Nd}$; their possible opposite signs do not change the norm. Consequently

$$\mathcal{A}_{\tau_i}(\sigma) \le 27\,3^{4\sigma} Z_{\mathrm{SF}}(\sigma)\zeta_K^{(\lambda)}(3\sigma - \tfrac{1}{2}), \qquad i = 1, 2. \tag{264}$$

In the stated range $z \le 1/9$, $3^{2+4\sigma} \le 3^7$, $3^{5/2+3\sigma} \le 3^6$, and $27\,3^{4\sigma} \le 3^8$. The brace in (263) is at most $3^8/2$. Finally $Z_{\mathrm{SF}}(\sigma) \le \zeta_K^{(\lambda)}(\sigma)$. The actual coefficient convention and all unit phases therefore give (262). □

For later comparison, an affine state $F_j(Az+B,|A|v)$, $A\neq 0$, has index and coefficient

$$\nu = A\mu, \qquad a_{j,A,B}(\nu) = |A|\, a_j(\mu) e(\mathrm{Tr}(\mu B)). \tag{265}$$

Its absolute Dirichlet norm is exactly $|A|^{1-2\sigma}\mathcal{A}_j(\sigma)$. For the ramified diagonal $A=\lambda^{-2B_0}$, this factor is $3^{B_0(2\sigma-1)}$. Commuting a moving translation past that diagonal changes $U(-\lambda^4 c/r)$ to $U(-\lambda^{4+2B_0}c/r)$. Together with $\nu=\lambda^{-2B_0}\mu$ its phase is still $e(-\mathrm{Tr}(\lambda^4\mu c/r))$. This explains why an omitted ramified scale changes actual phases, even though the eventual absolute bound is fixed-data controlled.

51.3 Exact rational reduction and the actual Kubota multiplier

For $a=A/C\in K$ reduced, choose $B_0,D_0\in\mathcal{O}$ with $AD_0-B_0C=1$ and put $\gamma_a=\left(\begin{smallmatrix}A & B_0\\ C & D_0\end{smallmatrix}\right)$. The Euclidean property of $\mathcal{O}$ supplies this choice. Let

$$U(x)=\begin{pmatrix}1 & x\\ 0 & 1\end{pmatrix}, \qquad W_C=\begin{pmatrix}0 & -1/C\\ C & 0\end{pmatrix}.$$

Integral a is included by $C=1$. Direct multiplication proves

$$\gamma_a^{-1}U(a)=\begin{pmatrix}D_0 & 1/C\\ -C & 0\end{pmatrix}=U(-D_0/C)W_{-C}. \tag{266}$$

Choose j with $\delta_a=\gamma_a\gamma_j^{-1}\in\Gamma_2$. Then the exact function identity is

$$\theta(U(a)w)=\chi(\delta_a)F_j(U(-D_0/C)W_{-C}w). \tag{267}$$

It uses known integral cusp functions rather than an unconstructed new cusp coefficient family.

The phase is computable, not an unspecified unit set to one. Find a rational integral $\sigma\in\mathrm{SL}_2(\mathbb{Z})$ with the same reduction modulo 3 as δ_a. Then $\delta'=\sigma^{-1}\delta_a\in\Gamma_1(3)$. If $\delta'=\left(\begin{smallmatrix}a' & b'\\ c' & d'\end{smallmatrix}\right)$, the source rule is

$$\chi(\delta_a)=\chi(\delta')=\begin{cases}(c'/a')_3, & c'\neq 0,\\ 1, & c'=0.\end{cases} \tag{268}$$

Here $a'\equiv 1 \pmod 3$ is nonzero and coprime to c'. The index j can be found by reducing modulo 3 and testing the 24-element rational subgroup against the 27 fixed cosets. Equivalently, a finite traversal of elementary generators constructs representatives. A representative differing from the published one changes its function by the fixed multiplier just computed. Its Fourier coefficients and constants carry that multiplier too.

One can also display the fixed ramified states explicitly. If a moving good level r survives the fixed denominator reduction, write $C=c_0r$, with c_0 in a fixed finite set and $(c_0,r)=1$. The matrix identity

$$U(A/C)W_r=\gamma_a U(-D_0/C)\,\mathrm{diag}(r/C,C/r)$$

and $D_0=br+9c_0d$ give the fixed state $F_j|U(-b/c_0)\,\mathrm{diag}(c_0^{-1},c_0)$ followed by the moving translation $U(-9c_0^2d/r)$. The general formula (267) also handles cancelled or nonunit fixed Fourier residues directly; this special display does not impose a primitive Fourier-support condition on them.

51.4 Radial cusp constants, angular derivatives and the Mellin Jacobian

Write $w = (z, v)$. The action in (267) is

$$z' = -D_0/C - \frac{\bar{z}}{C^2(|z|^2 + v^2)}, \qquad v' = \frac{v}{NC(|z|^2 + v^2)}. \tag{269}$$

The pulled-back constant of F_j is therefore

$$\chi(\delta_a)\sigma_j(NC)^{-2/3}v^{2/3}(|z|^2 + v^2)^{-2/3}. \tag{270}$$

It is radial in z. Every nonzero pure Wirtinger derivative ∂_z^m or $\partial_{\bar{z}}^m$ at $z = 0$ kills it. At large v the original translated theta constant is also independent of z. The known nonconstant lattice support gives exponential decay after angular differentiation at both ends, with the reciprocal-end rate depending on NC.

For completeness, the zeroth-angular constants have not disappeared. For a finite combination $\sum_a c(a)\theta(z + a, v)$ put

$$C_\infty = \frac{3^{5/2}}{2}\sum_a c(a), \qquad C_0^* = \sum_a c(a)\chi(\delta_a)\sigma_{j(a)}(NC_a)^{-2/3}. \tag{271}$$

Its regularized Mellin transform with weight v^{2s-2} has residues

$$\operatorname*{Res}_{s=1/6} \Lambda(s, 0) = -C_\infty/2, \qquad \operatorname*{Res}_{s=5/6} \Lambda(s, 0) = C_0^*/2. \tag{272}$$

These signs follow by integrating the powers at infinity and zero. A projector with zero sum kills C_∞, but need not kill C_0^*. No finite noncubic label implies such a cancellation.

For $\ell > 0$ set $\mathcal{D}_\ell = (2\pi i)^{-\ell}\partial_z^\ell$; for $\ell < 0$ use $\mathcal{D}_\ell = (2\pi i)^{-|\ell|}\partial_{\bar{z}}^{|\ell|}$. The angular Mellin transform is

$$\Lambda_a(s, \ell) = \int_0^\infty \mathcal{D}_\ell[\theta(z + a, v)]_{z=0} v^{2s+|\ell|-2}\, dv. \tag{273}$$

For $\ell \neq 0$ it is entire by the endpoint bounds just proved.

Proposition 40 (Exact angular rational-cusp identity)**.** *For every reduced rational translate and $\ell \neq 0$,*

$$\Lambda_a(s, \ell) = \chi(\delta_a)(-1)^{|\ell|}(C/\bar{C})^\ell(NC)^{1-2s}\Lambda_{j,-D_0/C}(1 - s, -\ell). \tag{274}$$

The right Mellin transform uses the actual known function $F_j(z - D_0/C, v)$.

Proof. For $\ell > 0$, set $\bar{z} = 0$ in the full-coordinate chain rule. Then z' and v' in (269) are constant as functions of z, while $\bar{z}'$ is affine with derivative $-1/(\bar{C}^2v^2)$. Thus the pure derivative has exactly the factor $\chi(\delta_a)(-1)^\ell\bar{C}^{-2\ell}v^{-2\ell}$ and the opposite angular derivative of F_j at $(-D_0/C, 1/(NCv))$. There are no lower angular or radial terms. The negative sign case is the conjugate-coordinate calculation.

In the integral substitute $y = 1/(NCv)$. Its Jacobian and power weight give $(NC)^{-2s+\ell+1}y^{\ell-2s}\, dy$. Combining the factor $\bar{C}^{-2\ell}$ produces precisely $(C/\bar{C})^\ell(NC)^{1-2s}$. The exponent $1 - 2s$, rather than $-2s$, is essential. This proves (274) for both signs. □

If the real part of the dual Dirichlet variable exceeds one, its series is

$$\sum_{\mu\neq 0} a_j(\mu)e(-\operatorname{Tr}(\mu D_0/C))(\mu/|\mu|)^{-\ell}(N\mu)^{-(1-s)}. \tag{275}$$

The additive and angular phases remain literal; their absolute values are one only after this exact expression has been established. Proposition 39 controls the series.

51.5 The integral-index primary projector and raw cube completion

Let κ be any one fixed residue member on the primary section. Use a fixed modulus containing its actual conductor and all original fixed exclusions. Separate its λ-part κ_λ and away-λ part κ_0. Choose $f \geq 2$ containing the λ-conductor, and put

$$\Phi(x) = \kappa_\lambda(x)\mathbf{1}_{x\equiv 1 \bmod \lambda^2} \quad (x \bmod \lambda^f), \qquad \Phi(x) = 0 \text{ on nonunits.} \tag{276}$$

The entire theta support has $\lambda^3\mu \in \mathcal{O}$. The correct argument is this integral index, not $\lambda^4\mu$ for a ramified character.

From the two τ families above, $\Phi(\lambda^3\mu)$ selects exactly

$$\mathcal{S} = \{\lambda^{-3}cd^3 : c, d \equiv 1 \pmod 3,\ c \text{ squarefree}\}. \tag{277}$$

Indeed the first family has positive λ-valuation after multiplication by λ^3, as do the second-family terms with $n \geq 1$. They are killed as nonunits. At $n = 0$ the negative branch is -1 modulo 3, and only the positive branch is 1. This proves the unit-state selection as well as its valuation condition.

For a moving good squarefree level $\mathfrak{t}$, let $\psi_\mathfrak{t}$ be either cubic orientation on its literal residue modulus. Let F_0 be the fixed away-λ raw modulus. The product $\psi_M(x) = \kappa_0(x)\psi_\mathfrak{t}(x)$ is a bounded periodic function on the valid period $M = F_0\mathfrak{t}$. This period is not asserted to be its primitive conductor. It is allowed to be imprimitive and to have overlapping $F_0, \mathfrak{t}$. No value $\kappa(\mathfrak{t})$ is divided out.

Define the coefficient-projected theta function using $\Phi(\lambda^3\mu)\psi_M(\lambda^3\mu)$. Its infinite constant is zero. On $\mathcal{S}$, $\tau(\mu) = 3^{5/2}\widetilde{g}(c)\sqrt{Nd}$. Multiplicativity on the primary section gives the exact cube identity

$$\Phi(\lambda^3\mu)\psi_M(\lambda^3\mu) = \kappa(c)\kappa(d)^3\psi_\mathfrak{t}(c)\mathbf{1}_{(d,\mathfrak{t})=1}. \tag{278}$$

Write

$$\mathcal{R}_{\kappa,\mathfrak{t},\ell}(s) = \sum_{c \text{ primary,SF}} \frac{\widetilde{g}(c)\kappa(c)\psi_\mathfrak{t}(c)\Theta_\ell(c)}{(Nc)^s}, \tag{279}$$

with every fixed raw zero in κ. Then initially $\Re s > 1$,

$$D_{\Phi,\psi}(s,\ell) = 3^{5/2+3s}i^\ell L^{\mathrm{raw}}(3s - \tfrac{1}{2}, \kappa^3\Theta_{3\ell})\mathcal{R}_{\kappa,\mathfrak{t},\ell}(s). \tag{280}$$

The factor i^ℓ is the actual angle of λ^{-3}. The cube series has denominator $(Nd)^{3s-1/2}$, and L^{raw} retains all original exclusions above $\lambda F_0\mathfrak{t}$, even after primitive conductor drop. There is no $\kappa(\lambda)$ factor or untwisted-zeta substitution. Conjugating the whole construction gives the actual conjugate Gauss root and opposite angular/cubic orientations.

The full ramified projector has an exact finite translation formula. The trace dual of $\mathcal{O}$ is $\lambda^{-1}\mathcal{O}$, so

$$\begin{aligned} b_h &= 3^{-f}\sum_{x \bmod \lambda^f} \Phi(x)e(-\operatorname{Tr}(xh/\lambda^{f+1})), \\ \Phi(x) &= \sum_{h \bmod \lambda^f} b_h e(\operatorname{Tr}(xh/\lambda^{f+1})), \qquad \sum_h |b_h| \leq 3^{f/2}. \end{aligned} \tag{281}$$

For the away period use normalized Fourier coefficients q_u with positive inversion sign and pairing $e(\operatorname{Tr}(xu/\alpha_M))$. Since M avoids λ, that pairing is perfect, and

$$\psi_M(x) = \sum_{u \bmod M} q_u e(\operatorname{Tr}(xu/\alpha_M)), \qquad \sum_u |q_u| \leq \sqrt{NM}. \tag{282}$$

Both mass bounds are Parseval followed by Cauchy. They hold for arbitrary bounded periodic functions; no primitive Fourier-support condition is inserted.

Consequently the exact projected function is a finite sum of $\theta(z + a_{h,u}, v)$ with

$$a_{h,u} = h\lambda^{2-f} + \lambda^3 u/\alpha_M, \qquad \sum_{h,u} |b_h q_u| \le 3^{f/2}\sqrt{NM}. \tag{283}$$

The inverse different explains the exponent $2 - f$. Every raw zero is reconstructed by this full sum. Each reduced denominator $C_{h,u}$ divides a fixed λ-power times M; hence $1 \le NC_{h,u} \ll_{\text{fixed}} N\mathfrak{t}$. Nonunit Fourier residues may lower that denominator. They remain terms of the exact transform, rather than being discarded.

51.6 The coarse dual bound, safe convexity and literal Perron endpoints

Let $R = N\mathfrak{t}$, and set

$$E(s) = L^{\text{raw}}(3s - \tfrac{1}{2}, \kappa^3\Theta_{3\ell})\mathcal{R}_{\kappa,\mathfrak{t},\ell}(s).$$

For $\ell \neq 0$ the finite sum of entire angular Mellins is entire. The source Fourier–Mellin calculation has factor

$$\gamma_\ell(s) = \frac{1}{4}(2\pi)^{-2s-|\ell|}\Gamma(s + \tfrac{|\ell|}{2} - \tfrac{1}{6})\Gamma(s + \tfrac{|\ell|}{2} + \tfrac{1}{6}). \tag{284}$$

Thus division by γ_ℓ, whose reciprocal is entire, and the nonzero elementary factor in (280) makes E entire as well.

On $s = -\eta + i\tau$, $0 < \eta < 1/12$, (274), (262) and (283) give

$$\begin{aligned} \sum_{h,u} |b_h q_u|(NC_{h,u})^{1+2\eta}\mathcal{A}_{j(h,u)}(1+\eta) &\ll_{\text{fixed},\eta} R^{3/2+2\eta}, \\ |E(-\eta + i\tau)| &\ll_{\text{fixed},\ell,\eta} R^{3/2+O(\eta)}(1+|\tau|)^{2+O(\eta)}. \end{aligned} \tag{285}$$

The height exponent follows from the quotient of the two Gamma products. The exact multiplier $\chi(\delta_{h,u})(-1)^{|\ell|}(C_{h,u}/\bar{C}_{h,u})^\ell$ has been written before bounding its modulus. This is a conservative $R^{3/2+\epsilon}$ cost. No refined $R^{1/2+\epsilon}$ Ramanujan estimate is used.

The line $\Re s = 1 + \eta$ is absolutely bounded uniformly in R, by the original Gauss and cube coefficients. For a requested final $\epsilon > 0$, choose $\eta > 0$ sufficiently small in terms of ϵ; the constants may depend on that choice. The finite-order continuation from the theta Mellin representation permits the damped strip maximum principle. Applying it to the entire numerator, then allocating small exponents, gives

$$|E(\sigma + i\tau)| \ll_{\text{fixed},\ell,\epsilon} (R(2+|\tau|))^\epsilon R^{\frac{3}{2}(1-\sigma)}(1+|\tau|)^{2(1-\sigma)} \tag{286}$$

on the safe half-strip $1/2 + \eta \le \sigma \le 1 + \eta$. Only there divide by L^{raw}: the real part of $3s - \frac{1}{2}$ exceeds one, so its inverse is bounded by absolutely convergent Euler products, uniformly in all raw omissions. No inverse-L continuation through unknown zeros or zero-free strip is asserted. Equation (286) therefore bounds $\mathcal{R}$ on that half-strip.

Here are the endpoint details needed for Theorem 37. Replace each real endpoint by the half-integer threshold selecting exactly the same integer norms. All ideals in a repeated norm fibre are counted. Such a fibre has divisor-bounded multiplicity in this quadratic field. For an unnormalized sharp prefix at length T, truncated Perron at height $\mathcal{H} \ge 1$ has error

$$O_\epsilon\big(T^\epsilon[1 + (T/\mathcal{H})\log(2T)]\big). \tag{287}$$

This follows by summing $\min(1, T/(\mathcal{H}|n-y|))$ over integer distances from the half-integer endpoint y. No norm tie receives a half-weight.

At $\sigma = 1/2 + \eta$ the vertical contribution is bounded by

$$(RT)^{\epsilon} T^{1/2} R^{3/4} \mathcal{H}.$$

If $T \geq R^{3/2}$, choose $\mathcal{H} = (T/R^{3/2})^{1/4}$. On horizontal sides the bound is at most $T/\mathcal{H}$, up to allocated subpowers, since $R^{3/2}\mathcal{H}^2 \leq T$. The vertical integral, truncation and horizontal contributions are then $O((RT)^{\epsilon} T^{3/4} R^{3/8})$. An interval is a difference of two prefixes. If $T < R^{3/2}$, the elementary $O(T)$ count is the better branch.

Finally apply partial summation to the same $T^{-1/2} W(x) x^{-1/2}$, $x = N\mathfrak{a}/T$. Its supremum and total variation are bounded by $O(T^{-1/2} V_1(W))$. This proves (254) with the exact profile. For $1/2 \leq T < 1$ only the unit can occur; its coefficient and profile are literal and satisfy the same bound with a fixed endpoint constant. For $T \leq 1/2$ there is no contribution. This extension uses no analytic theorem at a fictitious length.

51.7 The actual annular coefficient and exact eleven-factor expansion

The opened coefficient $B_U = \mu A_U$ in (255) is asserted on the actual annulus $N\mathfrak{m} > U$. It is distinct from original C_U. No identity for $B_U(\mathcal{O})$ is needed or asserted. Unit factors within the factorization of an annular column are essential and stay in every convolution below.

Let $\mathbf{1}$ be the constant-one function on good ideals, $\delta_{\mathcal{O}}$ the ideal-convolution identity, and $M_z(\mathfrak{a}) = \mu(\mathfrak{a}) \mathbf{1}_{N\mathfrak{a} \leq z_0}$, $z_0 = (2D)^{1/5}$. The exact good-ideal convolution identity is

$$\mu = \sum_{j=1}^{5} (-1)^{j-1} \binom{5}{j} \mathbf{1}^{*(j-1)} * M_z^{*j} + \mathbf{1}^{*4} * (\mu - M_z)^{*5}. \tag{288}$$

Indeed $\mathbf{1} * \mu = \delta_{\mathcal{O}}$, so expand $\mu * [\delta_{\mathcal{O}} - (\delta_{\mathcal{O}} - \mathbf{1} * M_z)^{*5}]$; also $\delta_{\mathcal{O}} - \mathbf{1} * M_z = \mathbf{1} * (\mu - M_z)$. Each factor of $\mu - M_z$ has norm strictly above z_0. The last term in (288) is supported strictly above $z_0^5 = 2D$ and vanishes at every residual $\mathfrak{c}$ under consideration.

The j-th term has j truncated Möbius factors and $j-1$ unrestricted coefficient-one factors. Together with the two original short factors it has $2j+1 \leq 11$ labels. The $j = 1$ term and unit labels remain. Restricting the total column to squarefree support forces all factors to be squarefree and pairwise coprime.

Use independent dyads $[M_i, 2M_i)$, $M_i \geq 1$, with the original own-factor cutoffs at U, z_0. If the box meets the annulus,

$$D/2^{11} \leq P_{\mathrm{box}} := \prod_i M_i \leq 2D. \tag{289}$$

Choose the partition from these scales only. If some $M_i \in [P_{\mathrm{box}}^{1/3}, P_{\mathrm{box}}^{2/3}]$, it is a balanced subset. If all $M_i < P_{\mathrm{box}}^{1/3}$, multiply successive scales until first reaching $P_{\mathrm{box}}^{1/3}$; the result is below $P_{\mathrm{box}}^{2/3}$. If no balanced subset exists, a unique scale exceeds $P_{\mathrm{box}}^{2/3}$; two such scales would have product above P_{box}. The product of all remaining scales is below $P_{\mathrm{box}}^{1/3}$.

At $D \geq 2^{64}$,

$$P_{\mathrm{box}}^{2/3} \geq D^{2/3}/2^{22/3} > \sqrt{D} \geq U, \qquad P_{\mathrm{box}}^{2/3} > (2D)^{1/5}.$$

Therefore an unbalanced unique long factor is genuinely one of the unrestricted coefficient-one factors. It cannot be an original short factor or a truncated Möbius factor. The actual whole prefix

$\mathfrak{p}$, the product of all remaining factors, satisfies $N\mathfrak{p} \le 2^{10} P_{\mathrm{box}}^{1/3} \ll D^{1/3}$. The $j = 1$ term has no unrestricted factor and is automatically balanced; its unit assignments have not been discarded.

For a fixed box define each macro coefficient by summing all its signed assignments within its assigned factor group. On a squarefree ideal the number of at most eleven labelled factorizations is at most $11^{\omega(\mathfrak{a})}$. Fixed binomial factors and dyadic logarithms thus give a uniform D^{ϵ} pointwise bound, with every own cutoff retained. There is no separate count over $N\mathfrak{a}, N\mathfrak{b} \le U$.

51.8 Whole-factor cubic matrices and the large-prefix part

A balanced partition gives two physical squarefree macro ideals $\mathfrak{u}, \mathfrak{v}$, each with scale in fixed multiples of $[D^{1/3}, D^{2/3}]$ and product comparable to D. Their own coefficients contain the original cutoffs and signs. Cubic Gauss CRT gives on coprime products

$$\overline{G_{\mathfrak{u}\mathfrak{v}}^{\mathrm{act}}} = \overline{G_{\mathfrak{u}}^{\mathrm{act}}}\,\overline{G_{\mathfrak{v}}^{\mathrm{act}}}\,\chi_{3,\mathfrak{u}}(\alpha_{\mathfrak{v}}). \tag{290}$$

The raw matrix entry is zero at a collision, so summing all squarefree macro pairs implements the original squarefree product. No primitive root at an overlapping product is invented.

The external physical cubic norm is

$$\big\|[\chi_{3,\mathfrak{u}}(\alpha_{\mathfrak{v}})]_{N\mathfrak{u}\asymp P,\, N\mathfrak{v}\asymp L}\big\|_{\mathrm{op}} \ll_{\epsilon} (PL)^{\epsilon}\sqrt{P + L + (PL)^{2/3}}. \tag{291}$$

The rows and columns are actual good squarefree ideals. All raw fixed/growing multiplicative masks are coefficient diagonals in this norm, after finite-selector expansion. Ideal counting gives squared coefficient norms D^{ϵ} when each coefficient is weighted by the inverse square root norm. Separate the one occurrence of the common W by

$$W(N\mathfrak{u}N\mathfrak{v}/D) = \frac{1}{2\pi}\int_{\mathbb{R}} \widehat{W}(t) D^{-it} (N\mathfrak{u})^{it} (N\mathfrak{v})^{it}\, dt.$$

The powers are unitary diagonals, costing $\mathcal{M}(W)$ once. Balanced boxes therefore cost $\mathcal{M}(W) D^{1/3+\epsilon}$.

For a Type-I box write $\mathfrak{m} = \mathfrak{p}\mathfrak{v}$, where $\mathfrak{v}$ has coefficient one and the intrinsic $H_{\mathfrak{p}}$ contains all remaining assignments. Then $N\mathfrak{p} \ll D^{1/3}$. The portion $N\mathfrak{p} > Z$, with its literal own-variable indicator, has dyadic bound

$$\mathcal{M}(W) D^{\epsilon}\sqrt{P + D/P + D^{2/3}},$$

and hence total bound

$$\mathcal{M}(W) D^{\epsilon}[D^{1/3} + \sqrt{D/Z}]. \tag{292}$$

Macro own-factor constraints remain coefficients. A chosen equivalent row-dependent norm interval costs the reviewed logarithmic interval-Schur factor, preserving complete norm fibres. It is not assigned operator norm one.

51.9 The full original opposite-column mask in the small-prefix row

Freeze the original $\mathfrak{n}$, the exact $\mathfrak{k}, \mathfrak{s}, \mathfrak{w}, \mathfrak{e}$, and a finite source component before estimating a Type-I row. The original pair mask is exactly

$$\mathbf{1}_{(\mathfrak{p}\mathfrak{v},\mathfrak{n})=1} = \mathbf{1}_{(\mathfrak{p},\mathfrak{n})=1}\mathbf{1}_{(\mathfrak{v},\mathfrak{n})=1}. \tag{293}$$

The first factor stays in the actual prefix coefficient, along with all original prefix/$\mathfrak{k}, \mathfrak{s}, \mathfrak{w}$ collisions, fixed numerator zeros and prefix phases. The intrinsic arithmetic $H_{\mathfrak{p}}$ still depends only on its own

factor assignments. Multiplication by these outer factors preserves its divisor bound. A forbidden prefix is zero before a row bound is applied.

The remaining long-variable deletion is

$$Q = \mathrm{rad}_{\mathrm{good}}(\mathfrak{p}\mathfrak{n}\mathfrak{k}\mathfrak{s}\mathfrak{w}).$$

It is not a substitute for the first factor in (293), and the growing $\mathfrak{n}$ is not part of a fixed character conductor. On the nonzero support $\mathfrak{p}, \mathfrak{k}$ are squarefree and coprime, by the original $\eta(\mathfrak{k})^2$ zero. The physical cubic frequency is $\mathfrak{t} = \mathfrak{p}\mathfrak{k}$.

The exact deletion formula is

$$\mathcal{F}_\psi(T; \mathfrak{t}; Q, I) = \sum_{\substack{\mathfrak{h} \mid Q \\ (\mathfrak{h}, \mathfrak{t}) = 1}} \frac{\mu(\mathfrak{h})\psi(\mathfrak{h})\overline{G_{\mathfrak{h}}^{\mathrm{act}}}\chi_{3,\mathfrak{h}}(\alpha_{\mathfrak{t}})}{\sqrt{N\mathfrak{h}}} \cdot \mathcal{F}_\psi(T/N\mathfrak{h}; \mathfrak{h}\mathfrak{t}; 1, I/N\mathfrak{h}). \tag{294}$$

Insert $\sum_{\mathfrak{h} \mid (Q, \mathfrak{v})} \mu(\mathfrak{h})$ and set $\mathfrak{v} = \mathfrak{h}\mathfrak{a}$. Gauss CRT and reciprocity combine the actual crosses into $\chi_{3,\mathfrak{a}}(\alpha_{\mathfrak{h}\mathfrak{t}})$. This raw entry is zero at $(\mathfrak{a}, \mathfrak{h}) > 1$ and encodes the squarefree-product condition. The outer raw entry kills $(\mathfrak{h}, \mathfrak{t}) > 1$. No zero is inverted, including if a numerator exponent becomes divisible by three. Original interval endpoints divide literally by $N\mathfrak{h}$.

Here $T = D/N\mathfrak{p}$, and the exact product is

$$(T/N\mathfrak{h})N(\mathfrak{h}\mathfrak{t}) = DN\mathfrak{k}. \tag{295}$$

There is no common-long-divisor decomposition between the two columns in this proof. The other original column remains fixed.

Using the sharp bound and the original prefix factor gives (258). For $T/N\mathfrak{h} \in [1/2, 1)$ use the actual unit contribution; smaller arguments contribute zero, with the profile endpoint rule. Large frequencies are allowed by the proven all-level sharp bound, with its trivial fallback. Thus

$$\sum_{\mathfrak{h} \mid Q} (N\mathfrak{h})^{-3/8} \le \tau(Q) \ll_\epsilon (NQ)^\epsilon.$$

On the exact support, $N\mathfrak{n} \le 2D$, $N\mathfrak{s} \le 2V$, $N\mathfrak{w} < 64$, so $NQ \ll DZVN\mathfrak{k}$. Summing $(N\mathfrak{p})^{-3/8}$ over $N\mathfrak{p} \le Z$ by ideal counting proves the small-prefix column bound

$$V_1(W)(DXA_1)^\epsilon D^{1/4}(N\mathfrak{k})^{3/8+\epsilon} Z^{5/8}. \tag{296}$$

Every prefix zero was retained until that positive estimate. No coherence or orientation mean has been removed.

51.10 The exact discrete prefix and the final normalization

Combine the balanced, large-prefix and small-prefix bounds, and put

$$Z = D^{2/9}, \qquad \sqrt{D/Z} = D^{1/4}Z^{5/8} = D^{7/18}. \tag{297}$$

The untouched original A_U column has the elementary bound $B_W D^{1/2+\epsilon}$. Apply it after the left estimate and use the original A_1/D once. All finite source/ray/primitive $\mathfrak{e}/\mathfrak{w}$ coefficient masses stay fixed, and the actual $\mathfrak{s}$ family is counted by $O(V)$ only at the end. The scalar g_0 is bounded and is never divided out.

The original finite $\mathfrak{k}$ sum is unchanged in the identity. Its small-prefix positive majorant is

$$\sum_{\substack{\mathfrak{k}\ \text{good,SF}\\ N\mathfrak{k}\le K_{\mathrm{cut}}}} (N\mathfrak{k})^{-13/8+\epsilon} \le \zeta_K(13/8-\epsilon), \qquad \epsilon<5/8. \tag{298}$$

Product and divisor losses can be allocated smaller exponents first; their $N\mathfrak{k}$ factors are absorbed by this convergent series. Balanced and large-prefix pieces use the easier exponent 2. This is a bound on a finite sum, not an Euler replacement of the model.

For the optional original unequal-norm projection, a fixed rational norm has $O_\epsilon(D^\epsilon)$ ideal coordinates. The actual A_U and every fixed HB component are divisor-bounded. Removing that same fixed-$\mathfrak{n}$ fibre costs $O(B_W D^{-1/2+\epsilon})$ in the left column, and therefore $O(B_W^2 A_1 V D^{-1+\epsilon})$ in the normalized form. Conjugate ideals remain different, and a norm inequality has not been made multiplicative.

This proves the leading-model part of (256). For the original exact prefix, all its Bessel arguments are at most $\Lambda = A_1 X/D^2$, and

$$|\dot G(x) - g_0| \le 3C_0^{-3/2}x.$$

There are $O(D^2)$ annular column pairs and $A_1/(D\sqrt{F_\Omega}) \le A_1/D^2$. The coefficient product is $O(B_W^2 D^\epsilon)$, the original finite $\mathfrak{k}$ absolute budget is bounded by $\sum(N\mathfrak{k})^{-2}$, and the actual sixth-part band has $O(V)$ ideals. Only in this remainder are these absolute counts taken. They give (250) directly for A_U, A_U. No A_U-to-A_Y error, principal density or scalar ray-count split is introduced.

Finally the exact powers are

$$\begin{aligned} &3/4+1/16-1/9 = 101/144, && 3/4-101/144 = 7/144,\\ &3/4+1/16-1/6 = 31/48, && 13/16+(3/4+3/8-2) = -1/16. \end{aligned} \tag{299}$$

The same norm-fibre correction has power $-3/16$. This completes the pure finite-prefix theorem with its once-only normalization and all original finite cutoffs.

51.11 Verification scope and the remaining arithmetic comparison

The accepted evidence includes exact rational cusp matrices and finite quotient/coset calculations, angular-coordinate and Mellin checks, literal primary-support cases, affine coefficient transport, and synthetic scale-composition/exponent checks. The typed exponent grid is a formal subfamily; it is not exhaustive physical dyad coverage after the constants in (289). The proof uses the separate exact $D \ge 2^{64}$ inequalities for those constants. None of these checks is actual-family energy data, training, asymptotics or a kernel-checked Lean proof.

All original sources, primary versions, raw outputs and failed preliminary harnesses remain preserved. The new theorem is a prose proof with the identified external theta and cubic-norm inputs. Its domain is the representation-fixed pure $\mathcal{T}_X$ prefix in (249). The infinite dual complement, other pure sectors, mixed and neighboring blocks, short/principal assembly, independent positive arithmetic determinant and zero-location comparisons remain open. The Dirichlet, Hecke, Dedekind, elliptic and general Artin scopes retain their separate hypotheses; no Artin holomorphy, desired positivity, zero location or GRH premise has been inserted. The whole signed and global unsigned comparisons have not changed.

52 Paired cutoffs, exact divisor signs and proved regimes

The comparison in this section concerns a specified part of the signed coprime-column form for the original unweighted, nonprincipal, two-pivot Eisenstein rows. We use the physical ideal norms

and the original ordered cutoff. Let D, U, H be the column, cutoff and capacity parameters, with $U^2 \le D \le 7U^2/2$ and $A_1 A_2^2 \ll H \le D$. The first two row ideals have dyadic norms A_1, A_2, $R = A_1 A_2$, and the actual primitive conductor of row u has norm $Q_u \asymp R$. The finite bad-prime set S, row type, and twist are fixed. The same smooth complex profile W supported on $[1, 2]$ is used on both sides. Ideal labels, including conjugate primes of the same norm, are retained individually.

52.1 The selected pairing

For each good squarefree intermediate Möbius divisor d, put $E = Nd$. Let V_d be its original row amplitude and $\mathcal{E}_d$ the equal-norm part of the original middle Poisson window. We study selections of

$$\frac{2}{D} \operatorname{Re} \sum_d \mu(d) \langle \mathcal{E}_d, V_d - \mathcal{E}_d \rangle. \tag{300}$$

Here the counting inner product is linear in the first argument. This is one term in the expansion of the original signed norm. In particular the intermediate d is different from the gcd of the final column pair.

On an original allocation let $z = e_0 z_0$ with $J \le Nz < 2J$, let $e_1 \mid \operatorname{rad}(Sd)$ be the other deletion label, and put $h = re_0$, with r the original squarefree short divisor. Set

$$X = \frac{D}{E(Nr)^2}, \qquad X' = \frac{X}{(Ne_0)^2}, \qquad M = \frac{R\, Ne_0\, Ne_1\, J}{X}.$$

The original middle window has $X' > B_*$ and $K_0 < Nk \le MD^\eta$, for its specified B_*, K_0, η. The primal test discards the guaranteed zero cells $X/((Ne_0)(Ne_1)J) \le 1/2$, and the original weights retain any further zeros. An allocation with $X' \le B_*$ belongs to the direct component. Its re-expressed Poisson sum stays in the complement in (300).

On the equality side write $z_{0,0} = g\ell$ and $k_0 = g\bar{\ell}$, so $Nz_{0,0} = Nk_0$. On the partner side the selected changes in the good exponents consist of erasing the squarefree cutoff part x, erasing a part of the good conjugate overlap, and inserting allowed positive increments at other good primes outside d. A conjugate erasure removes the full original exponent, which can be two. At an insertion prime the old residue plus the increment is at most five. Changes at the good primes of d keep their exact new residues. For fixed other labels write $z_{0,1} = z_* x$ and $k = Ik_*$. The physical insertion ideal I keeps each exponent $1, \dots, 5$ and is disjoint from the selected erasure supports and d. Thus $Nx \asymp B$ and $NI \asymp K$ are physical lengths. The two varying raw phases are $\psi_u(x)\overline{\psi_u(I)}$, or their conjugates if the pairing is conjugated. The fixed row factors retain the union of deleted primes, the literal section, the twist and the primitive S-values.

The subtraction on the partner imposes $Nz_* Nx = Nk_* NI$ only in its original middle window. At fixed x this fixes at most one integer norm for I, and the reverse count fixes at most one norm for x. An integer norm supports at most divisor-many ideals on the polynomial clip. With the physical inverse $1/Nx$, the paired count therefore costs $\min(1, K/B)$ up to subpowers. The two half radial powers leave

$$(B/K)^{1/2} \min(1, K/B) = \sqrt{\min(B/K, K/B)}.$$

The remaining harmonic weights give $O(R/E)$ for this subtraction at each d, with the inherited profile and subpower costs. Both original endpoints remain in the count. Equal norms in the direct and other components remain in the pairing.

52.2 Target bounds on the selected cross

The bounds have the common cost $O_{\epsilon,S,\nu,u_0}((DH)^{\epsilon}\mathcal{M}_{24}(W)^2)$, where $\mathcal{M}_{24}$ controls the indicated smooth profile derivatives. The live partner primal test supplies $E \ll D/B$ and partner radial scale $y_1' \gg Nk/R \gg K/R$. The proof first sums each actual column array. It then uses the row comparisons and the original label sums.

At the mixed scales $A_1 \asymp D^{5/12}$, $A_2 \asymp D^{1/6}$, the selected normalized cross has upper power $57/80$ over all physical ratios and all d. If $H \ll D^{3/4}$, a perturbed primitive sharp moment gives the target power $3/4$ for the same selected cross on

$$\frac{7}{24} \le \log_D A_1 \le \frac{439}{912}, \qquad A_1 A_2^2 \ll H.$$

The second pivot need only satisfy the stated original admissibility and capacity. The conductor comparison uses the actual R.

At the pure critical scales $A_1 \asymp R \asymp D^{3/4}$, $A_2 \asymp 1$, the target power $3/4$ holds over all d in each of the following physical sectors:

- $K \gtrsim D^{1089/1420}$, with any erasure length B.
- $B \gtrsim D^{1355/1482}$, with any insertion length K.
- $B \gtrsim R$ and $K \lesssim D^{921/1652}$.

The flexible central windows give further short and long threshold curves. There are also separate comparisons for a small fixed residual label product and for a small d-prefix. The comparable pure all-divisor comparison has power $55/68$, and the target holds on $Nd \le D^{9/68}$ in that sector. Each assertion concerns the selected changes described above. The remaining pure rectangles and the coherently averaged signed divisor sum are open.

52.3 Ordered assignment signs

Fix an ordered factorization $\zeta = \zeta_1\zeta_2$ of the actual squarefree cutoff label $\zeta = e_0 z_0$ on either side, with $(d,\zeta) = 1$ and $N\zeta_i \le U$. Write $T_i = U/N\zeta_i$. The assignment sum multiplying $\mu(\zeta)$ is

$$F_d = \sum_{\substack{a,b \mid d,\ (a,b)=1 \\ Na \le T_1,\ Nb \le T_2}} \mu(a)\mu(b).$$

Complete the b-sum and then subtract its excluded terms. With $m = d/b$ in the subtraction, squarefreeness gives the exact identity

$$\frac{F_d}{\mu(d)} = \mathbf{1}_{E \le T_1} + \mathbf{1}_{E \le T_2} - 1 - \sum_{\substack{m \mid d \\ T_1 < Nm < E/T_2}} \sum_{\substack{a \mid m \\ Na \le T_1}} \mu(m/a). \tag{301}$$

The complete b-sum forces $a = d$. The unit $m = 1$ gives the indicator at T_2, while the nonunit terms with $Nm \le T_1$ vanish by the full inner Möbius sum. Both inequalities of the remaining band are strict. Below the product cap the band is empty. The three possible assignment values are $\mu(d), 0, -\mu(d)$ according as both, just one, or neither individual capacity reaches E.

Fix the row, both ordered factorizations and the primal and dual labels. Select both e_1 labels to be fixed S-supported ideals, take the good changed-residue ideal on d to be the unit, and require d to avoid the full fixed prime union of the row and the primal and dual labels. When $EN\zeta \le U^2$ on both

sides, their capacity step functions multiply the original outer Möbius sign. The fixed-field deleted Möbius estimate supplies subpower cancellation for the resulting fixed-label scalar on $Nd \asymp Y$ with $Y \geq D^{\kappa}$ for a fixed $\kappa > 0$. In the pure comparable block, a physical Nx-interval of width $O(BD^{-1/8})$ gives the target after the full divisor sum by a separate sparse erased-column estimate. This includes the narrow interval immediately above the product cap on the actual partner allocation. The wide correction and the remaining coherent signed estimate require additional arguments. The whole signed and unsigned powers 13/12 and 9/8, the independent determinant comparison, and the scope-specific zero-location obligations remain as in the preceding sections.

53 Sharp insertion and ordered cutoff estimates

Throughout this section, x denotes the erased squarefree cutoff factor, I the positive insertion ideal, and B, K their respective physical dyadic lengths. The exact partner norms are $Nz = Nz_*Nx$ and $N\mathfrak{k} = Nk_*NI$. Here $\mathfrak{k}$ is the original dual ideal, and k_* is its fixed residual factor in a given cell. Write

$$E = Nd, \qquad A = A_1, \qquad R = A_1A_2, \qquad a = \log_D A, \quad r = \log_D R, \quad b = \log_D B, \quad k = \log_D K.$$

At the critical pure scales $a = r = 3/4$ at the level of powers, and at the critical mixed scales $a = 5/12,\ r = 7/12$. Fixed comparison constants in the scales are allowed throughout. A bound in this section includes the common factor $(DH)^{\epsilon}\mathcal{M}_{24}(W)^2$ for the paired cross, with the original fixed-data constants. In row moments the appropriate one-profile factor is understood separately.

The estimates concern the selected pairing with both original erasures, all positive off-d increases whose sum with the old residue is at most five, and the exact new residues on d. The old cutoff and conjugate erasures, including old exponent two in the conjugate erasure, remain in the phase. The cutoff is $Z = U^2$, with $U^2 \leq D \leq 7U^2/2$. The rows are the original unweighted literal nonprincipal rows. We keep the common W, fixed S, ν, u_0, literal ray, capacity and squarefree selectors, and the actual primitive $Q_u \asymp R$, Gauss root and raw zeros. The physical sieve is the order-six application in Section 45. Its exponent at column length D^t, relative to A, is

$$\Delta_a(t) = \max\bigl(0, t - a, t/3 - a/6, 5t/6 - 2a/3\bigr). \tag{302}$$

Its growth slope is at most one. A live subunit dyad is treated at constant scale.

53.1 The middle subtraction and scalar comparisons

First remove the negligible inherited tail $N\mathfrak{k} > D^5$. Its absolute estimate applies to each selection below. Fix the original allocation and dyad on the partner side. The compact logarithmic separation of the partner profile uses $\log(NI/K) - \log(Nx/B) - \log(Q_u/R)$. Its weighted Fourier absolute integral, with the extra factor $1 + |\xi|$ in log frequency ξ, is bounded by a constant times $\mathcal{M}_{24}(W)(1 + y_1')^{-4}$ at the actual dyadic profile scale y_1'. This follows by extending the profile on a fixed-width log interval and integrating by parts. The phase factors then separate on the erased array, insertion array and row. There is no condition on B/K in this separation. Sharp subinterval partial summation removes the insertion norm twist with the erased array and every other row multiplier fixed.

For the original middle equality, the equation $Nz_*Nx = Nk_*NI$ fixes at most one integer norm on either side given the other ideal. There are divisor-many ideals of each integer norm on the clipped support. This counts the entire insertion ideal, including its powerful part. The physical $1/Nx$ yields weighted mass $O((DH)^{\epsilon}\min(1, K/B))$. The two half radial powers supply $(B/K)^{1/2}$,

giving the factor $\sqrt{\min(B/K, K/B)} \le 1$. The remaining label sums give $O(R/E)$. Apply this subtraction only in the original middle window with its exact two endpoints. Equal norms in the re-expressed direct, low and far pieces remain there.

The original live primal test gives

$$E \ll D/B, \qquad D/(EE_1(Nh_1)^2) > Nz/4, \qquad y_1' \gg N\mathfrak{k}/R \gg K/R. \tag{303}$$

For each side $i = 0, 1$ write $h_i = r_i e_{0,i}$ for the original grouped square-divisor label and $E_i = Ne_{1,i}$ for its deletion norm. The good part of $e_{1,i}$ is a divisor of d, and its other part is supported on S. Write $h_0 = fa_{\mathrm{id}},\ h_1 = fb_{\mathrm{id}}$ with $f = (h_0, h_1)$. The three ideals $f, a_{\mathrm{id}}, b_{\mathrm{id}}$ are pairwise coprime and squarefree. The ideal b_{id} is distinguished from the length exponent b. On the equality side $z_{0,0} = g\ell,\ \mathfrak{k}_0 = g\bar{\ell}$. For the good part $\bar{\ell}_{\mathrm{g}}$ of the conjugate define

$$j = (\bar{\ell}_{\mathrm{g}}, d), \quad c = (z_*, a_{\mathrm{id}}, \bar{\ell}_{\mathrm{g}}), \quad a_0 = a_{\mathrm{id}}/c, \quad \ell_0 = \ell/(\ell, z_* b_{\mathrm{id}}).$$

The good conjugate erasure support $\mathfrak{q}_\ell$ is a squarefree divisor of $\bar{\ell}_{\mathrm{g}}$ avoiding $a_{\mathrm{id}}, b_{\mathrm{id}}, d$. Its full erased factor is $\mathfrak{e}_\ell = \prod_{p|\mathfrak{q}_\ell} p^{1+\mathbf{1}_{p|z_*}}$. Let $\widetilde{q}_d$ be the actual sixth-free new representative at the good primes of d. With k_S the original S-part and v_6 the full sixth-power repetition, the exact partner residual is

$$\begin{aligned} k_* = k_S v_6^6 \frac{z_* a_{\mathrm{id}}^4 b_{\mathrm{id}}^2 \ell_0^6}{\ell\, c^6} \\ \times \frac{(\bar{\ell}_{\mathrm{g}}/j)\widetilde{q}_d}{\mathfrak{e}_\ell}. \end{aligned} \tag{304}$$

The quotient follows from the good exponent difference of the two primitive phases. Away from d, before changing any residue, that difference has representative exponent

$$v_p(z_{0,1}) + 4v_p(a_{\mathrm{id}}) + 2v_p(b_{\mathrm{id}}) - v_p(\ell) + v_p(\bar{\ell}_{\mathrm{g}}) \pmod 6.$$

At ℓ_0 one adds six to the negative exponent, and at c one subtracts the overflow six. The selected cutoff erasure removes x from this product. The selected conjugate erasure divides by $\mathfrak{e}_\ell$, and the d-part uses $\widetilde{q}_d$ in place of the matched residue. The positive increases multiply this exact residual by I. The norm and raw value of v_6^6 are retained, including any overlap with the other factors. On the original admissible pairs these operations give the literal product $\psi_u(x)\overline{\psi_u(I)}$ times the fixed row phases, with the original raw zeros.

Let D^ρ be the product of the two row norms relative to $R\sqrt{K/B}$, after the projectors and powerful sums. The full cutoff inverse norm is included in these norms. For the exact paired scalar, start with

$$\frac{D}{E^2 E_0 E_1\, Ng\, N\ell\, Nz_*(Nf)^4(Nc)^2(Na_0)^2(Nb_{\mathrm{id}})^2}.$$

Up to the fixed dyadic constants its profile scales are

$$\begin{aligned} y_0 &= \frac{D}{EE_0 R(Nf\, Nc\, Na_0)^2}, \\ y_1' &\asymp \frac{D}{EE_1 R}\frac{(Na_0)^4}{(Nf\, Nc)^2}(N\ell_0 N v_6)^6 \frac{N\widetilde{q}_d}{Nj}\frac{Nk_S K}{B\, N\mathfrak{e}_\ell}. \end{aligned}$$

The actual profile uses Q_u, the physical ideal norms and the original fixed S-factors. Multiplying by $y_0^{-\alpha}(y_1')^{-\theta}$ gives the following weights before the relative row moment:

$$\begin{aligned} &D^{1-\alpha-\theta} R^{\alpha+\theta} E^{\alpha+\theta-2} E_0^{\alpha-1} E_1^{\theta-1} (B/K)^\theta \\ &\quad \times \frac{(Nf)^{2\alpha+2\theta-4}(Nc)^{2\alpha+2\theta-2}(Na_0)^{2\alpha-4\theta-2}}{Ng\, N\ell\, Nz_*(Nb_{\mathrm{id}})^2 (N\ell_0 N v_6)^{6\theta}(Nk_S)^\theta}(N\mathfrak{e}_\ell)^\theta (Nj/N\widetilde{q}_d)^\theta. \end{aligned} \tag{305}$$

The scalar is accompanied by the retained decay $(1+K/R)^{-(4-\theta)}$ from (303). Every absolute comparison follows the exact pairing with those labels.

In the harmonic ℓ-count the fixed $\bar{j} \mid \ell$ pays Nj to power one. A conjugate erasure prime p uses a distinct $\bar{p} \mid \ell$. Its conjugate lies in exactly one of $\ell_0, z_*, b_{\mathrm{id}}$, paying a second inverse norm. In the old exponent-two case, the original $p \mid z_*$ pays a third inverse norm. Thus the resulting local exponents are at most $-2+\theta$ and $-3+2\theta$, respectively. When Nc or its half power is needed, charge $c \mid z_*$. This charge is disjoint from both conjugate charges since $\mathfrak{q}_\ell \perp a_{\mathrm{id}}$, its conjugates lie in ℓ, and $c \mid a_{\mathrm{id}}, \bar{\ell}_{\mathrm{g}}$. The other sums below are harmonic on polynomial support, convergent inverse-power sums, or fixed-divisor choices.

For $1/6 \le \theta \le 1$, take $\alpha = 1-\theta$. The harmonic comparison at each d is

$$\frac{R}{E} D^{\rho} (B/K)^{\theta-1/2} (1+K/R)^{-(4-\theta)}. \tag{306}$$

The two three-half choices give the comparisons

$$\begin{aligned} &\frac{R^{3/2} D^{\rho}}{\sqrt{D}\sqrt{E}}, && (\alpha,\theta) = (1,1/2), \\ &\frac{R^{3/2} D^{\rho} (K/B)^{1/6}}{\sqrt{D} E^{1/3}}, && (\alpha,\theta) = (7/6,1/3), \end{aligned}$$

with the corresponding retained decays. In the second choice $E_0^{1/6} \ll_S E^{1/6}$. With $(\alpha,\theta) = (1,1/4)$, the E_0 cancels, the weights in $f, a_0, \ell_0 v_6$ have exponents $-3/2, -1, -3/2$, and the disjoint charge on c pays its half power. The resulting comparison is

$$\frac{R^{5/4} D^{\rho}}{D^{1/4} E^{3/4}} (K/B)^{1/4} (1+K/R)^{-15/4}. \tag{307}$$

By the first inequality in (303), summing the $E^{-3/4}$ costs $O((D/B)^{1/4})$ on its live range. The two three-half choices on $E \le D^e$ cost $D^{e/2}$ and $D^{2e/3}$, respectively.

The higher choice $(\alpha,\theta) = (5/2,1)$ leaves $(Nf)^3 (Nc)^5$ in (305). The exact primal restriction gives

$$(Nf)^2 (N b_{\mathrm{id}})^2 N z_* N x \ll D/(E E_1), \qquad (Nf\, Nc)^4 \ll (D/(E E_1 B))^4.$$

Write the remaining numerator as $(Nf)^{-1} Nc\, (Nf\, Nc)^4$ and pay the first two factors by the preceding harmonic counts. The additional E^{-4} and $E_0 \ll_S E$ give the harmonic comparison

$$\frac{R^{7/2} D^{3/2+\rho}}{B^4 E} (B/K)^{1/2} (1+K/R)^{-3}. \tag{308}$$

There is also a fixed-label selection for $(\alpha,\theta) = (1,0)$. For $p \mid a_0$, the original eligibility excludes b_{id}, $\ell, d, x, \mathfrak{q}_\ell$, and at most one of $z_*, \bar{\ell}_{\mathrm{g}}$ contains p. The exact residual has valuation four or five at this prime before adding the full sixth repetition. Thus $a_0^4 v_6^6 \mid k_*$. On $(N a_0)^4 (N v_6)^6 \ll D^{\xi}$, the joint count of a_0, v_6 is $O(D^{\xi/4})$, by summing $(N v_6)^{-3/2}$. Formula (305) then gives $R D^{\rho+\xi/4} (K/B)^{1/2}/E$. This selection is sufficient when $N k_* \ll D^{\xi}$. The comparison uses the full residual ideal, with its S-part and repetitions.

53.2 The sharp moment and the grouped sparse tail

Write $I = I_1 P$ uniquely, where I_1 is squarefree and the nonzero valuations of P are $2,3,4,5$. We have $I_1 \perp P$. Fix P at first. The common-prime projector for x, I_1 uses $s_1 \mid (x, I_1)$, with $\mu(s_1)$

and the two raw values on s_1. Require s_1 itself to satisfy the original fixed prime eligibility of I_1, including $s_1 \perp P$, before stripping it. The residual I_1' avoids s_1 by its squarefreeness, along with the original forbidden union. Keep the original fixed P-deletion on the erased array. After stripping s_1, the squared erased mass is $O((DH)^{\epsilon}/(BNs_1))$, including its actual $1/Nx$. The insertion length is $K/(NPNs_1)$. A dyad $NP \asymp M$ has $O(\sqrt{M})$ powerful ideals, by the unique representation $P = a_{\mathrm{p}}^2 b_{\mathrm{p}}^3$ with b_{p} squarefree and the convergent $\sum_{b_{\mathrm{p}}} (Nb_{\mathrm{p}})^{-3/2}$. Overlap is allowed at exponent five. This pays the $M^{-1/2}$ shortening in the product of the two moments.

For the sharp insertion bound at length L, expand the fixed deletion and squarefree selectors in the order $I_1' = e(r')^2 m$, where $e = e_f e_P$ divides the full forbidden union and r' avoids that union. All Möbius signs and fixed raw factors are retained. At fixed P, e, r' the remaining sharp interval is an unweighted m-interval after the exact partial summation. Its endpoints can be treated as sharp disk endpoints. Smooth a disk of norm length L at radial width proportional to $\delta\sqrt{L}$. The normalized shell row norm is bounded by a subpower times

$$D^{\Delta_a(l)/2}\big(\sqrt{\delta L} + L^{1/6}\big), \qquad l = \log_D L.$$

To see the additive term, split each ideal uniquely as $\mathfrak{s}\, b_0 v^6$ with $\mathfrak{s}$ supported on S and b_0 good sixth-free. At fixed $\mathfrak{s}, v$, the shell contains divisor-many ideals per integer norm and has norm width $O(\delta L/(N\mathfrak{s}(Nv)^6))$. Apply the physical sieve to b_0. The square roots of the width weights sum absolutely. The constant atoms sum over $N\mathfrak{s}(Nv)^6 \ll L$, giving $L^{1/6}$ up to subpowers. This also accounts for open or closed endpoints.

The smoothed primitive Poisson transform has the original $Q_u^{-1/2}$ and primitive Gauss root. On a dual dyad T put $Y = \sqrt{LT/R}$. The disk transform is L times the Fourier transform of the unit disk's chord length $2\sqrt{1-v^2}$ at frequency comparable to Y. Integrating once by parts, separating the two endpoint intervals of width $1/Y$, and integrating by parts again on the interior gives $O(L(1+Y)^{-3/2})$. The endpoint derivative integral is $O(Y^{-1/2})$; the interior boundary and second derivative terms are $O(Y^{1/2})$ before the second division by Y. The same argument with the chord length multiplied by v bounds the frequency derivative, so its logarithmic derivative costs one power of $1+Y$. The convolution bump and its log derivative provide $(1+\delta Y)^{-M_0}$ for arbitrarily large fixed M_0. Consequently the cutoff log function m and m' have L^2 bounds $O(L(1+Y)^{-3/2}(1+\delta Y)^{-M_0})$ and $O(L(1+Y)^{-1/2}(1+\delta Y)^{-M_0})$, respectively. Splitting the Fourier integral at $1 + \|m'\|_2/\|m\|_2$ and applying Cauchy and Plancherel gives absolute log Fourier mass $O(L(1+Y)^{-1}(1+\delta Y)^{-M_0})$. Multiplying by $\sqrt{T}/\sqrt{Q_u}$ yields

$$O_{M_0}\left(\sqrt{L}(1+\delta\sqrt{LT/R})^{-M_0}\right).$$

The zero frequency is absent by nonprincipality. With $\delta \asymp D^{\lambda-\Delta_a(l)}$ when $\Delta_a(l) > \lambda$, the effective dual exponent is

$$t = \max(0, r - l + 2(\Delta_a(l) - \lambda)). \tag{309}$$

The main shell term has row RMS $O(\sqrt{L}D^{\lambda/2})$. The additive term has the same bound whenever $\Delta_a(l) \le 2l/3$. The smoothed term has that bound if $\Delta_a(t) \le \lambda$. Above the cutoff the arbitrary smoothing decay pays the degree-one growth of the sieve and the divisor subpowers. The ordinary estimate applies directly if $\Delta_a(l) \le \lambda$.

In pure use $\lambda = 0$ for $l \le 4/5$. For $l \ge 3/8$ the convex function in (309) before truncation is at most $3/8$ by the endpoints. In mixed use $\lambda = 1/20$ for $l \le 5/8$. The smoothing branch requires $l > 43/120$, and the same convex comparison gives $t \le 43/120$. The additive term meets the stated criterion in both cases. The exact deletion e-sum has divisor cost at fixed P, the r'-sum is harmonic,

and the P-count pays the shortening above. The resulting sharp rectangle loss is $(\Delta_a(b)+\lambda)/2$. It uses the physical ratio $\sqrt{K/B}$ even when B and K are unequal. The ordinary two-moment alternative has loss $(\Delta_a(b)+\Delta_a(k))/2$.

The coupled comparison keeps a sharper tail for $Nr' \ge D^q$ at $NP \asymp M$. Freeze e_f, expand the remaining $(x,P)=1$ by $s_2 \mid (x,P)$, with $s_2 \perp s_1$, and fix the powered portion P_s on s_2. Set $P = P_s P'$ and group the insertion coefficients by the actual $P' e_P(r')^2 m$, of length $K' = K/(Ns_1\, Ne_f\, NP_s)$. The coefficients are divisor bounded and signed. They include the original full P-dyad, threshold, old residue test and clipped physical interval. For the sixth-free reduction of this grouped ideal, write it as $\mathfrak{s}\, b_0 v^6$, and split $P' = P_0 P_1,\ r' = r_0 r_1$ by the support of v. The original $r' \perp P$ gives $P_1 r_1^2 \mid b_0$. The two inverse-norm sums in the support count are bounded by

$$\min(1, (NP_s NP_0/M)^{1/2}), \qquad \min(1, Nr_0 D^{-q}).$$

The first is the powerful reciprocal dyad count and the second the inverse-square ideal tail. The resulting squared mass at fixed $\mathfrak{s}, v, P_0, r_0$ is at most their product times $K'/(N\mathfrak{s}(Nv)^6)$, up to subpowers. After square roots the v-sum converges, since $(NP_0)^{1/4}(Nr_0)^{1/2}(Nv)^{-3} \le (Nv)^{-5/4}$. The S-part sum also converges. Finally $NP_s \ge (Ns_2)^2$ makes the product of the insertion shortening and the erased s_2-shortening harmonic. Thus the relative norm saving in this branch is $M^{-1/4} D^{-q/2}$. It holds at the actual physical lengths for both ratio orientations.

53.3 Coupling the dual to the shorter ordered factor

Open x into its actual two ordered cutoff factors. Fix the assignments of the cutoff base and of s_1, and retain the individual cutoff bounds. Before expanding their remaining coprimality, both factors are individually squarefree and keep all their fixed deletions. Use the common squarefree w projector, with pre-stripping eligibility in both arrays and with both residual factors avoiding w after shortening. The fixed phase is $\mu(w)\psi_u(w)^2$. Sharp norm product tests separate with polylogarithmic Fourier mass on the clipped polynomial ranges, keeping their endpoints. The auxiliary primitive insertion dual has the same raw orientation as x. Multiply it by the shorter cutoff factor as an actual ideal before forming row norms. The grouped coefficients have divisor bound and no added coprimality with the auxiliary frequency. With shortened lengths X_s, X_l the physical product is $X_s X_l \asymp B/(Ns_1(Nw)^2)$. The actual cutoff inverse norm and the smoothed transform mass $\sqrt{L}$ therefore give the factor

$$\frac{\sqrt{K/B}}{Ns_1 Nw\, Nr'\sqrt{Ne}\sqrt{M}}.$$

The P-count pays the $M^{-1/2}$, and the s_1, w, r' sums remain harmonic.

For reference, set $n = \log_D Ns_1,\ j = \log_D Ne,\ p = \log_D Nr',\ v = \log_D M$, and write $\ell = k - v - n - j - 2p$ for the ordinary insertion log length. In each coupled application choose the relative target loss ρ, then use the sharp branch for large n and the extra deletion shortening $D^{-j/2}$ for large j. Put

$$\sigma = \Delta_a(b) + \Delta_a(k) - 2\rho, \qquad q = \max(0, \sigma - v/2).$$

The grouped branch $p \ge q$ pays the ordinary excess down to ρ. In the remaining branch take $\delta \asymp D^{-\beta}$. With $\beta = \max(0, \Delta_a(b-n) + \Delta_a(\ell) - j - 2\rho)$ the main shell meets the chosen loss directly. The additive term has exponent $(\Delta_a(b-n)+\Delta_a(\ell))/2 - \ell/3 - j/2$. The smoothed dual exponent is

$$t = \max(0, r - k + v + n + j + 2p + 2\beta).$$

Its two sieve losses are evaluated at $u_s + t, u_l$, where $u_i = \log_D X_i$, $u_s \le u_l \le 1/2$, and $u_s + u_l \le b - n$. The degree bound pays all larger dual dyads using the same smoothing decay. The following computations specify the small-p shell and cutoff choices in the regions of interest.

53.4 Pure central and comparable regions

Set $a = r = 3/4$. On $2/3 \le b \le a,\ 1/2 \le k \le 31/40$ the coupled loss is the smaller of $\Delta_a(b)/2$ and

$$\rho_c = \begin{cases} (25b - 2k)/204 - 7/272, & k \le a, \\ (25b + 22k)/204 - 31/272, & k \ge a. \end{cases}$$

To verify the second choice when it is smaller, put $j_0 = \Delta_a(b) - 2\rho_c,\ n_0 = 3j_0$. Use the sharp branches at $n \ge n_0$ or $j \ge j_0$, as described above. In the remaining small-p range choose

$$\beta = \sigma - v/3 - 2n/3 - 4j/3 - 2p/3, \qquad 6\rho_c = (a+b)/2 - k + 8\sigma/3.$$

Since $v + 2p \le 2\sigma$, the lower bounds for ℓ, β on the two k-pieces are, respectively,

	ℓ	β
$k \le a$	$47/68 - 9b/17 + 11k/51$	$73/408 - 9b/34 + 8k/153$
$k \ge a$	$13/34 - 9b/17 + 32k/51$	$-97/204 - 9b/34 + 283k/306.$

Their minima are $41/102, 1/153$ and $31/68, 1/51$. Thus $\ell > 3/8$ and $\beta > 0$. The lower slope $1/3$ gives $\Delta_a(\ell) \le \Delta_a(k) - (k - \ell)/3$; the main shell costs at most ρ_c. For $k \le a$ the additive excess is at most $-\ell/6 - n/6 - j/2$. For $k \ge a$ it is at most $13/96 - 31/204 < 0$.

The dual satisfies $t \le a - k + 8\sigma/3$. In the two k-pieces, $b/2 - k + 8\sigma/3$ is $(25b - 2k - 18)/34$ and $(25b + 22k - 36)/34$, respectively, both negative on the stated rectangle. Thus $u_s + t < a$. The geometric mean excess is bounded by

$$\tfrac{1}{6}\max(1/8, (a+b)/2 - k + 8\sigma/3) = \rho_c.$$

Indeed both sieve arguments use the lower slope, with $u_s \le b/2 - n/2,\ u_s + u_l \le b - n$ and $u_l \le 1/2$.

Write ρ for the smaller loss. Substituting it into the harmonic comparison at $\theta = 1/6$ gives the short target for $57/80 \le b \le a,\ k \le (172b + 21)/264$. The ordinary sharp loss already covers $k \le b/2 + 3/16$. At $\theta = 1$ and $k \ge a$, the retained decay gives the long target for $273/404 \le b \le a,\ k \ge (508b + 1743)/2768$. The sharp comparison covers $k \ge 5/8 + 4b/21$ through $4/5$, and the ordinary long insertion comparison decreases above $4/5$. The fixed product selection of Section 53.1 instead gives $k \le (308b + 21 - 204\xi)/400$ on the lower k-piece. The two divisor-prefix choices with partner exponents $1/2, 1/4$, retaining the actual partner decay, give target through

$$e \le \max\left\{\tfrac{1}{4} - 2\rho + 7(k-a)_+,\ \tfrac{1}{4} - 4\rho + b - k + 15(k-a)_+\right\}, \qquad Nd \le D^e.$$

For $B \asymp K$ and $b \le a$, the central formula gives $\rho \le 1/17$, with the sharp branch sufficient below its central rectangle. For comparable $b \ge a$, the shifted comparison below has loss at most $1/17 + (1/51 + 5/12)(b - a)$ until the long target applies. The partner decay in the half split pays this additional loss. The two half powers and the one-plus-half powers then give $D^{55/68+\epsilon}$ over all d and the target on $Nd \le D^{9/68}$, respectively. The latter follows from $5/8 + 1/17 + 9/136 = 3/4$.

53.5 Pure long insertion for every erasure length

Take $k = a + y,\ 0 \le y \le 1/56$. Work first at the reference erased exponent a, and put

$$\rho_0 = 1/17 + y/51, \quad n_0 = 6(1/16 - \rho_0), \quad j_0 = n_0/3, \quad \sigma = 9/68 + 27y/34.$$

The sharp and sparse branches use these reference thresholds, with the actual insertion length. Their erased ordinary excess increases by at most $5(b - a)_+/6$, before its square root. In the

remaining branch use $\beta = \Delta_a(a-n) + \Delta_a(\ell) - j - 2\rho_0$. Here $\ell \geq 31/68 - 22y/51 > 3/8$ and $\beta \geq 1/51 - 16y/153 > 0$. The main shell has reference loss ρ_0; the additive term has excess at most $1/8 + 5y/12 - (31/68 - 22y/51)/3 < 0$.

If $\ell \leq a$, then $\beta = \sigma - y/2 - v/3 - 2n/3 - 4j/3 - 2p/3$. If $\ell \geq a$, then $\beta = \sigma - 5v/6 - 7n/6 - 11j/6 - 5p/3$. The first gives $t \leq 8\sigma/3 - 2y$, and the second $t \leq 2\sigma - y \leq 8\sigma/3 - 2y = 6\rho_0$. At $b > a$, reduce u_s to at most $(a-n)/2$ and then u_l only enough to make the sum at most $a - n$. The reference pair remains ordered and nonnegative. The total reduction is at most $b - a$, either by the shorter-factor cap alone or by the original bound on the sum. The reference grouped length is below a, since $3/8 + 6\rho_0 < a$. Restoring the removed lengths adds at most $5(b-a)/12$ in the geometric mean, by the 5/6 Lipschitz bound through actual exponent one. Thus the relative loss is $\rho_0 + 5(b-a)_+/12$.

For $b \leq a$ use (306) at $\theta = 1$; for $b \geq a$ use (307) with the live divisor cap. The respective exponents after summation are at most

$$a + \rho_0 + (b-a)/2 - 7y/2, \qquad a + \rho_0 - (b-a)/12 - 7y/2.$$

Both meet target at $y \geq 6/355$. For $k \geq a + 1/56$ the sharp loss $\Delta_a(b)/2$ already suffices through 4/5, by the same two radial comparisons. Above 4/5 the ordinary insertion adds $\Delta_a(k)/2$; at the endpoint $1/16 + 1/12 - 7/40 < 0$, and the bound decreases thereafter. This gives $K \gtrsim D^{1089/1420}$ over all B and d.

53.6 Pure long erasure and the intermediate gap

For $a \leq b \leq 1,\ k \leq 4/5$, the sharp loss is $\rho_{\mathrm{sh}} = 5b/12 - 1/4$. The three-half choice with partner exponent 1/3, summed over $E \ll D/B$, has exponent $a + 5(1-b)/12 - (a-k)_+/6 - 7(k-a)_+/2$. The high radial comparison (308) has exponent $31/8 - 37b/12 - k/2 - 3(k-a)_+$. For $k \leq a$ the two reach target when

$$k \leq L_b = -7/4 + 5b/2 \quad \text{or} \quad k \geq H_b = 25/4 - 37b/6.$$

They overlap for $b \geq 12/13$. For $k \geq 4/5$ the ordinary insertion in the high comparison adds $\Delta_a(k)/2$. Its endpoint exponent $409/120 - 37b/12$ is below target for $b \geq 9/10$, and decreases with k. For the small residual product selection, the low endpoint is instead $1/2 + b/6 - \xi/2$; it overlaps the high endpoint at $b \geq 69/76 + 3\xi/38$.

For $1355/1482 \leq b \leq 12/13$, work in the gap $L_b \leq k \leq H_b$. Define

$$g = \min((k - L_b)/6, (H_b - k)/2), \qquad \rho = \rho_{\mathrm{sh}} - g, \qquad C_b = 17/4 - 4b.$$

The two losses agree at C_b, where $\rho_C = 3b/2 - 5/4$ is the minimum of ρ. It suffices to achieve the chosen ρ, by using the smaller radial comparison. Use $j_0 = \Delta_a(b) - 2\rho,\ n_0 = 6j_0/5,\ \sigma = j_0 + \Delta_a(k)$, and the exact nonnegative width exponent from Section 53.3. The two sharp branches and the grouped tail apply with these choices. The minimum of $k - 2\sigma - n_0 - j_0$ occurs at C_b and equals $(233b - 202)/30 > 1/3$; this bounds ℓ from below. The shell additive excess is at most $1/5 - 1/9 < \rho_C$.

When $\ell \geq 3/8$ and $\beta > 0$, the slopes on $b - n$ and on $[\ell, k]$ bound t by $t_1 = a - k + 8\sigma/3$. When $\beta = 0$, the bound is $t_2 = a - k + 2\sigma + n_0 + j_0$. For $\ell < 3/8,\ \beta > 0$ use $\beta = j_0 - 5n/6 - j$, which again gives $t \leq t_2$. Both upper expressions peak at C_b, as is seen from the slopes $1/3, -1$ of j_0. At C_b their difference is $(438b - 397)/180 \geq 0$, so $t \leq t_* = (95 - 96b)/18$.

For $u_s + u_l \leq b$, convexity and monotonicity of the sieve losses reduce the geometric mean maximum to $(u_s, u_l) = (b - 1/2, 1/2)$ and $(b/2, b/2)$. Indeed the maximizing u_l at fixed u_s is $\min(1/2, b - u_s)$. Before the breakpoint use monotonicity and after it use convexity. At the two

endpoints the shifted length is between a and one. The second endpoint is larger by $(1-b)/8$, with value

$$7b/24 + 5t_*/12 - 5/16 = 815/432 - 139b/72 \le \rho_C.$$

The last inequality is exactly $b \ge 1355/1482$. This closes the gap for every K after the two radial comparisons and the independent middle subtraction.

53.7 Pure short insertion over large erasures

For $b \ge a$ and $k < a$, the harmonic choice $\theta = 1/6$ and the quarter choice with the primal divisor cap meet target if the rectangle loss is at most

$$\rho = \max\{(b-k)/3,\ b/2 - k/4 - 3/16\}.$$

With the sharp loss this covers $b \le 3 - 4k$ or $b \ge 3k - 3/4$, which is exhaustive for $k \le 15/28$. For $15/28 \le k \le 921/1652$, work inside the remaining triangle and split it at $b = 9/8 - k/2$. Use the same j_0, n_0, σ and exact width as in the intermediate gap. Here

$$j_0 \le 3/236, \quad b - n_0 \ge 3117/4130 > a, \quad \ell \ge 3159/8260 > 3/8, \quad \rho \ge 1/16.$$

These bounds follow by the affine formulas on the two triangles. The main shell meets ρ, and its additive term is at most $(7/26 + 1/16)/2 - 1/8 < 1/16$. The same t_1, t_2 estimates apply, with $t_1 - t_2 \ge 67/4130$. Thus it is enough to evaluate $(\Delta_a(u_s + t_1) + \Delta_a(u_l))/2 - \rho$ at the two cutoff endpoints of the previous proof. On each triangle these expressions are convex. Their vertex values in the two endpoint orders are

k	b	$(b - 1/2, 1/2)$	$(b/2, b/2)$
15/28	6/7	−5/168	−1/48
921/1652	318/413	−83/9912	−83/9912
921/1652	699/826	−51/3304	0
921/1652	381/413	−787/19824	−85/2832.

The sharp branches, sparse tail and smoothed dual therefore meet ρ throughout. This proves the target for $K \lesssim D^{921/1652}$ and $B \gtrsim R$ over all d.

53.8 Mixed ratios and the critical-capacity strip

At $a_* = 5/12,\ r_* = 7/12$ use the mixed sharp insertion loss $\lambda/2 = 1/40$ through $k_0 = 5/8$. Denote Δ_{a_*} by Δ. For $k \ge b,\ k \le k_0$, the harmonic $\theta = 1$ comparison is at most $r_* + \Delta(k)/2 + 1/40 - 3(k - r_*)_+$, maximal at $k = r_*$. For $k \le b \le r_*$ use $\theta = 1/6$ and the sharp moment, with the same bound 57/80. For $k \le b,\ b \ge r_*$ and $k \le k_0$ the quarter comparison and live divisor cap give exponent

$$23/48 + \Delta(b)/2 + 1/40 + (k-b)/4 + (1-b)/4 - 15(k - r_*)_+/4.$$

At fixed b this is maximal at $k = r_*$. It equals 57/80 at $b = r_*$, decreases with slope $-1/12$ through 5/6, and is constant thereafter. For $k > k_0,\ k \ge b$ the ordinary $\theta = 1$ exponent is at most $r_* + \Delta(k) - 3(k - r_*)$, maximal at k_0 with value 101/144. For $k > k_0,\ k \le b$ use $\theta = 1/6$. The exponent decreases in k; at $k = k_0$ it is convex in b on $[k_0, 1]$, with endpoint values $2/3, 205/288$. Both are below 57/80. Adding the separate harmonic equality count gives 57/80 over all physical ratios and divisors.

Under $A_1A_2^2 \ll H \ll D^{3/4}$, the capacity bounds $r \le (3/4+a)/2$ at the level of exponents. The actual Q_u is still used in each moment. On the left, put $a = a_* - z,\ 0 \le z \le 1/8$. Use $\lambda = 1/20 + 3z/5$ through $k_0 = 5/8$. At $l = 43/120$ the three nonconstant monomials of $\Delta_a(l)$ are $-7/120 + z,\ 1/20 + z/6,\ 1/48 + 2z/3$, each at most λ. The smoothing branch is therefore longer than this l. The reference convex dual expression has endpoint maximum $31/90$, and the capacity and $\Delta_a(l) \le \Delta_{a_*}(l) + z$ give $t \le 31/90 + 3z/10$. The three nonconstant monomials at that upper bound are

$$-13/180 + 13z/10, \quad 49/1080 + 4z/15, \quad 1/108 + 11z/12,$$

each at most λ by the endpoints in z. The additive term is controlled since $3a \ge 7/8$. The erased and sharp insertion half-losses together increase by at most $4z/5$. In the harmonic branches the capacity prefactor saves $z/2$, while in the quarter branch it saves $5z/8$. In both ordinary branches $k > 5/8 > r_*$, so the active partner decay and capacity prefactor pay the increase of the two sieve factors. This gives $57/80 + 3z/10 \le 3/4$.

On the right take $a = a_* + x,\ 0 \le x \le 59/912$, and use $\lambda = 1/20 + x/4$ through $k_0 = 31/48$. On the smoothing branch $l > 43/120$. In (309) the possible r-increase of $x/2$ is paid by $2(\lambda - 1/20)$. The reference convex function on $[43/120, 31/48]$ has endpoint values $9/40, 43/120$. Thus $\Delta_a(t) \le 1/20$; the additive term is controlled by $k_0 < 3a$. In the ratio comparisons take the capacity upper bound $r_b = r_* + x/2$. The scalar estimates increase with r, while the sharp proof has already retained the actual conductor. In the sharp ranges the maximum is $r_b + (\Delta_a(r_b) + \lambda)/2 = 57/80 + x/2$. For the quarter range, after setting $k = r_b$, the b-dependent term $(\Delta_a(b) - b)/2$ is nonincreasing, giving the same maximum. In the ordinary ranges the $\theta = 1$ endpoint is $21/32 + 4x/3$. The $\theta = 1/6$ bound at $k = k_0$ has b-endpoint values $29/48 + 7x/4$ and $373/576 + 19x/12$; convexity bounds its interior. All are at most $3/4$, with the last endpoint attaining that value at $x = 59/912$. The union is the first-pivot strip $7/24 \le a \le 439/912$, subject to the original capacity and with no independent closeness requirement on A_2.

All comparisons here preserve the original $\mu(d)$ in the identity and use an absolute divisor bound at the end. The remaining pure rectangles, the coherent divisor sum, other signed components and the determinant and zero-location comparisons remain open.

54 Wide cutoff corrections and full-frequency compensation

Keep the pairing (300), the exact ordered cutoff identity (301), and all the original allocation, primitive-conductor and raw-zero conditions. At one live allocation put $h = re_0$, $E_1 = Ne_1$, and write the actual cutoff ideal as $\zeta = e_0 z_0 = \zeta_1\zeta_2$. Thus $T_i = U/N\zeta_i \ge 1$, $E = Nd$, and

$$\Omega = \frac{E}{T_1T_2} = \frac{EN\zeta}{U^2} < \frac{14}{E_1Nh}.$$

Let $\mathcal{P}$ consist of the good prime ideals of norms seven or thirteen, with the original S-deletion. These are at most four separate ideal labels. Put $s(E) = \mathbf{1}_{E\le T_1} + \mathbf{1}_{E\le T_2} - 1$. The exact wide correction is

$$\begin{aligned} F_d(T_1,T_2) &= \mu(d)\big(s(E) + C_d(T_1,T_2)\big), \\ C_d(T_1,T_2) &= \sum_{\substack{mn=d \\ Nm>T_1,\ Nn>T_2}} \left(1 - \sum_{\substack{p\in\mathcal{P} \\ p\mid m}} \mathbf{1}_{Nm>T_1Np}\right). \end{aligned} \tag{310}$$

The two factors are coprime because d is squarefree. The strict tests retain the original cutoff endpoints. For $E_1Nh \geq 2$ only the unweighted factorization count remains. For $E_1Nh \geq 14$ the correction vanishes.

At fixed row and fixed primal and dual labels, prescribe the exact intersection δ of d with the full forbidden prime union, including $\mathcal{P}$. In a divisor dyad $Y \leq Nd < 2Y$ put $Y' = Y/N\delta$. For fixed $\kappa > 0$ and $Y' \geq D^{\kappa}$, the fixed-label scalar after taking out its fixed coefficient and row phase has bound

$$\ll_{\kappa} \mathcal{M}_{24}(W)^2 \frac{4^{\omega(\delta)} \log^C(2D)}{Y\, N\delta} \exp(-c_{\kappa}\sqrt{\log(2Y')}). \tag{311}$$

The estimate includes both wide corrections, with the original full prime deletions. It gives subpower cancellation before a coherent average over the rows and other labels.

For each $p \in \mathcal{P}$, replacing the partner coefficient by its signed p-subtraction in (310) gives a zero primal amplitude $V_d^{[p]}$. The exact full-frequency complement is $V_d^{[p]} - \mathcal{E}_d^{[p]} = -\mathcal{E}_d^{[p]}$, where $\mathcal{E}_d^{[p]}$ is its middle equality projection. The original elementary equality norm gives a contribution of order at most $R(DH)^{\epsilon}\mathcal{M}_{24}(W)^2$ after the full divisor sum. This comparison requires all partner modes, including the re-expressed direct component.

In the pure comparable sector $A_1 \asymp R \asymp D^{3/4}$, $A_2 \asymp 1$, $B, K \asymp R$, the equality-side wide correction reaches $D^{3/4+\epsilon}$ on $Nd \leq D^{13/56}$ with the inherited profile cost. Its partner can carry the original coefficient, the step part, or the correction part. Thus this prefix includes their double correction. The remaining coherent wide sum and the other signed components remain open.

55 Derivations for the wide cutoff correction

55.1 The strict band and its prime terms

Use the original r, e_0, e_1, J, z_0 on each side. The primal-zero removal at the dyad $J \leq N\zeta < 2J$ implies $D/(EE_1(Nh)^2) > Nz_0/4$. Hence

$$\frac{EN\zeta}{U^2} < \frac{4D}{U^2}\frac{Ne_0}{E_1(Nh)^2} \leq \frac{14}{E_1Nh},$$

using $e_0 \mid h$ and $2D \leq 7U^2$. For a band label $T_1 < Nm < E/T_2$ in (301), $m \neq 1$, so the full inner Möbius sum vanishes. In its omitted part use $c = m/a$. The inner sum in (301) becomes

$$-\sum_{\substack{c \mid m \\ Nc < Nm/T_1}} \mu(c).$$

Here $1 < Nm/T_1 < \Omega < 14$. In the Eisenstein field the primes above two and three lie in S. The remaining inert primes have norm at least 25, and the only possible good prime norms below 14 are $7, 13$. Two good prime factors already have norm at least 49. Consequently this sum includes the unit and precisely the primes specified in (310). Substituting $n = d/m$ turns $Nm < E/T_2$ into $Nn > T_2$. The upper bound on Ω also proves the two stated conditions for vanishing of the prime term and the band.

If a p-subtraction is present, $Nm > T_1Np$ and $Nn > T_2$ give $EN\zeta > NpU^2 \geq 2D$. The original primal profile for that individual cutoff is

$$W\left(\frac{EE_1(Nh)^2Nz_0\,Nb}{D}\right), \qquad Nb \geq 1.$$

Its argument exceeds two, since $E_1(Nh)^2 \geq Ne_0$. The individual cutoff has zero primal contribution even if its dyad survived the original primal-zero removal. Its Poisson components are still present in the exact paired selection.

55.2 The fixed-intersection scalar estimate

Fix the actual row, primal and dual labels, ordered factorizations, original dyads, and both deletion allocations. Let $\mathcal{L}$ be the union of the primes of S, of the row, of all these labels, and of $\mathcal{P}$. Its radical has polynomial norm after the inherited $N\mathfrak{k} \leq D^5$ clip and the original support bounds. Choose the exact intersection of d with $\mathcal{L}$ as a squarefree ideal δ, subject to all required memberships and deletions. The good e_1 factors, conjugate d-overlaps, and actual new d-residue supports are fixed by these choices. Thus the original row phase is fixed on the free part of d, and an original zero stays zero. The free part avoids the full union $\mathcal{L}$.

In the product of the two copies of (310), their real factors $\mu(d)$ leave the original outer $\mu(d)$. The step-step term needs one ideal variable and a step-correction term uses $d = mn$. For two corrections, use the four cells of primes in d specified by their membership in the two factors m_0, m_1. Both m_i's and their complements are products of the appropriate cells. This is an exact squarefree factorization with pairwise coprime cells, including units. Expand the finite small-prime choices and assign the fixed primes of δ among the cells, with at most $4^{\omega(\delta)}$ assignments. The tests $p \mid m_i$ are thereby fixed. The cell ideals have form $\delta_j q_j$, and the sign is $\mu(\delta)\prod_j \mu(q_j)$.

Select the largest free cell by its norm, with deterministic index order on ties, and fix the other free cells. If its norm is t, then $t \geq (Y')^{1/4}$. Its interval lies on the scale $L = Y'/\prod_{\text{other}} Nq_j$, with $L \gg (Y')^{1/4}$ for every nonempty selection. Its deleted set $\mathcal{T}$ consists of $\mathcal{L}$ and the primes in the other free cells, and still has polynomial radical norm. Every factorization inequality is an interval test in t, since each m_i, n_i either includes the selected cell or is fixed. The tie order changes only an endpoint. The original primal and direct tests and the middle cutoffs also give finitely many interval tests. At fixed labels X' is proportional to $1/E$, the middle parameter M is proportional to E, and the integer norm equality is fixed. The original strict and nonstrict endpoints remain in the scalar sum.

For the fixed field $\mathbb{Q}(\omega)$, write $M_{\mathcal{T}}(x) = \sum_{Nc\leq x,\,(c,\mathcal{T})=1} \mu(c)$. The fixed-field deleted Möbius input used here is

$$|M_{\mathcal{T}}(x)| \ll_\kappa x \log^C(2D) \exp(-c_\kappa \sqrt{\log(2x)})$$

for $x \geq D^\kappa$ and polynomial radical norm. The underlying fixed-field zero-free input is Thorner–Zaman, *Algebra & Number Theory* 13 (2019), 1039–1068, Theorem 3.1, specialized to the trivial extension of this one fixed field. The classical truncated Perron shift for the reciprocal zeta function gives $M_\varnothing(x) \ll x\exp(-c\sqrt{\log(2x)})$. A possible fixed exceptional zero affects the fixed starting threshold. Artificial deletions are transferred exactly by

$$M_{\mathcal{T}}(x) = \sum_{\substack{Na\leq x\\ \operatorname{supp} a\subseteq\mathcal{T}}} M_\varnothing(x/Na).$$

All powers of the deleted primes occur in this convolution. Split at $Na = \sqrt{x}$. In the first range use the fixed-field bound. In the second use the elementary ideal count and $1/Na \leq x^{-\eta/2}(Na)^{-1+\eta}$, with $\eta \asymp 1/\sqrt{\log(2x)}$. The finite reciprocal Euler product at $1-\eta$ costs a power of $\log(2D)$. Indeed, if the deleted radical has norm $\ll D^{C_0}$, split its primes at $\log^2(2D)$. Below that threshold the ideal count bounds the reciprocal-power sum by $O((\log(2D))^{2\eta}\log\log(3D))$. Above it use the bound $O(\log(2D))$ for the sum of their log norms. For $x \geq D^\kappa$, these estimates give the claimed uniform logarithmic Euler factor and subpower decay. The underlying analytic zeta function stays fixed.

After factoring out the fixed scalar and phase, the weight is $E^{-2}G_W(c_0/E)\overline{G_W(c_1/E)}$ for fixed positive c_i. Its supremum plus integrated absolute derivative is $O(\mathcal{M}_{24}(W)^2Y^{-2})$ on each remaining interval. This follows from boundedness of the radial transform and its first log derivative, and the fixed dyadic range of E. Partial summation gives

$$\ll_\kappa \mathcal{M}_{24}(W)^2Y^{-2}L\log^C(2D)\exp(-c_\kappa\sqrt{\log(2Y')}).$$

Here $L \gg (Y')^{1/4}$ permits the deleted bound with a fixed smaller exponent in its polynomial lower cutoff. The remaining free cells sum harmonically by the factors $1/Nq_j$ in L. Together with the fixed-prime assignments this proves (311). It applies to a scalar with the row and all other ideal labels fixed.

55.3 The complete partner sum and the equality prefix

For each $p \in \mathcal{P}$, use the coefficient

$$-\mu(\zeta)\mu(d) \sum_{\substack{mn=d,\ p|m\\ Nm>T_1Np,\ Nn>T_2}} 1$$

on the partner and sum its original ordered factorizations. This coefficient is independent of the partner frequency and divisor bounded. The individual primal zero above proves $V_d^{[p]} = 0$. The exact original Poisson identity therefore gives the full complement $-\mathcal{E}_d^{[p]}$, with the same middle projection. Recombining all modes includes the re-expressed direct part and preserves every original raw value.

For the elementary equality bound, at fixed r, e_0, e_1, J put $L_0 = D/(EE_1(Nh)^2)$ and $J' = J/Ne_0$. The equal-norm ideals have $z_0 = g\ell$, $\mathfrak{k} = g\bar{\ell}$. There are $O(J'\log(2+2J'))$ choices on the dyad, up to the original constant-scale convention. Their norm Nz_0 appears inversely in the Poisson scalar, and the equal-norm profile scale is comparable to L_0/R. Thus the pointwise row bound is at most a subpower times $\mathcal{M}_{24}(W)L_0R^{-1/2}(1+L_0/R)^{-4}$. The $O(R)$ original rows give squared counting norm $O((DH)^\epsilon\mathcal{M}_{24}(W)^2RL_0)$. Summing the square roots over the original r, e_0, e_1, J uses the harmonic factors $1/(Nr\,Ne_0)$ and fixed divisor costs. After division by D, the squared norm is $O((DH)^\epsilon\mathcal{M}_{24}(W)^2R/E)$. The same proof handles the modified coefficient because it is divisor bounded and row independent at each allocation. Apply Cauchy to the original and modified middle projections and sum harmonically in d. This proves the target-size full-frequency compensation of Section 54.

For the pure comparable prefix, keep the correction on the equality side. Its coefficient is divisor bounded on its fixed labels. A partner correction, when included, is a coefficient on x alone after fixing the other labels. The insertion interval remains complete inside the original eligibility. The sharp insertion proof with powerful freezing and exact common-factor eligibility applies. Its relative loss is $D^{1/16}$ from the erased sieve factor $\Phi(A_1, B/Ns_1)/A_1 \ll D^{1/8}$. A nonzero equality correction implies $E_0Nh_0 < 14$, hence $E_0(Nh_0)^2 < 196$. Therefore the equality radial scale is $y_0 = D/(EE_0(Nh_0)^2R) \gg D/(ER)$. Use both half radial powers and keep the remaining equality-side $7/2$ decay. The fixed-d rectangle is bounded by a constant times

$$(DH)^\epsilon\mathcal{M}_{24}(W)^2\frac{R}{E}D^{1/16}\left(1+\frac{D}{ER}\right)^{-7/2}.$$

Count the exact original middle subtraction independently at the same inherited R/E cost with the common factors. For $E \le D^{13/56}$, $D/(ER) \gg D^{1/56}$ at the pure scales. The remaining decay pays

$D^{1/16}$ because $(7/2)(1/56) = 1/16$. Harmonic summation gives the stated prefix. The partner-only argument retains its own ratio $N\mathfrak{k}/Nz$, so the same automatic lower bound on its radial scale is unavailable. The coherent masked wide divisor sum remains open.

56 Joint moments for the wide cutoff pairing

We work at the pure scales $A = A_1 \asymp R \asymp D^{3/4}$, $A_2 \asymp 1$, with $U^2 \le D \le 7U^2/2$. The rows are the original unweighted nonprincipal rows, with fixed S, ν, u_0, their literal ray and capacity restrictions, and their primitive conductors $Q_u \asymp R$. Use one smooth complex W supported on $[1, 2]$, the full ordered cutoff, and the original raw zeros. A target bound means $O((DH)^{\epsilon}\mathcal{M}_{24}(W)^2 D^{3/4})$ for the specified contribution to the normalized pairing (300). Constants may depend on the fixed arithmetic data, scale constants and $\epsilon > 0$. Put $E = Nd$ and $e = \log_D E$. Here d is the intermediate divisor from the exact Möbius identity. All estimates include the original middle equality subtraction.

Theorem 41 (Full-frequency partner unit part)**.** *Take the unit part of C_d on the partner side of* (310)*, with all partner frequencies including the re-expressed direct term. A selection of the actual ordered pairs $mn = d$ may depend on d, m, n. Write $L_1 = U/Nm$, $L_2 = U/Nn$ and $l_s = \log_D \max(1, \min(L_1, L_2))$. The target holds over all $E \lesssim D^{1/6}$, and more generally on the dyads with $e \le 1/4$ and $l_s \ge 1/4 + e/2$ at the level of powers. The remaining upper cost at a single dyad with $e \le 1/4$, relative to R/E, has exponent $(1/4 + e/2 - l_s)_+/6$, where $x_+ = \max(0, x)$. The all-divisor bound for this component has power* $37/48$*.*

The low-divisor estimate uses the whole wide product support, even for balanced pairs with bounded nonunit d. Its gain comes from the forced radical of the actual product ideal in the joint positive norm. For instance, at $e = 1/4$ with the longer length of order $D^{1/2}$, the upper comparison still costs $D^{1/48}$ above R/E. This excess describes the estimate.

Theorem 42 (Selected partner frequencies)**.** *Let $T \le N\mathfrak{k} < 2T$ be an actual partner frequency dyad and write $t = \log_D T$. A selector of modulus at most one may depend on $\mathfrak{k}$, the actual equality allocation and the fixed partner labels. It is independent of the row and the two varying partner cutoff ideals. In the actual Poisson representation of the unit part, including the direct term and the corresponding selected middle subtraction, the target holds in the following regions:*

$$\begin{array}{c|l} 0 \le e \le 1/8 & t \le 15/32 + 3e/8 \\ 1/8 \le e \le 1/4 & t \le 33/64 \\ 1/4 \le e \le 5/16 & t \le 17/64 + e \\ 5/16 \le e < 27/56 & t \le 11/32 + 3e/4 \\ e \ge 27/56 & \textit{all frequency lengths.} \end{array} \tag{312}$$

Thus $T \lesssim D^{15/32}$ is sufficient over all intermediate divisors. Constant factors at the endpoints are allowed.

For selectors independent of the varying equality labels, further regions follow by grouping frequencies with prescribed ideal parts. If the rational content of $\mathfrak{k}$ is at least D^b, with $0 \le b \le 9/28$, then $T \lesssim D^{9/20+3b/5}$ suffices across all d. Here the content is the largest positive integer prime to the rational primes under S whose principal ideal divides $\mathfrak{k}$. At $E \gtrsim D^{1/4}$, content of order at least $D^{41/144}$ permits $T \lesssim D^{2/3}$. A good sixth-power base of norm at least of order $D^{1/16}$, or an S-part of norm at least of order $D^{3/10}$, suffices at every frequency length. These are conditions on the actual frequency, with all ideal labels and raw zeros retained.

Theorem 43 (Equality-wide correction)**.** *In the equality-side wide correction, write the equal-norm pair as $z_0 = g\ell,\ k = g\bar{\ell}$. Let ℓ_{g} retain the prime ideals of ℓ whose conjugates also avoid S. Against the full original partner complement, or its full-frequency step or wide part with its own subtraction, the target holds on $N\ell_{\mathrm{g}} \gtrsim D^{\theta(e)}$, where*

$$\theta(e) = \begin{cases} -17/12 + 37e/6, & 17/74 \le e \le 1/4, \\ 11/24 - 4e/3, & 1/4 \le e \le 11/32, \\ 0, & \text{otherwise.} \end{cases}$$

In the distinct pure comparable selected rectangles $B, K \asymp R$ of the preceding wide analysis, the equality-side prefix extends to $Nd \lesssim D^{111/476}$, with the original allocation tests and the same choices of partner coefficient.

The small-conjugate-support band, unbalanced middle divisors and small-part main frequency range leave contributions to the coherent signed problem. The full signed assembly and the independent determinant and zero-location arguments remain open.

57 Proofs of the joint cutoff estimates

Write $[F]_p = R^{-1/p}\|F\|_{\ell^p}$ for counting norms on the actual rows and $[F]_\infty = \|F\|_\infty$. Use the same convention after fixing the second pivot, since $A \asymp R$ and there are only boundedly many second pivots. The physical sieve factors are

$$\begin{aligned} \Phi_6(A,B) &= A + B + A^{5/6}B^{1/3} + A^{1/3}B^{5/6}, \\ \Phi_3(A,B) &= A + B + 2(AB)^{2/3}, \\ \Delta_6(x) &= \max(0, x - 3/4, x/3 - 1/8, 5x/6 - 1/2), \\ \Delta_3(x) &= \max(0, x - 3/4, 2x/3 - 1/4). \end{aligned}$$

The last two functions record the power of the corresponding sieve divided by A at $B = D^x$. We display bounds up to $(DH)^\epsilon$ and the appropriate W seminorms. The dyads, divisor multiplicities and harmonic norm sums have subpower cost. Shortened lengths are replaced by one only in the sieve factor, with a bounded nonempty subunit counted directly.

57.1 Exact cutoff arrays and the equality moment

For Theorem 41, fix the ordered $d = mn$. On the support $N\zeta_i > L_i$, the primal unused good squarefree ideal c has $Nc < 2D/(EN\zeta) < 7$. Thus the partner primal unit part is the sum with $c = 1$. Its coefficient is $\mu(d)\mathbf{1}_{(d,u)=1}$ times the product of the two Möbius and raw factors at ζ_1, ζ_2, with the common weight $W(EN\zeta_1 N\zeta_2/D)$. It keeps $(d, \zeta) = 1$, both original upper bounds $N\zeta_i \le U$, and $(\zeta_1, \zeta_2) = 1$. Each nonzero summand has $N\zeta_i < 7L_i$. Expand the common-prime test using $\sum_{w|(\zeta_1,\zeta_2)} \mu(w)$, with w good squarefree and coprime to d, and set $\zeta_i = wb_i$. The residual b_i avoid w, d. Keep $\mu(w)\psi_u(w)^2$ and the row deletion. Logarithmic Fourier separation of W gives two arrays B_s, B_l with lengths $M_i = L_i/Nw$, where s, l are the shorter and longer cutoffs. Their coefficients retain the exact endpoints. In particular their true masses multiply to $U/(\sqrt{E}Nw)$.

For the equality vector, at fixed e_1, g, ℓ and a dyad $P \le Nh < 2P$, group the actual $h = re_0$ coefficients over the permissible e_0-allocations. The row factor of the resulting sum F_P is $\psi_u(h)^2$. The coefficient is divisor bounded with the original middle and cutoff predicates. The external inverse norm is $D/(EE_1NgN\ell\sqrt{R})$, where $E_1 = Ne_1$ and the actual primitive factor $Q_u^{-1/2}$ is

bounded using $Q_u \asymp R$. The Gauss root and fixed raw factors are retained. Put $T_0 = \sqrt{D/(ER)}$, $T_1 = T_0/\sqrt{E_1}$. On a dyad $P \le Nh < 2P$, the normalized weight from the inverse square and the radial log separation is

$$P^{-2}\bigl(1 + (T_1/P)^2\bigr)^{-4}\mathcal{M}_{12}(W).$$

This keeps Q_u in the separated phase.

For a fixed integer $j \ge 2$, take the j-th power of the formed h-array and group by the product of its j ideals. The fixed factors on the second pivot and on S stay in the coefficients. In the variable first pivot the h-character is cubic. Write the grouped column as $b_0 q^3$, with b_0 cubefree, and freeze q with its raw multiplier $\psi_u(q)^6$. The squared mass on a fibre is bounded by $P^j/(Nq)^3$ up to divisor factors. The sum of the square roots converges in q, and the cubic comparison gives $[F_P]_{2j} \ll \sqrt{P}\,(\Phi_3(A, P^j)/A)^{1/(2j)}$. When $T_1 \ge 1$ the radial dyad sum is bounded at $P \asymp T_1$: below it use the eighth power in the weight, and above it the sieve factor grows at most as $P^{1/2}$ against $P^{-3/2}$. When $T_1 < 1$ use unit scale. For $e \le 1/4$, take $\tau = (1/4 - e)/2$; the replacement of T_1 by T_0 in the sieve costs only the monotone enlargement, and $E_1^{-1}T_1^{-3/2} = E_1^{-1/4}T_0^{-3/2}$. In the subunit case $E_1^{-1} \le T_0^{-2}$ when $T_0 \ge 1$. The equality bound after the g, ℓ harmonic sums is

$$[\mathcal{E}_d]_{2j} \ll \sqrt{D/E}\,D^{-\tau/2+\Delta_3(j\tau)/(2j)}(DH)^{\epsilon}\mathcal{M}_{12}(W). \tag{313}$$

The same dyad calculation with the elementary ideal count gives the pointwise factor

$$\sqrt{D/E}\min(1, \sqrt{D/(ER)}).$$

Set $l_i = \log_D \max(1, L_i)$. The cutoff second moment has excess $\Delta_6(l_i)/2$ over its square-root mass. For the fourth moment, square the actual cutoff sum. The grouped product is sixth-free, with exponent two allowed at a common prime, and has divisor-bounded multiplicity. Its fourth-root cost is $D^{\Delta_6(2l_i)/4}$. For $e \le 1/8$, pair the equality sixth moment with the shorter third and longer second moments. Interpolating the shorter between second and fourth with weight $2/3$ on the fourth yields

$$\frac{\Delta_6(l_l)}{2} + \frac{\Delta_6(l_s)}{6} + \frac{\Delta_6(2l_s)}{6} - \frac{\tau}{2} + \frac{\Delta_3(3\tau)}{6} = \frac{1}{48} - \frac{l_l - l_s}{12}.$$

Here $3/8 \le l_s \le l_l \le 1/2$, $l_s + l_l = 1 - e$ at the level of powers, and $\Delta_3(3\tau) = 0$. For $e \ge 1/8$ the pointwise equality bound and the two second moments have excess at most $1/48$, because $(l_s/3 - 1/8)_+ + (l_l/3 - 1/8)_+ \le 1/24$. The normalized masses and w-sum give R/E. The original $\mathcal{E}_d^{\mathrm{unit}}$ compensation has normalized energy of order at most R/E by the grouped elementary equality bound, and so does $\mathcal{E}_d$. Their inner product has the same upper bound. Summing over d harmonically proves the $37/48$ comparison.

57.2 A forced radical in the physical sieve

For an actual good sixth-free ideal $\mathfrak{b}$ of norm of order B, put $\mathfrak{b}^\sharp = \mathfrak{b}/\operatorname{rad}\mathfrak{b}$. If every column satisfies $N\operatorname{rad}(\mathfrak{b}^\sharp) \gtrsim P$, the physical positive comparison can use

$$\begin{aligned}\widetilde{\Psi}(A, B, P) = A + B + A^{5/6}B^{1/3}P^{-1/3} \\ + A^{1/6}BP^{-1/2} + A^{2/3}B^{1/2}P^{-1/2}.\end{aligned} \tag{314}$$

The factor multiplies the original squared coefficient mass, up to subpowers. Here is the support restriction inside the pinned de Faveri argument. In its Lemma 4.4 remove the outer common ideal of norm F. The remaining scales obey $B' \asymp B/F$ and $A/(CF) \le A' \le CA$ for a fixed constant

C, and the radical threshold can drop at most by F. A prime lost from that radical divides the removed ideal. The sharp part has norm at least the norm of its radical.

In the two positive norms of the proof of Lemma 5.1, fix the sharp label s_i and the internal Möbius divisor r, with squarefree part ru_i. The original coefficient already requires $\operatorname{rad}(s_i) \mid ru_i$. Set $q_i = \operatorname{rad}(s_i)/(\operatorname{rad}(s_i), r)$, so $u_i = q_i v_i$. The v_i are squarefree and avoid r, q_i. The new label is in bijection with the original label at these fixed indices. Keep the original denominator at rq_iv_i, the ray predicates, and the raw factor χ_{q_i} at the grouped row variable. Multiplicativity gives the same zeros, including when a row prime divides q_i. This factor is a contraction in the formed norm. The original squared coefficient sum is unchanged by the reindexing.

If P_i and C_i are the norm scales of s_i and its radical, the row cap is still $K_i \lesssim (B')^2/(A'P_i)$. The column cap is now $B'/(P_iC_i)$, since $N(rq_i) \geq N \operatorname{rad}(s_i)$. Both P_i, C_i are at least of order $\max(1, P/F)$. Insert the four terms of the sixth-power comparison from de Faveri's Theorem 1.1 and Lemma 6.1, $K + K^{1/6}L + K^{5/6}L^{1/3} + K^{1/3}L^{5/6}$. The two factors in its positive Cauchy estimate carry the prefactor $A'\sqrt{P_1P_2}/B'$. Consequently the four monomials before taking their geometric mean are bounded by

$$\begin{aligned} B', \quad & (A')^{5/6}(B')^{1/3}P_i^{-1/6}C_i^{-1}, \\ & (A')^{1/6}B'P_i^{-1/6}C_i^{-1/3}, \\ & (A')^{2/3}(B')^{1/2}P_i^{-1/6}C_i^{-5/6}. \end{aligned}$$

Put $G = \max(1, P/F)$. The resulting factors satisfy

$$\begin{aligned} F^{-1/3}G^{-7/6} &\leq P^{-1/3}, \\ F^{-1}G^{-1/2} &\leq P^{-1/2}, \\ F^{-1/2}G^{-1} &\leq P^{-1/2}. \end{aligned}$$

These follow by considering $F \leq P$ and $F \geq P$. The source diagonal term gives A. This proves (314) with its full masks.

For the pure joint norm apply Cauchy with the longer cutoff as one factor. The other norm, of $F_P\overline{B_s}$, equals that of F_PB_s by pointwise modulus. The exact signed pairing keeps its original conjugation. The latter positive norm groups actual h^2b_s injectively: exponents $0, 1, 2, 3$ determine membership of h, b_s at each prime. The radical of its sharp part is h. The length is of order $P^2 \max(1, M_s)$ and the squared mass is bounded by PM_s up to subpowers. The square root of (314) grows at most linearly in P at fixed M_s; the inverse square and square-root mass supply $P^{-3/2}$. Thus the same radial maximum at T_0, with the E_1-factor as above, applies. For $e \leq 1/4$ the joint excess over the square-root mass is $\frac{1}{2}(l_s/3 + \tau/3 - 1/8)_+$; the B term is below A in power, and the other two terms in (314) have no larger exponent. With the longer second moment the total relative power is

$$\tfrac{1}{2}(l_s/3 + \tau/3 - 1/8)_+ + \tfrac{1}{2}(l_l/3 - 1/8)_+ - \tau/2 = \tfrac{1}{6}(1/4 + e/2 - l_s)_+.$$

The equality follows from $l_l \geq 3/8$ and $l_s + l_l = 1 - e$. If $e \leq 1/6$, then $l_s \geq 1/2 - e \geq 1/4 + e/2$. The middle compensation still costs R/E, proving the low-divisor and length claims of Theorem 41.

57.3 The selected Poisson product

For Theorem 42, keep the actual selected partner frequencies in the Poisson expression. The live dyad gives $N\zeta_i < 14L_i$, and $E_1Nh < 14$ on the partner wide support. Keep the ordered e_0-assignment, the common-prime projector and all shortened residual deletions. The actual partner frequency weight separates in logarithmic variables on a frequency dyad, with absolute Fourier mass bounded

by $\mathcal{M}_{12}(W)(1+T/R)^{-4}$. The primitive factor remains $Q_u^{-1/2}$ and the factor from the re-expressed direct term remains in the array.

At fixed equality g, ℓ, e_1 and $P \le Nh_0 < 2P$, the raw equality phase times the conjugated partner dual phase is $\psi_u(h_0^2 \mathfrak{k})$. The row-independent selector is included in the coefficient at fixed $h_0, \mathfrak{k}$ by summing the exact equality allocations. It can give a nonseparable coefficient, with divisor-bounded multiplicity and all the original signs. Write the grouped product uniquely as $\mathfrak{s} b_0 v^6$, with $\mathfrak{s}$ supported on S, v good and b_0 good sixth-free. Factor $h_0 = h_{\mathrm{in}} h_{\mathrm{out}}$ by support in v. Then $h_{\mathrm{in}} \mid v$ and $h_{\mathrm{out}}^2 \mid b_0$, even when the actual frequency meets h_0. At fixed $\mathfrak{s}, v, h_{\mathrm{in}}$, count h_{out} at length P/Nh_{in} and divide the physical residual length by $(Nh_{\mathrm{out}})^2$. The squared coefficient mass is bounded by

$$(DH)^{\epsilon} \frac{PT\, Nh_{\mathrm{in}}}{N\mathfrak{s}\,(Nv)^6}.$$

The exact multipliers at $\mathfrak{s}, v$ remain in the positive comparison. The sums of the square-root weights converge, since the sum over $h_{\mathrm{in}} \mid v$ costs at most $(Nv)^{1/2+\epsilon}$. Also $N \operatorname{rad}(b_0^{\sharp}) \ge P/Nv$. Apply (314) with this threshold, truncated below by one, and with the true length at most of order $P^2T/(N\mathfrak{s}(Nv)^6)$. Each inverse threshold exponent β is at most six times its length exponent α. Thus the sieve may use the nominal arguments P^2T, P after this shortening. The radial sum at fixed T still has its maximum at the equality scale, as the sieve square root grows at most linearly in P.

Put $\Gamma_4 = (\Delta_6(2l_s) + \Delta_6(2l_l))/4$. The two fourth moments of the cutoffs give this excess, with their true shortened masses. Before the sieve factors, the normalized mass at fixed labels and radial dyad, omitting the harmless fixed shortening and the harmonic common-prime factor, is

$$\frac{\sqrt{L_1 L_2}\sqrt{T}\, P^{-3/2}}{EE_1 NgN\ell}.$$

The factor $P^{-3/2}$ includes the equality inverse square and its square-root mass. The outer normalization by D cancels the equality prefactor D; conversion of the counting norm back to the row sum supplies R against the two factors of size $Q_u^{-1/2}$. For $e \le 1/4$ and the indicated frequency ranges above $t = 3/8$, the relative joint sieve excess is $(t/3 + \tau/3 - 1/8)/2$. Adding the physical frequency factor, the cutoff moments and the equality radial gain gives $2t/3 - 7/16 + \Gamma_4 - \tau/3$. For $e \le 1/8$, $\Gamma_4 = 1/6 - 5e/12$. For $1/8 \le e \le 1/4$, convexity and the individual length caps give $\Gamma_4 \le 13/96 - e/6$. These are the first two branches of (312); smaller t follows by monotonicity of the upper comparison.

For $e \ge 1/4$, use the unit-scale equality dyad and the unshortened frequency sieve at T. The total power above R/E is

$$t/2 - 1/4 - e/2 + \Gamma + \Delta_6(t)/2.$$

For the third and fourth branches one may take $\Gamma = \Gamma_4$, bounded by $13/96 - e/6$ through $e = 5/16$ and by $1/12$ thereafter. For $e \ge 27/56$ use the lesser of Γ_4 and $\Gamma_2 = l_s/2 + \Delta_6(l_l)/2$; the latter uses the shorter pointwise bound and longer second moment. Under $l_s + l_l \le 29/56$ and $l_l \le 1/2$, the maximum of this minimum is $3/56$. Indeed l_l may be increased to its permitted endpoint; the two affine bounds cross at $l_s = 5/56,\ l_l = 3/7$. For $l_s \le 1/56$ the pointwise option is smaller than $3/56$, and for $l_s \ge 1/7$ convexity bounds the fourth-moment option by its values at $1/7, 29/112$. At $t = 3/4$ the resulting power is at most $3/16 - e/2 + 3/56 \le 0$. Above R, the remaining growth has slope at most one and the factor $(1+T/R)^{-4}$ gives a summable tail. For equality-dependent selectors, sum the actual equality label vectors only after using each label's own selected middle equality subtraction. Each of these subtractions has the elementary normalized energy R/E. The sum of the

equality-label pointwise majorants is $O((DH)^{\epsilon}\mathcal{M}_{12}(W)\sqrt{D/E})$, by the original harmonic g, ℓ sums and the inverse-square radial bound. Cauchy at each label gives the required total compensation of order R/E. The divisor sum is harmonic.

57.4 Sparse frequency and equality support

For a frequency-only selector, rational content at least D^b leaves $O(TD^{-b})$ ideals on a dyad. Divide each frequency by its rational principal factor and sum the inverse squares of the possible contents. For $j = 1, 2$, power the formed frequency array j times before taking its second moment. Grouping by product gives squared mass at most $(DH)^{\epsilon}(TD^{-b})^j$ by the divisor bound. At fixed S-part $\mathfrak{s}$, a dyad Q of the sixth-power base q has two bounds for the sum of square-root masses. They are $T^{j/2}Q^{1/2}D^{-jb/2}$ by Cauchy over disjoint fibres, and $T^{j/2}Q^{-2}$ by the individual shortened length and ideal count. The sieve length is at most of order $\max(1, T^j/Q^6)$. On writing $x = \log_D Q/j$, the bound is convex up to $b/5$ and nonincreasing afterwards. Its maximum is bounded by the larger value at $0, b/5$. Thus the $2j$-norm of the frequency array has excess

$$\lambda_j = \max\left\{\frac{\Delta_6(jt)}{2j} - \frac{b}{2},\ \frac{\Delta_6(j(t-6b/5)_+)}{2j} - \frac{2b}{5}\right\}.$$

The two cutoff moments and (313) give the following upper powers above R/E:

$$\begin{cases} t/2 - 13/48 + \lambda_1/3 + 2\lambda_2/3, & e \le 1/8, \\ t/2 - 9/32 + 5\lambda_1/6 + \lambda_2/6, & 1/8 \le e \le 1/4, \\ t/2 - 11/32 + \lambda_2, & e \ge 1/4. \end{cases}$$

In the first two lines use the equality sixth and twenty-fourth moments respectively and interpolate the frequency second and fourth moments. The last line uses the equality supremum, the longer second and shorter fourth moments. At $t = 9/20 + 3b/5$, the first branch of λ_j dominates for both j, since the lost sieve growth across the length shortening is at least $b/10$. The first excess is zero, and the other two are nonpositive. The last branch at $t = 2/3,\ b = 41/144$ is zero. The fixed S-part and sixth-power multipliers retain their exact raw values on each fibre.

If the frequency sixth-power base is at least D^a, the product of j such bases divides the grouped base, and summing its inverse-cube weights gives D^{-2ja}. The replacement λ_j is $-2a + \Delta_6(j(t - 6a)_+)/(2j)$. If the S-part has norm at least D^{σ}, its product has norm at least $D^{j\sigma}$. Summing the inverse square roots over this fixed prime support costs $D^{-j\sigma/2+\epsilon}$; use the Euler product with any fixed positive exponent below $1/2$. Now $\lambda_j = -\sigma/2 + \Delta_6(j(t-\sigma)_+)/(2j)$. At $t = 3/4$ these comparisons meet the target with $a = 1/16$ and $\sigma = 3/10$. Monotonicity gives smaller t and the original radial decay sums the upper tail. All comparisons keep the corresponding middle subtraction.

For the equality-wide sector, fix $g, h, e_0, e_1, \ell/\ell_{\mathrm{g}}$. The last factor lies on finitely many conjugates of primes in S, and its phase is retained. The live dyad gives $E_1 Nh < 14$. On the remaining conjugate-good part the actual phase is $\nu(\ell_{\mathrm{g}})\overline{\nu(\bar{\ell}_{\mathrm{g}})}(Nu/\ell_{\mathrm{g}})_6$, including zeros at either orientation. On $T \le N\ell_{\mathrm{g}} < 2T$ the rational conductor is $N\ell_{\mathrm{g}}$ and the squared coefficient mass is of order at most $1/T$ up to subpowers. Fixing the second pivot, the oriented rational Gram bound has sieve factor $A + T^{5/3}$ in the integer variable Nv_1; equal-norm row multiplicity has divisor cost. The sum over g is harmonic. Including the alternative direct count gives normalized equality energy at most

$$\frac{D}{E^2}\left(1 + \frac{D}{ER}\right)^{-8} \min\left(1, \frac{A + T^{5/3}}{AT}\right)(DH)^{\epsilon}\mathcal{M}_{24}(W)^2.$$

For the full partner complement the four upper powers of normalized energy are $3/4 - e,\ 1 - 2e,\ 23/24 - 4e/3,\ 13/12 - 11e/6$. For its wide coefficient the same bound follows by the primal unit reduction, the exact small-prime compensation and the physical sieve at D/E; for the step use the difference of the full and wide coefficients. Cauchy requires the extra equality gain $\max(0, P(e) - 1/2 - 8(1/4 - e)_+)$, where $P(e)$ is the maximum of the four displayed powers. This gives $\theta(e)$ in Theorem 43. The actual support $ET \ll D$ ensures that the $T^{2/3}/A$ term does not obstruct this threshold. The harmonic divisor sum finishes the full-complement bound.

For the distinct comparable selected rectangles, retain the extra equality factor $(1+D/(ER))^{-7/2}$ after the original two half radial powers. The coupled cutoff estimate on the same insertion rectangles costs $D^{1/17}$ relative to R/E. This estimate accepts the original coefficient or its step and wide parts. At fixed ordered m, n and, for a subtraction, fixed p, the additional wide tests are single-factor lower endpoints on the two ordered arrays. The step tests are single-factor upper endpoints. The exact common-prime projectors and actual insertion dual retain these endpoints. Since $(7/2)(1/4 - 111/476) = 1/17$, the equality radial factor pays that cost on the stated prefix. The selected middle subtraction has its separate R/E bound. This proves the second assertion of Theorem 43.

Classical reference NIST Digital Library of Mathematical Functions, Section25.15, Dirichlet L-functions, equations 25.15.5–6; primitive functional equation and finite Gauss sums (snapshot pinned 4 November 2026). The normalization and explicit finite Gauss sums used here are stated above.