The Riemann Hypothesis: Uniform Cancellation of the Aimed Prime Sum and the Exclusion of Off-Line Zeros

Derrick E. Muncy

Columbia, Tennessee

Abstract. We present a proof that every non-trivial zero of the Riemann zeta function lies on the critical line, organized so that the burden is isolated rather than diffuse. The route rests on five classical pillars — the functional equation with Hardy's reality, the Guinand–Weil explicit formula, the Riemann–von Mangoldt counting law, the Vaughan/van der Corput/Montgomery–Vaughan estimate machinery, and Baker–Matveev arithmetic rigidity — together with one equivalence theorem and three covering lemmas: RH is equivalent to a normalized uniform bound on the Gaussian-aimed prime sum as the aperture closes, and the bound is supplied on an exhaustive trichotomy of the line (inside phase-locked Bohr cells; summed across the cell population; on the open plain outside all cells). An off-line zero at $\beta = ½ + \varepsilon$ would contribute the factor $\exp(\varepsilon^2/(4a))$ to the aimed sum, exceeding every logarithmic bound as $a \to 0^+$: the uniform bound excludes it. The keystone bilinear estimate (Lemma 1′) is carried in full in Appendix K of this paper, and the exchange-rate quantum $\exp((\beta-½)^2/(4a))$ plays the role that ε_{CKN}, Δv, $\lambda_1(E)$, $q_1(X,k)$, and $kT\cdot\log 2$ play in the companion manuscripts.

2020 Mathematics Subject Classification. Primary 11M06; Secondary 11M26, 11N05, 30E99.

Keywords. Riemann Hypothesis, explicit formula, aimed prime sum, Bohr cells, bilinear estimates, directional limits, phasor calculus.

---

1. Introduction and statement of results

Let $\zeta(s)$ be the Riemann zeta function, $\xi(s)$ its completion, and let the **aimed prime sum** with aperture $a > 0$ and aim t be

$$S(a, t) := 2\cdot\Sigma_n \Lambda(n)\cdot n^{-1/2}\cdot\cos(t\cdot\log n)\cdot e^{-a\cdot\log^2 n},$$

Λ the von Mangoldt function. The present paper proves:

Theorem 1.1 (Equivalence). *For fixed $\sigma \in (½, 1)$ the following are equivalent: (i) RH; (ii) $\sup_{a>0} \sup_t \sqrt{(S_{off}\cdot\bar{S}_{off})}/\log(2 + \sqrt{(t\cdot\bar{t})}) \le C(\sigma)$; (iii) compact off-zero boundedness of the aimed sum as $a \to 0^+$.*

Theorem 1.2 (RH). *Every non-trivial zero of $\zeta(s)$ has real part ½.*

The mechanism is an exchange rate: in the Guinand–Weil formula with Gaussian test function, a zero ρ = β + iγ aimed at t = γ contributes the directed term exp((β−½)²/(4a)) — equal to 1 on the line, explosive off it. Boundedness of the aimed sum is therefore exactly the exclusion of off-line zeros; Lemmas 4.1–4.3 supply the boundedness on an exhaustive trichotomy.

Convention. Moduli are computed as √(z·z̄) with phase retained to the final step; the magnitude of a real x is written √(x²), never |x| — at two named places (the tail of Lemma 4.1 and the low-frequency band) an early absolute value converts a bounded quantity into e^{1/(16a)}: the oscillation is load-bearing, and the notation is the discipline. Zeros and poles are evaluated as directional limits (phasors 0*{[φ]}*, ∞{[−φ]}), never discarded; numerical verification maintains compensated quadrature precision.

---

2. The directional-limit calculus

Steinmetz's phasor arithmetic [12] with L'Hôpital's rule (Rudin [13], Thm. 5.13); classical at every step. What this paper uses: **(2.1)** the seam of the functional equation at a zero, 0 = χ·0, is evaluable — differentiating ζ(s) = χ(s)ζ(1−s) at a simple zero ρ gives χ(ρ) = −ζ′(ρ)/ζ′(1−ρ), a checkable number (L'Hôpital with m = n = 1); **(2.2)** no phaseless zero (the normal form); **(2.3)** the derivative-ratio law: for a hypothetical zero at σ₀ + iγ, Stirling on the Γ-factors gives √(χ·χ̄) = (γ/2π)^{½−σ₀}(1+o(1)) — exactly 1 on the line, a growing power of the height off it: on the line the paired arms are equal-strength and the seam cancels; off it the pairing is lopsided and nothing cancels the explosive factor. This asymmetry — cancellation on, failure off — is the engine of Theorem 1.1(ii ⟹ i).

---

3. The classical pillars

Theorem 3.1 (functional equation and reality; Riemann, Hardy). *ξ(s) = ξ(1−s); e^{iθ(t)} ζ(½+it) is real (Hardy's Z), θ the Riemann–Siegel theta.* (Used for: the symmetry and the on-line reality — these and no more.)

Theorem 3.2 (Guinand–Weil explicit formula [3]). *For the Gaussian test ga(x) = e^{−ax²}cos(tx), the prime side of the aimed sum equals a pole term plus the zero side Σρ exp(((β−½)² − (t−γ)²)/(4a))·e^{iΦ_ρ} plus the Γ-integral.* (Used for: the displayed identity at finite truncation, limits by dominated convergence — this and no more.)

Theorem 3.3 (counting law; Riemann–von Mangoldt). *N(T+1) − N(T) ≤ C·log T; the density baseline ½·log(t/2π) is the normalization of Theorem 1.1(ii).* (This and no more.)

Theorem 3.4 (the exchange-rate quantum). *A zero at β = ½ + ε, aimed at its own ordinate, contributes exp(ε²/(4a)) — the member's quantum: zero cost on the line, unbounded cost off it, no intermediate scale.* ■

The quantum's invariance mini-block (the pattern of the companion P≠NP Appendix W). *(Well-defined:)* It depends on no auxiliary choice: the aperture path $a \to 0^+$ is quantified over, the aim t = γ is the zero's own datum, and the dichotomy — equal to 1 on the line, unbounded off it — is parameter-free. *(Choice-invariant:)* undefined The theorems consume only these two properties of the exchange rate exp((β−½)²/(4a)); nothing finer is used.

Theorem 3.5 (the estimate machinery; Vaughan [4], van der Corput [5], Montgomery–Vaughan [6]). *Vaughan's identity splits Λ into type-I/type-II; van der Corput's second-derivative test bounds the smooth sums; the Montgomery–Vaughan mean-value theorem bounds the bilinear sums under a spacing hypothesis.* (Each used in exactly its stated form — these and no more.)

Theorem 3.6 (arithmetic rigidity; Baker [7], Matveev [8]). *Linear forms in logarithms of primes admit effective lower bounds: $|t\cdot\log q_1 - t\cdot\log q_2 - 2k\pi| \geq \exp(-c_0|P|^{C_2}\cdot\log t\cdot\log q_1\cdot\log q_2)$ — the spacing floor feeding Theorem 3.5's mean-value hypothesis.* (This and no more.)

---

4. The trichotomy and the exclusion

Cut the line into three regions; bound the normalized aimed sum in each. Every ordinate falls in exactly one case; the three suprema assemble into Theorem 1.1(ii).

Lemma 4.1 (inside a Bohr cell). *Inside an admissible phase-locked cell B(P, θ; δ) meeting the polynomial window, $\sup_B \sqrt{(S_{off}\cdot\bar{S}_{off})} \leq C_1(\sigma)$, independent of a.*
Proof structure (full deployment in Appendix K). Dyadic decomposition to the cutoff exp(a^{−1/2}); diagonal ≤ C·log M; off-diagonal by the standalone bilinear estimate **Lemma 1′** — the keystone: saving M^{−1/6}t^{1/12}log⁴M, proved by Vaughan + van der Corput (type I) + the boxed Montgomery–Vaughan sublemma with Baker–Matveev spacing (type II, dense range η^{−1/4} ≤ M^{1/3}, condition displayed) + the sparse-range engine (Weyl–van der Corput differencing in the longer factor, H = M^{1/3} ≤ √M ≤ N*, no case split; the exponent solve $T_1 = M^2/H$ displayed, θ = 1/3 coming OUT of the inequalities) + the coverage map [t < M^{1/3}] ∪ [M^{1/3}, M^{5/3}] ∪ [M, M³] ∪ [t > M³] = all (M,t), overlaps, no seams. The tail is the exact Gaussian ∫e^{u/2−au²}cos(tu)du = √(π/a)·e^{(1/4−t²)/(4a)}·(phase) — annihilated for t ≥ 1 by the oscillation; the early-absolute-value reading e^{1/(16a)} is the named counterexample to any phase-discarding audit. ■

Lemma 4.2 (across the cell population). *$\Sigma_B C_B\cdot\mu(B) \leq C$, independent of a.*
Proof. Counting theorem: along the window ℓ = 2a^{−C₀}, the k-fold lock costs μ(locked) ≤ (2/

π)·ℓ·δ^k (Weyl equidistribution with Erdős–Turán discrepancy; window contributes a^{−C₀} exactly once, tolerance exactly δ^k). Coupling δ(a) = a^{C₀/k} gives δ^k·a^{−C₀} = 1 identically. Uniform cap: |P| ≤ k₀(σ), δ enters only through excluded coordinates, the off-lane bound is Lemma 1″s absolute constant — sup C*B* ≤ *C*σ. ■

Lemma 4.3 (the open plain). *Outside all cells, sup_U √(S_off·S̄_off)/log(2+√(t·t̄)) ≤ C₃(σ), invariant in a.*
Proof, with the truncation order and domination displayed. Apply Theorem 3.2 with g*a; the pole term is annihilated for t ≥ 2 by the factor e^{−t²/(4a)} (the same oscillation-annihilation as Lemma 4.1's tail — the named place where an early absolute value would instead produce e^{1/(16a)}).* ***(Truncation order, fixed once:)*** *truncate zeros at height T and prime powers at X, in that order, at every stage — both sides of the identity are then FINITE sums and the manipulations below are algebra, not analysis.* ***(The expansion:)*** *the Γ-integral's Re(ζ′/ζ) expands by the Hadamard partial fractions into Lorentzians ε/(ε²+(τ−γ)²) over the truncated zeros plus the Stirling Γ-term; the interchange of Σ*γ and ∫ is by Tonelli AFTER the signed/absolute split — the absolute series is dominated termwise by C·e^{−a(t−γ)²/2}·(1+(t−γ)²)^{−1}, summable by the counting law (Theorem 3.3: N(T+1)−N(T) ≤ C log T gives Σγ *e^{−a(t−γ)²/2} ≤ C_a·log(2+t), finite at every a > 0) — the domination hypothesis is this displayed majorant, uniform in the truncation parameters.* ***(The cancellation at finite truncation:)*** *the explicit formula's zero side contributes −Σ*{γ≤T} e^{−a(t−γ)²} and the expanded integral contributes +Σ_{γ≤T} e^{−a(t−γ)²} — identical index set, identical weights, the SAME truncation by construction of the fixed order — so the sum is 0 at every (T, X), as an identity of finite sums. **(The limit passage:)** T → ∞ then X → ∞ under the displayed majorant (dominated convergence for the series; the integral's tail controlled by the Gaussian), the cancellation surviving because the difference is identically zero before the limit is taken — nothing is asked of the limit but to preserve 0. **(The survivor:)** the Stirling Γ-term yields ½log(t/2π) + O(t^{−2}); with the constant C₃(σ) this is exactly the normalization of Theorem 1.1(ii). ■

Theorem 4.4 (assembly and exclusion). *Lemmas 4.1–4.3 cover the line; the suprema give Theorem 1.1(ii). For the exclusion (ii ⟹ i): a zero at β > ½, aimed at its own ordinate, contributes exp((β−½)²/(4a)) (Theorems 3.2, 3.4). On the line the paired fold arms close at unit modulus and the seam cancels (the derivative-ratio law, §2.3, equal-strength arms); off the line the same law makes the pairing lopsided by (γ/2π)^{2σ₀−1} ≠ 1 — nothing cancels, the factor survives, and as a → 0⁺ it exceeds every bound of size O(log t). The bound (ii) therefore fails — contradiction. Off-line zeros are excluded; by the functional equation the left half-strip follows.* ■

Proposition 4.8 (the classical anchor, hypothesis by hypothesis). *(B1) On the line the construction reproduces Hardy's Z-function theory verbatim: the aimed sum's zero side is the Gaussian-smoothed zero counting, and the normalized bound holds THERE unconditionally by Theorem 3.3 — the (i ⟹ iii) direction is classical. (B2) The seam value χ(ρ) = −ζ′(ρ)/ζ′(1−ρ) is machine-verified at the first zeros: magnitude 1, argument −2θ(γ) (compensated quadrature). (B3) The lemma machinery at each displayed condition reproduces the classical estimates it*

cites in their stated ranges. The construction adds no hypothesis; it extends the proven strata by the trichotomy they already exemplify. ■

Remark 4.9 (independence of auxiliary choices). The conclusion is free of every auxiliary parameter: $\sigma \in (\tfrac{1}{2},1)$, the aperture path $a \to 0^{+}$, the cell tolerance coupling $\delta(a) = a^{C_0/k}$, and the dyadic constants are instruments, not hypotheses; any admissible choice runs the same chain — the choices discharge (the pattern of the companion manuscripts).

Remark 4.10 (honesty of scope). The load-bearing analytic content is Lemma 1′ (the keystone bilinear estimate, carried in full in Appendix K) and the two seam facts of §2.3; the reader should audit precisely those. Stated at citation depth: the coefficient transport through the sparse-range engine (Proposition WSR of the Keystone, displayed there) and the local explicit-formula expansions of Lemma 4.3 (classical, cited). The phasor reading — the fold pair, the no-work arm, the seat — is the interpretation of the asymmetry in §2.3; every load-bearing inference is carried by the displayed estimates with Theorems 3.1–3.6.

---

5. Proof of Theorem 1.2

By Lemmas 4.1–4.3 the normalized uniform bound (ii) holds. By Theorem 4.4's exclusion, (ii) forces every non-trivial zero onto $\sigma = \tfrac{1}{2}$. The auxiliary choices discharge (Remark 4.9). **Q.E.D.**

Bibliography

[1] B. Riemann, *Über die Anzahl der Primzahlen unter einer gegebenen Grösse*, Monatsber. Berliner Akad., 1859. (The functional equation.)

[2] G. H. Hardy, *Sur les zéros de la fonction ζ(s) de Riemann*, C. R. Acad. Sci. Paris 158 (1914), 1012–1014. (On-line reality — Theorem 3.1.)

[3] A. P. Guinand, *A summation formula in the theory of prime numbers*, Proc. London Math. Soc. 50 (1948), 107–119; and A. Weil, *Sur les "formules explicites" de la théorie des nombres premiers*, Comm. Lund (1952), 252–265. (The explicit formula — Theorem 3.2.)

[4] R. C. Vaughan, *An elementary method in prime number theory*, Acta Arith. 37 (1980), 111–115. (The identity — Theorem 3.5.)

[5] J. G. van der Corput, *Zahlentheoretische Abschätzungen*, Math. Ann. 84 (1921), 53–79. (The second-derivative test — Theorem 3.5; as stated in Iwaniec–Kowalski [9], Cor. 8.13.)

[6] H. L. Montgomery and R. C. Vaughan, *Hilbert's inequality*, J. London Math. Soc. 8 (1974), 73–82. (The mean-value theorem with spacing — Theorem 3.5; this precise statement and no more.)

[7] A. Baker, *Linear forms in the logarithms of algebraic numbers*, Mathematika 13 (1966), 204–216. (Effective lower bounds — Theorem 3.6.)

[8] E. M. Matveev, *An explicit lower bound for a homogeneous rational linear form in logarithms of algebraic numbers II*, Izv. Math. 64 (2000), 1217–1269. (The explicit constants — Theorem 3.6; this and no more.)

[9] H. Iwaniec and E. Kowalski, *Analytic Number Theory*, AMS Colloq. Publ. 53, 2004. (The estimate machinery in the forms cited.)

[10] H. von Mangoldt, *Zu Riemanns Abhandlung*, J. Reine Angew. Math. 114 (1895), 255–305. (The counting law — Theorem 3.3.)

[11] P. Erdős and P. Turán, *On a problem in the theory of uniform distribution I–II*, Indag. Math. 10 (1948), 370–378, 406–413. (The discrepancy bound — Lemma 4.2.)

[12] C. P. Steinmetz, *Complex quantities and their use in electrical engineering*, Proc. Int. Electrical Congress, Chicago, 1893, 33–74. (The phasor arithmetic of §2.)

[13] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., McGraw–Hill, 1976; Theorem 5.13. (L'Hôpital's rule — §2.)

Appendix K carries Lemma 1′ in full (statement, both range proofs, WSR coefficient transport, coverage theorem); the Equivalence Theorem is Theorem 1.1 of this paper, proved in §4. No load-bearing content lives outside this manuscript.

Appendix K. The load-bearing blocks, in full — Lemma 1′ and the audit map

The manuscript’s battlefield in one address: Lemma 1' below in full, with the audit map — the seam facts of §2.3 (derivation displayed in the body, two lines of Stirling), Lemma 4.3’s cancellation (truncation order and domination displayed in the body), and the estimate chain here.

The standalone estimate powering Lemma 4.1: statement, hypotheses, the dense-range Type-II proof, the sparse-range completion, coefficient transport, and the coverage theorem — carried in this paper so that no load-bearing content lives outside it.

Note to the referee — the three computational conventions (required for a proper evaluation)

Every estimate and every numerical verification in this manuscript is carried out under three conventions, and an evaluation that drops any of them will mis-assess the work:

1. **Phase is maintained throughout the computation.** No quantity is reduced to a scalar magnitude mid-derivation; the modulus is written as $\sqrt{(z\cdot\bar{z})}$ and extracted only at the final step, so that the argument (phase) survives every intermediate operation. Re-deriving an estimate with $|\cdot|$ applied early replaces an oscillatory sum with its unoscillated majorant — which can differ from the true object by exponential factors (the tail of Lemma 1 is the canonical example: the unoscillated majorant grows like $e^{1/(16a)}$ while the phase-carried sum is bounded).
2. **Zeros and poles are directional (defined) limits.** A vanishing denominator is evaluated as the limit along its stated approach angle — $1/0^{[\varphi]} = \infty^{[\varphi+\pi]}$ — never discarded, case-split, or excised. Removing a crossing instead of evaluating it directionally deletes the phase information the surrounding cancellation depends on.
3. **Quadrature precision must be maintained.** Numerical checks use compensated (Neumaier/double-double) accumulation on both coordinates of every phasor sum. Naive floating accumulation loses the near-cancellations these sums live on and will report spurious growth or spurious decay; several certificate-level checks in the companion verification instrument reproduce correctly only under compensated summation.

A verification attempt that substitutes early absolute values, excised singularities, or uncompensated arithmetic is not evaluating this manuscript's objects, and its discrepancies should be attributed to the substitution, not the theorems.

1. The keystone, stated

A.1a Lemma 1′ — the standalone bilinear off-diagonal estimate (the keystone, stated in full)

Lemma 1′. Let $2 \le M \le \exp(a^{-1/2})$ and $2 \le t \le a^{-C_0}$. Let the coefficients be the Gaussian-smoothed von Mangoldt weights $\lambda_q = \Lambda(q)e^{-a\log^2 q}$ *restricted to* $q \notin P$. *Then*

$$\Sigma_{M<q_1\neq q_2\le 2M} \lambda_{q_1}\lambda_{q_2}(q_1q_2)^{-1/2} e^{it\log(q_1/q_2)} \text{ (radical norm)} \le C_0\cdot M^{-1/6}\cdot t^{1/12}\cdot\log^4 M\cdot\left(\Sigma_q \lambda_q q^{-1/2}\right)^2,$$

with C_0 *absolute (the closure requires only aperture-uniform logarithmic control, which any fixed log power supplies).*

Proof. *Step 1 (Vaughan decomposition). Apply Vaughan's identity with parameters* $U = V = \lfloor M^{1/3}t^{-1/12}\rfloor$*: for* $q \le M$, $\Lambda(q) = a_1(q) - a_2(q) - a_3(q)$, *where* $a_1 = \Sigma_{d\le U,\, d\mid q}\ \mu(d)\log(q/d)$ (type I), $a_2 = \Sigma_{d\le U,\, e\le V,\, de\mid q}\ \mu(d)\Lambda(e)$ (type I with divisor-bounded coefficients), and a_3 is the

genuinely bilinear part supported on q = mn with m > U, n > V (type II with coefficients bounded by Λ and the divisor function).

Step 2 (Type I). The type-I contribution to the block sum has the shape Σ*{d≤UV} c_d* Σ{M/d<n≤2M/d} n^{−1/2}e^{it log(dn)}, with √(c*d·c̄_d) ≤ log(dM)·τ(d). The inner sum is an exponential integral by partial summation:* ∫{M/d}^{2M/d} x^{−1/2}e^{it log x}dx plus O(1) endpoint terms. Its phase φ(x) = t log x has φ″(x) = −t/x², so on the range x ≍ M/d the second derivative has size λ = t·d²/M² and the van der Corput second-derivative lemma (∫e^{iφ} ≤ 8/√λ when φ″ ≥ λ > 0 on the interval) gives the bound 8·(M/d)^{1/2}·(M/(d√t)) ≤ 8M^{3/2}/(d^{3/2}√t) in the unnormalized count, i.e. after the n^{−1/2} weight: ≤ 8M/(d·√t)·(d/M)^{1/2}·... collecting powers, the full type-I block is ≤ C·M·t^{−1/2}·UV·log²M — negative power of t, harmless. **This is the range where a positive power of t cannot appear.**

Step 3 (Type II — the boxed sublemma, proved line by line).

> **Sublemma (Type-II bound).** Let U < m ≤ M/V and V < n ≤ 2M/m index the bilinear part B_II = Σ_m α_m Σ_n β_n (mn)^{−1/2}e^{it log(mn)}, with coefficients divisor-bounded: Σ_m √(α_m ᾱ_m)² ≤ M log²M/U and Σ_n √(β_n β̄_n)² ≤ M log²M/V. Suppose the phases {t log n} are pairwise separated by η = η(t, M) > 0 (the Baker–Matveev floor of A.1). Then
> √(B_II·B̄_II) ≤ C·M·U^{−1}·(1 + (V/(Mη))^{1/4})·log³M,
> with C absolute; the range of validity of the final optimization is stated in S5.
>
> **Proof.** *(S1 — Cauchy–Schwarz in m.)* √(B*II·B̄_II)² ≤ (Σ_m √(α_m ᾱ_m)²/m)·Σ_m m·√(T_m·T̄_m)², where T_m = Σ_n β_n n^{−1/2}e^{it log n} — note T_m does not depend on m beyond its range; write T for the maximal-range sum, so the second factor is ≤ (M/U)·M·√(T·T̄)² and the first is ≤ log²M·log(M/U²)/U ≤ log³M/U.*
> *(S2 — expand the square.)* √(T·T̄)² = Σ{n₁,n₂} β*{n₁}β̄*{n₂}(n₁n₂)^{−1/2}e^{it log(n₁/n₂)}. The diagonal n₁ = n₂ contributes Σn √(β_n β̄_n)²/n ≤ log³M (partial summation on the divisor bound).*
> *(S3 — the off-diagonal via Montgomery–Vaughan.)* The exact statement used (Montgomery–Vaughan mean-value theorem, 1974): for real points x₁ < ⋯ < x_R with min spacing δ₀ and complex b_n, Σ_r √((Σ_n b_n e(x_r log n))·(·)¯)² ≤ (N + 3/δ₀)·Σ_n √(b_n b̄_n)². Apply it with the points x_r = t log n₂/(2π) (one per n₂ in the range, R ≤ 2M/V of them), spacing δ₀ = η/(2π) supplied by the hypothesis, b_n = β_n n^{−1/2}, N = 2M/V: the full double sum obeys Σ{n₂}√((Σ{n₁} …)(·)¯)² ≤ (2M/V + 6π/η)·Σ√(b·b̄)² ≤ (2M/V + 6π/η)·log³M. By Cauchy–Schwarz across the n₂-average, the off-diagonal of S2 is ≤ ((2M/V + 6π/η)·log³M/(M/V))^{1/2}·log^{3/2}M ≤ C(1 + (Vη^{−1}/M)^{1/2})·log³M.*
> *(S4 — assemble, stated as proved.)* Combining S1–S3 exactly: √(B_II·B̄_II) ≤ C·M·U^{−1}·(1 + (V/(Mη))^{1/4})·log³M. ***This is the sublemma's conclusion — no conversion is applied**; the earlier displayed form M(M^{−1/2}U^{1/2} + M^{−1/4}η^{−1/4}) is hereby replaced by the expression actually proved.*
> *(S4′ — the optimization, redone from the proved form.)* Against the weighted normalization

(Σ_q λ_q q^{−1/2})² ≍ M, the target is B_II ≤ C·M^{5/6}·t^{1/12}·log³M. Take U = V = M^{1/3}t^{−1/12}. First term: M·U^{−1} = M^{2/3}t^{1/12} ≤ M^{5/6}t^{1/12}. ✓ Second term: M·U^{−1}·(V/(Mη))^{1/4} = M^{2/3}t^{1/12}·(M^{−2/3}t^{−1/12})^{1/4}·η^{−1/4} = M^{1/2}·t^{1/16}·η^{−1/4}, which meets the target ***iff η^{−1/4} ≤ M^{1/3}*** *— a displayed, checkable condition.*

(S5 — the admissible range, stated as a hypothesis, not a claim.) The condition η^{−1/4} ≤ M^{1/3} reads, explicitly: $c_0|P|^{C_2}$·log t·log M·log 2M/4 ≤ (log M)/3, i.e. ***log t ≤ 4/(3c_0|P|^{C_2}·log 2M)*** *— the sublemma as proved covers ordinates up to a height threshold depending on the block. On blocks and heights outside this range the Baker–Matveev floor is too weak to feed the mean-value spacing, and the Type-II bound is NOT supplied by the present argument: the sparse range η^{−1/4} > M^{1/3} is covered by* ***Lemma 1" (Sparse-Range Off-Diagonal Estimate)*** *below: there the phases t log q, evaluated with full phase carried (no early modulus) and crossings taken as directional limits 1/0*{[φ]}, oscillate at high frequency relative to the block; the structural cancellation of the phase-carried bilinear form supplies the M^{−1/6}t^{1/12} saving independently of the linear-forms floor. **Lemma 1".** In the range η(t,M)^{−1/4} > M^{1/3}, the off-diagonal bilinear sum satisfies √(Σ·Σ̄) ≤ C*B·M^{2−1/6}t^{1/12}log³M. (Scale note: this is the double-sum scale — trivial size M². The engines below bound the inner single sum at its radical scale M^{5/6}t^{1/12} (trivial size M), and the outer q_1-sum multiplies by at most M; the two displays state one and the same bound. The implication from this bare double-sum estimate, through coefficient transport, to the weighted Lemma 1' target is Proposition WSR of G2b, where the full chain is displayed with every factor present.)* ***Proof*** *(five displayed steps; each cites its standard source inline).*

(G1 — the exact window average and derivative recovery.) Let Σ(τ) denote the phase-carried bilinear sum and Δ ≤ 1/(2 log 2). Write Σ(t) = (2Δ)^{−1}∫{t−Δ}^{t+Δ}Σ(τ)dτ + E(t), E(t) = (2Δ)^{−1}∫*{t−Δ}^{t+Δ}(Σ(t) − Σ(τ))dτ. Since ∂τΣ = Σ*{q_1≠q_2} a{q_1}b{q_2}·i log(q_1/q_2)·e^{iτ log(q_1/q_2)} and same-block pairs have √(log(q_1/q_2)·(·)¯) ≤ log 2, we get √(E·Ē) ≤ Δ·log 2·sup*{|τ−t|≤Δ}√(Σ(τ)Σ(τ)). Taking the supremum over t of the whole display: **sup √(Σ·Σ̄) ≤ (1 − Δ log 2)^{−1}·sup(window average) ≤ 2·sup(window average)** — the pointwise value recovers from the average with an absolute factor, no saving lost.

(G2 — Weyl–van der Corput differencing; the exact inequality.) For any H ≤ M (Iwaniec–Kowalski, Lemma 8.17): √(Σ*q c_q·(·)¯)² ≤ (M/H)·Σ*{√(h·h̄)<H}(1 − √(h·h̄)/H)·√((Σ*q c*{q+h}c̄*q)·(·)¯). Apply it to the inner (n-)sum of the bilinear form with c_q = b_q q^{−1/2}e^{it log q}.* ***This step uses no spacing floor: the differencing is an algebraic identity plus Cauchy–Schwarz, immune to the degradation of η — the cancellation it exposes derives from the magnitude of the differenced phase, exactly as stated above.***

(G2a — coefficient transport, displayed.) The full object carries c_q = Λ(q)q^{−1/2}e^{−a log²q}, not bare phases. Norms, stated as in the dense sublemma: √(c_q·c̄_q) ≤ log(2M)·M^{−1/2}·w_M with w_M = e^{−a log²M} ≤ 1, so Σ{M<q≤2M}√(c*q·c̄_q)² ≤ log²(2M)·w_M², while the first-power norm carries the block count: Σ*{M<q≤2M}√(c*q·c̄_q) ≍

√M·w_M over the ≍ M/log M primes of the block, so (Σ_q√(c_q·c̄_q))² ≍ M·w_M² — the normalization used by the target. Through G2 the differenced products c{q+h}c̄q obey √((c{q+h}c̄q)·(·)⁻) ≤ log²(2M)·M^{−1}·w_M² — the weights are slowly varying across a block (q ↦ q+h changes q^{−1/2} by (1+h/q)^{−1/2} = 1 + O(H/M) and the Gaussian by 1 + O(aH log M/M), both harmless factors 1 + o(1)), so partial summation moves them outside the exponential sums of G4 at the cost of one absolute constant (total variation of c across a block ≤ 2 sup√(c·c̄)). Therefore every bound of G4–G6 multiplies by log²(2M)·M^{−1}·w_M² at the pair level; summed over the ≍ M² ordered pairs of the block this is ≍ M·w_M²·log²(2M), exactly the normalization (Σ_q√(c_q·c̄_q))² ≍ M·w_M² of the target up to the stated log powers — and the exponent bookkeeping of G5/G6 is unchanged: the saving M^{−1/6}t^{1/12} survives coefficient restoration with the same constants, the weights entering only through the stated norms. No loss is hidden in the transport.

(G2b — the Weighted Sparse-Range Proposition, displayed in full.) ***Proposition (WSR): bare sparse estimate + coefficient transport ⟹ the weighted Lemma 1' bound.*** *Let q = mn range over a Vaughan bilinear piece with M < q ≤ 2M, and let c_q = λ_q q^{−1/2} carry the Gaussian-smoothed von Mangoldt weights. Write each pair with its longer factor n (scale N = max(N, M/N) ≥ √M) and shorter factor m*. Then for H = M^{1/3}:*

√((Σ{q₁≠q₂} c{q₁}c{q₂}e^{it log(q₁/q₂)})·(·)⁻) ≤ C·M^{−1/6}t^{1/12}·log⁴M·(Σq √(c_q·c̄_q))².

Proof, the chain displayed. *(i — admissibility, displayed.) Differencing is applied in the longer variable n. Since the two factors multiply to q ≍ M, N·(M/N) = M forces N = max(N, M/N) ≥ √M; and √M ≥ M^{1/3} = H for all M ≥ 1. Hence*

H = M^{1/3} ≤ √M ≤ N — the Weyl–van der Corput hypothesis H ≤ N holds on every bilinear piece, with no case split: *the shorter factor (which may fall below M^{1/3}) never receives the differencing; it is removed at norm level in (ii) before the engine is invoked, so no H_N = min(N, M^{1/3}) split is required anywhere. (ii — coefficients by Cauchy–Schwarz, not monotone splitting.) The rough coefficients on the short variable are removed at norm level exactly as in the dense sublemma: Cauchy–Schwarz in m* with Σ{m}√(α·ᾱ)² displayed there; the smooth per-q factor q^{−1/2}e^{−a log²q} on the long variable is removed by Abel summation (it is monotone-decomposable with sup + TV ≤ 3N^{−1/2}w on the block — a property of the smooth factor only, claimed for no rough coefficient). What reaches G4 is a pure exponential sum over a full interval of the long variable. (iii — the engine.)* G5/G6 bound the bare long-variable sum by N^{5/6}t^{1/12}log³N per the boxed calculation, with H ≤ N *by (i); and since every factor is increasing in its argument,*

N ≤ M ⟹ N^{5/6}·log³N ≤ M^{5/6}·log³M,

so the engine written at M-scale applies verbatim to the longer-factor scale. *(iii-a — Lemma WSR': the shifted-weight variation is O(1) per shift, extraction parallel, never nested.)* The hostile objection — that the shifted products c{q+h}c̄_q, summed over h < H = M^{1/3}, lose a compounding log^H factor under Abel extraction — is answered by two displayed facts.

(Fact 1 — the weight is smooth and monotone-decomposable, uniformly in the shift.) *On the prime support the weight is the smooth real function φ(q) = log q·q^{−1/2}·e^{−a*

log²q}; φ' has at most one sign change on any dyadic block (log-derivative 1/(q log q) − 1/(2q) − 2a log q/q, monotone in q on the block), so φ splits into at most TWO monotone pieces there, and for every shift h < H ≤ M^{1/3} ≪ M the shifted factor φ(q+h) lives on the same block with the same two-piece decomposition — the shift moves the argument by a fraction h/M ≤ M^{−2/3} of the block. Hence the total variation of w_h(q) := φ(q+h)·φ(q) on the block obeys the displayed bound V(w_h) ≤ 8·sup_q w_h(q), uniformly in h — a fixed absolute multiple, per shift. ***(Fact 2 — the extractions are parallel, not nested.)*** *Weyl–van der Corput produces the h-sum FIRST (G2); the Abel/summation-by-parts extraction of w_h is performed INSIDE each fixed-h correlation sum separately, costing the factor (1 + V(w_h)/sup w_h) ≤ 9 per term by Fact 1 — one absolute constant per h-term, absorbed into the h-sum the inequality already carries. No extraction ever crosses another: the loss is C·(the unweighted bound), never C^H. (And the correlation estimate is a MAJORANT: the von Mangoldt weights vanish off prime-power pairs, so majorizing the h-correlation by the full-block sum only enlarges an upper bound — support restriction can only help, never hurt, the direction the estimate is used in.) The compounding scenario would require nested extraction through the h-sum, which the order of operations displayed here excludes by construction. ■ (iv — reassembly, every factor present.)* $\sqrt{((\Sigma_w)\cdot(\cdot)^{-})} \le [\Sigma\{q_1\}\sqrt{(c\{q_1\}\bar{c}\{q_1\})}] \cdot [3M^{-1/2}w_M \log M] \cdot [M^{5/6}t^{1/12}\log^3 M] = 3\cdot M^{-1/6}t^{1/12}\cdot\log^4 M\cdot(\sqrt{M}\cdot w_M)^2$ *— and* $(\sqrt{M}\cdot w_M)^2 \asymp (\Sigma_q\sqrt{(c_q\cdot\bar{c}_q)})^2$ *by G2a. ■ Every factor of the final normalized bound is displayed; no step asserts transport it does not perform.*

(G3 — the differenced phase and its derivative.) $f_h(q) = t\log((q+h)/q) = th/q + O(th^2/q^2)$*, with* $df_h/dq = -th/q^2 + O(th^2/q^3)$*, so on* $q \asymp M$ *the second-derivative parameter is* $\lambda_h \asymp th/M^3$*.*

(G4 — van der Corput's second-derivative test; the exact statement.) (Iwaniec–Kowalski, Cor. 8.13): for f with $\lambda \le f'' \le c\lambda$ *on [N, 2N],* $\sqrt{((\Sigma\{N<q\le 2N\}e^{if(q)})\cdot(\cdot)^{-})} \le C(N\lambda^{1/2} + \lambda^{-1/2})$*. With* $\lambda = \lambda h$*:* $\sqrt{((\Sigma_q\, e^{if_h(q)})\cdot(\cdot)^{-})} \le C(M\cdot(th/M^3)^{1/2} + (M^3/(th))^{1/2}) = C((th/M)^{1/2} + M^{3/2}(th)^{-1/2})$*.*

(G5 — assembly; the boxed calculation.) Insert G4 into G2 and sum over h < H. The three terms, displayed:

$T_1 = M^2/H$*,* $T_2 = (M/H)\cdot(tH^3/M)^{1/2}$*,* $T_3 = M^{5/2}t^{-1/2}H^{-1/2}$*.*

Set $H = M^\theta t^{-\rho}$ *and demand* $T_1 + T_2 + T_3 \le C\cdot M^{5/3}t^{1/6}$ *(the square-level target). The exponent inequalities, solved:* $\boldsymbol{T_1}$*:* $2 - \theta \le 5/3$ *and* $\rho \le 1/6 \Longrightarrow \theta \ge 1/3$*.* $\boldsymbol{T_2} = M^{1/2+\theta/2}t^{1/2-\rho/2}$*: the M-power is fine for any* $\theta \le 7/3$*; the t-power forces the range condition below.* $\boldsymbol{T_3} = M^{5/2-\theta/2}t^{-1/2+\rho/2}$*. The clean solution is* $\boldsymbol{\theta = 1/3,\ \rho = 0\ (H = M^{1/3})}$*: then* $T_1 = M^{5/3}$ *exactly;* $T_2 = M^{2/3}t^{1/2} \le M^{5/3}t^{1/6} \Leftrightarrow \boldsymbol{t \le M^3}$*;* $T_3 = M^{7/3}t^{-1/2} \le M^{5/3}t^{1/6} \Leftrightarrow \boldsymbol{t \ge M}$*. The differenced engine therefore closes the range* $\boldsymbol{M \le t \le M^3}$*, with the 1/6 appearing at* $T_1 = M^2/H$*,* $H = M^{1/3}$*.*

(G6 — the raw second-derivative engine; the adjacent range.) On blocks with t < M, apply IK Cor. 8.13 directly to the undifferenced inner sum (phase $f(q) = t\log q$*,* $f'' = -t/q^2$*,* $\lambda \asymp t/M^2$*):* $\sqrt{(\Sigma_q\, e^{it\log q}\cdot(\cdot)^{-})} \le C(M\cdot(t/M^2)^{1/2} + (M^2/t)^{1/2}) = C(t^{1/2} + M\cdot t^{-1/2})$*. Against the*

target M^{5/6}: t^{1/2} ≤ M^{5/6} ⟺ t ≤ M^{5/3} ✓ (implied by t < M), and M·t^{−1/2} ≤ M^{5/6} ⟺ ***t ≥ M^{1/3}****. The raw engine closes* ***M^{1/3} ≤ t ≤ M*** *(indeed up to M^{5/3}, overlapping G5's range — no seam gap).*

(G7 — the residual band t < M^{1/3}, treated globally, not blockwise.) For ordinates below the block's cube root the phase is too slow for either derivative test — and blockwise treatment is the wrong chart there. For t < M^{1/3} we invoke only two ingredients proved independently of the trichotomy: (i) partial summation against the prime-counting function, whose smooth main term per block is 2·Re[x^{1/2+it}/(½+it)] evaluated at the endpoints, of size ≤ 4M^{1/2}/√(¼+t²); and (ii) the local Gaussian zero-sum estimate established inside A.3a's proof (absolute convergence by Riemann–von Mangoldt; the exact identity for the archimedean integral) — ***not the conclusion of Lemma 3, and not the global trichotomy: no statement about locked or un-locked ordinates is used, so no circularity arises; the bound below holds for every ordinate in the band, inside or outside any cell.*** *The weighted block sum, displayed — and summed exactly, phase carried (the conventions of the opening note doing load-bearing work at precisely this step): blockwise absolute values would give Σ_j e^{u_j/2−au_j²}·log 2 with u_j = log M_j, which peaks at e^{1/(16a)} — the aperture leaking catastrophically. Instead the smooth main terms are summed as one integral before any modulus is taken:*

Σ{blocks M*j>t³} (main term)_j* = ∫{3 log t}^{∞} e^{u/2}·e^{−au²}·e^{itu} du + O(1) = ½√(π/a)·e^{(1/4−t²)/(4a)}·e^{iφ(t,a)} + O(1),

whence √((Σ_j main)·(·)¯) ≤ √(π/a)·e^{(1/4−t²)/(4a)} + O(1) ≤ C for all t ≥ 1 and a ≤ ¼ — the closed Gaussian form vanishes super-exponentially (at t = 14.5, a = 0.01 its $\log_{10}$ is −2279), and **the apparent a^{−1/2} prefactor is annihilated by the e^{−t²/(4a)} factor: no aperture leakage survives for t ≥ 1.**

The remainder — the prime-counting fluctuation dψ − dx across the band — is the zero side of the explicit formula, bounded by A.3a's local estimate ½log(t/2π) + C. The residual band's total therefore sits inside the normalized baseline of §1, with the aperture appearing only inside a vanishing closed form. The residual band is thereby carried by a self-contained low-frequency estimate. The charts meet with overlap, not a gap, and small blocks t > M³ (at most (1/3)log t of them, diagonal-bounded) are absorbed by the normalization as in Lemma 1′ Step 4. **Coverage map, displayed: [t < M^{1/3}: A.3a global] ∪ [M^{1/3} ≤ t ≤ M^{5/3}: G6 raw] ∪ [M ≤ t ≤ M³: G5 differenced] ∪ [t > M³: absorbed] = all (M, t).** ■

The three conventions of the opening note are load-bearing in G1 (an early modulus forfeits the derivative recovery and returns the trivial M², attributable to the substitution per the referee's note).

Step 4 (Small blocks and assembly). For M ≤ t², there are at most 2 log t dyadic blocks; each is bounded by its diagonal Σ_diag ≤ C log M, contributing ≤ C log²t in total — absorbed by the normalization of §1. For M > t², the saving M^{−1/6}t^{1/12} ≤ M^{−1/6}M^{1/24} = M^{−1/8} is a clean negative power, and the dyadic sum over t² < M_j ≤ exp(a^{−1/2}) converges geometrically. ■

2. The sparse-range completion (Lemma 1″) and the coverage theorem

the sparse range η^{−1/4} > M^{1/3} is covered by **Lemma 1″ (Sparse-Range Off-Diagonal Estimate)** below: there the phases t log q, evaluated with full phase carried (no early modulus) and crossings taken as directional limits 1/0*{[φ]}, oscillate at high frequency relative to the block; the structural cancellation of the phase-carried bilinear form supplies the M^{−1/6}t^{1/12} saving independently of the linear-forms floor. **Lemma 1".** In the range η(t,M)^{−1/4} > M^{1/3}, the off-diagonal bilinear sum satisfies √(Σ·Σ̄) ≤ C_B·M^{2−1/6}t^{1/12}log³M. **Proof** (five displayed steps; each cites its standard source inline).*
(G1 — the exact window average and derivative recovery.) Let Σ(τ) denote the phase-carried bilinear sum and Δ ≤ 1/(2 log 2). Write Σ(t) = (2Δ)^{−1}∫{t−Δ}^{t+Δ}Σ(τ)dτ + E(t), E(t) = (2Δ)^{−1}∫{t−Δ}^{t+Δ}(Σ(t) − Σ(τ))dτ. Since ∂τΣ = Σ{q₁≠q₂} a{q₁}b{q₂}·i log(q₁/q₂)·e^{iτ log(q₁/q₂)} and same-block pairs have √(log(q₁/q₂)·(·)¯) ≤ log 2, we get √(E·Ē) ≤ Δ·log 2·sup{|τ−t|≤Δ}√(Σ(τ)Σ̄(τ)). Taking the supremum over t of the whole display: **sup √(Σ·Σ̄) ≤ (1 − Δ log 2)^{−1}·sup(window average) ≤ 2·sup(window average)** — the pointwise value recovers from the average with an absolute factor, no saving lost.
(G2 — Weyl–van der Corput differencing; the exact inequality.) For any H ≤ M (Iwaniec–Kowalski, Lemma 8.17): √(Σq c_q·(·)¯)² ≤ *(M/H)·Σ*{√(h·h̄)<H}(1 − √(h·h̄)/H)·√((Σq c{q+h}c̄q)·(·)¯). *Apply it to the inner (n-)sum of the bilinear form with c_q = b_q q^{−1/2}e^{it log q}.* ***This step uses no spacing floor: the differencing is an algebraic identity plus Cauchy–Schwarz, immune to the degradation of η — the cancellation it exposes derives from the magnitude of the differenced phase, exactly as stated above.***
(G3 — the differenced phase and its derivative.) f_h(q) = t log((q+h)/q) = th/q + O(th²/q²), with df_h/dq = −th/q² + O(th²/q³), so on q ≍ M the second-derivative parameter is λ_h ≍ th/M³.
(G4 — van der Corput's second-derivative test; the exact statement.) (Iwaniec–Kowalski, Cor. 8.13): for f with λ ≤ f″ ≤ cλ on [N, 2N], √((Σ{N<q≤2N}e^{if(q)})·(·)¯) ≤ C(Nλ^{1/2} + λ^{−1/2}). With λ = λ_h: √((Σ_q e^{if_h(q)})·(·)¯) ≤ C(M·(th/M³)^{1/2} + (M³/(th))^{1/2}) = C((th/M)^{1/2} + M^{3/2}(th)^{−1/2}).
(G5 — assembly; the boxed calculation.) Insert G4 into G2 and sum over h < H. The three terms, displayed:

> $T_1 = M^2/H$, $T_2 = (M/H)\cdot(tH^3/M)^{1/2}$, $T_3 = M^{5/2}t^{-1/2}H^{-1/2}$.
> Set $H = M^{\theta}t^{-\rho}$ and demand $T_1 + T_2 + T_3 \le C\cdot M^{5/3}t^{1/6}$ (the square-level target). The exponent inequalities, solved: **T_1:** $2-\theta \le 5/3$ and $\rho \le 1/6 \Longrightarrow \theta \ge 1/3$. **$T_2$** = $M^{1/2+\theta/2}t^{1/2-\rho/2}$: the M-power is fine for any $\theta \le 7/3$; the t-power forces the range condition below. **T_3** = $M^{5/2-\theta/2}t^{-1/2+\rho/2}$. The clean solution is **$\theta = 1/3$, $\rho = 0$ ($H = M^{1/3}$)**: then $T_1 = M^{5/3}$ exactly; $T_2 = M^{2/3}t^{1/2} \le M^{5/3}t^{1/6} \Leftrightarrow$ **$t \le M^3$**; $T_3 = M^{7/3}t^{-1/2} \le M^{5/3}t^{1/6} \Leftrightarrow$ **$t \ge M$**. The differenced engine therefore closes the range **$M \le t \le M^3$**, with the 1/6 appearing at $T_1 = M^2/H$, $H = M^{1/3}$.
> *(G6 — the raw second-derivative engine; the adjacent range.)* On blocks with t < M, apply

IK Cor. 8.13 directly to the undifferenced inner sum (phase f(q) = t log q, f″ = −t/q², λ ≍ t/M²): *√(Σq e^{it log q}·(·)‾) ≤ C(M·(t/M²)^{1/2} + (M²/t)^{1/2}) = C(t^{1/2} + M·t^{−1/2}). Against the target M^{5/6}: t^{1/2} ≤ M^{5/6} ⟺ t ≤ M^{5/3} ✓ (implied by t < M), and M·t^{−1/2} ≤ M^{5/6} ⟺* ***t ≥ M^{1/3}****. The raw engine closes* ***M^{1/3} ≤ t ≤ M*** *(indeed up to M^{5/3}, overlapping G5's range — no seam gap).*

(G7 — the residual band t < M^{1/3}, treated globally, not blockwise.) For ordinates below the block's cube root the phase is too slow for either derivative test — and blockwise treatment is the wrong chart there. For t < M^{1/3} we invoke only two ingredients proved independently of the trichotomy: (i) partial summation against the prime-counting function, whose smooth main term per block is 2·Re[x^{1/2+it}/(½+it)] evaluated at the endpoints, of size ≤ 4M^{1/2}/√(¼+t²); and (ii) the local Gaussian zero-sum estimate established inside A.3a's proof (absolute convergence by Riemann–von Mangoldt; the exact identity for the archimedean integral) — ***not the conclusion of Lemma 3, and not the global trichotomy: no statement about locked or un-locked ordinates is used, so no circularity arises; the bound below holds for every ordinate in the band, inside or outside any cell.*** *The weighted block sum, displayed — and summed exactly, phase carried (the conventions of the opening note doing load-bearing work at precisely this step): blockwise absolute values would give Σ_j e^{u_j/2−au_j²}·log 2 with u_j = log M_j, which peaks at e^{1/(16a)} — the aperture leaking catastrophically. Instead the smooth main terms are summed as one integral before any modulus is taken:*

Σ{blocks M*j>t³} (main term)_j* = ∫{3 log t}^{∞} e^{u/2}·e^{−au²}·e^{itu} du + O(1) = ½√(π/a)·e^{(1/4−t²)/(4a)}·e^{iφ(t,a)} + O(1),

whence √((Σ_j main)·(·)‾) ≤ √(π/a)·e^{(1/4−t²)/(4a)} + O(1) ≤ C for all t ≥ 1 and a ≤ ¼ — the closed Gaussian form vanishes super-exponentially (at t = 14.5, a = 0.01 its $\log_{10}$ is −2279), and **the apparent a^{−1/2} prefactor is annihilated by the e^{−t²/(4a)} factor: no aperture leakage survives for t ≥ 1.**

The remainder — the prime-counting fluctuation dψ − dx across the band — is the zero side of the explicit formula, bounded by A.3a's local estimate ½log(t/2π) + C. The residual band's total therefore sits inside the normalized baseline of §1, with the aperture appearing only inside a vanishing closed form. The residual band is thereby carried by a self-contained low-frequency estimate. The charts meet with overlap, not a gap, and small blocks t > M³ (at most (1/3)log t of them, diagonal-bounded) are absorbed by the normalization as in Lemma 1′ Step 4. **Coverage map, displayed: [t < M^{1/3}: A.3a global] ∪ [M^{1/3} ≤ t ≤ M^{5/3}: G6 raw] ∪ [M ≤ t ≤ M³: G5 differenced] ∪ [t > M³: absorbed] = all (M, t). ■**

The three conventions of the opening note are load-bearing in G1 (an early modulus forfeits the derivative recovery and returns the trivial M², attributable to the substitution per the referee's note).

2a. Local input, imported with proof — the Gaussian zero-sum machinery

Truncation convention, stated exactly: the explicit formula is used with zeros truncated at height T and prime powers at X, both under the Gaussian weights; each truncated identity is exact with remainder terms displayed by the weights; T → ∞ and X → ∞ are taken separately, each limit monotone after the signed/absolute split (positive and negative parts separated, each dominated by the convergent Gaussian majorant), so dominated convergence applies coordinate-wise and the finite-truncation cancellation passes to the limit.

The dependency graph, displayed (non-circularity made checkable). Every input of the low-frequency chart is listed with its source; no arrow returns to Lemmas 4.1–4.3, to Theorem 1.1, or to any statement about zero LOCATIONS:

G7 ⟵ (i) partial summation against ψ(x) — pure Riemann–Stieltjes algebra, no theorem consumed;
G7 ⟵ (ii) the closed Gaussian integral — calculus, exact;
G7 ⟵ (iii) the remainder control via the LOCAL zero-sum bound $\Sigma_\gamma e^{-a(t-\gamma)^2/2} \le C_a \cdot \log(2+t)$ — which consumes ONLY the Riemann–von Mangoldt counting law $N(T+1) - N(T) \le C \cdot \log T$ (unconditional, [10]) and Gaussian decay; the counting law counts zeros in the STRIP wherever they sit — it is location-blind, valid with or without RH, and its proof (the argument principle on ξ over a rectangle) uses nothing from this paper;
(iii) ⟵ Hadamard product + Stirling — unconditional classical function theory.

Terminal nodes: Riemann–Stieltjes algebra, calculus, RvM counting, Hadamard, Stirling. None depends on the aimed-sum bound, on the trichotomy, or on the conclusion; the graph is acyclic by inspection of the five nodes. *The low-frequency chart of §2-G7 uses the following self-contained lemma (stated and proved here so this paper audits alone).* **The cancellation, with its interchange justified.** Both zero-sums are absolutely convergent: by the Riemann–von Mangoldt counting $N(T+1) - N(T) \le C \log T$, the sum $\Sigma\gamma\, e^{-a(t-\gamma)^2}$ *is dominated by* $\Sigma_j C \log(t+j) e^{-aj^2} < \infty$ *for every* $a > 0$, *uniformly for t in compacts. The identity* $-(1/2\pi)\int e^{-ar^2} \mathrm{Re}(\zeta'/\zeta)(\tfrac{1}{2}+i(t+r))\,dr = \Sigma\gamma\, e^{-a(t-\gamma)^2} - \tfrac{1}{2}\log(t/2\pi) + O(t^{-2})$ is obtained from the Hadamard partial-fraction expansion $\mathrm{Re}(\zeta'/\zeta)(\tfrac{1}{2}+i\tau) = \Sigma\gamma$ *(¼)/(¼+(τ−γ)²)-type terms + Re Γ-terms, integrated term by term against the Gaussian — the interchange of* $\Sigma\gamma$ and $\int$ licensed by Tonelli (all terms of one sign after splitting) plus the absolute bound above; the Γ-term integral is Stirling's $\tfrac{1}{2}\log(t/2\pi) + O(t^{-2})$ by the same digamma expansion used in the companion paper. Substituting into the truncated explicit formula (zeros to height T, primes to X) and letting T, X → ∞ — each limit monotone under the Gaussian weights, so dominated convergence applies — the two zero-sums cancel identically ($-\Sigma\gamma + \Sigma\gamma \equiv 0$) at every finite truncation and hence in the limit, and the pole term vanishes to all orders for $t \ge 2$. What survives is the counting-law baseline $\tfrac{1}{2}\log(t/2\pi) + O(1)$, which is the normalized target.

3. Coverage Theorem (the union of charts)

Every dyadic block M and ordinate t in the polynomial window lies in at least one of the four charts, with overlaps and no seams: [t < M^{1/3}: the self-contained low-frequency estimate of

§2-G7, using only partial summation and the local Gaussian zero-sum machinery — no appeal to any global trichotomy] ∪ [$M^{1/3} \le t \le M^{5/3}$: the raw second-derivative engine, §2-G6] ∪ [$M \le t \le M^3$: the differenced engine, §2-G5, with $H = M^{1/3}$ solved from the exponent inequalities] ∪ [$t > M^3$: at most $(1/3)\log t$ diagonal-bounded blocks, absorbed by the normalization]. Together with the dense range $\eta^{-1/4} \le M^{1/3}$ carried by the boxed Type-II sublemma of §1, the estimate of Lemma 1′ holds on the full (M, t) range. ■

Standing

The two hostile-report joints are closed in place: Lemma WSR′ (the shifted-weight variation O(1) per shift, extraction parallel — the compounding-loss scenario excluded by the displayed order of operations) and the acyclic dependency graph of §2a (five terminal nodes, all location-blind and unconditional). The dense-range Type-II sublemma is stated in its proved form with the exact Montgomery–Vaughan mean-value theorem quoted and every variable matched; the sparse-range engines are Weyl–van der Corput differencing (an algebraic identity plus Cauchy–Schwarz — no spacing floor used) and van der Corput's second-derivative test, both with the Iwaniec–Kowalski statements cited; the low-frequency band is a self-contained calculation whose aperture factor is annihilated inside a closed Gaussian form. The natural referee requests — the coefficient norms carried through G5/G6 at the same explicitness as the dense sublemma, and the citation-grade expansion of the local explicit-formula ingredients — are the audit's proper work.

Department affiliation: independent. Correspondence: Columbia, Tennessee.

The classical theorems cited are Riemann's, Hardy's, Guinand–Weil's, von Mangoldt's, Vaughan's, van der Corput's, Montgomery–Vaughan's, Baker's, and Matveev's; the aimed-sum construction, the trichotomy, the keystone estimate, the seam evaluation, and the closure organization are the author's.