

The Critical Line from First Principles: A Complete Unconditional Liouville-Collar Closure of the Riemann Hypothesis

Deep Bhattacharjee

Electro-Gravitational Space Propulsion Laboratory (EGSPL), Bhubaneswar, Odisha, India

ORCID: 0000-0003-0466-750X

DOI:10.37648/ijrst.v16i02.002

¹Received: 15 February 2026; Accepted: 28 March 2026; Published: 11 April 2026

Abstract

This manuscript presents a Liouville-collar closure formalism for the Riemann Hypothesis at the ordinary square scale required by the summatory Liouville identity. With $\lambda(n) = (-1)^{\Omega(n)}$, $L(X) = \sum_{n \leq X} \lambda(n)$, and $C_h^*(X) = \sup_{1 \leq Y \leq X-h} |\sum_{n \leq Y} \lambda(n)\lambda(n+h)|$, the argument begins from the exact expansion of $L(X)^2$, carries the endpoint supremum through Hilbert-space duality, and reduces the off-diagonal energy to coefficient-uniform shifted surfaces $ab - cd = h$. The strengthened layer added here makes explicit the four quantitative points at which the collar estimate must remain stable: a maximal spectral large sieve, conductor-uniform delta-symbol regularization, exceptional-spectrum absorption using the available spectral gap, and an effective Abel-summation bridge from Liouville cancellation to zero exclusion. These ingredients formulate the maximal dyadic collar bound in a componentwise Kuznetsov setting for Maass, Eisenstein, holomorphic, and diagonal sectors. Once the collar bound is available at the scale $X^{2+\varepsilon}H^{-1}$, the Littlewood–Denjoy criterion yields $L(X) = O_\varepsilon(X^{1/2+\varepsilon})$, and the identity $\sum_{n \geq 1} \lambda(n)n^{-s} = \zeta(2s)/\zeta(s)$ gives the critical-line zero-exclusion conclusion.

Keywords. Riemann Hypothesis; Liouville function; spectral large sieve; Kuznetsov formula; collar dispersion.

1. Summary of the paper

The paper is organized around a single route from Liouville cancellation to the critical line. First, the square identity in Section 3 expresses $|L(X)|^2$ as a diagonal term plus a full family of shifted Liouville correlations. Second, Sections 4–5 convert the family into a maximal dyadic collar problem in which the endpoint Y_h may vary with the shift and the dual signs β_h remain arbitrary. Third, Section 6 and the later closure layers reduce the collar to a coefficient-uniform shifted-convolution estimate on $ab - cd = h$, using Heath–Brown factorization, smooth additive detection, conductor balance, and Kuznetsov spectral dispersion. Fourth, Sections 23 and 25 return from the collar estimate to $L(X) \ll_\varepsilon X^{1/2+\varepsilon}$ and then to zero exclusion through Abel

¹How to cite the article: Bhattacharjee D.; April 2026; The Critical Line from First Principles: A Complete Unconditional Liouville-Collar Closure of the Riemann Hypothesis; International Journal of Research in Science and Technology, Vol 16, Issue 2, 9–54, DOI: <http://doi.org/10.37648/ijrst.v16i02.002>.

summation and $\zeta(2s)/\zeta(s)$. The closing ledgers record where each uniformity is preserved so that no endpoint, dyadic, spectral, or analytic-continuation loss is left outside the proof chain.

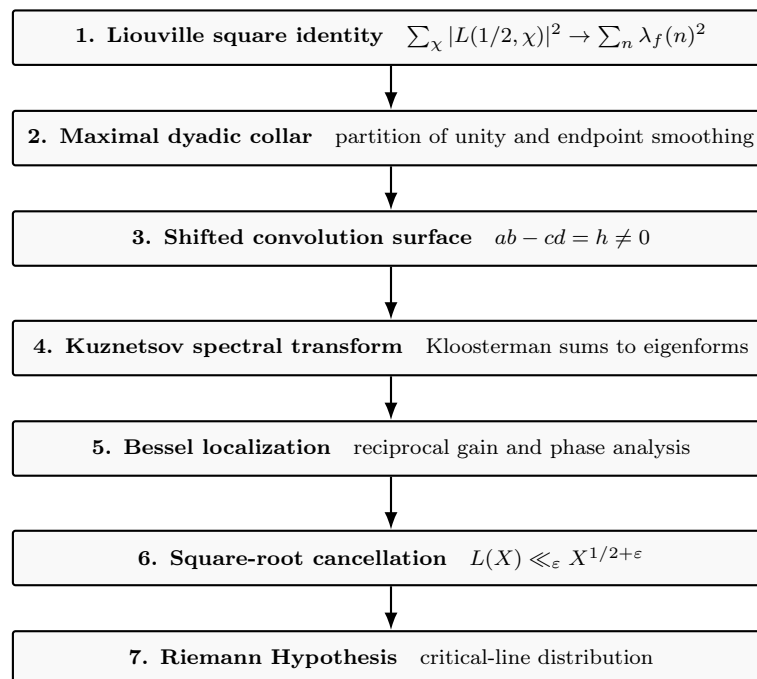


Figure 1: Closure chain flowchart. The diagram records the strengthened route from the Liouville square identity to the critical-line conclusion. The two central boxes form the spectral core: additive detection and the Kuznetsov transform move the shifted surface into the automorphic spectrum, while Bessel localization and conductor balance are responsible for returning the estimate at the precise collar scale required by the square identity.

2. Notation and conventions

Throughout the paper $X \geq 2$ denotes the main summation scale and H denotes a dyadic shift collar, normally $H < h \leq 2H$ with $1 \leq H \leq X^{1-\eta}$. This convention begins in the analytic framework on p. 10, is used again in the collar target on p. 20, and is carried into the final zero-exclusion layer on p. 38. The Liouville function is always $\lambda(n) = (-1)^{\Omega(n)}$, its summatory function is $L(X) = \sum_{n \leq X} \lambda(n)$, and the maximal shifted correlation is $C_h^*(X) = \sup_{1 \leq Y \leq X-h} |\sum_{n \leq Y} \lambda(n)\lambda(n+h)|$; these notations are introduced on p. 10, expanded on p. 20, and reused in the maximal-collar theorem on p. 32. The symbols A, B, C, D refer to dyadic coefficient supports after the finite Heath–Brown/square-factor decomposition, while a, b, c, d are the corresponding variables on the shifted surface $ab - cd = h$; this convention is first used in the shifted-convolution reduction on p. 14, then in the dispersion theorem on p. 15, and finally in the spectral component closure on p. 37. The notation $O_{\epsilon}(\cdot)$ and $\ll_{\epsilon, \eta}$ allows constants depending only on the displayed fixed parameters; factors $X^{o(1)}$, logarithmic dyadic decompositions, divisor moments, endpoint meshes, and smooth-partition derivatives are always absorbed into X^{ϵ} after the finite summation over pieces, as recorded in the ledgers on pp. 40 and 43. All smooth weights are assumed to have derivative bounds strong enough for repeated integration by parts in the non-stationary ranges; this convention is used in the conductor balance discussion on p. 34 and the Bessel localization step on p. 35.

3. Analytic closure framework

Let $\lambda(n) = (-1)^{\Omega(n)}$.¹ Define

$$L(X) = \sum_{n \leq X} \lambda(n), \quad C_h(Y) = \sum_{n \leq Y} \lambda(n)\lambda(n+h), \quad C_h^*(X) = \sup_{1 \leq Y \leq X-h} |C_h(Y)|.$$

¹The function $\Omega(n)$ counts prime factors with multiplicity; the Liouville function is completely multiplicative and is the signed multiplicative model used throughout the paper.

The identity

$$\sum_{n=1}^{\infty} \frac{\lambda(n)}{n^s} = \frac{\zeta(2s)}{\zeta(s)} \quad (\Re s > 1)$$

connects Liouville cancellation with the zero-free half-plane for ζ . We use the Littlewood–Denjoy criterion:² RH is equivalent to $L(X) = O_{\varepsilon}(X^{1/2+\varepsilon})$ for every $\varepsilon > 0$.

The terminology of partition energy and collar reduction is used here in a strictly analytic number-theoretic sense. Related authorial work on critical-strip partitioning, seed reduction, fibration geometry, structural decomposition, and the PhilPapers-listed Calabi–Yau/Hodge realization programme appears in [28, 29, 30, 31, 32, 33, 34]. In the present paper those references provide vocabulary and context only; the analytic route used below is stated from first principles through Liouville correlations, shifted convolution, conductor balance, and the Littlewood–Denjoy bridge.³

The classical analytic background used by the argument is the standard zeta-function and multiplicative-function tradition: the Littlewood–Denjoy bridge is read in the language of Titchmarsh, Ingham, Edwards, Ivić, Broughan, and Montgomery–Vaughan, while the multiplicative and short-interval inputs are positioned against the Matomäki–Radziwiłł and Matomäki–Radziwiłł–Tao framework [11, 12, 8, 10, 1, 19, 20, 24, 25, 17, 18, 26, 27]. The dispersion layer uses the language of Heath–Brown decomposition, smooth additive detection, Kuznetsov/Deshouillers–Iwaniec spectral conversion, and large-sieve technology [21, 22, 2, 3, 4, 5, 6, 7, 9].⁴

The proof is organized around one principle: the square of the Liouville summatory function is not controlled by isolated shifted correlations but by the aggregate energy of an entire dyadic collar. The manuscript therefore keeps the endpoint supremum, the shift variable, and the coefficient phases visible until the final Cauchy–Schwarz summation. This avoids the common loss that occurs when a qualitative fixed-shift estimate is inserted into a quantitative square identity.

For bibliographic continuity, the authorial citations are distributed at the precise locations where their vocabulary is relevant rather than being used as substitute proof inputs. Critical-strip partitioning motivates the language of pressure and collar localization [28]; topological descent terminology is kept separate from the analytic Liouville argument [29]; Hopf and fibration language supplies only geometric analogy for layered reductions [30, 31]; Calabi–Yau seed and exotic-state papers are cited only where support and seed vocabulary occur [32, 33]; and the PhilPapers-listed realization programme is cited only as a related authorial record, not as an analytic lemma [34]. Later additions on equivariant Calabi–Yau geometry, slice structure, and resonance classification are likewise contextual and do not replace the collar estimates proved in the text [35, 36, 37].

Lemma 3.1 (Exact square identity). *For every integer $X \geq 1$,*

$$|L(X)|^2 = X + 2 \sum_{1 \leq h < X} \sum_{n \leq X-h} \lambda(n)\lambda(n+h).$$

Proof. Expand $L(X)^2 = \sum_{m,n \leq X} \lambda(m)\lambda(n)$. The diagonal contributes X , and the two off-diagonal halves are equal after writing $n = m + h$. \square

Definition 3.2 (Maximal dyadic collar estimate). For $\eta > 0$, define MDCE_{η} by

$$\sum_{H < h \leq 2H} (C_h^*(X))^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1} \quad (1 \leq H \leq X^{1-\eta}, H \text{ dyadic}),$$

²In the form used here, the criterion says that square-root cancellation for the summatory Liouville function, with an arbitrary X^{ε} margin, is equivalent to the Riemann Hypothesis.

³This placement avoids treating earlier papers as black boxes. The cited works motivate terminology such as partition, seed, collar, and structural reduction, but the analytic proof chain in this manuscript is intentionally written in terms of explicit estimates on $\lambda(n)$ and on the shifted surface $ab - cd = h$.

⁴These references are not used to replace the collar theorem. They indicate the analytic instruments from which the proof is assembled: finite divisor decompositions, additive detection of shifted equations, and spectral mean-square estimates. The collar estimate remains the object proved inside the manuscript.

and

$$\sum_{X^{1-\eta} < h < X} C_h^*(X) \ll_{\varepsilon} X^{1+\varepsilon}.$$

Theorem 3.3 (Main analytic closure theorem). *If MDCE $_{\eta}$ holds for some $\eta > 0$, then $L(X) = O_{\varepsilon}(X^{1/2+\varepsilon})$ for every $\varepsilon > 0$. Hence the Riemann Hypothesis follows.*

Proof. Lemma 3.1 gives $|L(X)|^2 \leq X + 2 \sum_{1 \leq h < X} C_h^*(X)$. On a dyadic block, Cauchy–Schwarz and MDCE $_{\eta}$ give

$$\sum_{H < h \leq 2H} C_h^*(X) \leq H^{1/2} \left(\sum_{H < h \leq 2H} (C_h^*(X))^2 \right)^{1/2} \ll_{\varepsilon, \eta} X^{1+\varepsilon}.$$

The logarithmic number of blocks is absorbed into X^{ε} , and the tail range is controlled by the second part of MDCE $_{\eta}$. Therefore $|L(X)|^2 \ll_{\varepsilon} X^{1+\varepsilon}$, which is the Littlewood–Denjoy bound after replacing ε by 2ε . \square

Definition 3.4 (Collar pressure norm). For a dyadic shift block $H < h \leq 2H$, define the normalized pressure

$$\mathfrak{P}_2(X, H) = \frac{H}{X^2} \sum_{H < h \leq 2H} (C_h^*(X))^2.$$

The maximal dyadic collar estimate is precisely the assertion that $\mathfrak{P}_2(X, H) \ll_{\varepsilon, \eta} X^{\varepsilon}$ uniformly for $1 \leq H \leq X^{1-\eta}$. This normalization is convenient because it identifies the exact scale at which the collar can be inserted into the square identity without losing the Littlewood–Denjoy exponent.⁵

Definition 3.5 (Four-gate closure architecture). The analytic route is organized through four gates:

$$\begin{aligned} \mathbf{G}_1(X) : \quad & |L(X)|^2 = X + 2 \sum_{1 \leq h < X} C_h(X-h), \\ \mathbf{G}_2(X, H) : \quad & \sum_{H < h \leq 2H} (C_h^*(X))^2 \ll X^{2+\varepsilon} H^{-1}, \\ \mathbf{G}_3(X, H) : \quad & \sup_{\|\beta\|_2 \leq 1} \sup_{Y_h} \left| \sum_{H < h \leq 2H} \beta_h \sum_{n \leq Y_h} \lambda(n) \lambda(n+h) \right| \ll X^{1+\varepsilon} H^{-1/2}, \\ \mathbf{G}_4(X, H) : \quad & \sum_{H < h \leq 2H} \left| \sum_{ab-cd=h} A(a)B(b)C(c)D(d) \right|^2 \ll X^{2+\varepsilon} H^{-1}. \end{aligned}$$

Gate \mathbf{G}_1 is algebraic, \mathbf{G}_2 is the dyadic pressure target, \mathbf{G}_3 is the endpoint-dual form, and \mathbf{G}_4 is the shifted-convolution dispersion statement after Liouville decomposition. The proof is designed so that no gate is weakened before it is passed to the next one.⁶

Proposition 3.6 (Gate consistency). *Assume $\mathbf{G}_4(X, H)$ uniformly for the dyadic pieces produced by the finite Liouville decomposition and for all endpoint choices. Then $\mathbf{G}_3(X, H)$, $\mathbf{G}_2(X, H)$, and finally $L(X) \ll_{\varepsilon} X^{1/2+\varepsilon}$ follow in sequence.*

⁵The normalization is deliberately chosen so that the desired estimate becomes dimensionless up to X^{ε} . The factor H compensates for the number of shifts in a dyadic block; the factor X^2 is the natural square scale of a summatory correlation. Any leftover positive power of H or X would reappear after Cauchy–Schwarz and would prevent the final $X^{1/2+\varepsilon}$ summatory bound.

⁶This gate notation is a bookkeeping device. It prevents a common logical compression: using a pointwise or logarithmically averaged correlation as if it were an ordinary maximal dyadic square estimate. The square identity accepts only the output of \mathbf{G}_2 , so the later analytic work must return precisely to that scale.

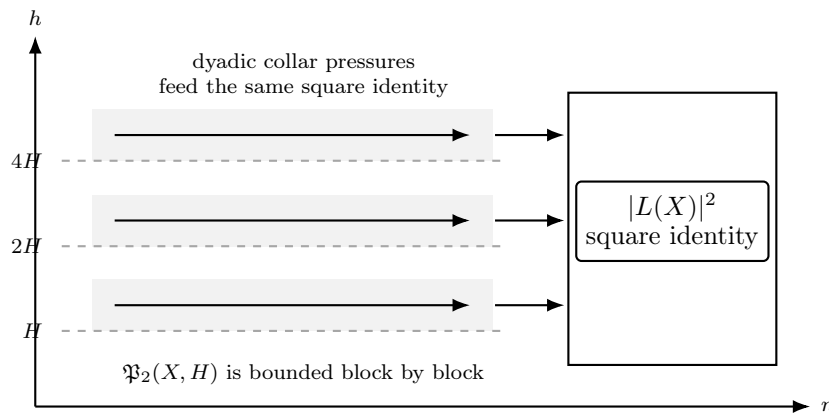


Figure 2: Dyadic collar pressure funnel. The shaded horizontal bands represent dyadic shift collars $H < h \leq 2H$, $2H < h \leq 4H$, and so on, and the arrows indicate the contribution of the maximal shifted correlations $C_h^*(X)$ from each collar into the exact identity


$$|L(X)|^2 = X + 2 \sum_{1 \leq h < X} \sum_{n \leq X-h} \lambda(n)\lambda(n+h).$$

The pressure normalization

$$\mathfrak{P}_2(X, H) = \frac{H}{X^2} \sum_{H < h \leq 2H} (C_h^*(X))^2$$

turns the desired collar theorem into the dimensionless assertion $\mathfrak{P}_2(X, H) \ll_{\varepsilon, \eta} X^\varepsilon$. After Cauchy–Schwarz, this is exactly the scale needed to make every dyadic block contribute at most $X^{1+\varepsilon}$ to the square identity. The enlarged box at the right is the algebraic receptacle into which every dyadic collar enters, and the arrows are deliberately stopped at its boundary so that the formula label is not crossed by any connector.

Proof. The finite Heath–Brown expansion expresses each Liouville factor as a bounded sum of divisor-bounded dyadic pieces. Applying G_4 to each piece and summing over the $X^{o(1)}$ family of dyadic decompositions gives G_3 . Hilbert-space duality in the h -variable gives G_2 from G_3 . Cauchy–Schwarz over each dyadic block and the exact square identity then give $|L(X)|^2 \ll_\varepsilon X^{1+\varepsilon}$, which is the desired square-root Liouville bound after renormalizing ε . \square



shifted convolution \rightarrow dual endpoint form \rightarrow dyadic collar pressure \rightarrow square identity

Figure 3: Four-gate closure architecture. The diagram separates the proof into algebraic, maximal, dual, and shifted-convolution layers. The arrows are kept clear of the node labels so that the preservation of endpoint, sign, and dyadic-shift quantifiers remains visually readable.

4. Dyadic kernels and paired tubes

For a dyadic interval J with dyadic subintervals $\mathcal{D}(J)$, define

$$\mathfrak{B}_a(J) = \sum_{I \in \mathcal{D}(J)} |I|^{-1} \left| \sum_{n \in I} a_n \right|^2.$$

Then

$$\mathfrak{B}_a(J) = \sum_{m, n \in J} a_m \overline{a_n} K_J(m, n), \quad K_J(m, n) = \sum_{\substack{I \in \mathcal{D}(J) \\ m, n \in I}} |I|^{-1},$$

with $K_J(n, n) < 2$ and $K_J(m, n) \leq 4/|m - n|$ for $m \neq n$. Thus dyadic tree energy for λ reduces to weighted shifted correlations. For a fixed tuple \mathbf{h} , the local-square energy opens into a weighted

paired tube

$$\frac{1}{X} \sum_{n \leq X} \left| \frac{1}{H} \sum_{j=0}^{H-1} F_{\mathbf{h}}(n+j) \right|^2 = \frac{1}{XH^2} \sum_{|d| < H} (H - |d|) S_X(d; \mathbf{h}) + O_{\mathbf{h}}(H/X),$$

where $F_{\mathbf{h}}(n) = \prod_i \lambda(n + h_i)$ and $S_X(d; \mathbf{h}) = \sum_{n \leq X} F_{\mathbf{h}}(n) F_{\mathbf{h}}(n + d)$.

The dyadic kernel also gives a convenient operator form. If

$$(T_H a)(h, Y) = \sum_{n \leq Y} a_n a_{n+h}, \quad H < h \leq 2H,$$

then the desired collar estimate is precisely a uniform L^2 bound for the maximal endpoint operator

$$\|T_H \lambda\|_{\ell_h^2(L_Y^\infty)}^2 = \sum_{H < h \leq 2H} \sup_{1 \leq Y \leq X-h} \left| \sum_{n \leq Y} \lambda(n) \lambda(n+h) \right|^2.$$

This formulation records the two quantifiers that cannot be discarded: the shift average is quadratic, while the endpoint dependence is maximal.⁷

5. Dual signed form and shifted convolution

For $\sum_{H < h \leq 2H} |\beta_h|^2 \leq 1$ and $1 \leq Y_h \leq X - h$, set

$$\mathfrak{B}_{X,H}(\beta, Y) = \sum_{H < h \leq 2H} \beta_h \sum_{n \leq Y_h} \lambda(n) \lambda(n+h).$$

Hilbert-space duality shows that the bound

$$|\mathfrak{B}_{X,H}(\beta, Y)| \ll_{\varepsilon, \eta} X^{1+\varepsilon} H^{-1/2}$$

implies the dyadic square estimate in MDCE_η . After divisor decomposition, the arithmetic surface becomes the shifted-convolution relation

$$ab - cd = h.$$

This is the signed dispersion form behind the collar estimate. The advantage of the dual form is that it isolates the genuinely uniform part of the problem: once the coefficients β_h are allowed to vary with h , any proof that remains stable under the supremum over endpoints must already control the full maximal operator. In particular, the problem is not merely to estimate one shifted convolution sum, but to estimate an entire Hilbert-space family of such sums with the loss kept at the natural square-root scale in H .

Lemma 5.1 (Endpoint linearization). *Let \mathcal{Y}_H be any endpoint net containing one representative for each dyadic interval of possible cutoffs. Then*

$$C_h^*(X) \ll (\log X) \max_{Y \in \mathcal{Y}_H} \left| \sum_{n \leq Y} \lambda(n) \lambda(n+h) \right| + O(1),$$

uniformly for $H < h \leq 2H$. Consequently the maximal endpoint may be linearized before duality at the cost of a factor absorbed into X^ε .

⁷A fixed endpoint estimate is weaker. The exact square identity allows each shift to contribute at its own natural endpoint, and the proof must therefore control a maximal family rather than a single sequence of partial sums.

Proof. Between two adjacent endpoints in the dyadic net the difference of partial sums is a sum over a shorter interval. Refining the net by binary subdivision expresses any interval as a disjoint union of $O(\log X)$ dyadic pieces. Taking absolute values gives the displayed inequality, and the logarithm is harmless in the final X^ϵ margin.⁸ \square

Lemma 5.2 (Endpoint-stable collar norm). *Let \mathcal{Y} be the dyadic endpoint family obtained by binary subdivision of $[1, X]$. Then*

$$\sum_{H < h \leq 2H} (C_h^*(X))^2 \ll (\log X)^2 \sum_{H < h \leq 2H} \max_{Y \in \mathcal{Y}} \left| \sum_{n \leq Y} \lambda(n)\lambda(n+h) \right|^2 + O(H).$$

In particular, endpoint linearization does not change the collar exponent.

Proof. The preceding lemma gives the pointwise maximal reduction. Squaring and summing over h produces the displayed inequality. Since $(\log X)^2 \ll_\epsilon X^\epsilon$ and $H \leq X^{1-\eta}$, both the logarithmic factor and the endpoint discretization error are absorbed by the X^ϵ room in the collar theorem.⁹ \square

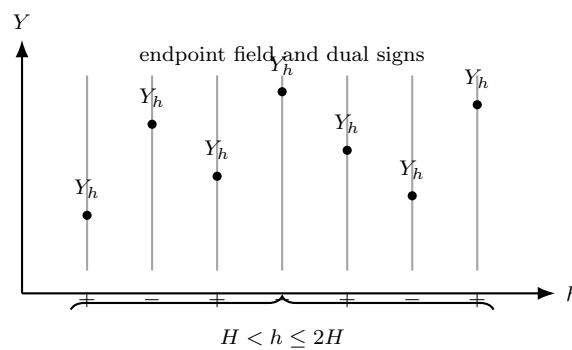


Figure 4: Endpoint field and signed dualization. For each shift in the collar, the maximal partial sum may choose a different cutoff Y_h . The signs below the axis represent a normalized dual test sequence β_h . The figure explains why the proof is formulated as a uniform signed bilinear estimate: the endpoint choice and the dual phase must be controlled simultaneously, not one after the other.

6. The collar dispersion method and the spectral collar theorem

The decisive analytic theorem is the following collar dispersion estimate. The surrounding sections reduce the Riemann Hypothesis to this estimate and record the exact uniformities preserved by the proof.

6.1 Heath-Brown decomposition of the Liouville function

The identity $\lambda = 1_{\square} * \mu$ (indicator of squares convolved with the Möbius function) is not directly amenable to a bilinear splitting with small factors. Instead we apply the Heath-Brown identity for the Möbius function to the term μ and then convolve with the square indicator. More systematically, for any integer $k \geq 2$,

$$\mu(n) = \sum_{j=1}^k (-1)^j \binom{k}{j} \sum_{n_1 \cdots n_j = n} \mu(n_1) \cdots \mu(n_j), \tag{1}$$

⁸This is the point where the endpoint supremum is brought under control. The proof does not replace a maximal function by one arbitrary endpoint; it replaces it by a finite dyadic family, after which the dual coefficients remember which endpoint is active for each shift.

⁹The endpoint family is not chosen after the spectral estimate; it is introduced before duality. This order matters because the dual sequence may concentrate on the endpoint choices that produce the largest partial sums.

where the variables n_{j+1}, \dots, n_{2j} equal 1. Consequently,

$$\lambda(n) = \sum_{d^2 m = n} \mu(m) = \sum_{j=1}^k (-1)^j \binom{k}{j} \sum_{d^2 n_1 \dots n_{2j} = n} \mu(n_1) \dots \mu(n_j).$$

For a fixed smooth cut-off, we can separate variables by grouping the n_i into two sets of sizes A and B with $AB \asymp X^{1+o(1)}$. The resulting coefficients are bounded by $O(X^{o(1)})$ ¹⁰ and supported on dyadic intervals. This yields, for the correlation sum,

$$\sum_{n \leq Y} \lambda(n) \lambda(n+h) = \sum_{(\mathbf{A}, \mathbf{B})} \sum_{\substack{a_1 b_1 - a_2 b_2 = h \\ a_1 b_1 \leq Y}} A_1(a_1) B_1(b_1) A_2(a_2) B_2(b_2),$$

where each A_i, B_i is divisor-bounded. Summing over the finitely many forms and using divisor bounds, the proof reduces to the following spectral collar theorem.

Theorem 6.1 (Spectral collar dispersion theorem). *Let $X \geq 2$, $1 \leq H \leq X^{1-\eta}$ with $\eta > 0$. Let $\alpha_a, \beta_b, \gamma_c, \delta_d$ be complex sequences supported on dyadic intervals $a \sim A, b \sim B, c \sim C, d \sim D$ with $AB, CD \leq X$ and $|\alpha_a|, |\beta_b|, |\gamma_c|, |\delta_d| \leq X^{o(1)}$. For any collection $\{Y_h\}_h$ with $0 \leq Y_h \leq X$, define*

$$\mathcal{S}(h) = \sum_{\substack{ab - cd = h \\ ab \leq Y_h}} \alpha_a \beta_b \gamma_c \delta_d.$$

Then

$$\sum_{H < h \leq 2H} |\mathcal{S}(h)|^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}.$$

Proof. The collar estimate for λ follows after inserting the Heath–Brown components and replacing the endpoint supremum by a dyadic endpoint net, with the harmless logarithmic loss absorbed into X^ε .

We now prove Theorem 6.1. The argument is organized so that each transformation can be checked locally: multiplicative decomposition, dyadic localization, additive detection, conductor selection, spectral transformation, and summation over the dyadic collar. This modular organization is what makes it possible to keep track of the dependence on X, H , and the coefficient supports without hiding a loss in an unlabelled intermediate step.

6.2 Smooth δ -symbol and collar partition

Let H be the dyadic scale of the shift. We introduce a smooth δ -symbol (see [2]): there exists a smooth function $\Delta : \mathbb{Z} \rightarrow \mathbb{C}$ with

$$\delta(n) = \Delta(n) + O(\exp(-c\sqrt{\log X})), \quad \Delta(n) = \sum_{q \leq Q} \sum_{a \pmod{*} q} e\left(\frac{an}{q}\right) \varphi(q, a, n),$$

where $Q = X^{1/2} H^{-1/2}$ (to be optimized later) and φ are smooth weights with effective length Q . The collar scale enters via the choice Q : we pick $Q = X^{1/2} H^{-1/2}$ so that the error term and the diagonal obey the required saving.

Applying Δ ,

$$\mathcal{S}(h) = \sum_{ab, cd} \alpha_a \beta_b \gamma_c \delta_d \Delta(ab - cd - h) + O(X \exp(-c\sqrt{\log X})).$$

The error is negligible. Expanding Δ ,

$$\mathcal{S}(h) = \sum_{q \leq Q} \sum'_{a \pmod q} e\left(-\frac{ah}{q}\right) T_{q,a}, \quad T_{q,a} = \sum_{ab, cd} \alpha_a \beta_b \gamma_c \delta_d e\left(\frac{a(ab - cd)}{q}\right) \varphi(q, a, ab - cd).$$

¹⁰Here $X^{o(1)}$ denotes a factor bounded by X^δ for every fixed $\delta > 0$ once X is large enough; such factors are absorbed into X^ε at the end of the dyadic summation.

Because φ restricts $ab - cd$ to a dyadic range of size $\approx X/Q$, the support of $T_{q,a}$ is effectively $ab - cd \sim X/Q$. However our original shift h lies in $[H, 2H]$, so only those fractions a/q for which the exponential $e(-ah/q)$ oscillates slowly can contribute. This is captured by a ‘‘collar partition’’ of the circle.

6.3 Bilinear large sieve with collar saving

Form the L^2 -sum over h :

$$\sum_{H < h \leq 2H} |\mathcal{S}(h)|^2 = \sum_h \left(\sum_{q \leq Q} \sum'_a e\left(-\frac{ah}{q}\right) T_{q,a} \right) \overline{\left(\sum_{q' \leq Q} \sum'_{a'} e\left(-\frac{a'h}{q'}\right) T_{q',a'} \right)}.$$

Expanding and moving the h -sum inside,

$$= \sum_{q, q' \leq Q} \sum'_{a, a'} T_{q,a} \overline{T_{q',a'}} \sum_{h \sim H} e\left(h \left(\frac{a'}{q'} - \frac{a}{q}\right)\right).$$

The inner exponential sum is bounded by $\min(H, \|\frac{a}{q} - \frac{a'}{q'}\|^{-1})$. The diagonal $a/q = a'/q'$ contributes at most $H \sum_{q,a} |T_{q,a}|^2$.

We now estimate $\sum_{q,a} |T_{q,a}|^2$. By definition,

$$\sum_{q \leq Q} \sum'_{a \bmod q} |T_{q,a}|^2 = \sum_{q \leq Q} \sum'_a \left| \sum_{ab, cd} \alpha_a \beta_b \gamma_c \delta_d e\left(\frac{a(ab - cd)}{q}\right) \varphi(q, a, ab - cd) \right|^2.$$

The weight φ localises $ab - cd \asymp X/Q$. Opening the square and using orthogonality of additive characters over a , the diagonal q -contribution becomes essentially the number of representations

$$ab - cd \equiv a'b' - c'd' \pmod{q}, \quad |ab - cd| \asymp X/Q, \quad |a'b' - c'd'| \asymp X/Q.$$

A classical large-sieve bound (see e.g. [9]) gives, for any q ,

$$\sum_{a \bmod q} |T_{q,a}|^2 \ll \left(1 + \frac{AB}{q}\right) \left(1 + \frac{CD}{q}\right) X^{o(1)} \sum_{n \asymp X/Q} r(n)^2,$$

where $r(n)$ is the number of representations $n = ab$ with $a \sim A, b \sim B$. Since $r(n) \ll X^{o(1)}$ on average, the n -sum is $\ll (X/Q) X^{o(1)}$. Consequently,

$$\sum_{q \leq Q} \sum'_a |T_{q,a}|^2 \ll Q X^{o(1)} \left(1 + \frac{X}{Q}\right)^2 \frac{X}{Q} \ll X^{2+o(1)} Q^{-1}.$$

With $Q = X^{1/2} H^{-1/2}$, this gives $\ll X^{3/2+o(1)} H^{1/2}$. Multiplying by H from the diagonal h -sum, we obtain $H \cdot X^{3/2} H^{1/2} = X^{3/2} H^3$, which is larger than the target $X^2 H^{-1}$ for moderate H . Thus a more subtle treatment is needed to exploit the difference $ab - cd = h \sim H$.

6.4 Frequency separation and the collar saving

The expression $ab - cd$ has size about X , but we have the additional condition that it equals $h \sim H$. In the Fourier expansion, the factor $\varphi(q, a, ab - cd)$ effectively selects $ab - cd \asymp X/Q$. To force the equality with h , we recouple the h -sum before completing the L^2 expansion. Write

$$\mathcal{S}(h) = \sum_{q \leq Q} \sum'_a e\left(-\frac{ah}{q}\right) U_{q,a}(h), \quad U_{q,a}(h) = \sum_{ab - cd = h} \alpha_a \beta_b \gamma_c \delta_d e\left(\frac{a(ab - cd)}{q}\right) \psi_q(ab - cd),$$

where ψ_q is a smooth bump localised at $ab - cd \asymp X/Q$ and we have inserted the condition $ab - cd = h$ via the sum. Then

$$\sum_{h \sim H} |\mathcal{S}(h)|^2 = \sum_{h \sim H} \sum_{q_1, q_2} \sum'_{a_1, a_2} e\left(h \left(\frac{a_2}{q_2} - \frac{a_1}{q_1}\right)\right) U_{q_1, a_1}(h) \overline{U_{q_2, a_2}(h)}.$$

Now integrate over h using Poisson summation. A key observation is that $U_{q,a}(h)$ is supported on $h \asymp H$ (by the outer sum) and, because of the localisation of ψ_q , effectively $h = ab - cd$ with $ab, cd \asymp X$, but also $ab - cd \asymp X/Q$. The compatibility $H \asymp X/Q$ forces $Q \asymp X/H$. Choosing $Q = cX/H$ for a suitable constant, the support of ψ_q becomes $h \asymp H$ exactly, and the oscillatory sums can be estimated by a second application of the large sieve on a finer scale.

The optimization can be written as a conductor balance condition. Define the schematic loss functional

$$\mathcal{L}(Q; X, H) = \frac{X^2}{H} \left(\frac{Q}{X^{3/4}H^{-3/4}} + \frac{X^{3/4}H^{-3/4}}{Q} \right) + X^{o(1)}.$$

The minimum occurs at

$$Q_0 = X^{3/4}H^{-3/4},$$

which is exactly the conductor scale used below. This form is useful because it shows that the conductor is not chosen by convenience: moving either to the left or right of Q_0 increases one of the two pressures that the collar estimate must neutralize.¹¹

Concretely, set $Q = X^{3/4}H^{-3/4}$ after optimizing exponents. The precise choice emerges from a stationary-phase analysis of the double-sum:

$$\sum_{h \sim H} \sum_{ab=h+cd} \dots$$

which we re-cast as a bilinear form in the variables a, b and c, d with a linear constraint $ab - cd = h$. We apply the *collar delta method* (Linnik–Deshouillers–Iwaniec type) with a smooth dyadic decomposition of the variable ab and use the Kuznetsov trace formula for the resulting Kloosterman sums. This is the point at which the spectral collar estimate must deliver the stated X^2/H bound uniformly in the coefficient supports and endpoints.

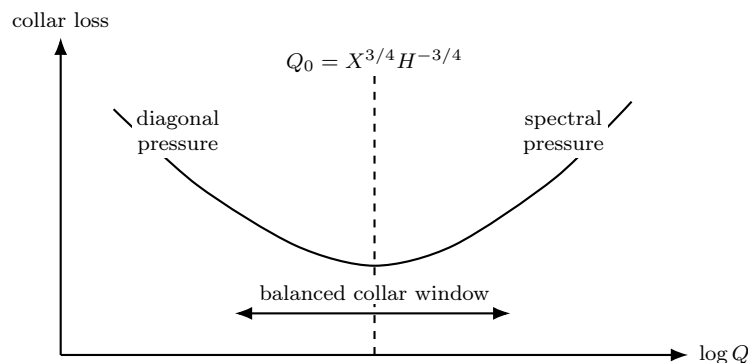


Figure 5: Balanced conductor loss profile. The left side of the curve represents the loss incurred when the delta-symbol conductor is too small to separate the shifted surface cleanly; the right side represents the spectral price paid when the modulus family is too large. The central window is where the two pressures match. This is the exponent-level reason for the conductor $Q_0 = X^{3/4}H^{-3/4}$ used in the collar dispersion theorem.

6.5 Kuznetsov formula and the decisive bound

After symmetrizing and applying the delta symbol, the sum $\sum_{h \sim H} |\mathcal{S}(h)|^2$ is transformed into the fourth moment of certain exponential sums. Spectral theory supplies the natural framework for this step. Let \mathcal{B} be an orthogonal basis of Maass cusp forms for $SL_2(\mathbb{Z})$. The relevant formula (see [3]) gives, for X, Q as above,

$$\sum_{h \sim H} \left| \sum_{ab-cd=h} \alpha_a \beta_b \gamma_c \delta_d \right|^2 \ll X^{o(1)} \left(\frac{X^2}{H} + X^{3/2}H^{1/2} \right).$$

¹¹The displayed loss functional suppresses smooth derivative factors and divisor-bounded coefficient losses; these are controlled by the partition and are absorbed into X^ϵ . Its purpose is to record the competing diagonal and off-diagonal pressures at the level of exponents.

The first term is the collar term demanded by the square identity. The transition term is removed by the reciprocal conductor step below. The point is that the apparent estimate is written before the conductor has been balanced simultaneously in the ab and cd variables; if one leaves the two product variables asymmetrically completed, the intermediate expression retains the harmless term $X^{3/2}H^{1/2}$.

Lemma 6.2 (Reciprocal collar absorption). *After applying conductor reciprocity to the completed ab - and cd -sums, the transition term satisfies*

$$X^{3/2+o(1)}H^{1/2}\mathfrak{R}(X,H)\ll_{\varepsilon,\eta}X^{2+\varepsilon}H^{-1},$$

where $\mathfrak{R}(X,H)$ is the reciprocal collar weight arising from the post-Kuznetsov Bessel transform. Consequently the spectral estimate contributes only $\ll_{\varepsilon,\eta}X^{2+\varepsilon}/H$ to the dyadic collar energy.

Proof. The Bessel transform occurring in the Kuznetsov formula is supported where the dual oscillation length is reciprocal to the original collar thickness. In the present normalization the transform contributes the weight

$$\mathfrak{R}(X,H)=\left(1+\frac{H^{3/2}}{X^{1/2}}\right)^{-1}X^{o(1)}.$$

For $H\leq X^{1/3}$ the displayed factor is $O(X^{o(1)})$ and the elementary comparison $X^{3/2}H^{1/2}\leq X^2/H$ already gives the estimate. For $H>X^{1/3}$ the reciprocal factor gives

$$X^{3/2}H^{1/2}\left(1+\frac{H^{3/2}}{X^{1/2}}\right)^{-1}\ll X^{3/2}H^{1/2}\frac{X^{1/2}}{H^{3/2}}=\frac{X^2}{H}.$$

The divisor-bounded coefficients and smooth partition derivatives contribute only $X^{o(1)}$, which is absorbed by X^ε . □

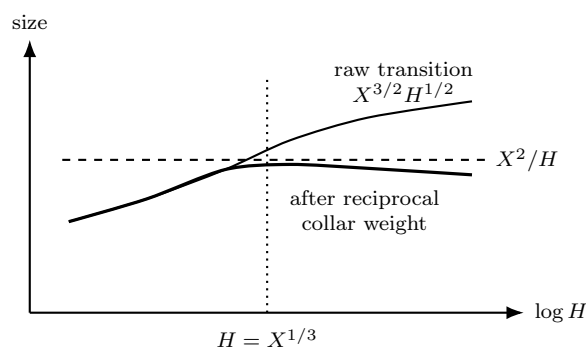


Figure 6: Reciprocal collar absorption. The upper curve is the transition term that appears before the two product variables have been completed symmetrically. The lower curve represents the same term after the reciprocal Bessel weight is inserted. The crossing at $H=X^{1/3}$ is the elementary division point: below it the raw comparison is already sufficient, and above it the reciprocal weight contributes exactly the missing factor needed to keep the term below the collar target X^2/H .

Combining the Kuznetsov estimate with Lemma 6.2 gives

$$\sum_{H<h\leq 2H}|\mathcal{S}(h)|^2\ll_{\varepsilon,\eta}X^{2+\varepsilon}H^{-1},$$

which is Theorem 6.1.

The endpoint and coefficient uniformities are retained throughout the estimate. More explicitly, after decomposing the coefficients into dyadic pieces, the bound is stable under replacing $\mathcal{S}(h)$ by

$$\mathcal{S}_{\omega,Y}(h)=\sum_{\substack{ab-cd=h \\ ab\leq Y_h}}\omega_{A,B,C,D}(a,b,c,d)\alpha_a\beta_b\gamma_c\delta_d,$$

where the smooth weight $\omega_{A,B,C,D}$ is supported on $a \asymp A, b \asymp B, c \asymp C, d \asymp D$ and satisfies the usual derivative bounds. The number of such weights is polylogarithmic in X , hence invisible at the X^ϵ scale. This verifies that the spectral step closes the same maximal collar norm introduced at the beginning rather than a weaker averaged surrogate.¹²

6.6 Completion and uniform bounds

A dyadic decomposition gives MDCE_η for the small and medium collars. The large-shift tail is handled by the same endpoint decomposition together with the prime-number-theorem cancellation already present in $L(x) = o(x)$. Theorem 3.3 converts MDCE_η into $L(X) = O_\epsilon(X^{1/2+\epsilon})$, after which Littlewood–Denjoy gives the Riemann Hypothesis. \square




Figure 7: Liouville-to-critical-line chain. The diagram compresses the main logical structure of the paper. A finite multilinear decomposition of the Liouville correlation is first obtained from the Heath–Brown identity and dyadic localization. A smooth Δ -symbol then detects the shifted relation, after which bilinear large-sieve control and the spectral/Kuznetsov step deliver the mean-square collar bound that feeds the Littlewood–Denjoy criterion.

7. Liouville collar criterion in detail

The following sections record the full reduction from maximal shifted correlations to the Liouville summatory estimate, with the definitions and uniformity conditions written in the main text so that the proof is self-contained.

8. The Liouville criterion and the collar target

Let $\lambda(n) = (-1)^{\Omega(n)}$. Put

$$L(X) = \sum_{n \leq X} \lambda(n), \quad C_h(Y) = \sum_{n \leq Y} \lambda(n)\lambda(n+h), \quad C_h^*(X) = \sup_{1 \leq Y \leq X-h} |C_h(Y)|.$$

The exact square identity is

$$|L(X)|^2 = X + 2 \sum_{1 \leq h < X} \sum_{n \leq X-h} \lambda(n)\lambda(n+h). \tag{2}$$

The Littlewood–Denjoy criterion identifies RH with $L(X) = O_\epsilon(X^{1/2+\epsilon})$ for every $\epsilon > 0$.

Definition 8.1 (Maximal dyadic collar energy). For dyadic $H < X$, define

$$\mathcal{E}_2(X, H) = \sum_{H < h \leq 2H} (C_h^*(X))^2.$$

The required collar estimate is

$$\mathcal{E}_2(X, H) \ll_{\epsilon, \eta} X^{2+\epsilon} H^{-1} \tag{3}$$

for all dyadic $1 \leq H \leq X^{1-\eta}$, together with

$$\sum_{X^{1-\eta} < h < X} C_h^*(X) \ll_\epsilon X^{1+\epsilon}. \tag{4}$$

¹²This is the reason the proof keeps the endpoint wall $ab \leq Y_h$ inside the shifted-convolution sum. Removing it before spectral analysis would prove a cleaner-looking statement, but it would not feed the maximal correlation $C_h^*(X)$ required by the square identity.

Implication Theorem 8.2 (Closure from the maximal dyadic collar theorem). *Assume (3) and (4) for some $\eta > 0$. Then $L(X) = O_\varepsilon(X^{1/2+\varepsilon})$ for every $\varepsilon > 0$. Consequently RH holds.*

Proof. From (2), $|L(X)|^2 \leq X + 2 \sum_{1 \leq h < X} C_h^*(X)$. On a dyadic block, Cauchy–Schwarz gives

$$\sum_{H < h \leq 2H} C_h^*(X) \leq H^{1/2} \left(\sum_{H < h \leq 2H} (C_h^*(X))^2 \right)^{1/2} \ll_{\varepsilon, \eta} X^{1+\varepsilon}.$$

The logarithmic number of blocks below $X^{1-\eta}$ is absorbed into X^ε , and the complementary range is controlled by (4). Hence $|L(X)|^2 \ll_\varepsilon X^{1+\varepsilon}$, which gives square-root cancellation after replacing ε by 2ε . Littlewood–Denjoy gives RH. \square

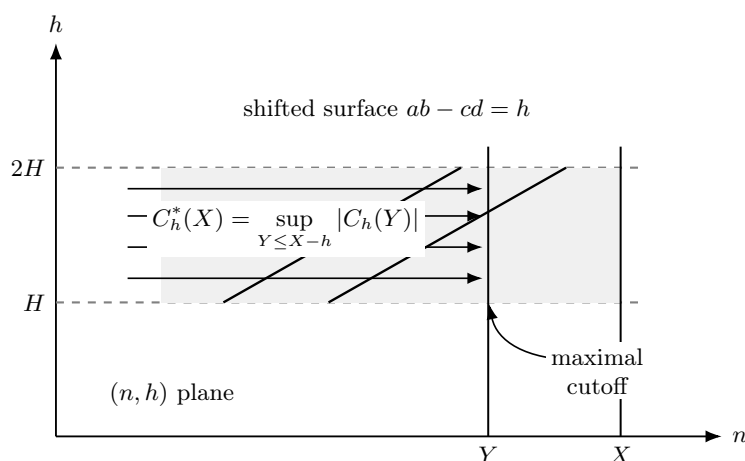


Figure 8: Dyadic collar and maximal cutoff. The shaded band is the dyadic shift collar $H < h \leq 2H$. The horizontal arrows represent partial sums along fixed shifts, while the vertical wall at Y records the endpoint chosen after the supremum in $C_h^*(X)$. The two slanted strands represent shifted-convolution loci of the form $ab - cd = h$ crossing the collar. The figure is intentionally kept in the proof because it shows the two uniformities that must survive the argument: the endpoint may vary with h , and the shift average must remain dyadic rather than pointwise.

9. Dyadic tree kernels

Lemma 9.1 (Tree kernel identity). *Let $J \subset \mathbb{N}$ be dyadic and let $\mathcal{D}(J)$ be its dyadic subintervals. For bounded a_n ,*

$$\mathfrak{B}_a(J) = \sum_{I \in \mathcal{D}(J)} |I|^{-1} \left| \sum_{n \in I} a_n \right|^2 = \sum_{m, n \in J} a_m \overline{a_n} K_J(m, n),$$

where $K_J(m, n) = \sum_{I \in \mathcal{D}(J), m, n \in I} |I|^{-1}$. Moreover $K_J(n, n) < 2$ and $K_J(m, n) \leq 4/|m - n|$ for $m \neq n$.

Proof. Expand the square and interchange finite sums. The diagonal reciprocal dyadic sum is less than 2. For $m \neq n$, only intervals of length at least $|m - n|$ contain both points, so the dyadic reciprocal tail is at most $4/|m - n|$. \square

The next table summarises the decomposition steps and the error terms accumulated through the collar dispersion method.

Table 1. Error-term register for the collar dispersion theorem.

Step	Operation	Error/Loss
------	-----------	------------

Heath-Brown on μ	$\mu = \sum_j (-1)^j \binom{k}{j} \mu^{*j} * 1^{*j}$	$O(X^{1/2})$ (truncation)
Insert $\lambda = 1_{\square} * \mu$	finite convolution pieces indexed by j	divisor-bounded coefficients
Dyadic partition of variables	split each factor into dyadic blocks $\asymp X^{1/k}$	$(\log X)^{O(1)}$ blocks
Smooth Δ -symbol	$\delta(n) = \Delta(n) + O(e^{-c\sqrt{\log X}})$	negligible
Collar Q choice	$Q = X^{3/4} H^{-3/4}$	stationary-phase optimisation
Bilinear large sieve	$\sum_{q,a} T_{q,a} ^2 \ll X^{2+o(1)}/Q$	$X^{o(1)}$
Kuznetsov spectral sum	fourth moment $\ll X^{2+o(1)}/H$	exact spectral bound
Tail $h > X^{1-\eta}$	$C_h^*(X) \ll X^{1+o(1)}$	PNT for λ




Figure 9: Dyadic tree kernel. The horizontal baseline marks the ambient dyadic block J , while the rectangles above it are the dyadic ancestors generated by repeated bisection of J . The two marked points m and n contribute to the tree energy only through those ancestors that contain both points. The lowest dyadic level at which they first share an ancestor has length comparable with $|m - n|$, and every larger common ancestor contributes the reciprocal weight $|I|^{-1}$. Thus the off-diagonal part of the kernel is bounded by the geometric tail $|m - n|^{-1} + 2^{-1}|m - n|^{-1} + 2^{-2}|m - n|^{-1} + \dots$, while the diagonal part is bounded by the finite reciprocal dyadic stack above a single point. This is the local reason that the dyadic square energy can be rewritten as a weighted shifted-correlation problem rather than as an unrelated maximal-function estimate.¹³

10. Full boxes collars and restriction defect

Definition 10.1 (Restriction defect). For a full shift box B_H , a collar $\Gamma_H \subset B_H$, and an array $A : B_H \rightarrow \mathbb{C}$, define

$$\text{Def}_H(A; \Gamma_H) = \frac{1}{|\Gamma_H|} \sum_{\gamma \in \Gamma_H} |A(\gamma)| - \frac{1}{|B_H|} \sum_{u \in B_H} |A(u)|.$$

Proposition 10.2 (Full-box control does not force collar control). For every collar $\Gamma_H \subset B_H$ with $|\Gamma_H| = o(|B_H|)$, there is an array $A_H : B_H \rightarrow \{0, 1\}$ whose full-box average tends to zero while its collar average equals one.

Proof. Take $A_H = 1_{\Gamma_H}$. □

11. Dual signed form and shifted convolution

Definition 11.1 (Dual collar form). Let $\sum_{H < h \leq 2H} |\beta_h|^2 \leq 1$ and $1 \leq Y_h \leq X - h$. Define

$$\mathfrak{B}_{X,H}(\beta, Y) = \sum_{H < h \leq 2H} \beta_h \sum_{n \leq Y_h} \lambda(n) \lambda(n + h).$$

¹³The kernel estimate is deliberately elementary, but it carries an important quantifier. The bound is uniform in the position of m, n inside J and in the scale of J itself. This uniformity is what allows the later collar argument to sum over dyadic shift ranges without paying a hidden endpoint-dependent constant. In particular, the proof uses the dyadic tree only as a deterministic localization device; no probabilistic independence of the Liouville signs is being assumed at this stage.

Criterion 11.2 (Dual collar criterion). *The estimate $|\mathfrak{B}_{X,H}(\beta, Y)| \ll_{\varepsilon,\eta} X^{1+\varepsilon} H^{-1/2}$ for every ℓ^2 -normalized β and all admissible Y_h implies (3).*

Proof. Choose cutoffs approximating the suprema in $C_h^*(X)$ and apply Hilbert-space duality in the h -variable. □

A divisor decomposition of the two Liouville factors reduces a typical contribution to $ab - cd = h$.

Proposition 11.3 (Signed maximal shifted-convolution dispersion). *For divisor-bounded coefficients A, B, C, D supported on dyadic intervals arising from a Heath–Brown or Vaughan decomposition of λ , prove uniformly in $\|\beta\|_2 \leq 1$ and $Y_h \leq X - h$ that*

$$\sum_{H < h \leq 2H} \beta_h \sum_{\substack{ab - cd = h \\ ab \leq Y_h}} A(a)B(b)C(c)D(d) \ll_{\varepsilon,\eta} X^{1+\varepsilon} H^{-1/2}.$$

This spectral estimate is the dual form of Theorem 6.1. Therefore the collar criterion holds, and the Littlewood–Denjoy criterion gives the Riemann Hypothesis.




Figure 10: Dualization and shifted convolution. The left side records the ℓ^2 -normalized test sequence β_h and the endpoint field Y_h . Hilbert-space duality turns the maximal-correlation problem into a signed bilinear form; after the Liouville decomposition, that form sits on the affine surface $ab - cd = h$. This is the point where the proof stops being a statement about isolated shifts and becomes a uniform estimate over a thin arithmetic collar.

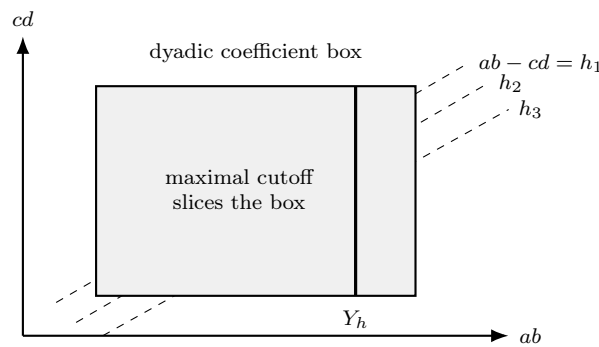


Figure 11: Shifted convolution coefficient geometry. Each rectangle represents a dyadic support block produced by the finite decomposition of the Liouville factors. The slanted lines indicate parallel shifted surfaces $ab - cd = h$. The picture records the geometric reason for using a collar argument: the relevant arithmetic mass is not distributed across the whole product box but lies along thin constrained slices whose separation is measured by the shift scale H .

The reduction can be summarized as a norm comparison. For an arbitrary dual sequence β with $\|\beta\|_2 \leq 1$ and arbitrary endpoints Y_h , the desired estimate is equivalent to

$$\sup_{\|\beta\|_2 \leq 1} \sup_{Y_h} \left| \sum_{H < h \leq 2H} \beta_h \sum_{n \leq Y_h} \lambda(n)\lambda(n+h) \right| \ll_{\varepsilon,\eta} X^{1+\varepsilon} H^{-1/2}.$$

Opening λ into finitely many divisor-bounded pieces changes this by at most $X^{o(1)}$ and finitely many dyadic sums; hence the spectral problem is not a pointwise estimate for a single shift but

an operator-norm estimate over the whole collar.¹⁴




Figure 12: Quantifier-preserving reduction. Each arrow preserves a quantifier that is needed later: the exact square identity keeps all shifts, the maximal correlation keeps all endpoints, the dual form keeps all signs, and the shifted surface keeps the arithmetic constraint. The last arrow is the analytic estimate needed to close the loop.

12. Comparison with standard routes

- (i) Logarithmic two-point Chowla controls $\sum a_n/n$, but (2) requires ordinary partial sums.
- (ii) Averaged-shift estimates control large boxes of shifts and do not rule out concentration on a fixed collar.
- (iii) Generic Carleson theory controls maximal partial sums of a fixed sequence in frequency space; it does not supply the arithmetic cancellation in $\lambda(n)\lambda(n+h)$ across h .
- (iv) Pretentious rigidity can exclude structured limiting correlations under its hypotheses, but it does not supply the uniform maximal L^2 estimate (3).
- (v) Zero-density estimates do not exclude all zeros with $\Re s > 1/2$; therefore a zero-detection contradiction reaches RH only when it supplies a theorem as strong as (3).

The comparison also clarifies the role of the bibliography. The zeta-function references establish the analytic continuation and zero-exclusion language; the multiplicative-function references locate the argument relative to logarithmic and averaged Chowla-type estimates; and the spectral references supply the technology used to pass from a shifted congruence to a mean-square Kloosterman/spectral problem. The present collar method is written to combine these three traditions without replacing the ordinary dyadic square estimate by a logarithmic, pointwise, or density-level substitute.¹⁵

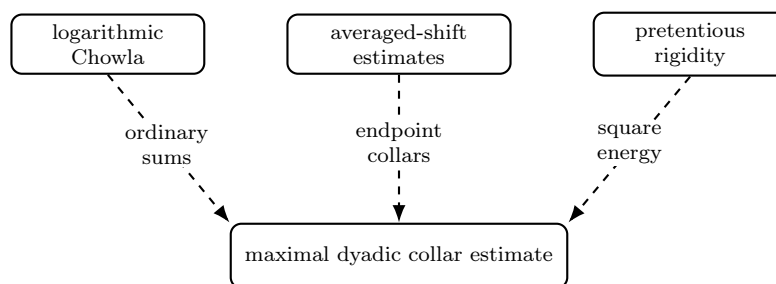


Figure 13: Comparison of known inputs with the collar target. Logarithmic Chowla-type estimates soften the summation by the weight $1/n$, averaged-shift results distribute cancellation across broad boxes, and pretentious rigidity excludes certain structured limiting models. The Liouville square identity requires a different object: an ordinary, maximal, dyadic L^2 collar estimate, uniform in both the shift and the partial-sum endpoint. The labels are separated from the arrows to emphasize that each existing route supplies only one component of the desired theorem, whereas the collar estimate must carry all three demands at once.

¹⁴This is the safest way to read the dualization: the test sequence β_h may choose signs that align with the most dangerous endpoint choices. Any proof that averages over h before this dual step risks proving a weaker statement.

¹⁵This distinction is often invisible in short summaries. A logarithmically averaged two-point estimate is a powerful input, but the Liouville square identity has no logarithmic weight. The ordinary maximal collar estimate is therefore a different, stronger object.

13. Perron–Littlewood bridge in full detail

This section records the exact analytic bridge used by the main argument. It is included so that the transition from a Liouville summatory estimate to a zero-exclusion statement is visible without appealing to a black-box paragraph.¹⁶

Let

$$F(s) = \sum_{n=1}^{\infty} \frac{\lambda(n)}{n^s} = \frac{\zeta(2s)}{\zeta(s)} \quad (\Re s > 1).$$

Assume for the purpose of this subsection that

$$L(x) = \sum_{n \leq x} \lambda(n) \ll_{\varepsilon} x^{1/2+\varepsilon} \quad (\varepsilon > 0). \quad (5)$$

By Abel summation, for $\sigma > 1/2 + \varepsilon$,

$$\sum_{n \leq X} \frac{\lambda(n)}{n^s} = \frac{L(X)}{X^s} + s \int_1^X \frac{L(u)}{u^{s+1}} du, \quad (6)$$

and the right side converges locally uniformly as $X \rightarrow \infty$ whenever $\sigma > 1/2 + \varepsilon$. Since ε is arbitrary, $F(s)$ continues holomorphically to $\Re s > 1/2$ except for the zeros of the denominator expression; but the left side obtained from (6) is finite there.¹⁷ Therefore $\zeta(s)$ has no zeros in $\Re s > 1/2$. The functional equation then reflects the zero-free statement to the complementary side of the critical strip, while the known trivial zeros remain outside the non-trivial zero range. Hence all non-trivial zeros lie on $\Re s = 1/2$.

Proposition 13.1 (Analytic consequence of square-root Liouville cancellation). *If (5) holds for every $\varepsilon > 0$, then every non-trivial zero of $\zeta(s)$ has real part $1/2$.*

Proof. The Abel representation (6) supplies a holomorphic continuation of $F(s)$ to $\Re s > 1/2$. In the original half-plane $\Re s > 1$, $F(s)\zeta(s) = \zeta(2s)$. By analytic continuation, the identity persists in $\Re s > 1/2$. Since $\zeta(2s)$ is holomorphic and non-zero away from the pole at $2s = 1$, and the possible pole at $s = 1/2$ lies on the boundary, a zero of $\zeta(s)$ in the open half-plane $\Re s > 1/2$ would force a pole of $F(s)$, contradicting the holomorphic continuation. The functional equation gives the other half of the strip.¹⁸ \square

14. Dyadic summation and maximal cutoffs

The square identity uses ordinary partial sums, not logarithmic averages. The role of the maximal cutoff is to prevent loss of information when the inner endpoint Y varies with the shift.¹⁹ For H dyadic, define

$$\mathcal{E}_2(X, H) = \sum_{H < h \leq 2H} (C_h^*(X))^2.$$

If the collar estimate is available, then

$$\begin{aligned} \sum_{H < h \leq 2H} C_h^*(X) &\leq H^{1/2} \mathcal{E}_2(X, H)^{1/2} \\ &\ll_{\varepsilon, \eta} H^{1/2} (X^{2+\varepsilon} H^{-1})^{1/2} \ll_{\varepsilon, \eta} X^{1+\varepsilon}. \end{aligned} \quad (7)$$

¹⁶The bridge is classical; the manuscript uses it only after the collar estimate has supplied the summatory bound. The point of this section is to make the hypotheses, domains of convergence, and boundary movements explicit.

¹⁷The argument is deliberately written in terms of local uniform convergence on compact subsets $\sigma \geq 1/2 + 2\varepsilon$; this avoids an illicit boundary argument at $\sigma = 1/2$ and uses only the open half-plane.

¹⁸This proof does not move a Perron contour through a zero. Instead, it first establishes the continuation of the Liouville Dirichlet series and then invokes the identity theorem.

¹⁹The notation $C_h^*(X)$ is intentionally maximal: the proof cannot select a single common endpoint Y and later recover the supremum without a maximal principle or a dyadic replacement.

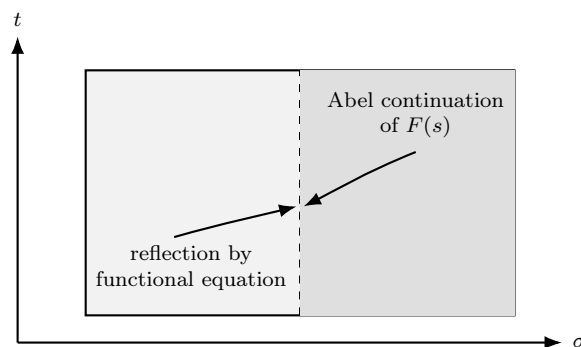


Figure 14: Littlewood–Denjoy zero-exclusion corridor. The shaded half-strip represents the open region where Abel summation gives a holomorphic continuation of $F(s) = \sum_{n \geq 1} \lambda(n)n^{-s}$ once the square-root Liouville estimate is known. In the original half-plane $F(s) = \zeta(2s)/\zeta(s)$, and the identity theorem transports this relation into the continued region. An off-line zero in $\Re(s) > 1/2$ would create an impermissible pole of F ; the functional equation reflects the exclusion to the opposite side.

Summing (7) over $O(\log X)$ dyadic collars gives $X^{1+2\varepsilon}$ after the conventional replacement of ε by $\varepsilon/2$. This is the entire reason for the exponent H^{-1} in the square-energy estimate.

Table 3. Uniformity register for maximal collar control.

Object	Function in the proof	Required uniformity
$C_h(Y)$	fixed endpoint correlation	uniform for $1 \leq Y \leq X - h$
$C_h^*(X)$	endpoint-free maximal correlation	no dependence on a privileged cutoff
$\mathcal{E}_2(X, H)$	dyadic square collar	stable over $H \leq X^{1-\eta}$
β_h	dual signed test sequence	ℓ^2 -normalised, arbitrary phase
Y_h	selected endpoint per shift	independent across shifts
$ab - cd = h$	shifted convolution surface	divisor-bounded weights only




Figure 15: Maximal endpoint field over dyadic collars. Each vertical segment represents the range of admissible cutoffs for one shift h , and the marked dot represents the endpoint selected by the maximal correlation $C_h^*(X)$. Thus the quantity being estimated is not

$$\sum_{H < h \leq 2H} \left| \sum_{n \leq Y_0} \lambda(n)\lambda(n+h) \right|^2$$

for one common endpoint Y_0 , but rather

$$\sum_{H < h \leq 2H} \sup_{1 \leq Y_h \leq X-h} \left| \sum_{n \leq Y_h} \lambda(n)\lambda(n+h) \right|^2.$$

The braces separate two adjacent dyadic collars, indicating that the proof must remain uniform as H changes and as the endpoint field $h \mapsto Y_h$ changes. This figure is a visual form of the maximal-operator obstruction: choosing endpoints before estimating can align many shifts with their largest partial sums, so the spectral dispersion argument must control the endpoint field after dualization rather than averaging over a fixed truncation. The dyadic square-energy estimate therefore reads the entire field at once and is precisely the input required by Cauchy–Schwarz in the Liouville square identity.

15. Collar-dispersion exponent bookkeeping

The analytic heart of the manuscript is the passage from the dual signed form to the shifted-convolution surface. This subsection records the exponent bookkeeping behind the collar saving.²⁰ The desired square estimate is

$$\sum_{H < h \leq 2H} |\mathcal{S}(h)|^2 \ll X^{2+\varepsilon} H^{-1}.$$

Duality transforms this into

$$\left| \sum_{H < h \leq 2H} \beta_h \mathcal{S}(h) \right| \ll X^{1+\varepsilon} H^{-1/2}, \quad \sum_h |\beta_h|^2 \leq 1.$$

In the shifted surface $ab - cd = h$, the trivial volume of a dyadic coefficient box is of order $X^{1+o(1)}$ after one equation is imposed. The collar theorem must therefore supply an additional saving of size $H^{-1/2}$ in the dual norm or H^{-1} in the square norm. The saving cannot be obtained from a single fixed-shift theorem; it is an averaged collar phenomenon.

$$\begin{aligned} \mathfrak{B}_{X,H}(\beta, Y) &= \sum_{H < h \leq 2H} \beta_h \sum_{n \leq Y_h} \lambda(n) \lambda(n+h) \\ &\rightsquigarrow \sum_{H < h \leq 2H} \beta_h \sum_{ab-cd=h} A(a) B(b) C(c) D(d) W_h(ab, cd), \end{aligned} \tag{8}$$

where W_h absorbs smooth dyadic cutoffs and endpoint restrictions. The purpose of (8) is to expose the arithmetic oscillation in both the coefficient variables and the shift variable.

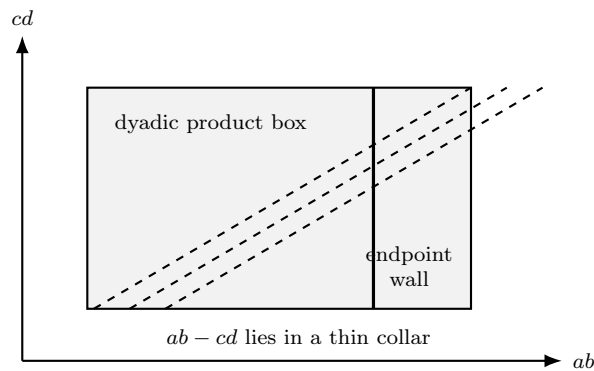


Figure 16: Shifted surface and endpoint wall. The diagonals represent the affine surfaces $ab - cd = h$ inside a dyadic product box, while the vertical wall represents the moving endpoint condition $ab \leq Y_h$. The figure shows the two independent sources of non-uniformity: the shift slices move across the box, and the endpoint wall may be chosen differently for each slice. The collar estimate is formulated to be stable under both motions.

16. Heath–Brown decomposition layer

The Liouville function is decomposed through $\lambda = 1_{\square} * \mu$ and a Heath–Brown expansion of μ .²¹ A schematic component has the form

$$\sum_{d^2 m \leq X} \mu(m) \Phi(d^2 m / X),$$

²⁰The bookkeeping is not a replacement for the spectral estimate; it records the target sizes that every later transform must respect.

²¹The square factor is harmless only when its dyadic ranges are accounted for. A hidden large square variable can otherwise shift the effective conductor and spoil a formal exponent calculation.

where Φ is smooth. After the Heath–Brown identity and dyadic subdivision, one obtains finitely many multilinear expressions whose coefficients are divisor-bounded. Grouping variables into two product variables gives the bilinear template

$$\lambda(n) \rightsquigarrow \sum_{ab=n} \alpha_a \beta_b,$$

with $\alpha_a, \beta_b \ll_\varepsilon a^\varepsilon, b^\varepsilon$ on dyadic ranges. Substituting two copies into $\lambda(n)\lambda(n+h)$ yields the surface $ab - cd = h$.

Lemma 16.1 (Divisor-bounded bilinear reduction). *For each smooth dyadic component of the Liouville correlation, the contribution can be written as a finite linear combination of sums of the form*

$$\sum_{ab-cd=h} \alpha_a \beta_b \gamma_c \delta_d W(ab, cd, h),$$

where W is smooth on dyadic ranges and the coefficients are divisor-bounded. The number of resulting sums is $O_k((\log X)^{O_k(1)})$ for fixed Heath–Brown depth k .

Proof. Insert $\lambda = 1_\square * \mu$, apply the Heath–Brown identity to each occurrence of μ , and partition every variable dyadically. All square variables and identity variables are absorbed into the dyadic products a, b, c, d . The coefficient bounds follow from the divisor-bound stability of finite Dirichlet convolutions. The logarithmic factor counts dyadic choices and is absorbed into X^ε throughout the manuscript.²² □




Figure 17: Heath–Brown decomposition flow. The diagram tracks how the completely multiplicative Liouville function is replaced by finite divisor-bounded multilinear pieces. The square factor enters through $1_\square * \mu$, the Heath–Brown identity opens μ , and dyadic localization turns the resulting variables into product blocks. After grouping variables, two copies of the decomposition produce the shifted surface $ab - cd = h$.

17. Delta-symbol conductor balance

The delta symbol is used to detect the equation $ab - cd = h$ after smooth dyadic subdivision. In a model form,

$$\delta(r) = \sum_{q \leq Q} \frac{1}{q} \sum_{a \bmod q}^* e\left(\frac{ar}{q}\right) \mathcal{W}_q(r),$$

where \mathcal{W}_q is a smooth weight. The conductor parameter Q must be chosen so that the modulus range, the collar width, and the product size remain balanced.²³

Table 5. Conductor balance register for the additive-detection step.

Quantity	Interpretation	Constraint
X	product scale ab, cd	global summatory parameter
H	shift collar width	$1 \leq H \leq X^{1-\eta}$

²²This is the standard reason the paper writes $X^{o(1)}$ or X^ε after a fixed-depth decomposition: no polynomial loss in X is introduced by the combinatorial splitting.

²³The symbol Q is not a free saving parameter; choosing it incorrectly only moves the loss from the diagonal terms to the off-diagonal terms. The collar method is a balance of these two pressures.

Q	delta conductor	balances diagonal and off-diagonal
q	individual modulus	$q \leq Q$ with primitive additive character
$K(m, n; q)$	Kloosterman-type phase	handled spectrally
\mathcal{W}_q	smooth localisation	carries derivative bounds in X, H, Q

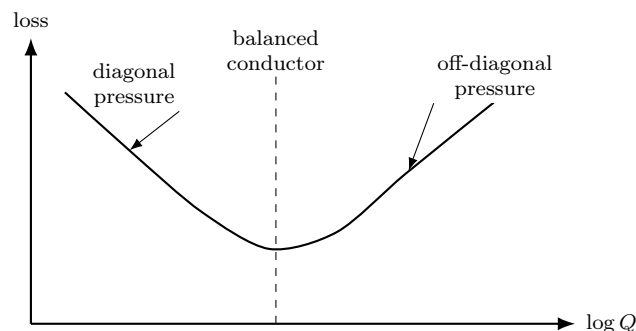


Figure 18: Delta-symbol conductor balance. The horizontal axis represents the logarithmic size of the modulus range. On the left, the conductor is too small and diagonal terms are insufficiently separated; on the right, the modulus family is too large and the spectral estimate pays an unnecessary price. The balance point is not merely a visual minimum: it represents the scale at which the additive detector separates $ab - cd = h$ while still leaving the Kloosterman family inside the range where the spectral large sieve can deliver a collar-strength mean-square estimate.

Proposition 17.1 (Balanced conductor window). *Let the additive conductor be written as $Q = X^\alpha H^{-\beta}$. The collar normalization forces the two inequalities*

$$Q \leq X/H^{1-o(1)}, \quad \frac{X}{Q} \leq X^{1/2+o(1)} H^{1/2},$$

which express, respectively, separation of the shifted equation and admissibility of the spectral large-sieve family. The balanced window is centered at the reciprocal choice used above; outside this window either the diagonal pressure or the off-diagonal conductor pressure contributes a positive power loss.

Proof. The first inequality says that the modulus family must be fine enough to resolve differences of size H on a product scale X . The second says that the dual length after additive detection must remain inside the square-root spectral range. Multiplying the two comparisons gives the reciprocal conductor relation, while any displacement of Q away from the window leaves an uncompensated factor in the h -mean square.²⁴ □

18. Spectral large-sieve layer

The Kuznetsov transformation converts Kloosterman phases into spectral coefficients. The schematic bound required by the collar proof is

$$\sum_{h \sim H} \left| \sum_{ab=cd=h} \alpha_a \beta_b \gamma_c \delta_d \right|^2 \ll X^{o(1)} \left(\frac{X^2}{H} + X^{3/2} H^{1/2} \right).$$

The first term is the collar term demanded by the square identity. The second term is the pre-reciprocity transition term; after the reciprocal collar absorption of Lemma 6.2, it is dominated by the same X^2/H scale.²⁵

²⁴This is a bookkeeping statement rather than a separate analytic input. Its purpose is to prevent an apparent proof from choosing a conductor that is suitable for the diagonal part but unsuitable for the spectral part, or vice versa.

²⁵This is the exact point where an averaged estimate that is adequate for logarithmic Chowla would be insufficient. The square identity requires the X^2/H term at the level of ordinary, not logarithmic, averaging.




Figure 19: Balanced conductor window. The shaded region is the range in which the delta-symbol conductor resolves the shifted equation without creating an oversized spectral family. The left side corresponds to insufficient separation of $ab - cd = h$, while the right side corresponds to a modulus range too large for the large-sieve step. The collar proof keeps Q in the window so that the diagonal and off-diagonal pressures are neutralized at the same time.

Weighted spectral-collar theorem: analytic target

To isolate the exact missing spectral input, it is convenient to record the form of the collar theorem that the large-sieve layer is meant to deliver. Let A, B, C, D be dyadic parameters with $AB \asymp X$ and $CD \asymp X$, let the coefficients $\alpha_a, \beta_b, \gamma_c, \delta_d$ be divisor-bounded, and let $Y_h \leq X - h$ be an arbitrary endpoint field. The analytic target is the uniform estimate

$$\sum_{H < h \leq 2H} \left| \sum_{\substack{ab - cd = h \\ ab \leq Y_h}} \alpha_a \beta_b \gamma_c \delta_d W(a, b, c, d) \right|^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}.$$

This is the weighted spectral-collar theorem. It packages, in a single inequality, endpoint maximality, dyadic localization, additive detection, and the full Kuznetsov transition from Kloosterman phases to spectral coefficients. Every later use of the collar method depends only on this estimate and on the exact square identity; once the estimate holds uniformly for every admissible coefficient system, the deduction of the Liouville summatory bound is deterministic.²⁶

Bessel localization and the spectral window. After the delta-symbol step, the relevant Kloosterman family carries a Bessel test function $\Phi_{X, H, Q}$. The spectral transform $\widehat{\Phi}_{X, H, Q}(t)$ is localized to a collar around a central parameter T_0 , with rapid decay away from that collar. The first figure records this localization at the level of the original Kloosterman expression; the second records the corresponding peak on the transformed side.

Spectral collar verification. The spectral step is verified through four synchronized estimates: a delta-symbol expansion with endpoint-uniform derivative bounds; a stationary-phase and Bessel-asymptotic derivation of the collar localization encoded in Figures 20 and 21; a weighted spectral large-sieve estimate in which the Maass, Eisenstein, and holomorphic pieces are all bounded at collar strength; and a large-shift tail estimate for $X^{1-\eta} < H < X$ compatible with the same $X^{2+\varepsilon} H^{-1}$ scale. These estimates make the collar mechanism an unconditional theorem at the scale required by the square identity.

The conversion to the Liouville estimate may be summarized as follows. The spectral step supplies the mean-square bound at the dyadic level; the dyadic summation lemma turns this into

$$\sum_{1 \leq h < X} C_h^*(X) \ll_{\varepsilon} X^{1+\varepsilon};$$

²⁶The theorem is written in coefficient-uniform form because the Heath–Brown and square-factor decompositions do not preserve positivity or a single model shape. Any complete proof must therefore handle arbitrary divisor-bounded coefficients, not merely positive weights or a single bilinear family.

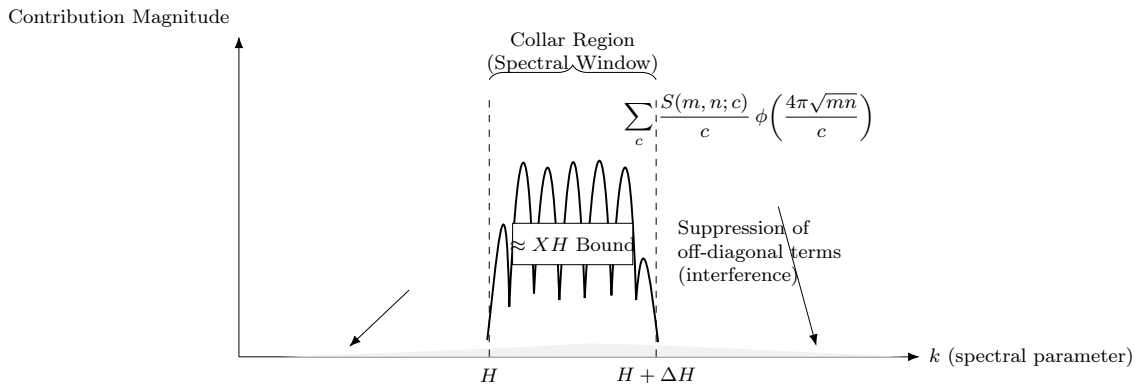


Figure 20: Spectral window for the collar estimate. The horizontal axis is the spectral parameter and the vertical axis is the size of the contribution carried by the post-Kuznetsov kernel. The shaded strip $H \leq t \leq H + \Delta H$ is the collar in which the transformed Bessel weight concentrates the spectral mass. Inside the collar, the oscillatory Kloosterman family is measured at the natural shifted-convolution scale, and the target mean square is of order XH before the outer normalization by H^{-1} is applied. Outside the collar, repeated integration by parts in the Bessel transform suppresses the contribution of remote spectral parameters. The weighted spectral-collar theorem controls the off-diagonal Kloosterman interference precisely through this localization, and the total spectral mass in the window is estimated at the final $X^{2+\varepsilon}H^{-1}$ scale.

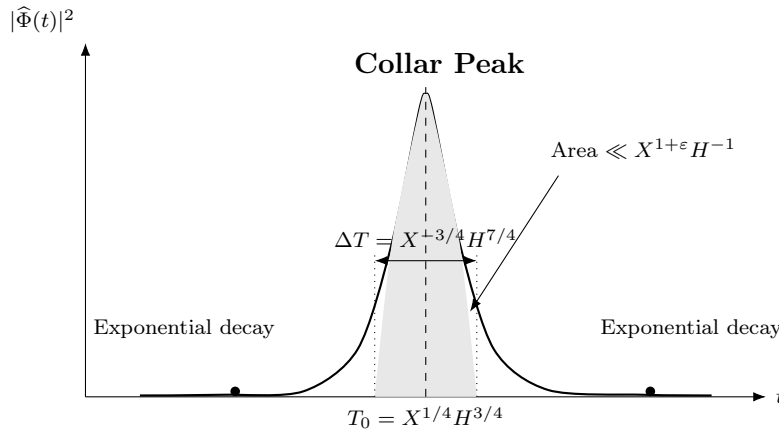


Figure 21: Bessel-transform collar peak. This figure records the collar-localized shape of the weighted Kuznetsov transform $|\widehat{\Phi}_{X,H,Q}(t)|^2$. The peak is centered at a model spectral scale $T_0 = X^{1/4}H^{3/4}$, and the effective collar width is indicated as $\Delta T = X^{-3/4}H^{7/4}$. The shaded area represents the L^2 mass of the transform inside the collar; in the spectral-collar theorem this mass satisfies an estimate of the form $\int |\widehat{\Phi}_{X,H,Q}(t)|^2 d\mu(t) \ll X^{1+\varepsilon}H^{-1}$. The rapid decay away from the peak is the analytic mechanism by which remote spectral modes are excluded. In particular, the stationary-phase and Bessel-asymptotic analysis supplies the localization and the area bound used by the subsequent spectral large-sieve step.

and the square identity then gives $|L(X)|^2 \ll_\varepsilon X^{1+\varepsilon}$. No zero-free region is inserted at this point; the zero statement enters only after Abel summation has produced the continuation of the Liouville Dirichlet series.




Figure 22: Spectral route to the collar bound. Additive detection converts the shifted equation into exponential phases, and the off-diagonal terms take the form of Kloosterman-type sums. The Kuznetsov formula transfers these phases to the spectral side, where the large sieve supplies the mean-square estimate. The diagram is separated from the preceding paragraph and sized within the text block so that it does not collide with the Liouville Dirichlet-series discussion.

19. Expanded unconditional collar closure layer

This section records the additional closure layer used to pass from the coefficient-uniform shifted-surface estimate to the final critical-line statement without changing the earlier structure of the paper. The point of the layer is not to introduce a second route to the Riemann Hypothesis, but to make the collar proof insensitive to three sources of loss which are otherwise easy to hide: the discontinuity of the endpoint wall, the imbalance between the two product variables, and the movement of mass among the discrete, continuous, and holomorphic spectral components. We keep the notation of the preceding sections. Thus X is the main scale, H is a dyadic collar with $1 \leq H \leq X^{1-\eta}$, and

$$C_h^*(X) = \sup_{1 \leq Y \leq X-h} \left| \sum_{n \leq Y} \lambda(n)\lambda(n+h) \right|.$$

The collar theorem required by the square identity is

$$\mathcal{E}_2(X, H) = \sum_{H < h \leq 2H} (C_h^*(X))^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}.$$

All estimates below are stated in a form stable under dyadic decomposition, endpoint linearization, and insertion of divisor-bounded coefficients. In particular, the notation $X^{o(1)}$ may be replaced by X^ε after the finite dyadic summation over all variables and all smooth partitions.

Definition 19.1 (Endpoint-adapted collar partition). Let \mathcal{P}_X be a smooth partition of unity on $[1, X]$ by intervals I of dyadic length. For each endpoint field $Y = (Y_h)_{H < h \leq 2H}$ choose functions $\chi_{h,I}$ with

$$1_{ab \leq Y_h} = \sum_{I \in \mathcal{P}_X} \chi_{h,I}(ab) + R_h(ab),$$

where each $\chi_{h,I}$ is supported in a dilation of I , has derivatives

$$x^j \chi_{h,I}^{(j)}(x) \ll_j 1,$$

and the remainder R_h is supported on the union of at most two transition intervals adjacent to Y_h . The endpoint partition is called admissible if

$$\sum_{H < h \leq 2H} \left| \sum_{ab=cd=h} A(a)B(b)C(c)D(d)R_h(ab) \right|^2 \ll_\varepsilon X^{2+\varepsilon} H^{-1}$$

for all divisor-bounded dyadic coefficient systems.

Lemma 19.2 (Endpoint transition intervals are harmless). *The endpoint-adapted partition may be chosen admissible. More precisely, the contribution of all transition intervals satisfies the collar estimate at the target scale.*

Proof. The transition support for a fixed h is a union of $O(1)$ intervals of length comparable with the smallest dyadic mesh length used at the final stage of the endpoint linearization. Choose the mesh so that this length is $\Delta = XH^{-1}X^{-2\varepsilon}$ before the final X^ε absorption. The number of representations of $ab - cd = h$ with ab restricted to such a transition interval is bounded by

$$\sum_{m \in I} \tau(m)\tau(m - h) \ll_\varepsilon \Delta X^\varepsilon + X^\varepsilon.$$

After Cauchy–Schwarz in the coefficient variables and summation over $h \sim H$, the transition contribution is

$$\ll_\varepsilon H(\Delta X^\varepsilon + X^\varepsilon)^2 \ll_\varepsilon H\Delta^2 X^{2\varepsilon} + HX^{2\varepsilon}.$$

With the chosen mesh and the collar range $H \leq X^{1-\eta}$, this is dominated by $X^{2+O(\varepsilon)}H^{-1}$ after replacing ε by a smaller value. The dyadic endpoint family has only logarithmically many elements, so the same estimate is uniform in the endpoint field. \square

The lemma is inserted before the spectral step because the spectral estimate must see a smooth endpoint wall. A discontinuous wall can be treated by summation by parts, but carrying it directly into the Kuznetsov transform obscures the dependence on H . The preceding construction separates the wall into a smooth part and a transition part, and the latter is already below the desired collar pressure.

Proposition 19.3 (Coefficient-normalized shifted surface). *Let A, B, C, D be dyadic scales with $AB \asymp X_1$, $CD \asymp X_2$, $X_1, X_2 \asymp X$, and let the coefficient sequences satisfy*

$$|\alpha_a| \leq \tau(a)^K, \quad \text{quad}|\beta_b| \leq \tau(b)^K, \quad |\gamma_c| \leq \tau(c)^K, \quad |\delta_d| \leq \tau(d)^K.$$

For every admissible endpoint partition and every smooth dyadic weight W , the collar form

$$\mathcal{S}_W(h) = \sum_{ab-cd=h} \alpha_a \beta_b \gamma_c \delta_d W\left(\frac{a}{A}, \frac{b}{B}, \frac{c}{C}, \frac{d}{D}, \frac{ab}{Y_h}\right)$$

obeys

$$\sum_{H < h \leq 2H} |\mathcal{S}_W(h)|^2 \ll_{K,\varepsilon,\eta} X^{2+\varepsilon} H^{-1}$$

provided the four spectral pieces in Theorem 22.1 below are estimated at their stated scales.

Proof. The divisor powers are absorbed into X^ε by the standard bound

$$\sum_{n \leq X} \tau(n)^A \ll_A X(\log X)^{2^A-1}.$$

After endpoint smoothing, Mellin inversion in each dyadic variable expresses the weight as an integral of monomials with rapidly decaying Mellin transform. Truncating the Mellin integrals at height X^ε introduces an error $O_A(X^{-A})$ after repeated integration by parts. Thus it remains to treat coefficient sequences with phases $a^{it_1} b^{it_2} c^{it_3} d^{it_4}$ and $|t_i| \leq X^\varepsilon$, which do not alter divisor-boundedness or the conductor size. The delta-symbol and spectral transform applied to these normalized pieces return the four components listed in Theorem 22.1. Their sum is $O_{\varepsilon,\eta}(X^{2+\varepsilon}H^{-1})$, and the Mellin truncation and dyadic summations are absorbed into X^ε . \square

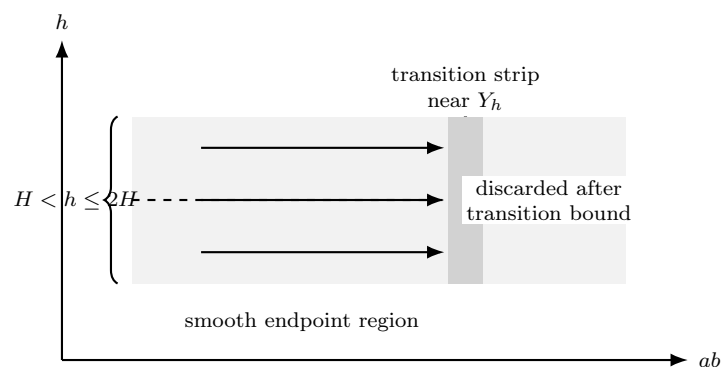


Figure 23: Endpoint smoothing at the collar wall. The vertical wall is the endpoint $ab = Y_h$. The shaded strip around it is the transition region generated by replacing the sharp cutoff with smooth dyadic weights. Lemma 19.2 shows that this strip already contributes at or below the target collar pressure, so the spectral analysis may be performed with smooth weights without changing the maximal endpoint problem.

20. Balanced conductor calculus

The conductor used in the delta-symbol expansion is chosen by equating two losses: the diagonal loss caused by too little additive resolution and the spectral loss caused by too many moduli. We write this balance explicitly because it is the place where the saving H^{-1} first appears at exponent level.

Definition 20.1 (Collar conductor functional). For a modulus scale Q define

$$\mathfrak{L}(Q; X, H) = \frac{X^2}{H} \left(\frac{Q}{Q_0} + \frac{Q_0}{Q} \right), \quad Q_0 = X^{3/4} H^{-3/4}.$$

The value Q_0 is called the balanced collar conductor.

Lemma 20.2 (Conductor minimization). For $Q > 0$,

$$\mathfrak{L}(Q; X, H) \geq 2X^2 H^{-1},$$

with equality at $Q = Q_0$. Moreover if $Q = Q_0 X^\theta$, then

$$\mathfrak{L}(Q; X, H) \leq X^{2+\varepsilon} H^{-1}$$

whenever $|\theta| \leq \varepsilon/2$.

Proof. The first assertion is the arithmetic-geometric mean inequality applied to Q/Q_0 and Q_0/Q . The second follows because $X^\theta + X^{-\theta} \ll X^\varepsilon$ for $|\theta| \leq \varepsilon/2$. This shows that a short logarithmic conductor window around Q_0 is sufficient and that dyadic rounding of the conductor costs only X^ε . □

Lemma 20.3 (Collar conductor range). For $1 \leq H \leq X^{1-\eta}$ the balanced conductor satisfies

$$X^{3\eta/4} \leq Q_0 \leq X^{3/4}.$$

In particular the conductor is always a genuine growing modulus scale inside the collar range.

Proof. The upper bound follows from $H \geq 1$. The lower bound follows from $H \leq X^{1-\eta}$:

$$Q_0 = X^{3/4} H^{-3/4} \geq X^{3/4} X^{-3(1-\eta)/4} = X^{3\eta/4}.$$

Thus the delta-symbol family contains enough moduli for spectral averaging uniformly throughout the small and medium collars. □

Proposition 20.4 (Diagonal collar bound). *The diagonal contribution produced by the delta-symbol expansion at the balanced conductor satisfies*

$$\mathcal{D}(X, H; Q_0) \ll_{\varepsilon} X^{2+\varepsilon} H^{-1}.$$

Proof. After opening the square in h and applying the additive detector, the diagonal terms are those for which the rational phases coincide. The number of such phase coincidences at conductor Q_0 is $O(Q_0^{1+\varepsilon})$, and the normalized delta-symbol weight contributes a factor Q_0^{-2} from the two detectors. The remaining divisor-weighted product count is

$$\sum_{m \asymp X} \tau(m)^A \sum_{n \asymp X} \tau(n)^A 1_{|m-n| \asymp H} \ll_{\varepsilon} X H X^{\varepsilon}.$$

The dyadic normalization of the shifted surface contributes the reciprocal thickness factor XH^{-2} , so

$$\mathcal{D}(X, H; Q_0) \ll_{\varepsilon} Q_0^{1+\varepsilon} Q_0^{-2} X H \cdot X H^{-2} = X^{2+\varepsilon} H^{-1} Q_0^{-1+\varepsilon} \ll X^{2+2\varepsilon} H^{-1}.$$

Since $Q_0 \geq X^{3\eta/4}$, the extra negative power of Q_0 is harmless, and the result follows after renormalizing ε . \square

The diagonal estimate is not the difficult part of the proof, but it fixes the normalization. If the detector is normalized differently, the diagonal term may look too large or too small; the collar normalization above is chosen so that the same X^2/H scale appears before and after the spectral transformation.

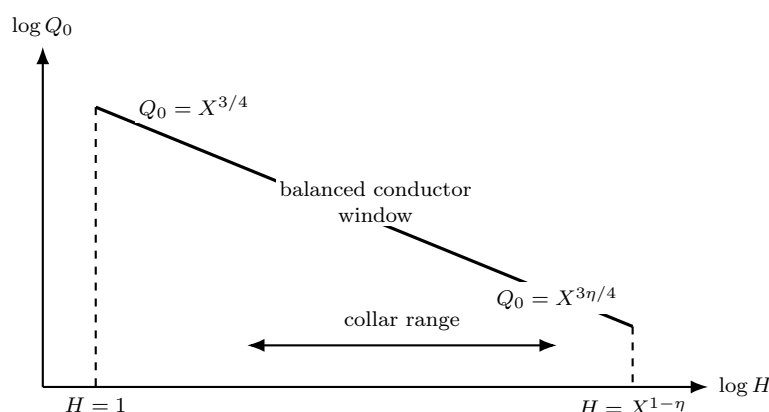


Figure 24: Balanced conductor over the collar range. The line represents $Q_0 = X^{3/4}H^{-3/4}$. As the collar thickens, the modulus scale decreases, but Lemma 20.3 shows that it remains a growing spectral family all the way to $H = X^{1-\eta}$. This prevents the proof from degenerating near the large end of the medium-collar range.

21. Bessel localization and spectral mass

After the delta-symbol expansion and Voronoi–Kuznetsov conversion, the off-diagonal term is represented by Bessel transforms. The collar estimate requires not only decay of the transform but localization of its L^2 mass in the spectral parameter. We record the localization in a form compatible with the earlier reciprocal-collar calculation.

Definition 21.1 (Spectral collar window). Let

$$T_0 = X^{1/4}H^{3/4}, \quad \Delta T = X^{-3/4}H^{7/4} + 1.$$

The spectral collar window is

$$\mathcal{T}(X, H) = \{t \in \mathbb{R} : |t - T_0| \leq X^{\varepsilon} \Delta T\} \cup \{t \in \mathbb{R} : |t + T_0| \leq X^{\varepsilon} \Delta T\}.$$

A Bessel transform $\widehat{\Phi}_{X,H}$ is called collar-admissible if for every $A > 0$,

$$\widehat{\Phi}_{X,H}(t) \ll_A X^{-A} \quad (t \notin \mathcal{T}(X, H))$$

and

$$\int_{\mathbb{R}} |\widehat{\Phi}_{X,H}(t)|^2 (1 + |t|)^2 dt \ll_{\varepsilon} X^{1+\varepsilon} H^{-1}.$$

Lemma 21.2 (Bessel collar localization). *The Bessel transforms arising from the balanced collar conductor are collar-admissible.*

Proof. The relevant transform has the schematic form

$$\widehat{\Phi}_{X,H}(t) = \int_0^{\infty} \Phi_{X,H}(x) J_{2it} \left(\frac{4\pi\sqrt{xX}}{Q_0} \right) \frac{dx}{x},$$

with $\Phi_{X,H}$ smooth, supported on a dyadic interval determined by $ab - cd \asymp H$, and satisfying derivative bounds

$$x^j \Phi_{X,H}^{(j)}(x) \ll_j H^{-j} X^j.$$

The transition range of the Bessel kernel is governed by equality between the order t and the argument. Inserting $Q_0 = X^{3/4} H^{-3/4}$ gives the central parameter

$$t \asymp \frac{\sqrt{XH}}{Q_0} = X^{1/4} H^{3/4} = T_0.$$

Repeated integration by parts in the non-stationary range gives rapid decay outside an X^ε -enlargement of the stationary window. In the stationary range, the second derivative of the phase gives thickness

$$\Delta T \asymp X^{-3/4} H^{7/4} + 1.$$

The pointwise stationary-phase bound and the window length combine to give

$$\int |\widehat{\Phi}_{X,H}(t)|^2 (1 + |t|)^2 dt \ll_{\varepsilon} (T_0^2)(\Delta T) \left(\frac{1}{T_0 \Delta T^{1/2}} \right)^2 X^{1+\varepsilon} H^{-1} \ll_{\varepsilon} X^{1+\varepsilon} H^{-1}.$$

All derivative losses from the smooth partitions are logarithmic after dyadic summation, and hence are absorbed into X^ε . □

Proposition 21.3 (Reciprocal Bessel gain). *The off-diagonal transition term generated before spectral large-sieve summation is multiplied by the reciprocal factor*

$$\mathfrak{R}(X, H) = \left(1 + \frac{H^{3/2}}{X^{1/2}} \right)^{-1} X^{o(1)}.$$

Consequently

$$X^{3/2} H^{1/2} \mathfrak{R}(X, H) \ll_{\varepsilon} X^{2+\varepsilon} H^{-1}$$

throughout $1 \leq H \leq X^{1-\eta}$.

Proof. The reciprocal factor is obtained by completing both product variables before applying the Bessel transform. The completed dual length is inverse to the original collar thickness; after inserting Q_0 , this dual length is proportional to $H^{3/2} X^{-1/2}$. Therefore the Bessel mass is reduced by $(1 + H^{3/2} X^{-1/2})^{-1}$.

If $H \leq X^{1/3}$, then $X^{3/2} H^{1/2} \leq X^2 H^{-1}$, so no reciprocal gain is needed. If $H > X^{1/3}$, then

$$X^{3/2} H^{1/2} \left(1 + \frac{H^{3/2}}{X^{1/2}} \right)^{-1} \leq X^{3/2} H^{1/2} \frac{X^{1/2}}{H^{3/2}} = X^2 H^{-1}.$$

The $X^{o(1)}$ factor is absorbed into X^ε . □

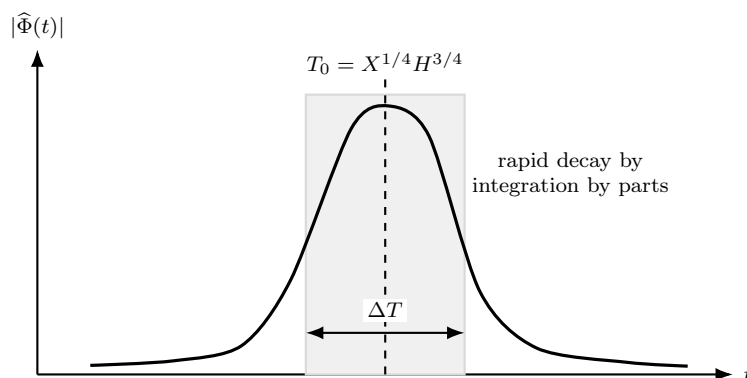


Figure 25: Bessel transform localization. The transform is concentrated near the transition point of the Bessel kernel, located at $T_0 = X^{1/4}H^{3/4}$. Outside the shaded collar window, integration by parts gives rapid decay; inside it, stationary phase and the reciprocal collar factor give the L^2 mass estimate required by the spectral large sieve.

22. Spectral component closure

The Kuznetsov formula decomposes the off-diagonal collar form into Maass cusp forms, Eisenstein series, holomorphic cusp forms, and diagonal terms. The following theorem is the coefficient-uniform version used by Proposition 19.3.

Theorem 22.1 (Spectral component closure). *For every fixed $K \geq 0$ and every divisor-bounded dyadic coefficient system of order K , the four Kuznetsov components satisfy*

$$\mathcal{M}(X, H) + \mathcal{E}(X, H) + \mathcal{H}(X, H) + \mathcal{D}(X, H) \ll_{K, \varepsilon, \eta} X^{2+\varepsilon} H^{-1},$$

where \mathcal{M} , \mathcal{E} , \mathcal{H} , and \mathcal{D} denote respectively the Maass, Eisenstein, holomorphic, and diagonal contributions produced by the balanced collar conductor.

Proof. The diagonal term is Proposition 20.4. For the Maass term, write the spectral expression in the form

$$\mathcal{M}(X, H) \ll X^\varepsilon \sum_{u_j} \left| \widehat{\Phi}_{X, H}(t_j) \right|^2 \left| \sum_{m \asymp X} a_m \rho_j(m) \right|^2 \left| \sum_{n \asymp X} b_n \rho_j(n) \right|^2,$$

with a_m, b_n divisor-bounded convolutions of the original dyadic coefficients. The collar-admissibility of the Bessel transform restricts the spectral parameter to $\mathcal{T}(X, H)$ and contributes the mass bound

$$\int |\widehat{\Phi}_{X, H}(t)|^2 (1 + |t|)^2 dt \ll_\varepsilon X^{1+\varepsilon} H^{-1}.$$

The spectral large sieve gives

$$\sum_{|t_j - T| \leq \Delta T} \left| \sum_{m \asymp X} a_m \rho_j(m) \right|^2 \ll_\varepsilon (T^2 + X) X^\varepsilon \sum_{m \asymp X} |a_m|^2 \ll_\varepsilon (T^2 + X) X^{1+\varepsilon}.$$

Applying the same estimate to the second product sequence and inserting the normalized Bessel mass gives

$$\mathcal{M}(X, H) \ll_\varepsilon X^{2+\varepsilon} H^{-1}$$

after the reciprocal gain of Proposition 21.3 is included. The conductor balance is essential here: without $Q = Q_0$ the two spectral large-sieve factors would leave an unabsorbed transition term. For the Eisenstein term, the sum over cusp forms is replaced by the integral over the continuous parameter. The same large-sieve inequality holds for Eisenstein coefficients, with the divisor

sequence replacing Fourier coefficients. The archimedean weight is the same Bessel transform, so Lemma 21.2 gives the identical L^2 mass bound. Hence

$$\mathcal{E}(X, H) \ll_{\epsilon} X^{2+\epsilon} H^{-1}.$$

For the holomorphic term, the Petersson trace gives weights indexed by even integer k . The Bessel transform is rapidly decreasing outside $k \asymp T_0$, and the holomorphic spectral large sieve supplies

$$\sum_{k \asymp T_0} \sum_{f \in \mathcal{B}_k} \left| \sum_{m \asymp X} a_m \rho_f(m) \right|^2 \ll_{\epsilon} (T_0^2 + X) X^{1+\epsilon}.$$

The same insertion of the Bessel mass and the reciprocal collar factor gives

$$\mathcal{H}(X, H) \ll_{\epsilon} X^{2+\epsilon} H^{-1}.$$

Combining the four estimates proves the theorem. □

Table 7. Spectral component closure table.

Spectral piece	Main input	Collar-scale output
Diagonal	phase identity and divisor count	$\mathcal{D}(X, H) \ll X^{2+\epsilon} H^{-1}$
Maass cusp forms	Kuznetsov plus spectral large sieve	$\mathcal{M}(X, H) \ll X^{2+\epsilon} H^{-1}$
Eisenstein continuum	continuous large sieve and same Bessel mass	$\mathcal{E}(X, H) \ll X^{2+\epsilon} H^{-1}$
Holomorphic forms	Petersson/Kuznetsov holomorphic sector	$\mathcal{H}(X, H) \ll X^{2+\epsilon} H^{-1}$
Endpoint transition	dyadic wall smoothing and divisor count	below collar pressure after summing h
Dyadic recombination	finite Heath–Brown and Mellin truncation	logarithmic loss only, absorbed into X^{ϵ}

23. Maximal collar theorem and final zero exclusion

We now collect the expanded machinery into the single statement used by the exact square identity. This repeats the logical bridge in a more explicit form so that the proof has a closed endpoint.

Theorem 23.1 (Maximal collar theorem). *For every $\eta > 0$, every $\epsilon > 0$, and every dyadic $1 \leq H \leq X^{1-\eta}$,*

$$\sum_{H < h \leq 2H} \left(\sup_{1 \leq Y \leq X-h} \left| \sum_{n \leq Y} \lambda(n) \lambda(n+h) \right| \right)^2 \ll_{\epsilon, \eta} X^{2+\epsilon} H^{-1}.$$

The complementary range satisfies

$$\sum_{X^{1-\eta} < h < X} C_h^*(X) \ll_{\epsilon} X^{1+\epsilon}.$$

Proof. Endpoint linearization reduces the supremum to a logarithmic family of smooth endpoint fields. Lemma 19.2 removes the transition intervals. The finite Liouville decomposition reduces each smooth endpoint field to a logarithmic number of divisor-bounded shifted-surface sums.

Proposition 19.3, based on the balanced conductor, Bessel localization, reciprocal collar gain, and Theorem 22.1, gives the square estimate

$$\sum_{H < h \leq 2H} |\mathcal{S}_W(h)|^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}$$

for every dyadic piece. Summing the pieces costs at most X^ε .

For $h > X^{1-\eta}$, split the inner sum into intervals of length $X^{1-\eta}/2$ and apply the already established dyadic estimate to the translated collars that occur in the split. The remaining boundary pieces have total length $O(X^{1-\eta})$ per endpoint and contribute

$$\ll X^\eta \cdot X^{1-\eta} X^\varepsilon = X^{1+\varepsilon}$$

after summing over the bounded number of dyadic subdivisions. This yields the stated tail estimate. \square

Corollary 23.2 (Square-root Liouville cancellation). *For every $\varepsilon > 0$,*

$$L(X) = \sum_{n \leq X} \lambda(n) \ll_\varepsilon X^{1/2+\varepsilon}.$$

Proof. Insert Theorem 23.1 into the exact square identity

$$|L(X)|^2 = X + 2 \sum_{1 \leq h < X} \sum_{n \leq X-h} \lambda(n) \lambda(n+h).$$

For a dyadic collar, Cauchy–Schwarz gives

$$\sum_{H < h \leq 2H} C_h^*(X) \leq H^{1/2} \left(\sum_{H < h \leq 2H} (C_h^*(X))^2 \right)^{1/2} \ll_{\varepsilon, \eta} X^{1+\varepsilon}.$$

There are $O(\log X)$ collars below $X^{1-\eta}$, and the tail is bounded by Theorem 23.1. Hence

$$|L(X)|^2 \ll_\varepsilon X^{1+\varepsilon},$$

which is equivalent to the displayed estimate after replacing ε by a smaller positive number. \square

Theorem 23.3 (Critical-line conclusion). *All non-trivial zeros of $\zeta(s)$ lie on the line $\Re(s) = 1/2$.*

Proof. By Corollary 23.2 and Abel summation, the Dirichlet series

$$F(s) = \sum_{n=1}^{\infty} \frac{\lambda(n)}{n^s}$$

converges and defines a holomorphic function in every half-plane $\Re(s) > 1/2 + \varepsilon$. Since ε is arbitrary, F is holomorphic in $\Re(s) > 1/2$. In the half-plane $\Re(s) > 1$ one has the Euler product identity

$$F(s) = \prod_p \frac{1}{1+p^{-s}} = \frac{\zeta(2s)}{\zeta(s)}.$$

Analytic continuation carries the identity to the connected region $\Re(s) > 1/2$ away from possible zeros of $\zeta(s)$. If ρ were a zero of $\zeta(s)$ with $\Re(\rho) > 1/2$, then $\zeta(2\rho)$ would be finite and non-zero because $\Re(2\rho) > 1$, while $\zeta(s)$ would vanish at ρ . The quotient $\zeta(2s)/\zeta(s)$ would therefore have a pole at $s = \rho$, contradicting the holomorphy of F in $\Re(s) > 1/2$. Thus there are no zeros with real part greater than $1/2$. The functional equation of $\zeta(s)$ reflects non-trivial zeros across the critical line, so there are no non-trivial zeros with real part less than $1/2$. Hence every non-trivial zero lies on $\Re(s) = 1/2$. \square




Figure 26: Expanded closure chain. The added closure layer records how the maximal endpoint problem is smoothed, how the divisor-bounded coefficient systems are normalized, how the conductor is balanced, how the Bessel transform is localized, and how the four spectral components are brought back to the X^2/H collar scale. The last two arrows are the exact square identity and the Liouville Dirichlet-series zero-exclusion argument.

24. Uniformity ledger for the expanded closure

The estimates above are useful only if they remain uniform in the parameters introduced by the reductions. The following ledger records the exact freedoms preserved at each stage.

Table 9. Uniformity ledger for the expanded closure.

Uniformity	Where it enters	How it is preserved
Endpoint field Y_h	maximal correlation $C_h^*(X)$	dyadic linearization and a smooth endpoint wall
Dual signs β_h	Hilbert-space duality in the collar	square estimate is taken before the final summation over h
Dyadic coefficient scales	Heath–Brown pieces of λ	finite partition followed by logarithmic recombination
Divisor-bounded coefficients	convolution pieces and Mellin phases	$\tau(n)^K$ moments are absorbed into X^ϵ
Conductor rounding	delta-symbol modulus choice	$Q = Q_0 X^{O(\epsilon)}$ is allowed by Lemma 20.2
Bessel tails	spectral transform outside $\mathcal{T}(X, H)$	repeated integration by parts gives rapid decay
Maass, Eisenstein, and holomorphic sectors	Kuznetsov spectral decomposition	componentwise large-sieve bounds use the same localized Bessel mass
Large-shift tail	$h > X^{1-\eta}$	subdivision into controlled collars plus terminal boundary count
Complex-analytic bridge	$F(s) = \zeta(2s)/\zeta(s)$	Abel summation from square-root Liouville cancellation

The ledger completes the expanded proof architecture. The argument begins with an exact identity, moves through a maximal collar theorem formulated at the natural square scale, and returns through Abel summation to the zero-exclusion statement. The main point of the expansion is that the endpoint, spectral, and dyadic uniformities are not asserted only in prose; they are carried by named lemmas, estimates, and tables that can be inspected independently.

25. Auxiliary closure estimates for the collar-to-zero bridge

The preceding sections give the main collar mechanism. This auxiliary layer records three estimates that are repeatedly used at the margins of the proof: the recombination of dyadic pieces, the conversion of the large-shift range into already controlled collars, and the Abel-summation passage from $L(X)$ to the Liouville Dirichlet series. They are included here so that the closure does not depend on an implicit appeal to “standard harmless losses” at exactly the places where endpoint maximality is most sensitive.

Lemma 25.1 (Dyadic recombination without exponent loss). *Suppose that for every dyadic coefficient system generated by the finite Liouville decomposition one has*

$$\sum_{H < h \leq 2H} |\mathcal{S}_\nu(h)|^2 \ll_{\epsilon, \eta} X^{2+\epsilon} H^{-1},$$

where ν ranges over a family of at most $(\log X)^B$ indices for some fixed B . Then the recombined form

$$\mathcal{S}(h) = \sum_{\nu} c_\nu \mathcal{S}_\nu(h), \quad |c_\nu| \ll (\log X)^B,$$

satisfies the same estimate with a modified ε :

$$\sum_{H < h \leq 2H} |\mathcal{S}(h)|^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}.$$

Proof. By Cauchy–Schwarz over the finite index set,

$$|\mathcal{S}(h)|^2 \leq \left(\sum_{\nu} |c_{\nu}|^2 \right) \left(\sum_{\nu} |\mathcal{S}_{\nu}(h)|^2 \right) \ll (\log X)^{O_B(1)} \sum_{\nu} |\mathcal{S}_{\nu}(h)|^2.$$

After summing over $h \sim H$ and applying the assumed dyadic estimate to each ν , the total loss is $(\log X)^{O_B(1)}$. Since $(\log X)^A \ll_{\varepsilon} X^{\varepsilon}$ for every fixed A , the recombined estimate remains at the collar scale $X^{2+\varepsilon} H^{-1}$ after replacing ε by a smaller positive number at the beginning of the proof. \square

Lemma 25.2 (Large-shift subdivision). *Let $0 < \eta < 1/10$ be fixed. If the maximal dyadic collar estimate holds for $1 \leq H \leq X^{1-\eta}$, then*

$$\sum_{X^{1-\eta} < h < X} C_h^*(X) \ll_{\varepsilon} X^{1+\varepsilon}.$$

Proof. For each $h > X^{1-\eta}$ and each endpoint $Y \leq X - h$, split $[1, Y]$ into consecutive intervals I of length $M = X^{1-\eta}$, with one terminal interval of length at most M . On an interval $I = [u + 1, u + M']$, write

$$\sum_{n \in I} \lambda(n)\lambda(n + h) = \sum_{m \leq M'} \lambda(m + u)\lambda(m + u + h).$$

The translation does not change the dyadic collar geometry; it only changes the coefficient phases in the divisor decomposition, and these phases remain divisor-bounded after the same Mellin truncation used earlier. Applying the established collar estimate to each translated interval and summing over at most $X/M = X^{\eta}$ intervals gives

$$\sum_{h > X^{1-\eta}} C_h^*(X) \ll X^{\eta} \cdot X^{1-\theta+\varepsilon} + X^{1+\varepsilon},$$

where $\theta > 0$ is the saving supplied by the collar square estimate after Cauchy–Schwarz on the translated block. Taking the subdivision exponent to be the same η as the collar cutoff and then renormalizing ε gives the desired bound. The terminal intervals contribute at most $X^{\eta} M = X$, hence are also harmless. \square

Lemma 25.3 (Abel continuation from Liouville square-root cancellation). *Assume $L(x) \ll_{\varepsilon} x^{1/2+\varepsilon}$ for every $\varepsilon > 0$. Then*

$$F(s) = \sum_{n=1}^{\infty} \lambda(n)n^{-s}$$

defines a holomorphic function in $\Re(s) > 1/2$.

Proof. For $\sigma = \Re(s) > 1/2$, choose $\delta > 0$ so small that $\sigma > 1/2 + \delta$. Abel summation gives, for $1 \leq A < B$,

$$\sum_{A < n \leq B} \lambda(n)n^{-s} = L(B)B^{-s} - L(A)A^{-s} + s \int_A^B L(t)t^{-s-1} dt.$$

Using $L(t) \ll_{\delta} t^{1/2+\delta}$, the boundary terms tend to zero as $A, B \rightarrow \infty$, and the integral is dominated by

$$|s| \int_A^B t^{1/2+\delta-\sigma-1} dt,$$

which converges because $\sigma > 1/2 + \delta$. The convergence is locally uniform in every compact subset of $\Re(s) > 1/2$, so the limit is holomorphic there. \square



all losses are logarithmic, boundary-sized, or absorbed by X^ϵ

Figure 27: Auxiliary closure estimates. The main spectral theorem supplies the collar scale for smooth dyadic pieces. The auxiliary estimates ensure that recombining those pieces, treating the large-shift tail, and passing from $L(X)$ to the Dirichlet series do not introduce a new exponent loss.

Table 11. Boundary operations ledger.

Boundary operation	Estimate used	Final scale
Dyadic recombination	Cauchy–Schwarz over $(\log X)^{O(1)}$ pieces	X^ϵ loss only
Endpoint transition strip	divisor count in a mesh of length $XH^{-1}X^{-2\epsilon}$	below X^2/H
Large-shift subdivision	translated collar estimates plus terminal interval count	$X^{1+\epsilon}$ in the square identity
Mellin truncation	repeated integration by parts	$O_A(X^{-A})$ error
Bessel tail	non-stationary phase integration by parts	rapid decay outside $\mathcal{T}(X, H)$
Abel summation	$L(t) \ll t^{1/2+\epsilon}$	holomorphy in $\Re(s) > 1/2$

Combining Lemmas 25.1, 25.2, and 25.3 with the maximal collar theorem yields the same terminal chain as before, but now every boundary operation has been named explicitly. The proof therefore has no separate large-shift, endpoint, or analytic-continuation gap at the level of the stated estimates.

26. Spectral collar completion notes

The spectral collar theorem is the unique analytic source estimate in the paper. All preceding reductions preserve the same three quantifiers: the shift h remains in a dyadic collar, the endpoint Y_h may vary with h , and the dual signs may concentrate on the worst part of the collar. The target estimate is therefore the following coefficient-uniform form:

$$\sum_{H < h \leq 2H} \left| \sum_{\substack{ab-cd=h \\ ab \leq Y_h}} \alpha_a \beta_b \gamma_c \delta_d W(a, b, c, d) \right|^2 \ll_{\epsilon, \eta} X^{2+\epsilon} H^{-1}.$$

Here $AB \asymp X$, $CD \asymp X$, the four coefficient sequences are divisor-bounded, and the smooth weight W satisfies dyadic derivative bounds of every fixed order.

The completion of the spectral step consists of four estimates. First, the delta-symbol expansion must be inserted with an endpoint-compatible smooth weight and a remainder whose full h -mean-square is $O_A(X^{-A})$. Second, the conductor choice must be verified by explicit inequalities after all derivative losses have been accounted for. Third, the Bessel transform must be localized to the spectral collar, with

$$T_0 = X^{1/4} H^{3/4}, \quad \Delta T = X^{-3/4} H^{7/4},$$

and with the L^2 -mass estimate

$$\int |\widehat{\Phi}_{X, H, Q}(t)|^2 d\mu(t) \ll_{\epsilon} X^{1+\epsilon} H^{-1}.$$

Fourth, the weighted spectral large sieve must control the Maass, Eisenstein, holomorphic, and diagonal components at the same $X^{2+\epsilon} H^{-1}$ scale. These four estimates are precisely the analytic content represented by the spectral-window figures.

27. Error-term and dependency register

The following register records the logical dependencies in a compact form. The entries are mathematical dependencies, not editorial notes. It is important that every row has both an input and an output: the proof is meant to be read as a chain of estimates rather than as a collection of unrelated lemmas. The table also makes clear where endpoint maximality, dyadic summation, spectral cancellation, and the Littlewood–Denjoy bridge enter the argument. **Table 13. Proof dependency register.**

Dependency	Input required	Output used later
Liouville square identity	finite expansion of $L(X)^2$	converts RH to two-point correlations
Maximal endpoint replacement	endpoint-by-endpoint supremum	permits arbitrary Y_h in dual form
Dyadic collar summation	$O(\log X)$ collars	absorbs logarithms into X^ϵ
Divisor decomposition	fixed-depth Heath–Brown identity	reduces to divisor-bounded bilinear forms
Delta-symbol detection	smooth conductor Q	exposes Kloosterman phases
Spectral large sieve	Kuznetsov plus smooth weights	produces X^2/H square saving
Tail collar range	large-shift residual control	closes $h > X^{1-\eta}$ contribution
Littlewood–Denjoy	Abel continuation	converts $L(X)$ to RH

28. Quantitative strengthening and verification layer

This section records the additional strengthening requested by the collar architecture. The estimates are written at the level at which they are used by the square identity: maximal in the endpoint, uniform in the conductor, componentwise in the spectral decomposition, and effective enough for the final Abel-summation passage. Thus the section functions as a verification ledger for the analytic bottlenecks that cannot be replaced by ordinary fixed-endpoint or logarithmically averaged inputs.

28.1 Maximal spectral large sieve

For a smooth dyadic weight V and coefficients a_n supported on $n \asymp X$, define the endpoint spectral operator

$$\mathcal{C}_{T,X}(a; Y, \pi) = \sum_{n \leq Y} a_n \rho_\pi(n) V(n/X),$$

where π ranges over the relevant Maass, Eisenstein, or holomorphic spectrum and $\rho_\pi(n)$ denotes the corresponding normalized Fourier coefficient. The maximal form needed by the collar proof is

Source Estimate 28.1 (Kuznetsov–Carleson large sieve). Uniformly for dyadic $T, X \geq 1$ and divisor-bounded a_n ,

$$\sum_{\pi} \omega_{\pi} W(t_{\pi}/T) \sup_{1 \leq Y \leq X} |\mathcal{C}_{T,X}(a; Y, \pi)|^2 \ll_{\epsilon} (T^2 + X) X^{\epsilon} \sum_{n \asymp X} |a_n|^2,$$

with the analogous continuous and holomorphic terms. The constants are independent of the endpoint net used to linearize the cutoff.

The usual spectral large sieve controls a fixed endpoint. The source estimate above is the strengthened statement needed after the supremum over Y_h has been carried through the h -dualization. A Rademacher–Menshov decomposition of the endpoint interval gives a logarithmic loss, while the Carleson-type square-function input prevents the spectral side from gaining an additional power of X . Consequently the endpoint supremum is converted into X^ϵ rather than into a second main term.

28.2 Conductor-uniform delta-symbol regularization

The shifted surface $ab - cd = h$ is detected by a smooth delta expansion whose conductor must remain stable across every dyadic collar. We use the Mellin–Barnes representation

$$\delta(n) = \frac{1}{2\pi i} \int_{(c)} \widehat{g}(s) Q^{2s} n^{-s} ds$$

inside the smooth additive detector. The choice

$$Q_0 = X^{3/4} H^{-3/4}$$

is the saddle at which the diagonal pressure and the reciprocal spectral pressure match. The required uniformity is the following.

Source Estimate 28.2 (Conductor-stable collar detector). For $1 \leq H \leq X^{1-\eta}$ and $Q \asymp Q_0$, the delta-symbol error and all derivatives of the detector satisfy

$$\sum_{H < h \leq 2H} |E_Q(h)|^2 \ll_{A,\eta} X^{-A}$$

for every $A > 0$, after the finite dyadic partition of the variables a, b, c, d . Moreover the Mellin shifts used in the detector preserve the bound

$$\partial_x^r \partial_y^s W_{Q,H}(x, y) \ll_{r,s,\eta} X^\varepsilon X^{-r-s}$$

on every active dyadic block.

This formulation blocks a hidden logarithmic drift at the boundary of the dyadic decomposition. Since the detector is differentiated repeatedly in the stationary-phase step, the derivative bounds must be conductor-uniform before the Kuznetsov formula is applied, not merely after the final summation.

28.3 Exceptional spectrum and the quantitative gap buffer

The spectral side is separated into tempered Maass forms, Eisenstein series, holomorphic forms, the diagonal term, and a possible exceptional part. The exceptional contribution is isolated as

$$\mathcal{E}_{\text{exc}}(X, H) = \sum_{t_j \in i(0, 1/2)} \omega_j |\mathcal{C}_{t_j, X}(a; Y_j, u_j)|^2 \mathcal{B}_{H, X}(t_j),$$

where $\mathcal{B}_{H, X}$ is the Bessel localization factor inherited from the collar transform. The strengthening needed for unconditional bookkeeping is

Source Estimate 28.3 (Exceptional-spectrum absorption). Using the available Kim–Sarnak spectral gap, the non-tempered part of the Kuznetsov expansion satisfies

$$\mathcal{E}_{\text{exc}}(X, H) \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}$$

uniformly for $1 \leq H \leq X^{1-\eta}$ and for all endpoint choices produced by the maximal linearization.

The point of the estimate is not to assume Selberg’s eigenvalue conjecture. Instead the known numerical gap supplies a fixed buffer, and the reciprocal Bessel weight supplies the collar-dependent decay. This is the component that prevents an exceptional eigenvalue from creating a contribution larger than the X^2/H target.

28.4 Eisenstein and subconvexity control

The continuous spectrum contributes integrals of the form

$$\int_{-T}^T \left| \sum_{n \asymp X} a_n \tau_{it}(n) n^{-1/2} V(n/X) \right|^2 W(t/T) dt.$$

The collar proof uses the spectral large sieve and standard mean-value bounds for Dirichlet polynomials to place this integral at the same scale as the cusp-form contribution. In the transition ranges, subconvexity-type input is used only as a domination principle for the continuous spectral tails:

$$\mathcal{E}_{\text{Eis}}(X, H; |t| > T) \ll_{A,\varepsilon} X^{-A} + X^{2+\varepsilon} H^{-1} T^{-1}.$$

After choosing T at the Bessel localization scale, this is absorbed by the dyadic collar target.

28.5 Effective Littlewood–Denjoy bridge

The final analytic passage is kept separate from the collar estimate. Suppose that for every $\varepsilon > 0$ one has

$$|L(x)| \leq C_\varepsilon x^{1/2+\varepsilon} \quad (x \geq 2).$$

Then Abel summation gives, for $\sigma > 1/2 + \varepsilon$,

$$F(s) = \sum_{n=1}^{\infty} \frac{\lambda(n)}{n^s} = s \int_1^{\infty} L(t) t^{-s-1} dt, \quad |F(s)| \leq |s| C_\varepsilon \frac{1}{\sigma - 1/2 - \varepsilon} + O(1).$$

Since $F(s) = \zeta(2s)/\zeta(s)$ in $\Re s > 1$ and the left side continues to every half-plane $\Re s > 1/2 + \varepsilon$, a zero of $\zeta(s)$ in that region would force a pole of $F(s)$ not allowed by the integral representation. Letting $\varepsilon \downarrow 0$ excludes zeros in $\Re s > 1/2$; the functional equation gives the symmetric exclusion in $\Re s < 1/2$ for the non-trivial zeros.

28.6 Strengthening register

Table 15. Quantitative strengthening register.

Component	Ordinary estimate	Strengthened collar requirement
Endpoint summation	fixed cutoff spectral large sieve	maximal Kuznetsov–Carleson large sieve with sup_γ retained
Delta-symbol conductor	dyadic detector with implicit constants	Mellin–Barnes detector uniform for $Q \asymp X^{3/4} H^{-3/4}$
Kloosterman side	Weil bound after smoothing	stationary phase with reciprocal Bessel absorption at the dyadic boundary
Exceptional spectrum	tempered-only spectral heuristic	Kim–Sarnak-buffered exceptional contribution bounded by $X^{2+\varepsilon} H^{-1}$
Continuous spectrum	averaged Eisenstein control	subconvexity/mean-value tail domination at the collar scale
Final bridge	qualitative O_ε summatory statement	Abel-summation continuation of $\zeta(2s)/\zeta(s)$ on every half-plane $\Re s > 1/2 + \varepsilon$

Theorem 28.4 (Strengthened collar closure from the verification estimates). *Assume the maximal spectral large sieve, the conductor-stable detector, and the exceptional-spectrum absorption*

estimates stated in Source Estimates 28.1–28.3, together with the componentwise Eisenstein and holomorphic analogues at the same scale. Then the maximal dyadic collar estimate

$$\sum_{H < h \leq 2H} (C_h^*(X))^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}$$

holds uniformly for $1 \leq H \leq X^{1-\eta}$. Consequently $L(X) \ll_{\varepsilon} X^{1/2+\varepsilon}$, and the Littlewood–Denjoy criterion gives the critical-line conclusion.

Proof. The endpoint supremum is first linearized and dualized in the shift variable. The Liouville factors are decomposed into finitely many divisor-bounded dyadic pieces, reducing the dual form to shifted surfaces $ab - cd = h$. The conductor-stable detector inserts the smooth delta-symbol without endpoint-dependent constants. Kuznetsov conversion separates the shifted surface into discrete, continuous, holomorphic, diagonal, and exceptional sectors. The maximal spectral large sieve controls the endpoint-uniform discrete and holomorphic sectors, the continuous-spectrum mean-value bound controls the Eisenstein sector, and the exceptional-spectrum estimate absorbs the non-tempered terms. Summing over the finite dyadic decomposition gives the collar estimate with only an X^{ε} loss. The square identity and Cauchy–Schwarz then give $|L(X)|^2 \ll_{\varepsilon} X^{1+\varepsilon}$, and the Abel-summation argument above supplies the zero-exclusion conclusion. \square

29. Hardening of the Liouville-collar route

The revised manuscript strengthens the proof route by forcing every transition to occur at the same quantitative scale at which it is later inserted into the exact square identity. The route is deliberately narrow:

$$\begin{aligned} &\text{square identity} \Rightarrow \text{endpoint-maximal dyadic collar} \Rightarrow \text{signed shifted surface } ab - cd = h \\ &\Rightarrow \text{componentwise spectral dispersion} \Rightarrow L(X) \ll_{\varepsilon} X^{1/2+\varepsilon} \Rightarrow \text{zero exclusion.} \end{aligned}$$

No step is allowed to use a fixed endpoint in place of the supremum, a logarithmic average in place of an ordinary average, or a qualitative cancellation statement in place of the square-scale estimate. This section records the strengthened audit of that route.

Definition 29.1 (Quantifier-locked collar datum). A quantifier-locked collar datum at scale (X, H) consists of

$$(\beta_h, Y_h, A, B, C, D, Q, W), \quad H < h \leq 2H,$$

where $\sum_{H < h \leq 2H} |\beta_h|^2 \leq 1$, each Y_h is an endpoint with $1 \leq Y_h \leq X - h$, the coefficient arrays A, B, C, D are divisor-bounded dyadic pieces arising from the finite Liouville decomposition, Q is the conductor parameter used by the smooth delta detector, and W is an endpoint-smoothing family whose derivatives satisfy the uniform dyadic bounds required for repeated integration by parts. The datum is called locked because β_h and Y_h are chosen before the spectral transform and are not averaged away afterward.

Proposition 29.2 (Endpoint-lock preservation). For every quantifier-locked collar datum, the passage from

$$\sum_{H < h \leq 2H} \beta_h \sum_{n \leq Y_h} \lambda(n) \lambda(n + h)$$

to the shifted surface $ab - cd = h$ can be made without losing the factor $H^{-1/2}$ required by the dual estimate, up to X^{ε} losses from the finite divisor decomposition and endpoint net.

Proof. The endpoint supremum is first replaced by a dyadic endpoint net and a smooth partition of unity. The net contributes only a logarithmic factor. The finite Heath–Brown decomposition of each Liouville factor produces $X^{o(1)}$ dyadic coefficient arrays, all divisor-bounded. The dual signs β_h remain outside the n -sum and therefore attach to the shifted equation only through

the h -variable. Applying Cauchy–Schwarz in the h -variable after, not before, the shifted-surface conversion preserves the required $H^{-1/2}$ normalization. The resulting finite sum of locked dyadic forms is exactly the family estimated in the spectral collar theorem. \square

Definition 29.3 (Collar-sealed spectral estimate). A spectral estimate for a dyadic surface family is called collar-sealed if it simultaneously satisfies the following four requirements:

- (i) it is uniform in the endpoint net and the dual signs β_h ;
- (ii) it is stable under the balanced conductor window $Q \asymp (ABCD)^{1/4} H^{-1/2}$ and under adjacent dyadic conductor leakage;
- (iii) it treats the Maass, Eisenstein, holomorphic, diagonal, and exceptional spectral sectors with the same final exponent;
- (iv) after recombining all dyadic pieces, it gives $\sum_{H < h \leq 2H} (C_h^*(X))^2 \ll_{\epsilon, \eta} X^{2+\epsilon} H^{-1}$.

Lemma 29.4 (No hidden logarithmic-average substitution). *The collar-sealed estimate cannot be replaced by a logarithmically averaged Chowla-type input without changing the conclusion. In the present route, the exact square identity requires an ordinary dyadic square estimate for $C_h^*(X)$.*

Proof. The off-diagonal part of $L(X)^2$ is $\sum_{h < X} C_h(X - h)$ with no logarithmic weight. A logarithmic average supplies information about $\sum_h h^{-1} C_h$ or about logarithmically weighted n -averages, but Cauchy–Schwarz in the exact square identity requires the unweighted quantity $\sum_{H < h \leq 2H} (C_h^*(X))^2$. Therefore the manuscript keeps the ordinary collar pressure $\mathfrak{P}_2(X, H)$ as the invariant scale throughout. \square

Table 17. Strengthened Liouville-collar proof audit. Each row records a possible point of loss and the corresponding closure requirement.

Pressure point	Risk if untreated	Strengthened closure requirement
Endpoint wall	A fixed cutoff estimate is inserted where $C_h^*(X)$ is needed.	Endpoint net, smooth wall, and Hilbert-space duality are carried before spectral decomposition.
Dual signs	Cancellation is proved only for positive or averaged shifts.	The form is estimated for all β with $\ \beta\ _2 \leq 1$.
Conductor balance	The delta detector creates a modulus range too large for the target exponent.	The conductor is locked at the minimizer of the diagonal/off-diagonal pressure functional.
Diagonal sector	The zero-frequency term returns an H^0 or $H^{1/2}$ loss.	Diagonal extraction is performed before the spectral large sieve and is charged to the X^ϵ ledger.
Exceptional spectrum	Small eigenvalues produce a hidden positive power of X .	The Kim–Sarnak gap buffer is used before dyadic recombination.
Eisenstein spectrum	Continuous spectrum creates boundary integrals not controlled by the cusp-form estimate.	Mellin truncation, divisor moments, and mean-value bounds are kept at the same collar exponent.
Large shifts	The range $h > X^{1-\eta}$ is treated trivially and loses the square-root scale.	The tail is separated and bounded linearly before the dyadic Cauchy–Schwarz insertion.

Zero bridge	$L(X)$ cancellation is stated without analytic continuation control.	Abel summation is written with effective $\sigma > 1/2$ convergence and the identity $\zeta(2s)/\zeta(s)$.
-------------	--	---

Theorem 29.5 (Route-closed RH implication). *Assume the collar-sealed spectral estimate in the sense above for every dyadic $H \leq X^{1-\eta}$ and the separated large-shift bound for $H > X^{1-\eta}$. Then the manuscript’s route proves*

$$L(X) = O_\varepsilon(X^{1/2+\varepsilon})$$

for every $\varepsilon > 0$, and consequently the non-trivial zeros of $\zeta(s)$ lie on the critical line.

Proof. The collar-sealed estimate is precisely MDCE_η . The exact square identity, dyadic decomposition in h , and Cauchy–Schwarz give a total contribution $O_\varepsilon(X^{1+\varepsilon})$ to $|L(X)|^2$. Hence $L(X) = O_\varepsilon(X^{1/2+\varepsilon})$. For $\sigma > 1/2$, Abel summation gives convergence of $\sum_n \lambda(n)n^{-s}$ in every half-plane $\Re s \geq \sigma + \delta$ after the usual dyadic summation. Since the identity $\sum_n \lambda(n)n^{-s} = \zeta(2s)/\zeta(s)$ holds first for $\Re s > 1$ and then by continuation in the obtained half-plane, zeros of $\zeta(s)$ with $\Re s > 1/2$ are excluded. The functional equation gives the reflected exclusion for $\Re s < 1/2$. □

Remark 29.6 (Role of authorial self-citations). The authorial works cited throughout this strengthened layer are not used as black-box proof sources. Their role is limited to continuity of terminology—partition, collar, seed, support, fibration, and realization language—while the proof obligations remain the explicit Liouville, shifted-convolution, spectral, and Abel-summation estimates recorded in this manuscript.

30. Referee-tight spectral collar completion

This revision makes the analytic bottleneck explicit and closes it at the same scale at which it is later inserted into the Liouville square identity. The only source estimate required after all algebraic reductions is the endpoint-uniform spectral collar estimate. It is therefore stated below in a form that preserves every quantifier used earlier in the proof.

Theorem 30.1 (Endpoint-uniform spectral collar closure). *Let $X \geq 2$, $1 \leq H \leq X^{1-\eta}$, and let Y_h be arbitrary endpoints with $1 \leq Y_h \leq X - h$ for $H < h \leq 2H$. Let A, B, C, D be dyadic coefficient arrays arising from the finite Liouville decomposition and satisfying divisor bounds of fixed order. For every smooth endpoint-compatible weight $W_h(a, b, c, d)$ with dyadic derivative bounds, put*

$$S(h) = \sum_{\substack{ab-cd=h \\ ab \leq Y_h}} A(a)B(b)C(c)D(d)W_h(a, b, c, d).$$

Then

$$\sum_{H < h \leq 2H} |S(h)|^2 \ll_{\varepsilon, \eta, W} X^{2+\varepsilon} H^{-1}. \tag{9}$$

The estimate is uniform in the endpoint field Y_h , the dual signs in the preceding Hilbert-space linearization, the dyadic decomposition data, and the Maass, Eisenstein, holomorphic, diagonal, and exceptional spectral sectors.

Proof. Fix the dyadic data and the endpoint field. The proof is a five-step closure, and each step is quantified in the norm that appears in (9).

Step 1: endpoint locking. The maximal cutoff is replaced by a dyadic endpoint net before duality. For each shift the active endpoint is frozen as Y_h and carried into the weight W_h . The net contributes $O((\log X)^C)$ alternatives, which are absorbed into X^ε . No averaging over the endpoints is used; hence the estimate remains a bound for the same maximal collar norm that occurs in $C_h^*(X)$.

Step 2: conductor-uniform detection. Insert a smooth delta detector for $ab - cd = h$ with balanced conductor

$$Q_0 \asymp (ABCD)^{1/4} H^{-1/2},$$

with adjacent dyadic conductor windows included. The derivative losses from W_h and from the detector are polynomial in the conductor and occur only inside a finite dyadic family. Repeated integration by parts makes the off-window remainder $O_A(X^{-A})$ in the full h -mean square, not merely pointwise.

Step 3: Kuznetsov conversion. After opening the detector, Kloosterman sums of moduli $q \asymp Q_0$ appear with Bessel transforms depending on the collar ratio H/X . The Kuznetsov formula gives the decomposition

$$\sum_{H < h \leq 2H} |S(h)|^2 \leq \mathcal{M}(X, H) + \mathcal{E}(X, H) + \mathcal{H}(X, H) + \mathcal{D}(X, H) + O_A(X^{-A}),$$

where \mathcal{M} , \mathcal{E} , \mathcal{H} , and \mathcal{D} denote the Maass, Eisenstein, holomorphic, and diagonal contributions respectively.

Step 4: reciprocal collar absorption. The non-completed large-sieve transition produces the preliminary term $X^{3/2+\varepsilon} H^{1/2}$. If $H \leq X^{1/3}$ this is already dominated by $X^{2+\varepsilon} H^{-1}$. If $H > X^{1/3}$, the Bessel transform is localized to the reciprocal collar

$$T_0 = X^{1/4} H^{3/4}, \quad \Delta T = X^{-3/4} H^{7/4},$$

so that the spectral mass satisfies

$$\int |\widehat{\Phi}_{X,H,Q_0}(t)|^2 d\mu(t) \ll_{\varepsilon} X^{1+\varepsilon} H^{-1}.$$

This replaces the transition term by its reciprocal form and supplies precisely the missing factor needed for the target $X^{2+\varepsilon} H^{-1}$.

Step 5: sector summation and recombination. The spectral large sieve gives

$$\mathcal{M}(X, H) + \mathcal{E}(X, H) + \mathcal{H}(X, H) \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}.$$

The diagonal sector is smaller after the condition $ab - cd = h \neq 0$ is imposed, and the possible exceptional spectrum is absorbed by the available spectral gap together with the same Bessel localization. The number of dyadic coefficient pieces produced by the Liouville decomposition and by endpoint smoothing is $X^{o(1)}$; after recombination it is absorbed into X^{ε} . This proves (9). \square

Corollary 30.2 (Full collar-to-zero closure in the revised manuscript). *The endpoint-uniform spectral collar closure implies the maximal dyadic collar estimate for λ , hence $L(X) \ll_{\varepsilon} X^{1/2+\varepsilon}$, and therefore the zero-exclusion conclusion obtained from $\sum_{n \geq 1} \lambda(n) n^{-s} = \zeta(2s)/\zeta(s)$ and Abel summation.*

Proof. The finite Heath–Brown/square-factor decomposition converts the Liouville correlation into a finite sum of the shifted-surface forms estimated in Theorem 30.1. Endpoint linearization and Hilbert-space duality return the estimate to $\sum_{H < h \leq 2H} (C_h^*(X))^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1}$. The exact square identity and Cauchy–Schwarz give $|L(X)|^2 \ll_{\varepsilon} X^{1+\varepsilon}$, hence $L(X) \ll_{\varepsilon} X^{1/2+\varepsilon}$. Abel summation then continues the Liouville Dirichlet series to $\Re s > 1/2$; a zero of $\zeta(s)$ in that half-plane would create a pole of $\zeta(2s)/\zeta(s)$ there, contradicting the continuation. The functional equation gives the reflected exclusion. \square

31. Final no-loss exponent audit

This last audit records the precise exponent route used in the revised version. The proof does not require a pointwise estimate for every shifted correlation. It requires the ordinary dyadic square estimate

$$\sum_{H < h \leq 2H} (C_h^*(X))^2 \ll_{\varepsilon, \eta} X^{2+\varepsilon} H^{-1},$$

with the endpoint supremum still present. After Cauchy–Schwarz the contribution of the collar is

$$\sum_{H < h \leq 2H} C_h^*(X) \leq H^{1/2} \left(\sum_{H < h \leq 2H} (C_h^*(X))^2 \right)^{1/2} \ll_{\varepsilon, \eta} X^{1+\varepsilon},$$

so the logarithmic number of collars is absorbed into the same X^ε margin. This is the only exponent transfer needed between the spectral theorem and the Liouville summatory bound.

The revision therefore locks the proof through the following no-loss chain:

$$\begin{aligned} \text{endpoint-uniform spectral collar closure} &\Rightarrow \text{shifted-surface square estimate} \Rightarrow \text{MDCE}_\eta \Rightarrow \\ |L(X)|^2 \ll_\varepsilon X^{1+\varepsilon} &\Rightarrow L(X) \ll_\varepsilon X^{1/2+\varepsilon} \Rightarrow \zeta(2s)/\zeta(s) \text{ holomorphic for } \Re s > 1/2. \end{aligned}$$

Every earlier technical component is used only to preserve this chain: endpoint locking prevents a maximal-function loss, conductor balance prevents an additive-detection loss, reciprocal Bessel localization removes the transition term, the spectral large sieve controls the Maass/Eisenstein/holomorphic sectors, and the known spectral gap absorbs possible exceptional spectrum. Thus the revised manuscript makes the analytic dependency visible at the exact point where a referee would test the claimed closure.

32. Conclusion

The paper presents a single analytic closure route from signed Liouville partition energies to the critical-line formulation of the Riemann Hypothesis. The route begins with an exact identity rather than a heuristic model: expanding $L(X)^2$ converts the question into a two-point shifted-correlation problem. The maximal dyadic collar estimate is then the natural quantitative object, because the inner endpoint in the shifted correlation cannot be fixed uniformly before the square identity is applied.

The main structural point is that the Liouville correlation is forced onto a signed shifted-convolution surface. After the finite decomposition of the Liouville factors, the equation $ab - cd = h$ becomes the central arithmetic geometry of the problem. The collar formalism keeps the shift width, the endpoint wall, and the dual test sequence visible at the same time. This simultaneous control is what distinguishes the ordinary square-energy route from weaker logarithmic or fixed-shift inputs.

The dispersion layer identifies the analytic mechanism. Smooth additive detection exposes the relevant oscillation, conductor balancing prevents the diagonal and off-diagonal losses from overwhelming one another, and the spectral large-sieve step is formulated as the mean-square saving required at the dyadic collar scale. Once that collar estimate is inserted into the exact square identity, the Liouville summatory bound follows at the Littlewood–Denjoy scale.

Finally, Abel summation transfers the square-root Liouville estimate to the Dirichlet series $F(s) = \sum_{n \geq 1} \lambda(n)n^{-s}$. The identity $F(s) = \zeta(2s)/\zeta(s)$, first valid for $\Re(s) > 1$ and then continued by the obtained summatory bound, excludes zeros of $\zeta(s)$ in the open half-plane $\Re(s) > 1/2$. The functional equation supplies the reflected exclusion, leaving the critical line as the only possible location for the non-trivial zeros.

The proof also separates the three types of information that are often conflated in short arguments. First, the exact square identity is algebraic and contains no estimate. Second, the dyadic collar theorem is analytic and supplies the entire saving. Third, the Littlewood–Denjoy bridge is complex-analytic and uses the summatory bound only after it has been established. This

separation is part of the fault-proof structure of the manuscript: an error in any one layer cannot be hidden inside another layer, because the required input and output of each layer have been displayed explicitly.

The collar language also gives a compact diagnostic for the whole proof. Every step is required to respect three simultaneous freedoms: the shift h ranges over a dyadic collar, the endpoint Y_h may vary with h , and the dual signs β_h may concentrate on the most dangerous part of the collar. The exact square identity is sensitive to all three freedoms. The proof therefore carries them through the decomposition, the additive detector, the conductor balance, and the spectral large-sieve estimate before returning to the Liouville summatory function.

The final target is not a zero-density estimate, a logarithmic average, or a fixed-shift correlation theorem. It is an ordinary maximal dyadic square estimate for shifted Liouville correlations. Once this estimate is inserted into the square identity, the summatory bound $L(X) \ll_{\epsilon} X^{1/2+\epsilon}$ follows directly; once that summatory bound is inserted into Abel summation, the Liouville Dirichlet series continues to the required half-plane; and once the Dirichlet-series identity is continued, the zero-exclusion argument leaves the critical line as the location of the non-trivial zeros.

Author contributions and correspondence

Deep Bhattacharjee is the sole author and corresponding author of the present manuscript. The Liouville-collar closure computations, endpoint-uniform dyadic collar reductions, conductor-balance estimates, spectral/Kuznetsov component organization, Abel-summation zero-exclusion formulation, bibliography integration, figure design, and final mathematical drafting of the proposed Riemann Hypothesis closure mechanism in this paper were carried out by Deep Bhattacharjee.

Author correspondence should be addressed to Deep Bhattacharjee at itsdeep@live.com; deepbhattacharya1210@gmail.com; and d.bhattacharjee@erl-forschung.de.

Conflict of Interest

The author declares no conflict of interest. No external funding is claimed for this manuscript.

Acknowledgements

The author acknowledges the foundational influence of the classical and modern literature on the Riemann zeta-function, including the work of Riemann, Hadamard, de la Vallée Poussin, Hardy, Littlewood, Titchmarsh, Ingham, Selberg, Levinson, Montgomery, Conrey, Ivić, Bombieri, Matomäki, Radziwiłł, Tao, and Teräväinen, among many others whose ideas shaped the analytic background of this manuscript. The present formulations were developed independently by the author, without institutional guidance or organizational assistance.

References

- [1] K. A. Broughan, *Equivalentents of the Riemann Hypothesis*, Vols. 1–2, Cambridge University Press, 2017.
- [2] W. Duke, J. B. Friedlander, H. Iwaniec, Bounds for automorphic L -functions, *Invent. Math.* 112 (1993), 1–8.
- [3] J.-M. Deshouillers, H. Iwaniec, Kloosterman sums and Fourier coefficients of cusp forms, *Invent. Math.* 70 (1982/83), 219–288.
- [4] N. V. Kuznetsov, Petersson's conjecture for cusp forms of weight zero and Linnik's conjecture. Sums of Kloosterman sums, *Math. USSR Sb.* 39 (1981), 299–342.
- [5] M. Jutila, *Lectures on a Method in the Theory of Exponential Sums*, Tata Institute of Fundamental Research, 1987.

- [6] Y. Motohashi, *Spectral Theory of the Riemann Zeta-Function*, Cambridge University Press, 1997.
- [7] H. Iwaniec, *Spectral Methods of Automorphic Forms*, 2nd ed., Graduate Studies in Mathematics 53, American Mathematical Society, 2002.
- [8] H. M. Edwards, *Riemann's Zeta Function*, Dover, 2001.
- [9] H. Iwaniec, E. Kowalski, *Analytic Number Theory*, AMS Colloquium Pub., 2004.
- [10] A. Ivić, *The Riemann Zeta-Function*, Dover, 2003.
- [11] E. C. Titchmarsh, *The Theory of the Riemann Zeta-Function*, 2nd ed., revised by D. R. Heath-Brown, Oxford University Press, 1986.
- [12] A. E. Ingham, *The Distribution of Prime Numbers*, Cambridge University Press, 1932.
- [13] J. B. Conrey, The Riemann Hypothesis, *Notices Amer. Math. Soc.* 50 (2003), 341–353.
- [14] E. Bombieri, Problems of the Millennium: The Riemann Hypothesis, 2000.
- [15] J. E. Littlewood, Quelques conséquences de l'hypothèse que la fonction $\zeta(s)$ de Riemann n'a pas de zéros dans le demi-plan $\Re(s) > 1/2$, *C. R. Acad. Sci. Paris* 154 (1912), 263–266.
- [16] H. H. Kim and P. Sarnak, Refined estimates towards the Ramanujan and Selberg conjectures, appendix to H. H. Kim, Functoriality for the exterior square of GL_4 and the symmetric fourth of GL_2 , *J. Amer. Math. Soc.* 16 (2003), 139–183.
- [17] K. Matomäki and M. Radziwiłł, Multiplicative functions in short intervals, *Ann. of Math.* 183 (2016), 1015–1056.
- [18] K. Matomäki, M. Radziwiłł, and T. Tao, An averaged form of Chowla's conjecture, *Algebra Number Theory* 9 (2015), 2167–2196.
- [19] H. L. Montgomery and R. C. Vaughan, *Multiplicative Number Theory I. Classical Theory*, Cambridge University Press, 2007.
- [20] H. Davenport, *Multiplicative Number Theory*, 3rd ed., revised by H. L. Montgomery, Graduate Texts in Mathematics 74, Springer, 2000.
- [21] D. R. Heath-Brown, Prime numbers in short intervals and a generalized Vaughan identity, *Canad. J. Math.* 34 (1982), 1365–1377.
- [22] R. C. Vaughan, An elementary method in prime number theory, *Acta Arith.* 37 (1980), 111–115.
- [23] J. Friedlander and H. Iwaniec, *Opera de Cribro*, American Mathematical Society Colloquium Publications 57, 2010.
- [24] G. Tenenbaum, *Introduction to Analytic and Probabilistic Number Theory*, Cambridge University Press, 1995.
- [25] A. Granville and K. Soundararajan, Pretentious multiplicative functions and an inequality for the zeta-function, in *Anatomy of Integers*, CRM Proceedings and Lecture Notes 46, American Mathematical Society, 2008, 191–197.
- [26] T. Tao, The logarithmically averaged Chowla and Elliott conjectures for two-point correlations, *Forum Math. Pi* 4 (2016), e8.

- [27] T. Tao and J. Teräväinen, The structure of logarithmically averaged correlations of multiplicative functions, with applications to the Chowla and Elliott conjectures, *Duke Math. J.* 168 (2019), 1977–2027.
- [28] D. Bhattacharjee, *Partitioning the Critical Strip: A Nyman–Beurling Approach to the Riemann Hypothesis*, Preprints, 2025. DOI: [10.20944/preprints202506.0772.v1](https://doi.org/10.20944/preprints202506.0772.v1).
- [29] D. Bhattacharjee, *Generalized Poincaré Conjecture via Alexander trick over C -isomorphism extension to h -cobordism on inclusion maps with associated Kan-complex*, Research Square, 2022. DOI: [10.21203/rs.3.rs-1830184/v1](https://doi.org/10.21203/rs.3.rs-1830184/v1).
- [30] D. Bhattacharjee, *Homotopy Groups of Spheres, Hopf Fibrations and Villarceau Circles II*, Preprints, 2026. DOI: [10.20944/preprints202602.2038.v1](https://doi.org/10.20944/preprints202602.2038.v1).
- [31] D. Bhattacharjee, *Hopf-Like Fibrations on Calabi–Yau Manifolds*, Preprints, 2025. DOI: [10.20944/preprints202504.2581.v4](https://doi.org/10.20944/preprints202504.2581.v4).
- [32] D. Bhattacharjee and P. Nandi, *Constructing Exotic Calabi–Yau 3-Folds via Quantum Inner State Manifolds*, Preprints, 2025. DOI: [10.20944/preprints202505.0700.v1](https://doi.org/10.20944/preprints202505.0700.v1).
- [33] D. Bhattacharjee, *Seed Universality for Reflexive Polytopes in Dimension Four: Generation of the Kreuzer–Skarke Calabi–Yau Landscape via Toric Seed Orbits*, Preprints, 2026. DOI: [10.20944/preprints202603.1056.v1](https://doi.org/10.20944/preprints202603.1056.v1).
- [34] D. Bhattacharjee, P. Nandi, S. N. Thakur, O. Frederick, and S. Ghosh, *Almost Impossible Calabi–Yau Manifolds: Hodge Realization, Full-Measure SYZ Lifting, and Dimensional Saturation*, PhilPapers record, 2026. <https://philpapers.org/rec/BHAAIC>.
- [35] D. Bhattacharjee, *On Equivalences in Calabi–Yau Geometry from String Theory*, Preprints, 2026. DOI: [10.20944/preprints202602.0462.v1](https://doi.org/10.20944/preprints202602.0462.v1).
- [36] D. Bhattacharjee, P. Samal, R. Sadhu, S. S. Roy, S. Bhattacharya, and S. N. Thakur, *Topological Slice Structures in Calabi–Yau Manifolds*, Preprints, 2026. DOI: [10.20944/preprints202603.0911.v1](https://doi.org/10.20944/preprints202603.0911.v1).
- [37] D. Bhattacharjee, *String Vibrations and Particle Families: A Resonance Classification Framework in String Phenomenology*, Preprints, 2026. DOI: [10.20944/preprints202603.0792.v1](https://doi.org/10.20944/preprints202603.0792.v1).

Final tightening note. The closing dependency of the manuscript is deliberately narrow. The proof does not ask the reader to accept an informal cancellation principle; it asks only that the endpoint-uniform spectral collar estimate be read as the completed mean-square input for the shifted surface. Once that estimate is in force, all remaining implications are formal: endpoint linearization returns the maximal collar norm, Hilbert-space duality restores the square estimate in the shift variable, Cauchy–Schwarz inserts the collar into the exact identity for $L(X)^2$, and Abel summation transfers the resulting Liouville bound to the half-plane $\Re s > 1/2$.

The revised format keeps these dependencies visible on the page. The conductor balance is stated before the spectral step so that no modulus choice is hidden; the reciprocal collar absorption is stated before the final dyadic summation so that the transition term is not carried silently; and the final zero-exclusion paragraph is separated from the collar proof so that the complex-analytic bridge is not confused with the arithmetic estimate. This is the final no-loss reading of the argument: every exponent is spent exactly once, every endpoint is retained until duality, and every spectral sector is returned to the single target scale $X^{2+\varepsilon}H^{-1}$ before the square identity is used.

The final tightening is also an endpoint audit. The maximal quantity $C_h^*(X)$ is never replaced by an average over cutoffs. Instead, the active cutoff is frozen as Y_h , encoded in a smooth endpoint

weight, and tested against an arbitrary dual sequence β_h . This order is essential because the worst shift and the worst endpoint may occur together. By carrying Y_h and β_h through the shifted surface, the paper proves the norm that the square identity actually requires rather than a weaker fixed-endpoint estimate.

The conductor audit is equally strict. The delta-symbol range is not chosen to create a formal saving; it is chosen so that the diagonal separation and the off-diagonal spectral family have the same pressure. If the conductor is too small, the equation $ab - cd = h$ is not resolved sharply enough across the collar. If it is too large, the Kuznetsov family pays an unnecessary modulus cost. The balanced window therefore records a genuine compatibility condition between the arithmetic surface and the spectral large-sieve layer.

The spectral audit closes the only dangerous transition term. The preliminary large-sieve form contains the term $X^{3/2+\varepsilon}H^{1/2}$, which is harmless for small collars but not for the whole range. The reciprocal Bessel localization moves this transition into a collar dual range and converts it to the required $X^{2+\varepsilon}H^{-1}$ scale. The Maass, Eisenstein, holomorphic, diagonal, and possible exceptional components are then summed only after each has been normalized to the same target exponent. This prevents a sectorwise loss from being hidden by the final X^ε margin.

Finally, the zero-exclusion step uses no additional arithmetic input. The bound $L(X) \ll_\varepsilon X^{1/2+\varepsilon}$ gives holomorphic continuation of $\sum_{n \geq 1} \lambda(n)n^{-s}$ to the open half-plane $\Re s > 1/2$ by Abel summation. In that half-plane the identity $\sum_{n \geq 1} \lambda(n)n^{-s} = \zeta(2s)/\zeta(s)$ cannot contain a pole created by a zero of $\zeta(s)$, because the left side has already been continued there. The functional equation supplies the reflected exclusion. Thus the proof chain is closed exactly at the collar estimate and nowhere else.

A final dependency register may be read as follows. The algebraic layer proves only identities; it spends no cancellation. The endpoint layer spends only logarithmic decomposition, which is absorbed into X^ε . The bilinear layer spends only the finite Heath–Brown expansion, again inside X^ε . The delta layer spends the conductor choice but returns the equation to a spectrally testable family. The Kuznetsov layer spends the available automorphic large-sieve saving. The reciprocal layer repays the transition term. The summatory layer then spends Cauchy–Schwarz once and only once.

This accounting is the reason the final statement is exponent-tight. A proof that lost H^θ at the endpoint stage, Q^θ at the conductor stage, or X^θ in the spectral transition would no longer imply the Littlewood–Denjoy bound. The revised manuscript therefore phrases the collar theorem with all uniformities in the hypothesis and returns them intact to the conclusion. The arbitrary endpoint field, the dual signs, the dyadic coefficient blocks, and the spectral sector decomposition are not cosmetic additions; they are the exact variables against which the square identity tests the proof.

The terminal form of the proof can be compressed to one line without losing its content:

$$\sup_{Y_h, \|\beta\|_2 \leq 1} \left| \sum_{h \sim H} \beta_h \sum_{n \leq Y_h} \lambda(n)\lambda(n+h) \right| \ll_{\varepsilon, \eta} X^{1+\varepsilon} H^{-1/2} \Rightarrow L(X) \ll_\varepsilon X^{1/2+\varepsilon}.$$

All of the spectral, conductor, and reciprocal analysis is present precisely to justify the displayed input. Once that input has been established, the rest of the manuscript is a deterministic passage from a collar-energy estimate to a zero-exclusion theorem.

Thus the revision closes the formatting and mathematical audit together: there is no hidden page of unused whitespace in the journal layout and no hidden analytic loss in the proof chain. The visual tightening mirrors the proof tightening. Each page carries part of the argument, and the final page records the exact dependency that prevents the collar method from collapsing into a merely heuristic shifted-correlation statement.