

A zero-free half-plane with boundary 874861/1000000

Hailey Collet
HaileyCollet@gmail.com

Abstract

We prove nonvanishing for finite-order Hecke L -functions of trivial infinite type over $\mathbb{Q}(\sqrt{-3})$ and for all ordinary Dirichlet L -functions in the strict half-plane $\Re s > 874861/1000000$, apart from principal poles. The proof extends the density range, derives a strictly interior local paired moment, and combines the physical probe with exact rational inequalities and a common-gap Mellin continuation. The accompanying Lean declarations discharge the analytic inputs using the pinned dependencies described in Appendix A. Verification status is recorded separately in Appendix A.4.

This paper gives a proof of the zero-free bound below and describes its Lean formalization. The source uses the pinned analytic results listed in Appendix A, with their arithmetic families, eligibility, masks, profiles and conductor conventions. The final Lean declarations have no additional analytic hypotheses. Appendix A.4 records the verification status separately from the mathematical argument.

Put

$$c = \frac{125139}{1000000}, \quad \theta = 1 - c = \frac{874861}{1000000}.$$

Theorem. Every finite-order Hecke character of trivial infinite type over $K = \mathbb{Q}(\omega)$, in any nonzero integral modulus, has nonzero ordinary L -function in $\Re s > \theta$, except at its principal pole at $s = 1$. Consequently, for every natural number $q \neq 0$, every complex Dirichlet character χ modulo q , and every $s \in \mathbb{C}$,

$$\theta < \Re s, \quad \neg(\chi = 1 \text{ and } s = 1) \implies L(s, \chi) \neq 0. \quad (\text{T})$$

There is no primitivity or reality restriction on χ . The assertion includes $\zeta(s)$ by taking modulus one. It makes no assertion on the boundary line, for nontrivial archimedean Hecke types, or with constants uniform in a growing conductor.

1. Conventions, fixed families and the base estimates

Let $\omega = e^{2\pi i/3}$, $\mathcal{O} = \mathbb{Z}[\omega]$, $\lambda = 1 + 2\omega$, $q_a = Na$, and $\alpha(a) = a/|a|$. Use

$$e(z) = \exp(2\pi i \operatorname{Tr}_{K/\mathbb{Q}}(z/\lambda)), \quad d\mu(z) = 2 \, dx \, dy / \sqrt{3},$$

and Fourier sign $e(-zy)$. Ideal sums use their primary generators congruent to one modulo $3\mathcal{O}$; element-row sums retain all six units. A fixed excluded set S contains primes above six, the target conductors, and a sufficiently large fixed small-prime cutoff. Every character power, including powers zero and six, is extended by zero at nonunits. Cancellation of a phase never removes such a zero.

For the sextic symbol put $\psi_u(n) = \chi_n(u)$. With either whole orientation and a fixed finite ray character ν , set $\varphi_u = \nu\psi_u^{\pm 1}$ and

$$M_u(D; W) = D^{-1/2} \sum_n \mu(n) \varphi_u(n) W(q_n/D), \quad S_u(N; W) = N^{-1/2} \sum_n \varphi_u(n) W(q_n/N). \quad (1.1)$$

The smooth profiles are common functions on fixed annuli. Bounded real norm powers and independently varying heights have finite smooth-seminorm and polynomial-height costs. Arithmetic coefficients and moving support are not profile seminorms.

For the positive-slot plain moment, fix the ray modulus, subgroup containing the global units, anchor η_0 , and the literal ideal $Q_* = \text{internalQ}(Q, \eta_0)$, nonzero, proper and contained in (72). Each slot is a canonical good-prime pool with an allowed relative ray character ν_i and the same annular profile W :

$$Q_{u,i} = Z^{-z_i/2} \sum_{p \in \mathcal{P}_i^{\text{can}}} \varphi_u(p) \nu_i(p) W(q_p/Z^{z_i}) (q_p/Z^{z_i})^{-\sigma_i + i\nu_i}, \quad Q_{u,I} = \prod_{i \in I} Q_{u,i}. \quad (1.2)$$

The finite slot list is fixed before $Z \rightarrow \infty$. A common polynomial-size puncture can depend on frozen outer labels but not on the row currently summed. The natural eligible rows have precisely the original non-FixedInducingRow gate attached to Q_* , not an arbitrary finite-group exclusion. The enlarged source state counts the row width plus all moving character support before cancellation; its additional common puncture is not added to that width. The local theorem below uses a fixed parent of moving width zero. Its descendants can be enlarged source states.

Define β_* as the supremum of $\frac{1}{2}$ and real parts of positive-real-part zeros of all primitive finite-order trivial-infinity Hecke functions, omitting the principal pole. The initial seven-eighths theorem gives $\beta_* \leq 7/8$. We import this initialization, not any later optimized endpoint. In the contradiction branch $\beta_* > \theta$, set

$$\mu = \beta_* - \theta > 0, \quad \kappa = 2\beta_* - 1 \in [7/10, 3/4]. \quad (1.3)$$

Here and below $\varepsilon > 0$ denotes a freely allocated small power loss. A bound with $\mathcal{H}_J$ includes finitely many smooth seminorms and a fixed polynomial in all independent input heights. Its degree may depend on the subsequently fixed arithmetic family; the real exponent saving will not.

Base estimates. The following are the actual pinned PR9/PR15 contracts, stated here to specify what is used. Their source locations and independence limits are in Appendix A.

1. The unmarked inverse moment is $\sum_{0 < q_u \leq H} |M_u(D)|^2 \ll H(HD)^\varepsilon$ for $H \geq D^{1+\vartheta}$, for every fixed $\vartheta > 0$, including sixth-power rows. The derived no-slot inverse moment on rows of norm $\ll U$, for $0 \leq r \leq 2$, has exponent

$$e_I(r) = \max\{1, (1 + 5r)/6\}. \quad (1.4)$$

It includes the common cutoff profile and its required scale derivatives.

2. The marked inverse moment with bounded row-independent slot coefficients has exponent $U^{1+\varepsilon}$ whenever

$$r + 2z_I \leq 1 - c_1, \quad 2r + 8z_I \leq 3 - c_2, \quad (1.5)$$

for fixed positive margins. Either complete orientation and every subcollection are allowed; no slot mesh is required here.

3. On the actual canonical positive family, the paired plain moment is

$$\sum_{u \text{ eligible}} |S_{u,1}(Z^{n_1}) S_{u,2}(Z^{n_2}) Q_{u,I}|^2 \ll Z^{M+\varepsilon} \mathcal{H}_J \quad \text{if } n_1 + n_2 + 6\kappa z_I \leq M. \quad (1.6)$$

Here M is row plus moving-character width, $\beta_* \geq 51/100$, and $2\beta_* - 1 \leq \kappa$. The mesh is chosen from bounded geometry and loss before the arithmetic family, slot count, profiles and live heights. The no-slot form has no plain-sum constraint for bounded $n_1, n_2 \geq 0$, and includes all nonprincipal inducing rows, including fixed-conductor ones. Two-variable scale recovery uses the four profile pairs formed from a profile and $-W/2 - yW'$. All descendant calls below use only (1.6) with coefficient $\frac{9}{2}$, or this unrestricted no-slot form.

4. For the fixed-dictionary all-free-row count, PR15's no-slot plain theorem gives

$$\sum_{u \text{ free}} |P_j(U^m)P_k(U^m)|^2 \ll (1+H)^J U^{\max(1,2m)+\varepsilon}, \quad m \leq 3/4, \quad j+k \leq 2. \quad (1.7)$$

Real powers lie in $[0, 1]$; H is arbitrary after its threshold. Finite exceptional free rows are restored by the absolute bound $|P_n| \ll \sqrt{U^m}$ with the specified profile norm. This theorem is different from the unrestricted eligible physical no-slot input in item 3.

5. The supported-witness contract applies to a finite nonprincipal natural dictionary of moduli at most U , a nonfloor buffered bin $b > 51/100$, and a fixed cutoff $t_* \in [1, 3/2]$. With an arbitrarily small tolerance ϵ_d , it supplies one actual zero, one dictionary label, a common inverse/plain frequency, and dyads satisfying

$$r \leq t_* + \epsilon_d, \quad r + m \geq t_* - \epsilon_d, \quad 0 \leq m \leq 1/2 + 75\epsilon_d, \quad r \geq t_* - 1/2 - 76\epsilon_d, \quad (1.8)$$

$$|M(U^r)|^2 \geq U^{(2b-1)r-2\epsilon_d}, \quad |P(U^m)|^2 \geq U^{(2b-1)m-2\epsilon_d}. \quad (1.9)$$

These are actual common-profile witnesses, not row-dependent artificial truncations. The saturation proof requires only $2b - 1 \geq 1/50$. Exact buffer and height choices are given in Lemma 2.1.

The arithmetic foundation of items 1–3 is the completed theta reflection and marked energy in PR9 B2. The completed index is cn^3 , with c squarefree but (c, n) unrestricted. Marks divide that entire index and may meet the row. Both collision masks, all cusp sources, fixed phases and natural zeros survive reflection. The unmarked energy estimate alone is not a replacement. Its physical corollary is

$$M + 2z \leq N \quad \implies \quad \sum_{q_m \ll Z^M} |\mathcal{T}(m)|^2 \ll Z^{M+(N+2z-2M)_+/4+\varepsilon}. \quad (1.10)$$

All exact Poisson identities used below are applied to the full signed coefficient before absolute values. Common Fourier measures are chosen before live labels. A tail is bounded on the original arithmetic expression using a fixed raw polynomial degree before a decay order is chosen. These are part of the imported contracts, not conclusions of the rational certificate.

2. Enlarged density and the strictly interior local moment

Lemma 2.1. Fixed-dictionary density

Fix a finite dictionary of natural characters with coefficients $c_{\eta_\lambda}(I)\mathfrak{s}_{m_0^{\delta}u}(I)$, $m_0 = 2(1 - \omega)$, and natural modulus norm at most $K_\lambda N(u)$. Let $N(Y, a, T)$ count free rows $N(u) \leq Y$ for which some member has a zero with real part at least a and height at most T . For every $\varepsilon > 0$, there exist C, J, Y_0 , before the live parameters a, T , such that

$$N(Y, a, T) \leq C(1+T)^J Y^{\rho(2a-1)+\varepsilon}, \quad \rho(\delta) = 1 - \frac{10\delta}{3(5-2\delta)}, \quad \frac{33}{40} \leq a \leq \frac{7}{8}. \quad (2.1)$$

The same bound after the standard sixth-power convolution applies to raw rows, retaining natural holes and orientation.

Proof. Put $K_0 = \max(1, \max K_\lambda)$ and $U = K_0 Y$. Use the natural modulus bound; an existential representative with the same ideal coefficients alone supplies no conductor bound. Remove the finite union of free fixed-inducing rows. Their primes belong to a fixed finite support and their valuations are at most five; units add a factor six. All remaining dictionary characters are nonprincipal. The removed free rows contribute $O(1)$. If the counted population is nonempty, $\beta_* \geq a > 51/100$, so the old plain input applies; otherwise there is nothing to prove.

Choose $0 < \epsilon_d \leq 1/1000$ and a moment loss $\epsilon_m > 0$. Obtain a common polynomial height degree J_0 for (1.4), (1.7), their real-power/scale derivatives and rowwise Sobolev bounds. The inverse and plain counting costs include respectively $12(1 + H)$ and $192(1 + H)$. Choose positive $e, k_{\text{aux}}, \eta_{\text{aux}}$ with

$$e < 1/1000, \quad k_{\text{aux}} \leq 1, \quad 288e + 8k_{\text{aux}} + 2\eta_{\text{aux}} \leq \epsilon_d/2. \quad (2.2)$$

Then choose $\tau_d > 0$ so that $\tau_d < 1/2$, $4\tau_d < \eta_{\text{aux}}$, and $(J_0 + 1)\tau_d$ fits the allotted loss. These choices are possible in this order and precede Y, a, T . The auxiliary k_{aux} is not the physical parameter κ .

Suppose $T \leq U^{\tau_d}$. The actual buffered-bin theorem, with $I = \lceil 49/(100e) \rceil$, partitions the population into finitely many bins, with $1 \leq i \leq I$ and a grid value b , satisfying

$$b \leq M(3iU^{\tau_d}) < b + e, \quad M(3(i+1)U^{\tau_d}) < b + 2e, \quad b > a - e, \quad b \leq 7/8. \quad (2.3)$$

Here $M(H)$, with the row u implicit, is the actual zero maximum of its fixed parent dictionary at height H , with floor .51. The buffer supplies the zero-free disks needed by the supported theorem. Since $a \geq .825$, the floor cannot occur and $\delta_b = 2b - 1 > .648 > 1/50$. Take

$$t_* = \frac{5}{5 - 2\delta_b} \in [1, 3/2] \quad (2.4)$$

before choosing witnesses on the whole bin. Apply (1.8)–(1.9) at normalized width one. A witness frequency is $2\pi\nu_u - \Im\rho_u$, bounded by $H_i = 2\pi U^{\tau_d/2} + 3iU^{\tau_d}$. The actual zero's real power is in $[0, 1]$. The inverse cutoff is the fixed dyadic profile multiplied by $V_{\leq}(U^r y/U^{t_*})$; its derivatives are uniformly controlled on a fixed annulus. Thus the actual no-slot count inputs apply for every witness: $0 \leq r < 2$ and $m \leq 23/40 < 3/4$.

Partition by the finite witness label and the two dyads, at cost $O((\log U)^2)$. Within such a part the inverse and squared plain-pair counts give

$$\begin{aligned} \#\mathcal{R} &\ll (1 + H_i)^{J_0+1} U^{eI(r) - \delta_b r + 2\epsilon_d + \epsilon_m}, \\ \#\mathcal{R} &\ll (1 + H_i)^{J_0+1} U^{\max(1, 2m) - 2\delta_b m + 4\epsilon_d + \epsilon_m}. \end{aligned} \quad (2.5)$$

For $r \leq 1$, $\max(1, 2m) \leq 1 + 150\epsilon_d$ and $\max(r, 2m) \geq 2(r + m)/3 \geq 2(t_* - \epsilon_d)/3$. Taking the better count bounds the exponent by $1 - (2/3)\delta_b t_* + 156\epsilon_d + \epsilon_m$. For $r > 1$, the first exponent is increasing in r , since $\delta_b \leq 3/4 < 5/6$; its maximum is at most $1/6 + (5/6 - \delta_b)t_* + 3\epsilon_d + \epsilon_m$. Both main expressions equal $\rho(\delta_b)$ by (2.4). No marked count or theorem restricted to $\delta \geq .7$ has been used.

On the interval in question $-\rho'(\delta) = 50/[3(5 - 2\delta)^2] \leq 200/147$. Therefore bin rounding costs less than $3e$. Allocate

$$156\epsilon_d + \epsilon_m + 3e + \text{dyadic loss} + (J_0 + 1)\tau_d < \varepsilon. \quad (2.6)$$

This proves the small-height bound, with thresholds before a, T . If $T > U^{\tau_d}$, elementary lattice counting gives $N \ll U$; choose an integer $J \geq 1/\tau_d$, so $U \leq T^J$. Since $\rho \geq 2/7 > 0$, this proves the remaining heights. Enlarge the constant for bounded Y .

Finally write a raw row as a sixth-power-free core times a sixth power. Summing the free estimate incurs $\sum_v N(v)^{-6\rho}$, convergent uniformly because $\rho \geq 2/7 > 1/6$. This also pays the infinitely many sixth-power multiples of the finite exceptional cores; they are not declared a finite raw population. Natural Euler holes add no positive-real-part zeros. Whole-character conjugation reflects zero heights and preserves the modulus. Thus the raw-row assertion holds in exactly the same fixed finite dictionary. At physical use choose a distinct external height exponent only after the returned J . $\square$

Theorem 2.2. Interior affine moment on the actual source

Put

$$\begin{aligned} q_0 &= 17/25, \quad \rho_0 = 103/273, \quad K_a = 34/21, \\ \Lambda &= 19312/4787, \quad B_0 = 14200/14361, \quad B = B_0 - 10^{-5}. \end{aligned} \quad (2.7)$$

In this theorem and its proof, set $M = 1$. For every $\varepsilon > 0$, the actual fixed parent of §1 satisfies

$$\sum_{u \text{ eligible}} |S_{u,1}(Z^{n_1})S_{u,2}(Z^{n_2})Q_{u,I}|^2 \ll Z^{M+\varepsilon}\mathcal{H}_J \quad (2.8)$$

when

$$.41M \leq n_1, n_2 \leq .44M, \quad .034M \leq z_I \leq .04M, \quad n_1 + n_2 + \Lambda z_I \leq BM. \quad (2.9)$$

One sufficiently fine common slot mesh works before the slot dictionary, live scales, masks and heights. The assertion includes the four plain derivative pairs. The width-one statement is the input used by the final assembly. We prove only this strict interior; no limiting closed-plane theorem is required.

Proof. Write $S_1 = (n_1 + n_2)/M$, $t = z_I/M$, $A = n_1 + n_2 + z_I$, and $g = (n_1 + n_2 + (9/2)z_I - M)_+$. Use the 600 original-slot cells

$$\begin{aligned} b_j &= .034 + j \cdot 10^{-5}, \quad g_j = B - 1 + (9/2 - \Lambda)b_j, \\ \bar{a}_j &= K_a g_j + \alpha, \quad A_j = B - (\Lambda - 1)b_{j-1}, \quad \alpha = 1/400000. \end{aligned} \quad (2.10)$$

On $b_{j-1} < t \leq b_j$, with $j = 1$ also at $t = .034$, fix amplifier width $a = M\bar{a}_j$. Then $g/M \leq g_j$, $A/M \leq A_j$, and every $g_j > 0$. This choice depends only on the original total slot width. It remains fixed in both equal-product rectangles, every child and every profile derivative. No affine hypothesis is imposed on descendants.

The original amplification theorem allows $\sigma = 3a > 0$ with sufficiently small mesh; it does not impose an extra smallness gate on σ . Its literal prime pool is $(Z^a/2, Z^a]$, with its original normalized average. Removing row/support/projector primes deletes only $O(\log Z)$ primes. The bounds

$$\begin{aligned} \bar{a}_{\min} &= \frac{902420521}{120632400000}, \quad \bar{a}_{\max} = \frac{289458497}{24126480000}, \\ A_{\max} &= \frac{1271830639}{1436100000}, \quad \frac{1}{4} - 7\bar{a}_{\max} = \frac{572201503}{3446640000} > 0 \end{aligned} \quad (2.11)$$

meet the original support gates. Choose the relative slot mesh below $\bar{a}_{\min}/10$ and the final loss budgets. After the arithmetic family and finite slot system are fixed, take finite maxima of all required orders and thresholds over these cells and four profile pairs before live data.

Comparison and first transform. Compare with plain lengths $M/4, n_1 + n_2 - M/4$, retaining the whole-product character, puncture, norm power, slot coefficients and profiles in the equal-product difference. Deletion/restoration of additional common masks has absolute mass $Z^{o(1)}$. Reflect only a comparison second plain longer than $M/2$, using its actual combined conductor. The reflected length is at most $M/2 + \xi$; lower annuli have summable central coefficients and upper tails have arbitrary decay. Every comparison daughter has old budget at most

$$M/4 + M/2 + (9/2)(.04M) + \xi = .93M + \xi < M. \quad (2.12)$$

Subunit retained scales use bounded dilation and empty slot products use the unrestricted nonprincipal no-slot input. The comparison is therefore bounded by the old estimate.

Enlarge the positive norm of the entire centered difference to the source smooth row ball. The first zero frequency is bounded by the original powerful-product count. On nonzero frequencies extract the complete common supports C, D . Let c, d, p be their logarithmic norms and the norm of $\text{rad}(CD)$; let R be the active common radical width and E the remaining common row-mask divisor width. Let s_0 be the first full coprimality-projector width. Keep its signs and labels inside the Hilbert-space Cauchy sum. First restrict to supported first-allocation labels, then sum the complete second-allocation cost for each such label. The zero test applies to the original full Fourier row before its ordinary and fixed-inducing parts are separated. Fourier data uses the actual cardinality of the surviving slot set; an ambient cardinality bound is used only for common loss rates and outside allocation factors. With $B_c = (3c - 5d - R)_+/6$, the required C-side estimate is

$$\sum_{\lambda} \mathcal{N}_{C,\lambda} \ll Z^{p+A-c+R+E+B_c+\varepsilon}, \quad (2.13)$$

and similarly after exchanging c, d . The outside power is $M - A - R/2 - E$. Joint Cauchy yields $M + p - (c + d)/2 + R/2 + (B_c + B_d)/2 \leq M$, because the original primewise balance is $B_c + B_d \leq c + d - 2p - R$. Per-label branches include $-s_0$ and pay the count $Z^{p+s_0+o(1)}$ once. The high branch below proves the joint estimate directly.

Actual amplification and second transform. Put

$$K_0 = 2A - c - d + R + E - M, \quad K_{\text{main}} = K_0 + (d - c)_+ + 6a + \omega, \quad (2.14)$$

where $\omega > 0$ pays retained-frequency and support-doubling padding. This is the actual smooth ball supplied by the retained-radius theorem. An arbitrarily smaller declared radius for the old ball would not suffice. For an extracted error with widths (w, w_o) , retain its actual radius $K_0 + \max(d - c - 2w + w_o, 0) + 3a + \delta_1 + \text{reserve}$. Negative logarithmic lengths are still positive physical radii; emptiness is tested on supported norms.

The exact second transform and full residual coprimality projector give genuine common-part data $c_2, d_2, b_2 = (c_2 + d_2)/2$, radical width p_2 , gcd width g_2 , and unit-branch radical width t_2 . The nominal main child width is

$$M' = M - (c - d)_+ - 6a - \omega - g_2 + t_2, \quad (A - c - c_2 - M') - (A - M) \leq 6a + \omega, \quad (2.15)$$

and its inner allowance is $M' + b_2 - p_2 + B_c + a$, at least $M' + a$. The genuine first projector divides both second common parts; the residual second projector has no such divisor assertion. Allocate the latter with its full inclusion-exclusion signs and both extraction scalars. A selected plain quotient remains unrestricted; do not add a new puncture to it. The actual nonexceptional child is precisely the canonical natural state with its original fixed ideal and natural holes. Cauchy and splitting into positive rectangle energies occur before bin restriction. The four clipped rectangle pairs retain the original scalar inequality paying the second projector; Appendix A states it.

Low bins. Partition within each fixed parent translate and the fixed finite dictionary provided by signed reconstruction. In a low bin the buffered upper zero parameter is at most $q_0 + O(e_{\text{bin}})$. The actual whole-character prime bound charges this squared slot cost. Greedily delete a prefix until the old $\frac{9}{2}$ capacity holds or the slots are exhausted. Only after paying this pointwise deletion drop the bin predicate and apply the old eligible moment. Exhaustion calls the unrestricted no-slot theorem. For actual width difference $d_* \leq \delta_2$, which may be negative, and $0 \leq r \leq 1$,

$$d_* + r(F - d_*)_+ \leq \delta_2 + rF_+. \quad (2.16)$$

Take $r = 2q_0/9$ and $F \leq g + 6a + \text{small}$. The main low reserve is at least

$$a(1 - 4q_0/3) - (2q_0/9)g \geq M(7/75)\alpha = 7M/30000000. \quad (2.17)$$

The $6a$ defect is paid even if the original $g = 0$. One greedy overshoot, crossing-bin rounding, finite aggregate support error and radius padding fit smaller fractions of this reserve. No bin indicator is inserted into either Poisson formula, and no improved child moment is invoked.

High bins. High bins have lower zero parameter $\delta \geq q_0$, hence zero real part at least .84, within Lemma 2.1. The literal signed row, including its fixed unit, is

$$w = u_0 e F_c G_c V h', \quad W = M - R - (c - d)_+ + L_c - 6a, \quad L_c = \log_Z N F_c. \quad (2.18)$$

The two active residues are $1 + t_p$ and $1 - t_p$ modulo six; residue zero retains a hole. If R_0 is canceled radical width, then $L_c \leq 5(R - R_0)$. The first projector and all genuine second-common primes divide w . Missing first radical width is at most $p - R - E + R_0$. The signed arbitrary-valuation reconstruction is divisor-bounded at fixed w and missing radical. It supplies the natural conductor $\ll_{\text{fixed}} Nw$ and a fixed finite dictionary; it does not enlarge that dictionary over live labels.

Lemma 2.1 and its raw lift therefore give the joint count

$$Z^{\rho(\delta)W + p - R - E + R_0 + \varepsilon} (1 + T_h)^J. \quad (2.19)$$

A slightly negative nominal W is completed at the actual supported nonnegative width and charged to padding; a genuinely negative supported upper bound is empty. A nonempty bounded range is never assigned a count less than one. Whole-character conjugation reflects height; conjugating only the fixed factor is not this identity.

Combining the pointwise plain and slot bounds, both extraction scalars, and outside powers, the excess over (2.13) is

$$M(\rho - 1) + \delta A + a(5 - 6\rho) + (1 - \rho)(c - d)_+ - \delta c - (1 + \rho)R - E - B_c + \rho L_c + R_0 + g_2 - t_2 - (1 + \delta)b_2 + \text{small}. \quad (2.20)$$

On $[17/25, 3/4]$, $\rho + \delta - 1 \geq 1/28$, $5\rho - 1 \geq 3/7$, $4\rho - 1 \geq 1/7$, and $4\rho - 1 - \delta \leq -1166/6825$. Using $\min(c, d) \geq p \geq R$ and $L_c \leq 5(R - R_0)$, the middle line is at most $(4\rho - 1 - \delta)p - E - B_c \leq 0$. The final line is at most $-\delta b_2 - t_2 \leq 0$. The second projector has been paid by its scalars, not recounted as a divisor of w .

It remains to bound

$$H_j(\delta) = 1 - \rho(\delta) - \delta A_j - \bar{a}_j(5 - 6\rho(\delta)). \quad (2.21)$$

Its derivative increases with δ and is bounded below by $4397519373071/15608701631250 > 0$, using (2.11). Thus the minimum is at q_0 . The identity

$$q_0(\Lambda - 1) = K_a(5 - 6\rho_0)(9/2 - \Lambda) \quad (2.22)$$

shows that its coefficient of b_j vanishes. Consequently all 600 cells have exactly the same high reserve:

$$H_j(q_0) = \frac{81379}{15925}(B_0 - B) - \frac{9877}{4787}10^{-5} - \frac{249}{91}\alpha = \frac{720486667}{30493190000000} > 0. \quad (2.23)$$

Bins crossing q_0 were assigned to the low branch, with their rounding paid. Using only the older density range $\delta \geq .7$ would leave a gap here.

Extracted errors, zeros and fixed-inducing children. A literal amplifier prime has logarithmic norm $\lambda_p \geq a - \log_Z 2 \geq 3a/4$ eventually. Its three extracted alternatives are $(w, w_o) = (\lambda_p, \lambda_p), (6\lambda_p, 0), (7\lambda_p, \lambda_p)$, with the original squared scalars $q_p^{-1}, (1 - q_p^{-1})^2, q_p^{-1}$. They use the unamplified error radii and global slot price $\frac{3}{4}$; (2.16) with deletion rate $\frac{1}{6}$ leaves reserves $a/4 - g/6, 13a/4 - g/6, 4a - g/6$. The prime average stays normalized. No low-bin heredity is used.

The first error-allocation weight is kept undivided. Division by the extracted child volume occurs only in the normalized error coefficient. The retained-budget inequality has the form

$$\text{tail} + \frac{56}{\#\mathcal{P}}(c_{\text{amp}} \|\mathcal{R}_{\text{main}}\| + \|\mathcal{R}_{\text{error}}\|).$$

Here $\mathcal{P}$ is the literal amplifier-prime pool and the two $\mathcal{R}$'s are the original retained main and error sums. The factor $56/\#\mathcal{P}$ is paid once; it is not introduced again when the normalized costs are contracted back to retained rows.

The estimate holds for every original radical-capped witness v , with the same constant and threshold before that witness and the moving source data. It does not choose a replacement witness. In the source sum, each first-allocation factor B_p is included once. A zero first coefficient contributes zero; for nonzero coefficients, the original full-zero filter is applied before the ordinary/exceptional split, so cancellation in the full source is respected.

The first zero frequency retains the powerful-product bound. The artificial zero Gauss row is added only in the smooth ball, remains unamplified, and carries its $-s_0$ allowance. The genuine second zero has excess at most $A - M - \min(c, d) - B_c + 5a + \text{small}$.

Fixed-inducing children retain the exact predicate, forced residues, masks and positive radius factor $K_{\text{phys}}^{5/6}$. The actual main-radius substitutions give $F_1 \geq 2c/3 - 4a - \text{small}$ and $F_2 \geq 2b_2/3$. With $v = c + \min(c_2, d_2)$,

$$\max_{v \geq 0} \{A - 5M/6 - 2v/3 - (M/4 - v)_+\} = A - M. \quad (2.24)$$

Thus main and error exceptional excesses are at most $A - M + 4a, A - M + 9a$. Leading lattice terms cancel because the two rectangles retain exactly the same whole character, mask, independent profile heights and whole-product norm powers. A small-power error cannot be substituted for this calculation.

The resulting normalized terminal reserves are

Branch	Reserve
low	$\frac{7}{30000000}$
high	$\frac{720486667}{30493190000000}$
first extracted error	$\frac{530959781}{482529600000}$

Branch	Reserve
sixth extracted error	$\frac{11360006033}{482529600000}$
seventh extracted error	$\frac{3516816899}{120632400000}$
second zero	$\frac{8754683599}{120632400000}$
main fixed-inducing	$\frac{127361803}{1507905000}$
error fixed-inducing	$\frac{988504553}{40210800000}$

These are whole-cell bounds. The low and high identities are (2.17), (2.23). The three error expressions are affine in b_j , and the last three are $1 - A_j - k\bar{a}_j$ for $k = 5, 4, 9$, also affine. Their interval endpoints give the displayed minima. Crucially A_j and $\bar{a}_j$ belong to the same cell. The frozen 600-cell receipt remains available, but the proof needs only these affine endpoint calculations.

For completeness, the formal assembly makes the small losses explicit. Write $\varepsilon_a > 0$ for the loss allotted to the affine estimate and put

$$g = \min \left\{ \frac{\varepsilon_a}{128}, \frac{7}{3840000000}, \frac{1}{128} \right\}, \quad L_2 = 3(N + 2).$$

The width-one construction uses the fixed source cap 130, mask bound 28, and dictionary degree 31. The cap 130 is unrelated to the affine-cone constant B in (2.7). Its common scalar rate is

$$\delta_f = \min \left\{ \frac{g}{71N + 41}, \frac{g}{9N + 164}, \frac{g}{128}, \frac{g}{2048(N + 1)} \right\} > 0. \quad (2.25)$$

Here N is the ambient slot count, while every Fourier and physical datum for a subset T still uses $\#T$. The density coordinate δ in (2.19) is distinct from this small scalar rate. The coarser error-dictionary rate is $g/(8(N + 1))$; it is not substituted for δ_f .

In the normalized error sum, the projector and shell costs contribute $\delta_f + 3L_2\delta_f$, the box count contributes another δ_f , and the physical source sum contributes $5\delta_f - 7/600000000$. The finite prime-ratio payment adds one more δ_f . Thus the complete exponent is

$$8\delta_f + 3L_2\delta_f - \frac{7}{600000000} \leq -\frac{7}{1200000000}. \quad (2.26)$$

This includes the box count as well as the original undivided first error weight. The corresponding Lean inequality is `geometry_affine_error_final_source_exponent`.

The mesh and g are chosen before the slot dictionary. After it is fixed, choose each subset's high-density degree J_T , then one positive geometric cutoff τ_g small enough for every T , including $2\tau_g(J_T + 1) \leq g/8$. The low-family degree is chosen after τ_g . The detector Ψ is introduced after this common cutoff; its finite profile controls, polynomial degrees and constants may then depend on Ψ and T , but precede Z , all moving source data and v . Common Fourier majorants precede the labels, and row-dependent profile parameters act as contractions inside the joint norm. These choices retain all independent heights. The two side estimates (2.13), their joint Cauchy bound, the comparisons and terminal terms prove (2.8). $\square$

3. The physical moment, all shortenings and the marked detector

Fix the rational geometry

$$\begin{aligned} h &= \frac{3969737987}{5000000000}, & \ell &= \frac{11}{3} - 4\theta = \frac{125417}{750000}, \\ x &= 1 + \ell - h = \frac{5599126039}{15000000000}, & y &= h - 2\ell = \frac{6892533961}{15000000000}, \\ \kappa_* &= \frac{374861}{500000}. \end{aligned} \quad (3.1)$$

Write $X = Z^x, Y = Z^y$. We have $0 < \ell < x < y < 2x, y < 1/2 < h < 2y, 1/6 < \ell < 1/5, h > 4\ell$, and $x + y = 1 - \ell$. Put $d_0 = 1/200$, $d_- = h - 1/1000$, $d_+ = h + 10^{-10} = 317579039/400000000$, and $t_0 = 1/1000$. The original positive slot widths ℓ_i will have sum ℓ .

Physical source and capacity. At row norm $U = Z^d$, an Euler free shift $h_0 = U^t$ gives a reflected plain scale at most U^n , where $n = 1 - y/d + t$. On $d \in [d_-, d_+]$, $0 \leq t \leq t_0$, use

$$z_A = \frac{B - 2n}{\Lambda}, \quad dz_A(0) = \frac{2y - (2 - B)d}{\Lambda}. \quad (3.2)$$

The intercept $1 - B$ in this formula is essential. The literal pair is the natural extended sextic source with SourcePrincipalS, whole-character conjugation, common profile $\overline{F.W}$, slot real power $\frac{33}{50}$, the required fixed inverse relative character, and independent heights $-\Im z_i$. The four derivative moments use this same family and slot subcollection. It is the `finite_physical_pair_positiveAt` contract, not a freely selected character family. Fixed-inducing free rows have bounded norm and leave growing dyads $d \geq d_0$.

Direct rational extrema give

$$\begin{aligned} \frac{41}{100} &< \frac{5001680000}{11894213961} \leq n \leq \frac{402287137237}{952737117000} < N := \frac{264}{625}, \\ \frac{17}{500} &< \frac{65807969436717121}{1839925920350400000} \leq z_A \leq \frac{841262571937768693}{22970106001483200000} < \frac{37}{1000} < \frac{1}{25}. \end{aligned} \quad (3.3)$$

The gaps at the upper bounds are respectively $\frac{745104919}{4763685585000}$ and $\frac{8631350117109707}{22970106001483200000}$. Choose scale boxes, annular padding and a fixed capacity decrement below these gaps and the later row saving. A Z -threshold of at least U_0^{200} covers all $d \geq d_0$; no limit $d \downarrow 0$ is used.

For an amplitude class, let $w_i = \ell_i/d$, $L = \ell/d$. A positive main slot has gain $g_i \in [0, \delta/2]$; an error or small-amplitude slot has gain zero. Define $q = \sum_i w_i g_i / L = \delta p$, $0 \leq p \leq 1/2$, keeping **all** widths in the denominator. The original whole-slot selector, with $w_i \leq \text{mesh}$, gives a single positive-main subcollection I such that

$$z_I \leq z_A, \quad |Q_{u,I}|^2 \geq U^{2qz_A - \delta \text{ mesh} - \text{small}}. \quad (3.4)$$

If the positive supply is exhausted, its total gain is $qL \geq qz_A$. Thus no lower bound on positive-main width and no zero-gain artificial factor is needed. This includes $q = 0$.

Every actual shortening has a paid moment:

- If $z_I \leq 17/500$, the old budget is $2N + (9/2)(17/500) = 1 - 11/5000$.
- If either plain width is at most $\frac{41}{100}$, its budget is $41/100 + N + (9/2)(37/1000) = 1 - 11/10000$.
- Otherwise both plain widths are in $[41/100, N]$, and $n_1 + n_2 + \Lambda z_I \leq 2n + \Lambda z_A = B$; apply Theorem 2.2.
- If I is empty, use the unrestricted old no-slot moment. It has no additional plain-sum gate.

In each case the selector still gives the target gain qz_A , even if z_I is shorter. The same I and atlas index work for the four derivative pairs. A fixed parent modulus gives effective width $1 + \nu$. Normalize before choosing the atlas: its two lower cutoffs are now $(17/500)(1 + \nu)$, $(41/100)(1 + \nu)$. The old complements lose at most $153\nu/1000$, $41\nu/100$; the shaped inequality relaxes to $B(1 + \nu)$. Choose ν below the common reserves before the target-specific threshold. Endpoint overlap is harmless.

Supply is verified without pruning the 600-cell library:

$$\begin{aligned} \ell - d_+ z_{\max} &= \frac{1269281592530944914352773973}{9188042400593280000000000000} > 0, & \frac{\ell}{d_+} - \frac{4}{21} &= \frac{44784748}{2223053273} > 0, \\ 2\ell - \frac{1}{3} &= \frac{139}{125000} > 0, & 4\ell - y &= \frac{1046942013}{5000000000} > 0. \end{aligned} \quad (3.5)$$

For $d < 1/2$, the original selector has $L \geq 2\ell > 1/3 > 1/(6\kappa)$. A convenience wrapper restricted to $d \geq 1/2$ is not invoked there.

Euler shifts. The exact physical Euler expansion has coefficients $c_u(h_0, h_r)$, support $h_0 h_r \ll Y$, and

$$\sum_{h_0, h_r} |c_u(h_0, h_r)| h_0^{\delta/2} \ll U^\varepsilon. \quad (3.6)$$

Its primitive conductor damping and repeated-power deficit are charged together once. Reflection has scale $\ll (U/Y)h_0$, and natural Euler restoration only shortens it. In particular $h_0 \ll Y$; neither squarefree row restriction nor an unweighted coefficient bound replaces (3.6).

For $t \leq t_0$, replacing $z_A(0)$ by $z_A(t)$ loses numerator gain tq/Λ , paid by $\delta t/2$, since $q \leq \delta/2$ and $\Lambda > 1$. For $t \geq t_0$, put $b = 2y/d - 1$ and use old capacity

$$z_{\text{old}}(t) = \max\{0, (b - 2t)/(9/2)\}, \quad S_b(t) = \frac{t}{2} - \frac{1}{4} \left(\frac{b - (1 - B)}{\Lambda} - z_{\text{old}}(t) \right). \quad (3.7)$$

A negative parenthesis is already favorable; otherwise (3.6) pays whenever $S_b \geq 0$. This continuous piecewise affine function has slopes $7/18, 1/2$. Here $b \leq 1890853961/11894213961 < 4/25$, and

$$S_b(t_0) \geq \frac{4035883}{69523200000} > 0. \quad (3.8)$$

Thus every larger shift is paid. At clipped capacity zero the actual plain lengths are at most $U^{1+o(1)}$, within the old unrestricted bounded-length no-slot contract. For $d < d_-$ retain the original old capacity and the same weighted shift argument.

The full physical detector. The high-bin split $\frac{17}{25}$ in Module 2 belongs to transformed rows. The physical detector still uses the entire rectangle $0 \leq \delta \leq 3/4$, $0 \leq p \leq 1/2$. Define, for the actual κ ,

$$C_\kappa(p) = \frac{2(1-p)(1-p/(3\kappa))}{3-p-2p/(3\kappa)}, \quad V = \frac{5}{6} - \delta, \quad R_\kappa = 1 - \delta + \frac{\delta C_\kappa V}{2(V + \delta C_\kappa)}. \quad (3.9)$$

To obtain its count, a saturated witness of cutoff $t \in [1, 3/2]$ has inverse and plain exponents r, m with $r + m \geq t$ up to the allocated losses. The short plain count decreases as m increases, so its worst admissible envelope substitutes $m = t - r$. Old slot capacities $(1 - r)/2$ and $(1 - 2m)/(6\kappa)$ in the short range give the two count exponents

$$A_I = 1 - q - (\delta - q)r, \quad A_P = 1 - \frac{q}{3\kappa} - 2 \left(\delta - \frac{q}{3\kappa} \right) m. \quad (3.10)$$

Their crossing is

$$r_*(t) = \frac{2(1 - p/(3\kappa))t + p/(3\kappa) - p}{3 - p - 2p/(3\kappa)}. \quad (3.11)$$

For the used range, $r_* \geq 13/21$ and $1/3 \leq t - r_* \leq 1/2$. At this crossing the short envelope is $1 - \delta + \delta C_\kappa(3/2 - t)$. The long inverse envelope from (1.4) is $1 - \delta + V(t - 1)$. Balancing them chooses

$$t = 1 + \frac{\delta C_\kappa}{2(V + \delta C_\kappa)}$$

and gives (3.9). The selected inverse capacity is at most $\frac{4}{21}$ and the plain capacity at most $\frac{1}{12}$, so (3.5) supplies them. A decrement $\nu_0 > 0$ of inverse capacity pays the two strict MI margins $2\nu_0$ and $8\nu_0 + (2r - 1) \geq 8\nu_0 + 5/21$. If a nominal capacity is smaller than the decrement, use the unmarked bound with an $O(\nu_0)$ loss; treat the plain decrement in the same way. The whole-slot selector, supported witness, common scale derivatives, frequency discretization and fixed finite dictionary therefore give $U^{R_\kappa + \varepsilon}(1 + T)^J$ rows. Floors are handled separately.

For clarity, the numerator for one fixed main/error term of the exact replacement (5.8) is the fully inverted physical object

$$\mathcal{N}_u(Y; s, z) = \frac{1}{2\pi i} \int_{(3)} Y^{w-1} \widehat{W}_1(w) L_K^S(w, \psi_u) \mathcal{H}_{\eta, u, Z}^{\text{term}}(s, w, z) dw. \quad (3.12)$$

The reciprocal $L_K^S(s, \overline{\eta\psi_u})^{-1}$ is independent of w and is bounded only after this inversion. On the residue-free detector contour $\Re s = (1 + \delta)/2 + \text{buffer}$, $\Re z = 17/50$, Hölder with the paired moment and row count gives

$$\sum_{\text{class}} |\mathcal{N}_u| \ll Y^{-1/2} Z^{\ell(17/50 - 1/2) + \ell q} U^{1/4 + 3R_\kappa/4 - qz_A/2 + \varepsilon} \mathcal{H}_J. \quad (3.13)$$

The exponent $-qz_A/2$ is the fourth-root contribution of (3.4). Expanding the physical weights in (3.13), the row contribution relative to the probe normalization has exponent $G_\kappa - k_{\text{probe}}$, where

$$G_\kappa(d, \delta, p) = \frac{1 + \delta}{2} + \frac{13h}{75} - \frac{y}{2} - \frac{\ell}{2} + \ell\delta p + d \left(\frac{1}{4} + \frac{3}{4}R_\kappa - \frac{17}{50} \right) - \frac{\delta p}{2} C_d. \quad (3.14)$$

Here $C_d = [2y - (2 - B)d]/\Lambda$ for the new rows and $C_d = \min(d, 2y - d)/(6\kappa)$ for old rows. This explicitly retains the fixed affine deficit, both plain factors, and the original slot gain.

4. Exact endpoint inequalities and supremum feedback

First set $\kappa = \kappa_* = 2\theta - 1$. Put

$$D = 3 - p - \frac{2p}{3\kappa_*}, \quad P = 2(1 - p) \left(1 - \frac{p}{3\kappa_*} \right), \quad J = VD + \delta P. \quad (4.1)$$

On the full detector rectangle, $V \geq 1/12$, $P \geq 0$, and $D \geq D(1/2) > 0$. Hence $J \geq D(1/2)/12 = 4622915/26989992 > 1/10$, and $JR = (1 - \delta)J + \delta PV/2$. For a proposed saving v , the polynomial

$$F_{d,v} = J \left[\theta - v - \frac{1 + \delta}{2} - \frac{13h}{75} + \frac{y}{2} + \frac{\ell}{2} - \ell\delta p - d \left(\frac{1}{4} - \frac{17}{50} \right) + \frac{\delta p C_d}{2} \right] - \frac{3d}{4} \left[(1 - \delta)J + \frac{\delta PV}{2} \right] = J(\theta - v - G_{\kappa_*}) \quad (4.2)$$

has degree at most $(2, 3)$ in (δ, p) .

Here is a finite exact proof on closed rectangles. On a box $[a, b] \times [c, d']$, expand $F(a + (b - a)u, c + (d' - c)v') = \sum f_{rs} u^r (v')^s$. Its tensor Bernstein coefficients of degree $(2, 3)$ are

$$\beta_{ij} = \sum_{\substack{r \leq i \\ s \leq j}} f_{rs} \frac{\binom{i}{r} \binom{j}{s}}{\binom{2}{r} \binom{3}{s}}. \quad (4.3)$$

The corresponding basis functions are nonnegative and sum to one. Positive coefficients consequently prove positivity everywhere, including the boundary, rather than just at selected evaluation points.

For the new row $d = h, v = 10^{-6}$ and old row $d = d_-, v = 10^{-5}$, use the same six boxes:

$$\begin{aligned} [0, 3/8] \times [0, 1/2], \quad [3/8, 3/4] \times [0, 1/4], \quad [3/8, 9/16] \times [1/4, 3/8], \\ [3/8, 15/32] \times [3/8, 1/2], \quad [15/32, 9/16] \times [3/8, 1/2], \quad [9/16, 3/4] \times [1/4, 1/2]. \end{aligned} \quad (4.4)$$

They partition the rectangle with disjoint interiors. For the old row $d = 1/2, v = 10^{-3}$, one whole rectangle suffices. Formula (4.3), applied to (4.2), gives the following minima among all 12 coefficients per box:

Row	Boxes	Minimum coefficient
new $h, 10^{-6}$	6	$\frac{113219985368027151827923}{800610394374144000000000000000}$
old $d_-, 10^{-5}$	6	$\frac{8604603547887123299947}{388511823018700800000000000000}$
old $1/2, 10^{-3}$	1	$\frac{130800958135492294036529}{2529373847778000000000000000}$

All 156 coefficients, the input power coefficients, and exact box bounds are retained in the pinned endpoint receipt. These values are exact rationals and are independently reconstructed by the bounded check accompanying this manuscript.

The old exponent is affine for $d \in [1/2, d_-]$. At the junction only,

$$C_{d_-}^{\text{new}} - C_{d_-}^{\text{old}} = \frac{33165294566691730080019}{3257692034400000000000000000} > 0. \quad (4.5)$$

Thus the old junction bound also bounds the new exponent there. The new exponent is affine on $[d_-, h]$, so its two endpoints suffice. The capacity difference has negative slope; no upward monotonicity is claimed. Since $R \leq 1$, the new exponent has d -slope at most $33/50 + (3/8)(2 - B)/(2\Lambda) < 2$. Extending to d_+ leaves saving $\frac{4999}{50000000000}$.

Below $\frac{1}{2}$, the old expression is convex in d , because $\min(d, 2y - d)$ is concave and is multiplied by $-\delta p/(12\kappa_*)$. At the lower endpoint its elementary bound has reserve

$$\theta - \left(\frac{7}{8} - \frac{49h}{150} + \frac{7\ell}{8} + \frac{33d_0}{50} \right) = \frac{82198036363}{750000000000} > \frac{1}{10}. \quad (4.6)$$

This and the half-row certificate cover $[d_0, 1/2]$. Floor classes have their separate elementary reserve

$$\theta - \left(\frac{51}{100} + \frac{h}{3} + \frac{51\ell}{100} \right) = \frac{223923613}{15000000000} > \frac{1}{100}, \quad (4.7)$$

with room for the d_+ extension. They do not require a nonfloor witness.

Finally return to the actual $\kappa = 2\beta_* - 1$. On $7/10 \leq \kappa \leq 3/4$,

$$0 \leq \partial_\kappa C_\kappa = \frac{6p(1-p)^2}{(9\kappa - 3\kappa p - 2p)^2} \leq \frac{1}{18}, \quad 0 \leq \partial_\kappa R = \frac{\delta V^2}{2(V + \delta C)^2} \leq \frac{5}{12}. \quad (4.8)$$

For the first bound the denominator before squaring is at least $\frac{17}{4}$ and $6p(1-p)^2 \leq 8/9$, so it is at most $128/2601 < 1/18$. The second follows already from $\delta/2 \leq 3/8 < 5/12$. It follows that

$$\partial_\kappa G_{\text{new}} \leq \frac{5d_+}{288} = \frac{317579039}{23040000000} < \frac{1}{50}. \quad (4.9)$$

For the old exponent retain also the derivative of $1/(6\kappa)$. Using $\delta p \leq 3/8$, $\min(d, 2y - d) \leq 1/2$, $d \leq 1$, gives

$$\partial_\kappa G_{\text{old}} \leq \frac{695}{14112} < \frac{1}{8}. \quad (4.10)$$

Since $\kappa - \kappa_* = 2\mu$, the exponents rise by at most $\mu/25$ and $\mu/4$, whereas β_* rises by μ . Every strict row gap thus persists at the actual supremum parameter. The old capacity decreases with κ , so the junction comparison persists as well. No numerical theorem at a previous improved boundary is a premise of this argument.

5. The physical probe, the common gap and ordinary L -functions

Definition and physical bound. For a primitive target η , let b_* be the product of fixed excluded prime generators. Choose the primitive residue character $\xi \bmod b_*$ locally quadratic at odd primes and cubic at norm four, so ξ^6 is principal on units. Set $\tau = q_{b_*}^{-1/2} g_\xi(b_*, 1)$ and $\Xi(A) = \xi(A) \chi_A(b_*)$. The completed theta coefficients $\gamma_2(c), G(A), R(A, a)$ are exactly the normalized coefficients in Appendix A; their indices, phases and zeros are not free parameters. For nonnegative nonzero smooth annular W_0, W_1 , define

$$T_{m,a,\eta}(t) = \sum_{c \text{ sf}, n} \gamma_2(c) \overline{\alpha(A) G(A)} \eta(A) \Xi(A)^{-1} R(A, a) \chi_A(m) q_c^{-1/2-t} q_n^{-1-3t}, \quad A = cn^3, \quad (5.1)$$

$$I_\eta(X, Y, Z) = \frac{1}{Y} \sum_a \frac{W_1(q_a/Y) \chi_a(b_*)}{\tau \xi(a) \sqrt{q_{b_*} q_a X}} \sum_m \xi(m) q_a^{-1/2} g_{\chi_a}(a, -m) W_0\left(\frac{q_m}{q_{b_*} q_a X}\right) \cdot \frac{1}{2\pi i} \int_{(4)} Z^t e^{t^2} T_{m,a,\eta}(t) dt. \quad (5.2)$$

Thus the logarithmic Gaussian is $V(r) = e^{-(\log r)^2/4}/(2\sqrt{\pi})$ with transform e^{t^2} .

Let $I_{\eta;D}$ insert the mark $D \mid cn^3$, including collisions between D, c and n . Choose disjoint canonical good-prime slots $P_i = Z^{\ell_i}$, with the profiles w_i obtained from one common annular profile on those slots as required in §1. The exact compensator is

$$I_\eta^{\text{comp}}(Z) = \sum_{(p_i)} \prod_i w_i(q_{p_i}/P_i) \sum_{J \subseteq \{1, \dots, r\}} (-1)^{|J|} q_{p_J}^{-3/2} \overline{\eta(p_{J^c})} I_{\eta; p_{J^c}}(X/q_{p_J}, Y/q_{p_J}, Z q_{p_{J^c}}). \quad (5.3)$$

All good primes are retained. In particular the coefficient contains $\bar{\eta}$, not η .

The physical signed Gram bound concerns the literal coefficient

$$A_{m,\sigma,v} = Y^{-3/2} \sum_{a \in \sigma} \frac{\chi_a(b_*)}{\tau \xi(a)} W_1(q_a/Y) (q_a/Y)^{-1+iv} g_{\chi_a}(a, -m).$$

For $Q = q_{b_*} XY$, $P_a = Y^2/Q \geq 1$, it is

$$\sum_{q_m \ll Q} |A_{m,\sigma,v}|^2 \ll (1 + |v|)^J \frac{Q}{Y} (1 + P_a^{1/6} + P_a^2/Y) Z^\varepsilon. \quad (5.4)$$

Indeed full row Poisson leaves prefactor Q/Y^3 , diagonal $\mathbf{1}_{u=v} \phi(u)$, and the complete local projector before column separation. The complete column mean vanishes unless every good valuation outside the shared gcd is divisible by six. Mean-zero column Poisson contributes $(Q/Y) P_a^2/Y$; the retained sixth-power exception contributes $(Q/Y) P_a^{1/6}$. This proof uses the displayed signed coefficients and does not extend to arbitrary column weights.

Combine (5.4) with the completed marked-energy estimate in Module 1. For $M = x + y$, its exact physical consequence is

$$|I_\eta^{\text{comp}}| \ll Z^{x/2+(y-x)/12+(1+3\ell-2M)_+/8+\varepsilon}, \quad (5.5)$$

provided $0 < \ell < x \leq y$, $M + \ell \leq 1$, and $y - 2x + \ell < 0$. A rescaled subset of width d has $M' = M - 2d$, $z' = \ell - d$, $N' = 1 + \ell - d$. Combining its outside coefficient and tuple count adds

$-d + (1 + 3\ell - 2M + d)_+/8$, a decreasing function of d . This explains why no stronger, separate energy bound on each rescaled subset is required.

For (3.1),

$$y - 2x + \ell = -\frac{599126039}{5000000000}, \quad 1 + 3\ell - 2(x + y) = -\frac{24583}{150000}. \quad (5.6)$$

Set $k = k_{\text{probe}} = 5/6 - x/3 - \ell/6$. Then $(5x + y)/12 = \theta - k$.

Exact principal signal. Row Poisson and $s = t + 1 - z$ give the absolutely convergent identity

$$I_\eta = \frac{1}{(2\pi i)^3} \int_{(2)} \int_{(3)} \int_{(3)} X^{1/2-z} Z^{s+z-1} Y^{w-1} e^{(s+z-1)^2} \mathcal{M}(z) \widehat{W}_1(w) \cdot \sum_u^{(6)} q_u^{-z} \overline{\xi(u)} \frac{\zeta_K^S(6z) L_K^S(w, \psi_u)}{L_K^S(s, \eta \overline{\psi_u})} H_{\eta, u}(s, w, z) dz dw ds. \quad (5.7)$$

The sum is over sixth-power-free nonzero element rows, including all units. The excluded fixed primitive character removes the zero Fourier row, and factoring a frequency as ua^6 supplies $\zeta_K^S(6z)$. The complete six-row local table in Appendix A defines $H_{\eta, u}$ and the quotient-free replacement G_p . Compensation replaces H by

$$\mathcal{H}_{\eta, u, Z} = \sum_{(p_i)} \prod_i w_i(Q_i/P_i) Q_i^{z-1} G_{p_i} \prod_{p \notin S \cup \{p_i\}} H_p, \quad G_p = -\overline{\psi_u(p)} H_p + \Delta_p. \quad (5.8)$$

This equality holds also at row primes. It never divides by H_p or discards a collision.

The products converge normally for $\Re s \geq .51, \Re z \geq .34, \Re w \geq -.01, \Re(s + w) \geq 1 + \epsilon_1$. The principal box is $\Re s \geq .87, \Re w \geq 3/4, \Re z \geq 1/8$. The unseparated correction has a fixed raw polynomial bound on these boxes. Its special bounds on $(\beta_* + e, 1/2, .34)$ and $(2, 2, v)$ are $q_u^\epsilon Z^{.34\ell}$ and $q_u^\epsilon Z^{v\ell}$, respectively.

Only $u = 1$ has principal numerator. A nonunit sixth-power-free row has nontrivial tame inertia; for a unit $u \neq 1$, the primary CRT class $n \equiv -2 - 3\omega \pmod{36\mathcal{O}}$ has $q_n \equiv 7 \pmod{36}$, hence $\chi_n(u) = u \neq 1$. At $w = 1, z = 1/6$, write $t = Q^{-1}$, $D = \eta(p)Q^{-s}$, $R = \overline{\alpha(p)}^6 \eta(p)^6 Q^{3-6s}$. The principal factor is

$$H_{\eta, p}(s) = \frac{1-t^2}{1-R} \left(1 + \frac{t(D-R)}{1-D} \right). \quad (5.9)$$

Enlarge the target's fixed cutoff so that $|H_\eta(s) - 1| \leq 1/2$ for $\Re s \geq .87$. On this residue locus only, $G_p/H_p = -1 + O(Q^{-.87})$. Define

$$c_S = \frac{(\text{Res}_{s=1} \zeta_K^S(s))^2}{6} \widehat{W}_1(1) \mathcal{M}(1/6) > 0, \quad A(Z) = (-1)^r Z^{-\ell/6} \prod_i \sum_p w_i(Q/P_i) Q^{-5/6}. \quad (5.10)$$

The radial normalization is exact:

$$\mathcal{M}(\sigma) = (2\pi/\sqrt{3})^{1-2\sigma} \frac{\Gamma(\sigma)}{\Gamma(1-\sigma)} \int_0^\infty W_0(r) r^{-\sigma} dr > 0 \quad (0 < \sigma < 1). \quad (5.11)$$

Fixed-field prime ideal counting gives $|A(Z)| \asymp (\log Z)^{-r}$ and, for fixed $0 < \mu_P < .87 \min_i \ell_i$,

$$\mathcal{H}_{\eta, 1, Z}(s, 1, 1/6) = H_\eta(s) Z^{\ell/6} A(Z) (1 + O(Z^{-\mu_P})). \quad (5.12)$$

Both residue scalars are present in (5.10). Put

$$J_\eta(Z) = \frac{I_\eta^{\text{comp}}(Z)}{c_S A(Z)}, \quad f_\eta(Z) = \frac{1}{2\pi i} \int_{(2)} Z^{s-k} e^{(s-5/6)^2} \frac{H_\eta(s)}{L_K^S(s, \eta)} ds. \quad (5.13)$$

Equations (5.5)–(5.6) imply $J_\eta(Z) \ll_\eta Z^{\theta-k+\epsilon}$.

All complementary terms. We prove a target-independent power improvement between the two functions in (5.13). For $u = 1$, move first to $(\Re s, \Re w, \Re z) = (\beta_* + e, 1 + e, 1/6 + e)$. Move w to $\frac{3}{4}$, then z in its residue to $\frac{1}{8}$. The only numerator poles crossed are $w = 1, z = 1/6$. The double residue is f_η plus the error from (5.12). Use the quotient-free correction on the full principal box for the other two remainders; do not use the residue approximation there. Their relative savings below $\beta_* - k$ are $y/4, h/24, \mu_P$. For example the smaller source contour checks have reserves

$$\frac{y}{20} - \frac{1+h}{1000} - \frac{1}{4000} = \frac{313967484089}{15000000000000} > 0, \quad \frac{h}{600} - \frac{1}{1000} - \frac{1}{4000} = \frac{219737987}{3000000000000} > 0. \quad (5.14)$$

For retained nonprincipal rows $Z^{d_0} \leq q_u \leq Z^{d_+}$, first move z to $\frac{17}{50}$ with $s = 2, w = 3$, then s to its global zero-free line, then the buffered retained segment to $\Re s = (1 + \delta)/2 + \text{buffer}$. Outside the retained heights keep the reciprocal on the global line. Perform the complete w inversion at $\Re w = 3$. Equations (3.13)–(3.14) and Module 4 pay every nonfloor amplitude class; (4.7) pays floors. No reciprocal zero or nonprincipal numerator pole is crossed.

For $q_u \leq Z^{d_0}$, including the finitely many nonprincipal units, use $(\beta_* + e, 1/2, 17/50)$. The ordinary functional equation and convexity give the sufficient numerator bound $L_K^S(1/2 + it, \psi_u) \ll U^{3/5+\varepsilon}(1 + |t|)^A$. The raw correction bound, ideal counting and $q_u^{-17/50}$ leave relative exponent

$$\frac{13h}{75} - \frac{y}{2} + \frac{63d_0}{50} = -\frac{21458387121}{2500000000000} < -\frac{1}{20}. \quad (5.15)$$

These rows need no detector and no artificial nonprincipal assumption on the genuine principal row, which was already separated.

For $q_u > Z^{h+\zeta}$, $\zeta = 10^{-10}$, keep the original absolute integration order and take $(\Re s, \Re w, \Re z) = (2, 2, v)$, $v > 2$. With $B_{\text{raw}} = x/2 + 1 + y$, the raw dyad is at most

$$Z^{B_{\text{raw}} + hv + \varepsilon} U^{1 + \varepsilon - v},$$

so its sum is bounded by

$$Z^{B_{\text{raw}} + (h+\zeta)(1+\varepsilon) - \zeta v + \varepsilon}. \quad (5.16)$$

A sufficiently large fixed v , chosen after ζ , pays any required negative power. This is an absolute tail argument, not an extension of the central row estimate.

The independent frequencies are $\Im(s + z), \Im z, \Im w$. The Gaussian controls the first; the two Mellin transforms have arbitrary polynomial decay in the others on the fixed real strips. Every joining segment keeps its reciprocal inside the buffer or on a global zero-free line. The unseparated correction has raw bound $Z^{B'} U^\varepsilon$ with B' fixed before the final tail order. Consequently discarded heights give $O_N(Z^{B''} T^{-N})$. The reflected profiles are common: near zero their Mellin shift crosses only the simple gamma pole and gives $O(1)$, so their annular decomposition has summable normalized coefficients. Complete that decomposition and its scale supremum before the moment. These steps do not differentiate arithmetic indicators or equate independent heights.

Order of choices. First fix the contradiction gap μ , the geometry and strict rational reserves. Choose the moment loss, capacity decrement, amplitude mesh and parent-width tolerance from these reserves. The old and new moment meshes are uniform in the later fixed arithmetic family. Choose a finite positive slot system of total width ℓ , sufficiently fine on all $d \geq d_0$, and freeze it. Its smallest width fixes $\mu_P > 0$. Choose real buffers and one $\Delta > 0$, smaller than fixed fractions of every source and row reserve, $y/4, h/24, \mu_P$, and $\mu/4$.

Only now fix the primitive target, its S , finite dictionaries and profiles. Within the local moment construction, follow the order after (2.26): high-density degrees, common geometric cutoff τ_g , low-family degrees, then the detector and its profile controls. The internal density exponent of Lemma 2.1 is chosen before its returned polynomial height degree.

For the final Mellin truncation choose a separate positive external exponent τ , small enough that every resulting factor $(1 + Z^\tau)^J$ consumes less than the reserved common gap and the smallest-slot and d_0 restrictions hold. This final tail cutoff is distinct from τ_g . Choose its tail order

afterward, then take the finite maximum of target-dependent thresholds. Constants and thresholds can depend on the target, but the frozen slots and Δ cannot. Thus

$$J_\eta(Z) - f_\eta(Z) \ll_\eta Z^{\beta_* - k - \Delta}. \quad (5.17)$$

Reduce Δ if needed while retaining $\Delta < \mu/4$, and take the loss in the physical probe bound below $\mu/4$. We obtain

$$f_\eta(Z) \ll_\eta Z^{\beta_* - k - \Delta} \quad (Z \rightarrow \infty). \quad (5.18)$$

For $0 < Z \leq 1$, move the defining line in (5.13) to any fixed $A > 2$. Absolute Euler convergence and the Gaussian give $f_\eta(Z) = O_{\eta,A}(Z^{A-k})$.

Continuation and transfer. The ordinary Mellin integral

$$\mathfrak{F}_\eta(s) = \int_0^\infty f_\eta(Z) Z^{-s+k} \frac{dZ}{Z} \quad (5.19)$$

is therefore holomorphic on $\Re s > \beta_* - \Delta$. On $\Re s > 1$, Fourier inversion in $\log Z$, with both sides integrable on a fixed vertical line, identifies it exactly as

$$\mathfrak{F}_\eta(s) = e^{(s-5/6)^2} \frac{H_\eta(s)}{L_K^S(s, \eta)}. \quad (5.20)$$

Since $\beta_* - \Delta > \theta > .87$, the numerator never vanishes. If $\beta_* > \theta$, the definition of supremum supplies a primitive zero with real part greater than $\beta_* - \Delta/2$. Choose its target only after the common Δ , then invoke that target's allowed constants and threshold. The right side of (5.20) has a pole at that zero, contrary to (5.19). No supremum attainment or conductor-uniform constant was assumed. Hence $\beta_* \leq \theta$.

For an imprimitive Hecke character induced from η^* ,

$$L_m(s, \eta) = L(s, \eta^*) \prod_{\mathfrak{p} | m, \mathfrak{p} \nmid \eta^*} (1 - \eta^*(\mathfrak{p})(N\mathfrak{p})^{-s}). \quad (5.21)$$

Every displayed Euler factor is nonzero for $\Re s > 0$. Classical Hecke nonvanishing on $\Re s = 1$, and the Euler product above it, complete the Hecke assertion with its principal-pole exception.

For any ordinary complex Dirichlet character $\chi \bmod q$, $q \neq 0$, use the finite-order trivial-infinity norm lift η_χ . With the original imprimitive and ramified factors retained, its exact identity away from $s = 0, 1$ is

$$L_K(s, \eta_\chi) = L(s, \chi) L(s, \chi\chi_{-3}). \quad (5.22)$$

On $\theta < \Re s < 1$ both factors are finite and the product is nonzero, so $L(s, \chi) \neq 0$. On $\Re s \geq 1$ apply classical ordinary Dirichlet nonvanishing directly, including nonprincipal χ at $s = 1$. The sole exception is $\chi = 1, s = 1$. For modulus one the function is ζ . This proves (T), using the pinned analytic results specified in the appendix. $\square$

Appendix A. Analytic dependencies, source contracts and attribution

This appendix specifies the analytic dependencies of the proof at $c = 125139/1000000, \theta = 874861/1000000$. The imported theorems fix the families, hypotheses, profiles, marks, conductor conventions and parameter ranges. Their hypotheses are discharged in the final assembly; no moment or energy estimate remains as a premise of the endpoint. The prose describes the used results rather than reproducing each upstream proof. Exact source and verification records accompany the formalization.

A.1. Exact base and definition contacts

Use PR9 head `58344dfdbe756cf2f743da1908ffe0582418cf5b`, its manuscript [the pinned manuscript](#), and source archive [source-public.tar.gz](#). The manuscript SHA256 is `2b4c67c3e817d7854ee57e451cea1789c8a4a1d03a72ddb83b980421167096d8`; Git blob `ae5d0488c155d35f10b39038f7b9b9f2d5e00071`. The archive SHA256 recorded by the intake is `23171d67353bc555b48be2e4a77ff7cbcd559489fdf03ba665b9dae637ea1ef5`. The archive pin records inherited provenance; it does not establish a fresh full source replay.

Use Gao PR15 head `ca182b70ac2aee2d053d679b4ddf1233883f9eef`, directory [proofs/gao-local-moment-20261011-v3](#). Selected written files are bound by exact byte hashes and Git blobs in the accompanying source provenance. No mutable branch tip is an input.

The mathematical dependency graph is:

$$\begin{aligned}
& \text{arithmetic} + \text{completed reflection} \\
& \longrightarrow \text{UM, MI, old MP/MP0, signed physical Gram,} \\
& \quad \text{UM} + \text{all-free no-slot plain} + \text{actual witness} \\
& \quad \longrightarrow D^+ \\
& \quad \longrightarrow \text{local paired plain moment,} \\
& \text{local moment} + \text{literal physical bridge} + \text{weighted shifts} \\
& \quad \longrightarrow \text{all-row rational inequalities} \\
& \quad \longrightarrow \text{one common-gap Mellin probe} \\
& \longrightarrow \text{Hecke zero exclusion} \longrightarrow \text{ordinary Dirichlet zero exclusion.}
\end{aligned}$$

The Lean assembly implements this graph through the original retained source sums. `GeometryAffineMainUniform` combines the complete low and high families; `GeometrySourceMainPayment` specializes them to the weighted source. `GeometryAffineErrorFamily` combines the actual per-source menu, joint normalized transport and scalar aggregate; `GeometryAffineErrorCompletePayment` and `GeometrySourceErrorPayment` return the original retained error estimate. `GeometryAffinePaidComponents` and `GeometrySourceBudget` combine these estimates, and `GeometryChosenSetup` fixes the mesh before the chosen source data. Finally `Cycle61.FinalAssembly.chosen_affine_input` supplies the affine input and `beta_1e_rational` proves the exact supremum bound.

The final weighted source uses the same fixed principal source character for its seed and anchor, $\eta_{\text{seed}} = \eta_0$, and the internal ideal $Q_* = \text{internalQ}(Q, \eta_0)$. These source choices are distinct from the target character of the zero-exclusion theorem. A radial support bound is replaced by $b_\Phi = \max(2, b_\Phi^{\text{initial}})$ before the common geometric cutoff and detector are chosen; this bound is separate from the source cap $M_{\text{cap}} = 2$.

The public comparison declarations are `Cycle61Verification.allDirichlet`, `zeta` and `allHecke`. They retain the strict half-plane and principal-pole exceptions. `zetaTotal` separately records Mathlib’s total-function convention at the pole; the ordinary analytic theorem above remains an off-pole assertion.

The following are the imports, rather than numerical ancestor conclusions.

ID	Exact contract used	Source contact
F0	Eisenstein conventions, all finite-order trivial-infinity characters, primitive zero supremum with sentinel $\frac{1}{2}$, initial bound $\beta_* \leq 7/8$; ordinary zero-free boundary at $\Re s = 1$	PR9 T0, T2–T4, B0.1

ID	Exact contract used	Source contact
F1	Extended sextic symbol at all powers, fixed ray reciprocity and Gauss phases; full-modulus finite Poisson and CRT at shared primes	PR9 B1, especially B1.1–B1.18; B4
F2	Completed theta reflection, all cusp sources, whole-index marks, both collision masks, common Fourier profiles and whole-series tails; marked energy implication $M + 2z \leq N \Rightarrow Z^{M+(N+2z-2M)_+/4+\varepsilon}$	PR9 B2.1–B2.41; B0.26
F3	UM, MI, old positive MP with actual natural eligibility and coefficient 6κ , unrestricted nonprincipal MP0; moving support before cancellation; independent heights and four derivative profile pairs	PR9 B0.1, B3.1–B3.8, B3.4 natural-source contract
W	Supported simultaneous inverse/plain witnesses with $b > .51$, $t_* \in [1, 3/2]$, actual nonfloor zero, buffered finite dictionary, upper plain support $m \leq 1/2 + 75\epsilon_d$, no-slot counts and common profile discretization	PR9 source contacts DetectorSupportedWitness, DetectorSimultaneousSupport, DetectorWitnessRows; count/profile contacts below
G0	Fourier masks kept inside joint outer labels, fixed translated natural dictionary, modulus $Nf \leq KNu$, signed reconstructed row $u_0 e F_e G_e V h'$ with holes	PR15 FourierMaskTransport, FixedDictionaryBridge, SignedReconstruction
G1	All-free no-slot paired plain exponent $\max(1, 2m)$, $m \leq 3/4$, all heights, real powers $[0, 1]$, finite fixed-inducing rows restored absolutely	PR15 PlainUnmarkedDensity, with PR9 natural-source exceptions
G2	Actual first amplifier data, two transforms, original radii, all error/zero/fixed-inducing terms, joint high labels	PR15 PhysicalLocalAssembly; PR9 B3.5–B3.8 and the three exact source contacts below
P	Literal physical source, common subcollection and four plain derivatives, old-only downward-scale complements, whole-slot exhaustion, primitive/natural reflection with combined conductor and weighted Euler mass	PR15 RestrictedCanonicalSource, PhysicalScaleShortening, EndpointAssembly; PR9 B0.37–B0.44, B4
R	Signed physical Gram, exact compensator and six-row Euler table, principal residue, global/buffered reciprocal bounds, complete three-frequency tails, common-gap Mellin return	PR9 B0.3–B0.5 and B0.7
T	Natural-to-primitive Euler deletion, Hecke family identity, exact norm-lift ordinary product, ordinary $\Re s \geq 1$ nonvanishing	PR9 T2–T4 and declaration contacts below

The new proof replaces only the density range, local numerical region and the ensuing physical/rational constants. In particular it derives D+ on [33/40, 7/8] rather than treating PR15's narrower stated interval as if it already covered this range. In transformed high bins it uses $\delta \geq 17/25$; the physical detector remains on its full rectangle.

Arithmetic definitions needed to read the physical probe

For a squarefree primary n ,

$$\gamma_j(n) = q_n^{-1/2} \sum_{a \bmod n} \chi_n(a)^j e(a/n), \quad a_\nu(n) = \overline{\alpha(n)} \gamma_2(n) \nu(n).$$

For odd $a + b\omega$, the pinned convention is

$$\Gamma_q(a + b\omega) = \frac{1 + i^{-b} + i^a + i^{b-a}}{2}, \quad G(c) = \overline{\chi_c(4)} \Gamma_q(c), \quad R(v, w) = \frac{\Gamma_q(vw)}{\Gamma_q(v) \Gamma_q(w)}.$$

The finite identities include

$$\begin{aligned} \chi_b(a) &= R(a, b) \chi_a(b), \quad \gamma_2(n)^3 = \mu(n) \alpha(n), \quad \gamma_1(n) \gamma_2(n) = \mu(n) \alpha(n) G(n), \\ \mu(n) \gamma_{-1}(n) &= \chi_n(-1) G(n)^{-1} \overline{\alpha(n)} \gamma_2(n). \end{aligned}$$

The fixed phases need not have order dividing six. These formulas and the primary-generator convention determine (5.1); they are not independent choices fixed by the source definitions and checked endpoint adapters.

The completed theta function is

$$\mathcal{C}_V(N; m, \nu) = \sum_{c \text{ sf}, n} \frac{\overline{\alpha(cn^3)} \gamma_2(c) \nu(cn^3) \chi_{cn^3}(m)}{\sqrt{q_c} q_n} V(q_c q_n^3 / N).$$

Its unmarked energy is $(HN)^\varepsilon (H + H^2/N)$. Its marked energy is a separate imported assertion F2: a divisibility mark acts on all of cn^3 , and c, n , marks and rows may share primes. The full source retains all nonzero element sectors and all cusp terms.

Literal local Euler data

For $p \notin S$, $Q = Np$, $j = v_p(u) \in \{0, \dots, 5\}$, put

$$\begin{aligned} \rho &= \chi_p(u/p^j), \quad \mathfrak{a} = \overline{\alpha(p)}^3 \eta(p)^3, \quad \mathfrak{b} = \overline{\alpha(p)}^2 \eta(p)^2 \overline{\chi_p(4)}, \quad \mathfrak{b}^3 = \mathfrak{a}^2, \\ V &= Q^{-6z}, \quad R = \mathfrak{a}^2 Q^{4-6s-6z}, \quad W^\circ = \rho Q^{-w}, \quad W = \chi_p(u) Q^{-w}, \quad D = \eta(p) \overline{\chi_p(u)} Q^{-s}. \end{aligned}$$

Thus $W = D = 0$ at a row prime. Define

$$\begin{aligned} P_p^* &= \frac{1}{1-R} \left\{ \frac{R(1-Q^{-1}) - \eta(p)(Q-1)Q^{-s-w} V^{\mathbf{1}_{j \leq 1}}}{1-V} + J_j \right\}, \quad P_p = (1-V)^{-1} + \mathbf{1}_{j=0} \frac{W^\circ}{1-W^\circ} + P_p^*, \\ H_p &= P_p(1-V)(1-W)/(1-D). \end{aligned}$$

j	J_j
0	$-\eta(p) \rho^{-1} Q^{-s} + W^\circ R$
1	$\eta(p) Q^{-s-w}$
2	$\mathfrak{a} \rho^{-3} Q^{3/2-3s}$
3	$-\eta(p) \mathfrak{b} \rho^{-2} Q^{2-3s-w} + \mathfrak{b}^2 \rho^{-4} Q^{2-4s}$
4	$-\eta(p) \mathfrak{a} \rho^{-3} Q^{5/2-4s-w}$
5	$-\mathfrak{a}^2 Q^{3-6s}$

With $c_p = \overline{\eta(p)} Q^s$, $h_p = Q^{-w}$, the quotient-free compensator is

$$G_p = \frac{(c_p - h_p)(1-V)(1-W)P_p^* - h_p(1-VW)}{1-D}.$$

This is the source definition used in (5.8). No inverse of H_p is introduced. Set $\Delta_p = G_p + \overline{\psi_u(p)}H_p$ even when $p \mid u$. The polynomials in Q^{-w} have degrees $\deg H_p \leq 2, \deg \Delta_p \leq 3$, one smaller at row primes. Their denominators are independent of w . At a row prime of valuation j , assign $\min(k, j-1)$ powers of a monomial shift Q^k to the powerful part $q_u/q_{\text{rad } u}$; assign the rest to h_0 . The exact inequalities B0.42–B0.44 yield the weighted mass (3.6). For example the apparent Q coefficient at $j \geq 2$ is canceled by central normalization, and the quadratic error at $j = 2$ is $Q^{1-a}Q^{-3/2}Q^{(2a-1)/2} = Q^{-1}$. This is why treating all rows as squarefree or charging the deficit twice would change the proof.

Primitive reflection uses $\mathfrak{c}_u = 3N\mathfrak{f}_u/(2\pi)^2 \asymp_S q_{\text{rad } u}$ and

$$\widehat{W}_1^\vee(v) = \widehat{W}_1(1-v) \frac{\Gamma(v)}{\Gamma(1-v)}.$$

Delete redundant Euler factors finitely before reflection and restore their natural zero masks afterward. The restoring coefficients have fixed bounded absolute mass and shorten the scale. The combined conductor—not its primitive factor alone—is required in F1/P.

Second projector and the four clipped rectangle pairs

The source normalization used after (2.15) is PR9 B3.65–B3.69. The residual projector is the full signed sum $\sum_{t|(a,b)} \mu(t)$. Allocate a prime of t by $\mathbf{1}_{p|\prod_i n_i} = \sum_{\emptyset \neq J} (-1)^{|J|+1} \prod_{i \in J} \mathbf{1}_{p|n_i}$. A selected plain becomes pl_i with unrestricted quotient l_i ; both rectangle scales divide by Np . A selected prime slot freezes, and selecting two different primes in one prime slot is zero.

For side j , let α_j be its pre-allocation total log length, e_j its entrance-slot annular discrepancy, and let r_j be the sum of the selected plain-prime log norms and newly frozen slot widths. The new frozen-slot discrepancy is ω_j . On lower-endpoint divisor dyads $Z^{t-} \leq Nt < 2Z^{t-}$,

$$r_j + \omega_j \geq t_-, \quad |\omega_j| \leq \theta_N, \quad |e_j + \omega_j| \leq \theta_N.$$

The raw scalar is $Z^{-r_j/2}$ relative to the pre-normalizer, and the raw post product has length $\alpha_j + e_j - r_j$. For one actual retained rectangle with plain lengths x_i , let $y_i = x_i - a_{j,i}$, $\pi_0 = \sum_i (-x_i)_+$, $\pi = \sum_i (-y_i)_+$. An unsupported rectangle is zero; otherwise $0 \leq \pi_0 \leq \pi \leq \theta_N$. Its clipped reduction and total are

$$r_{\text{clip},j} = r_j - (\pi - \pi_0) \geq 0, \quad A_{\text{clip},j} = \alpha_j + e_j - r_j + \pi.$$

The exact scalar is $Z^{(-r_{\text{clip},j} + e_j + \pi_0)/2}$. Consequently the divisor count and both extraction scalars satisfy

$$t_- - \frac{r_{\text{clip},1} + r_{\text{clip},2}}{2} + \frac{e_1 + \pi_{0,1} + e_2 + \pi_{0,2}}{2} \leq 2\theta_N.$$

This applies separately to the four positive rectangle pairs after Cauchy. No common clipped reduction is assigned to a signed difference. In high bins the residual pointwise factor $-\delta(r_{\text{clip},1} + r_{\text{clip},2})/2$ is favorable. The second projector is therefore paid without pretending it divides the reconstructed row. Source-native fixed-inducing and nonexceptional predicates remain attached to every such child.

Source declarations and independent source pin limits

These existing declaration contacts specify the imported mathematical contracts. Listing them does not assert a fresh compilation, replay or independent dependency audit.

- `FirstAmplificationChoice.eventually_original_amplification`, source lines 712–740: SHA256 9676fcf5ed9200b782ce1565ed9a3dddcba7779743ef55d081d6457d69018f5f, Git blob 1e7a23d2f1319bd7a025a0ee8148a577c685b52d.

- `AmplifiedRetainedRadius.eventually_original_amplification_common`, lines 220–255: SHA256 `8ab0f7421c9355a43124887413fa7080494929064ab5ac04b67cdf97c36aece1`, blob `f6d0c98a4212c85ac9de0c0ce7de35f519383d89`.
- `SecondExceptionalFixedQSource.original_exceptional_energy`, lines 61–80: SHA256 `49d2a891448583561e4f607a615bd86524c3a7ddbece00a20dcdfb6103a7c62c`, blob `286754d49563f233397d28b24f5e5138ef1027f2`.

The supported-witness contacts are `DetectorSupportedWitness`, `DetectorSimultaneousSupport` and `DetectorWitnessRows`. Support and counting contacts are `DetectorZeros`, `DetectorSaturation`, `DetectorDyadicActual`, `DetectorNoSlotInverseCount`, `DetectorNoSlotPlainCount` and `DetectorProfilesUniform`. The first three have direct fixed-head source receipts; the remaining support modules are matched intake evidence. Those receipts do not establish a new independent replay of their complete dependencies. Exact file bindings are recorded in the accompanying source provenance.

Physical contacts are `finite_physical_pair_positiveAt`, `Physical.CertifiedZero.finite_pair_zeroAt`, `PhysicalMomentInput.source_physical_slot_conj` and `reflected_pair_norm_eq`, `weighted_source_keep_iff`, `weighted_source_exceptional_norm`, and `SlotSelection.zero_error_whole_slots`. They fix the actual source, orientation, eligibility, zero-slot range and selector quantifiers. `PhysicalMomentInput` SHA256 is `05c498a791d60edfd5f600e2fa16c51d0012345aaca2d48de6cbe2bc53620170`; `Physical-SourceState` is `2244f9fb94cf780737d291bf1872b5bba232f1d1f879a3c59996f0d34e7a9026`; `NaturalFixedRaySourceFreeExceptional` is `4dbb1822fa52e5743cd4a43d77baa773d9c299e554d13496f5a1c9d6d4318b35`.

The required definition/transfer names are `HeckeFamily.Character`, `HeckeZeroSupremum.beta`; `HeckeDirichlet.character`, `modulus`, `idealCoeff_character`, `LFunction_eq_dirichlet_product`; `ShortDraftHeckeBridge.baseChangeChar`; `CenteredMomentNaturalPrimitive.polynomial_natural_deletion` and `exists_natural_primitive`; `DirichletCharacter.LFunction_ne_zero_of_one_le_re` and `LFunction_modOne_eq`. `Cycle25.PrimitiveProbeContract`, `UniformCommonProbe` and `beta_le_of_common_probe` in `Cycle25/Assembly/Continuation.lean` expose the already variable-boundary common-gap logic. They do not, by themselves, establish this new numerical endpoint. Likewise old `Cycle25.dirichlet_nonzero/dirichlet_nonzero_catalogue` are not reused as new endpoint declarations.

A.2. Rational evidence and proof scope

The source-range extension and endpoint have separate source and rational-check receipts. These receipts establish their stated comparisons; they do not automatically verify every restatement or imported analytic dependency in this paper.

The source architecture is unchanged. $D+$ is newly derived from the actual supported witness and old no-slot estimates on $[\cdot 825, \cdot 875]$. The new constants are $q_0 = .68$, $\Lambda = 19312/4787$, $B = 14200/14361 - 10^{-5}$. Physical geometry is (3.1). The local proof has constant low and high reserves by the exact affine identity (2.22). Every other displayed atlas reserve is affine in its cell endpoint. Thus a short algebraic argument replaces a 600-case enumeration as a premise, while all 600 source checks remain immutable evidence. The separate 156-coefficient endpoint certificate remains necessary in this version.

The helper `rational-check.py` independently expands (4.2) with rational polynomial arithmetic, transforms it by (4.3), and compares every coefficient to the frozen receipt, including zero coefficients and closed-box cover. It also checks the constant-reserve identity, physical extrema, complementary/shift reserves and feedback bounds. It uses Python’s exact Fraction arithmetic and hashing only. It neither imports nor executes the author’s or critic’s helpers. A matching certificate proves only its stated finite inequalities. The analytic lemmas, source dictionary and quantifier order are separate mathematical dependencies.

Numerical endpoint theorems of C52, C54, C55, C56 and C58, PR9’s later quartic optimization, and the older PR15 endpoint are not premises. The initial seven-eighths theorem remains F0. The closed-plane every-epsilon extension is unnecessary: the chosen physical application has strictly interior padding paid in Module 3. We retain all 600 cells; the c5-only first250-cell simplification is invalid for this geometry and is not used. No new 300-cell pruning theorem is needed.

A.3. Primary imports, scope and attribution

The upstream analytic import list is explicit:

- [Dunn–Radziwiłł](#), [arXiv:2109.07463v3](#), sections 5 and Appendix A, for unconditional cubic-theta and cusp material. The conditional prime asymptotic is not used.
- [Goldmakher–Louvel](#), [arXiv:1112.1642v2](#), Definition 1 and Theorem 1.1, for the quadratic Hecke-family sieve. PR9 B2 verifies the actual primary square classes, conductor $a\lambda^{e(a)}$, product conductor, fixed reciprocity bicharacter and transpose application; both original masks are retained.
- [Heath-Brown](#), [Kummer’s conjecture for cubic Gauss sums](#), Theorem 2, also restated as [Dunn–Radziwiłł](#) Theorem 9.1, for the cubic sieve. Squarefree means squarefree ideals in $\mathcal{O}$, not squarefree rational norms; rational prime factors are not silently excluded.
- [Gao–Zhao](#), [arXiv:1707.00091v3](#), equation (1.1), for the stated primitive Hecke functional-equation normalization, as matched to the pinned reflection declarations.
- [Thorner–Zaman](#), [arXiv:1803.02823v3](#), Theorem 1.1 specialized to the identity extension of the fixed Eisenstein field, for annular prime ideal counting. No uniformity in a growing extension is inferred.

These are primary-source contacts of imported mathematics, not new claims that each entire external paper has been independently rechecked for this proof. The retained source audits identify the used statements and applicability limits. The frozen source closure records the used dependency graph. Its complete rebuild and independent replay are verification obligations distinct from matching source bytes; their status is recorded in A.4.

Credit remains with the OpenAI initial seven-eighths framework at `fd4aeeb2ee4fc729c18d98444fed42fd0529eeeb`; the AI-produced PR9 completed theta/quartic, variable moment, detector and source architecture developed within the research program directed by Hailey Collet; Akash Levy’s weighted numerator work at `2fd60c0926b66ea18d7436f5ed55250fd006ab5d`; Gao Zhongpai’s AI-assisted PR15 local construction and source bridges; and the subsequent QRH source/endpoint extensions and independent reviews. This selected assembly does not erase their contributions or relabel imported proofs as newly originated here. Preserve PR15’s Apache-2.0 LICENSE, NOTICE.md, UPSTREAM_NOTICE.md and the upstream notices in any derivative artifact.

A.4. Formal statements and verification

The development build checks `Cycle61.FinalAssembly.beta_le_rational` and all four declarations in `Cycle61.Verification` with only `propext`, `Classical.choice` and `Quot.sound`. The rational bound is unconditional. The auxiliary lower-bound gate $51/100 \leq \beta_*$ is used only in the contradiction branch $\theta < \beta_*$, where it follows from the exact value of θ ; it is not a hypothesis of the final theorem.

The Dirichlet statement uses the ordinary `Mathlib.DirichletCharacter.LFunction`, every nonzero natural modulus, and every complex character. The Hecke statement uses `OAI.SevenEighths.HeckeFamily.Character` and its `LFunction`, with principal residue equal to one excluded at $s = 1$. The adapters retain imprimitive Euler factors, the ramified prime three and the norm-lift product’s domain. The proof also retains literal eligibility and masks, common profiles, old-only

descendants, the amplifier and original retained-radius quantifiers, physical complements and weighted shifts, and the common gap chosen before the target family.

The verification source is SHA256 b966ba23494d14cd0c0cfb220e8e07864646ef9d3a66539b8933d1f3e2c48ae1; the unconditional assembly is 96fc445a406dc2ef3726ec666b743f2bc71ecdb8b45a195d12d48af0479c3825. The complete clean rebuild of the frozen 3,695-module source closure passed on 13 October 2026, using Lean 4.34.1 with no submitted objects reused and no frozen sources changed. Each final comparison export reports only the three standard axioms. For the original source, the literal statements `allDirichlet`, `zeta` and `allHecke` and their exported proof closures passed the Comparator and independent Lean/NanoDa/con-ron replay on 13 October 2026. Dependency seals agreed before and after the run. `zetaTotal` is the separate fourth development export. The transformed submission also passed a fresh 194-module compilation, the canonical all-Dirichlet Comparator and all three kernels. Both local replays used a recorded native judge adapter because this host refused nested namespaces; the kernels retained their own sandboxes. Protected exact-head catalogue admission requires a separate tracker run on the submitted commit. The theorem’s boundary is $\theta = .874861$, with $c = .125139$; no claim about the current catalogue frontier is made here.

The submission uses the dependency-ordered, section-isolated packaging approach of [nielstron’s PR26](#). It maps 772 local modules into 193 bundles plus the literal `Candidate.bound` wrapper, retaining the originals and a reproducible transformation map. The map records import, scope, name-qualification and token-spacing changes. In the bundled copy, `actual_main_low_all_height_splice` has a local elaboration heartbeat allowance of 40,000 instead of 20,000; its statement and proof body are unchanged. Protected replay time limits are unchanged. The transformed package was checked separately from the original sources.

Authorship and AI assistance

Hailey Collet initiated and directed the research program, selected targets and sources, proposed questions, reconciled assumptions, and directed the assembly of compatible proof components. She brought Gao’s refinement into the program and required its contribution to remain explicit. Her authorship credits research conception, direction and source integration; it does not attribute the detailed mathematical derivations to her personally.

The detailed derivations and substantial drafting were produced with AI assistance. Earlier work in the research program was developed mainly with GPT6-Pro and non-Pro Astra, with contributions from GPT 6.1 Sol, Grok 4.7 and Gemini 3.1 Pro. Some research prompts were also AI-written and selected or adapted by the author. This is the author’s account of the tools used across the program, rather than a per-message record or an assignment of specific models to every later stage. It does not certify mathematical verification.

Copyright Hailey Collet. This paper is offered under Apache-2.0, with the source attributions and upstream notices retained. Hailey Collet is the sole human author.