A Proof of the Binary Goldbach Conjecture for All Sufficiently Large Even Integers

Denis Saltykov (ds1678@gmail.com)

29 May 2026

Status: manuscript-MD publication overview.

This manuscript layer is prepared for LaTeX generation. The mathematical source is the lemma and external-input layer. The present file fixes the article order and records the exact files that a converter should include.

0.1 Main Theorems

0.1.1 Theorem 1.1 (Sufficiently large binary Goldbach)

For every sufficiently large even integer N, there exist primes p_1, p_2 such that

$$N = p_1 + p_2.$$

0.1.2 Theorem 1.2 (Weighted Goldbach asymptotic)

For every sufficiently large even integer N,

$$R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2) = \mathfrak{S}(N)N + o(N),$$

where

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p\mid N\\ p>2}} \frac{p-1}{p-2}, \qquad C_2 = \prod_{p>2}\left(1 - \frac{1}{(p-1)^2}\right) > 0.$$

For even N, $\mathfrak{S}(N) \geq 2C_2 > 0$. Theorem 1.1 follows from Theorem 1.2 by prime-power removal and positivity.

0.2 Canonical Build Order

The future LaTeX converter should read `build_order.md`; it should not scan the directory recursively.

0.3 Source Policy

Proof prose refers to logical objects: Theorem 1.2, Lemma B1, Lemma F3T, Lemma E10X, X9L-GT, X10, and so on. File paths appear only in Appendix G and in workflow metadata. Mathematical corrections must be made first in the proof-source layer and then propagated back into this manuscript layer.

1 Abstract and Front Matter

Author: Denis Saltykov.

Affiliation: independent researcher.

Email: ds1678@gmail.com.

Date: 29 May 2026.

MSC 2020: primary 11P32; secondary 11N05, 11N36, 11L05, 11B30.

Keywords: binary Goldbach problem, von Mangoldt correlations, Heath–Brown identity, Kloosterman fractions, Liouville orthogonality, local density, routing lemma.

1.1 Abstract

We prove that every sufficiently large even integer is a sum of two primes. The proof establishes the weighted asymptotic

$$\sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2) = \mathfrak{S}(N)N + o(N)$$

for even N, with the classical Goldbach singular series $\mathfrak{S}(N)$. The argument begins with a fixed-depth Heath–Brown decomposition of both von Mangoldt factors and routes every resulting B1-origin atom into one of five terminal classes: Edge, CKP, GoodAWACK, LongAP/Local, or LocalDiag. Edge terms are estimated by strict savings, CKP terms by a Duke–Friedlander–Iwaniec Kloosterman-fraction input after a smooth-weight derivative verification, GoodAWACK terms by a TC1 global testing route and a finite-grammar closure, and local terms by explicit tagged Λ_Q-local projections. A global error-budget lemma verifies that all terminal errors are simultaneously summable. Prime-power removal and positivity of the singular series then give a genuine prime representation.

2 Introduction and Relation to Known Work

The binary Goldbach problem asks whether every even integer greater than 2 is a sum of two primes. The classical circle method proves strong averaged statements and, beginning with Vinogradov, proves ternary Goldbach for sufficiently large odd integers. The completed weak Goldbach theorem removes the remaining finite range in the ternary problem. The binary problem is more delicate because the expected main term has only one degree of freedom, and minor-arc cancellation must interact with arithmetic local factors without losing the main term.

The present proof is organized around a finite decomposition-and-routing architecture. A fixed-depth Heath–Brown identity expands the two von Mangoldt factors into finitely many typed product variables. Each product block is then partitioned by deterministic routing operations. The routing theorem is designed so that every terminal atom belongs to a class with a specified analytic or local mechanism.

The five terminal classes are:

1. Edge, where strict size, conductor, boundary, or square-divisor savings give $o(N)$;
2. CKP, where a balanced bilinear Kloosterman-fraction structure is present;
3. GoodAWACK, where the remaining Branch B atoms are handled by the TC1 global testing route and the E10Y-certified finite GoodAWACK routing grammar;

4. LongAP/Local, where no nonlocal arithmetic coefficient survives and the term satisfies the explicit H4 tagged local projection condition;
5. LocalDiag, where forced local dependence is admitted only when it is the same tagged LPI local projection later assembled by H4.

Two features distinguish the route from a direct circle-method argument. First, the proof does not try to estimate all descendants with one universal analytic theorem. It proves a complete finite classification and uses different tools on the different terminal classes. Second, the proof avoids the earlier inverse-Gowers route X8. The GoodAWACK branch is instead closed by the combination of a TC1 global testing dichotomy and a structural finite-grammar theorem for rank-dropping affine regroupings.

The most delicate points are the completeness of the terminal routing, the CKP/X10 smooth-weight match, the X16/Shiu carrier-slice estimate used inside the TC1 route, the GoodAWACK finite-grammar closure, and the local-main reconstruction. These points are isolated as named lemmas and appendices rather than hidden in the final assembly.

3 Main Theorems and Proof Strategy

3.1 Weighted Form

Let

$$R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2).$$

The central theorem is Theorem 1.2:

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$$

for sufficiently large even N.

The singular series is

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p\mid N\\ p>2}} \frac{p-1}{p-2}, \qquad C_2 = \prod_{p>2}\left(1 - \frac{1}{(p-1)^2}\right).$$

Since $\mathfrak{S}(N) \geq 2C_2 > 0$ for even N, the weighted asymptotic is eventually positive.

3.2 Decomposition and Routing

Lemma B1 gives an exact finite Heath–Brown decomposition of both factors in $R_\Lambda(N)$. Lemmas B3, F3, F4, E5, and F3T convert every B1-origin atom into a finite tagged sum of terminal atoms in exactly the five terminal classes

Edge, CKP, GoodAWACK, LongAP/Local, LocalDiag.

The tagged partition is exact; overlaps between visual algebraic shapes do not create double counting because the routing history is part of the terminal tag.

3.3 Terminal Estimates

The terminal estimates are:

1. Edge terms first satisfy one of the strict C1P saving predicates; C1A verifies admission of all active Edge inputs, and C1 proves their total contribution is $o(N)$.
2. CKP terms equal the explicit H4 tagged local projection plus $o(N)$ by G8a, CKPD, X10, and B1LD.
3. GoodAWACK terms are $o(N)$ by E10L, TNG, X9L-GT, X16BRS/X16C, E10Y, E10M, E10X, and E10K.
4. LongAP/Local terms satisfy the H4 tagged admission condition plus $o(N)$ by D1 and H4.
5. LocalDiag terms are admitted only when they are tagged LPI local projections later assembled by H4.

The global error-budget lemma GEB proves that these estimates remain $o(N)$ after all dyadic and routing summations.

3.4 Assembly

The local projections admitted by D1, G8a/B1LD, and LocalDiag are all passed through H4. H4 uses the explicit Λ_Q-model, tagged linearity over the B1/F3 partition, no double counting, and the CRT local factor computation to prove

$$M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).$$

Combining this with the terminal error bounds gives Theorem 1.2. Lemmas G2, G1, and G0H then remove nontrivial prime powers and convert positivity into a prime representation.

4 Notation, Parameters, and Error Budget

Throughout, N is an even integer tending to infinity and

$$L = \log N.$$

All implicit constants may depend on the fixed decomposition depth and on the fixed smooth cutoffs, but not on N.

4.1 Fixed Depth and Dyadic Partitions

The Heath–Brown depth is fixed once and for all. We use a parameter J_0 large enough for the global hierarchy. The explicit consistency witness used in the parameter register is

$$J_0 = 20, \qquad \eta = \frac{1}{40}, \qquad \theta = \frac{1}{4000}.$$

The witness is not optimized. Its role is to show that all inequalities in the parameter hierarchy can be satisfied simultaneously.

Every variable produced by the B1 decomposition is smoothly dyadically localized. For fixed J_0, the total number of active dyadic and routing cells is $L^{O(1)}$.

4.2 Terminal Classes

The proof uses the following terminal classes.

Class	Meaning	Output
Edge	strict saving, boundary, short volume, square-divisor, or conductor loss	$o(N)$
CKP	balanced Kloosterman-fraction branch	local projection $+o(N)$
GoodAWACK	Branch B affine/global-testing branch	$o(N)$
LongAP/Local	long local arithmetic progression branch	local projection $+o(N)$
LocalDiag	forced local dependence or diagonal branch	local projection

4.3 Global Error Budget

GEB records the summability principle used throughout the proof. If a per-terminal estimate gains either a fixed power of N or a sufficiently large negative power of L, then the $L^{O(1)}$ terminal multiplicity is harmless.

The constants are chosen in the following order:

1. fix $\theta \ll \eta$;
2. choose $J_0 \geq J_*(\eta)$;
3. fix the routing multiplicity constant $C_0(J_0)$;
4. fix the Edge and divisor losses;
5. choose the X16/Shiu logarithmic floor exponent;
6. choose the TC1 near-global modulus and AP exponents;
7. choose CKP high-frequency and DFI derivative thresholds.

With this order, later logarithmic exponents can always dominate earlier polylogarithmic losses. GEB is invoked only as a bookkeeping lemma; it does not replace the branch estimates.

5 External Inputs

The proof uses four external inputs, each in a fixed stated form.

5.1 X1: Heath–Brown Identity

X1 is the fixed-depth Heath–Brown identity used to decompose each von Mangoldt factor. The active formulation is the one needed by B1: after choosing J_0, Λ is represented by a finite linear combination of typed divisor-convolution factors with controlled coefficients and smooth dyadic localization.

5.2 X9L-GT: Near-Global Liouville/AP Orthogonality

X9L-GT is the Davenport/AP input used only after the TC1 route has produced near-global active B1-origin coarea tests. It is not invoked on arbitrary shifted short intervals. TTH supplies the near-global length condition

$$H \geq X(\log X)^{-B_\kappa}.$$

The TC1 testing family has only polylogarithmic modulus and AP complexity, so the Davenport saving can be chosen to dominate the recorded losses.

5.3 X10: DFI Kloosterman-Fraction Estimate

X10 is the Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction estimate in the smooth weighted form required by CKP. The CKP branch uses X10 only after G1a–G3a reduce the nonzero-frequency terms to the DFI form and CKPD verifies the actual two-variable smooth weight.

No later Kloosterman-fraction strengthening is an active dependency.

5.4 X16: Shiu/AP Divisor-Average Input

X16 is the Shiu-type arithmetic-progression divisor-average input used in the BRS carrier-slice step of the TC1 route. X16C proves the active normalized carrier estimate, including the divisor-function second moment and the logarithmic floor needed by GEB.

5.5 Citation Boundary

Each external theorem is invoked only through its named stated interface: X1, X9L-GT, X10, and X16. The manuscript does not appeal to informal variants of these results.

6 Heath–Brown Decomposition

6.1 B1 Blocks

Lemma B1 applies the fixed-depth Heath–Brown identity to both copies of Λ in $R_\Lambda(N)$. It writes $R_\Lambda(N)$ as a finite sum of B1 blocks. Each B1 block consists of two finite lists of product variables, smooth dyadic cutoffs, and elementary coefficient types

$$\mu \cdot 1_{\leq y}, \qquad 1, \qquad \log .$$

For fixed J_0, the number of elementary variables in a B1 block is bounded, and the number of dyadic cells is $L^{O(1)}$.

6.2 Exactness

The B1 decomposition is an identity before estimates are applied. Smooth dyadic partitions are inserted as exact partitions of unity up to boundary terms already routed to Edge. Thus no main term is lost at the decomposition stage.

6.3 Structural Output

The output of B1 is a structural finite-convolution problem together with the equation inherited from $n_1 + n_2 = N$. The routing layer decides which grouped variables, divisor relations, local congruences, and oscillatory features survive to terminal atoms.

The subsequent proof uses only B1-origin atoms. This origin condition is important in the GoodAWACK finite-grammar argument: formal affine systems that do not arise from the B1/B3/F3/F4 grammar are not terminal descendants of the proof.

7 Routing Exhaustion

7.1 Finite Grouping

Lemma B3 supplies a finite set of product-grouping candidates for each B1 block. Candidate labels may overlap at this preliminary stage; they are not yet terminal classes.

7.2 Authoritative Routing Operations

Lemma F3 supplies the complete routing operations for actual F3 atoms. These include controlled CRT absorption, ordinary divisor decisions through F4, square-divisor routing, grouping selection or elimination, terminal Edge detection, terminal LocalDiag detection, and final terminal labelling.

Lemma E5 is used only as content stability for transports generated by the routing grammar. It is not an independent generator of terminal affine systems.

7.3 Exhaustion Theorem

Lemma F3T proves the finite routing exhaustion theorem:

$$\mathrm{B1} + \mathrm{B3} + \mathrm{F3} + \mathrm{F4} + \mathrm{E5} \Longrightarrow \{\mathrm{Edge}, \mathrm{CKP}, \mathrm{GoodAWACK}, \mathrm{LongAP/Local}, \mathrm{LocalDiag}\}.$$

More precisely, each B1-origin atom is partitioned into a finite disjoint sum of tagged terminal atoms in exactly those five classes. No sixth terminal class remains.

The proof uses a deterministic routing order and a well-founded routing measure $\mathfrak{M}^\sharp$. Nonterminal operations strictly decrease this measure, while terminal rows assign one of the five tags. If a cell visually satisfies more than one terminal predicate, the tag records the first applicable class in the deterministic order. Therefore visual overlap does not imply double counting.

7.4 Terminal Interfaces

Once F3T has assigned terminal tags, the estimates are supplied by the corresponding branch lemmas:

$$\mathrm{Edge} \to C1P/C1A/C1, \quad \mathrm{CKP} \to G8a, \quad \mathrm{GoodAWACK} \to E10L, \quad \mathrm{LongAP/Local} \to D1/H4, \quad \mathrm{LocalDiag} \to H$$

8 Edge Estimates

8.1 Edge Predicates

Edge terms are terminal cells carrying a strict saving predicate. Lemma C1P defines these predicates before any late branch estimate is invoked. The active predicates include boundary or short-volume loss, Type I saving, large gcd/content, square-divisor tails, high Fourier tails, and small-conductor layers.

Lemma C1 proves that any cell satisfying one of the strict C1P Edge predicates contributes $o(N)$ after its allowed coefficient and polylogarithmic losses.

8.2 Admission

Lemma C1A proves the complementary admission statement needed by the manuscript: every active terminal atom routed into Edge satisfies one of the strict C1P predicates. Thus Edge is not a residual label. It is a verified saving class.

The admission ledger treats the Edge sources from F3/F4, CKP excluded ranges, BRS carrier slices, square-divisor routing, boundary terms, and high-frequency tails. Each source is paired with a saving estimate and a summability check.

8.3 Summation

For fixed J_0, the number of Edge cells is $L^{O(1)}$. C1 supplies either a fixed power saving in N or a logarithmic saving chosen larger than the routing multiplicity. GEB therefore gives

$$R_{\mathrm{Edge}}(N) = o(N).$$

9 LongAP/Local Terms

9.1 Terminal LongAP/Local Atoms

A LongAP/Local atom is a terminal B1-origin cell satisfying the intrinsic F3P LongAP/Local predicate: its long-variable coefficients lie in the controlled local coefficient algebra on a long arithmetic-progression fibre. It is not allowed to contain a surviving nonlocal μ-, λ-, Fourier-, Kloosterman-, or nilsequence-type coefficient.

9.2 Exclusion of Nonlocal Coefficients

Lemma D1.2A proves that the F3P LongAP/Local predicate has the advertised consequence inside the routed B1/F3 partition: no nonlocal arithmetic coefficient survives in a terminal LongAP/Local atom. The proof combines the positive F3P predicate with the finite F3/F4 routing alternatives:

1. a surviving oscillatory or Liouville/Mobius-type factor routes to GoodAWACK or CKP;
2. a strict saving predicate routes to Edge;
3. forced local dependence routes to LocalDiag;
4. only local AP/congruence data remains in LongAP/Local.

Thus D1 may replace F3P-LongAP/Local atoms by the explicit H4 tagged local projections without discarding hidden arithmetic oscillation.

9.3 Output

Lemma D1 gives, for each LongAP/Local terminal cell,

$$R_{\text{LongAP/Local}}(N) = M_{\text{LongAP/Local}}(N) + o(N),$$

where the main term is the H4 Λ_Q-projection of the same tagged B1/F3 cell. Boundary and smoothing errors are Edge-admitted.

10 The CKP Branch

10.1 CKP Claim

The CKP branch proves

$$R_{\text{CKP}}(N) = M_{\text{CKP}}(N) + o(N),$$

where $M_{\text{CKP}}(N)$ is the explicit H4 tagged local projection obtained by the same Λ_Q-replacement used in Appendix E.

10.2 Internal Reduction

The CKP proof follows the chain

$$G1a + G2a + G3a + CKPD + G4a/X10 + X10\text{-}ER + B1LD \Longrightarrow G8a.$$

Lemma G1a performs the gcd splitting $u = ga$, $u' = gq$, with $(a,q) = 1$. Lemma G2a performs the smooth AP/Fourier expansion. The zero-frequency term is local. Lemma G3a writes each nonzero central frequency as a bilinear Kloosterman-fraction sum

$$\sum_{\substack{a\sim A_g,\ q\sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q)W_{g,h}(a,q)e\left(\frac{hN_g\overline{a}}{q}\right).$$

10.3 X10 Match

X10 is invoked only for the central nonzero-frequency range. CKPD proves that the actual two-variable Fourier weight $W_{g,h}(a,q)$, including the dependence $z(a,q,y) = (N_g - ay)/q$, satisfies the smooth derivative hypotheses of the DFI Kloosterman-fraction theorem with only polylogarithmic loss.

Excluded ranges are not sent to X10. The X10ER routing statement sends them to Edge through C1P/C1A/C1, to local zero-frequency terms, or to the CKP auxiliary exclusions recorded in G8a.

10.4 Output

The nonzero-frequency CKP contribution is $o(N)$. The zero-frequency term is identified by B1LD with the local model used by H4. Therefore G8a supplies the CKP input needed by I1.

11 The GoodAWACK Branch

11.1 Branch B Theorem

Lemma E10L proves that the total terminal GoodAWACK contribution is $o(N)$. The proof has two components:

1. the TC1 global testing route;
2. the HighTC finite-grammar closure.

11.2 TC1 Global Testing

The TC1 route does not choose an arbitrary shifted short interval. It tests only structural B1-origin coarea families whose cells have not already been routed away. Theorem TNG-A gives the single near-global-or-routed interface:

every unrouted TC1 test is near-global, or routed away before X9L-GT.

In the near-global alternative the test satisfies

$$H \geq X(\log X)^{-B_\kappa}.$$

Thus X9L-GT applies in its Davenport/AP form with polylogarithmic modulus and AP complexity. The resulting orthogonality estimate is summable by GEB.

In the routed alternative, TTD, ROC, BRS, and X16BRS/X16C send the test to Edge, LongAP/Local, CKP, LocalDiag, or zero before X9L-GT is invoked.

11.3 Finite GoodAWACK Grammar

The HighTC branch is closed by the finite combinatorial grammar package E10Y/E10X/E10M/E10K. E10Y proves that every actual terminal GoodAWACK skeleton is generated from B1/B3 grouped cells by the listed B1/B3/F3/F4/E5 grammar: fixing/projection, controlled CRT restriction, fixed divisor quotient, F4-tagged variable quotient, local/diagonal/gcd dependence, CKP-balanced structure, Edge-type saving or boundary routing, full-rank affine regrouping, post-terminal slicing after terminal vectors are fixed, E5 auxiliary inheritance, and final terminal labelling. E5 is used only as content stability and is not an independent terminal generator.

E10X proves by induction on the E10Y-certified finite grammar that every rank-dropping affine operation created along a derivation carries an allowed origin tag. E10S and E10S-MECH are non-logical source-maintenance and reproducibility records for the maintained Branch B source list; they are not premises of the induction. E10M then proves that no untagged rank-dropping AFF occurrence survives in an actual terminal GoodAWACK skeleton. E10K converts this into AFF-origin completeness, and E10X eliminates the FreeAffineHighTC residual.

The formal $4AP$-like family $Y_i = x + ir$ remains useful as an interface test, but E10X proves that it has no untagged actual terminal occurrence in the finite routing grammar.

11.4 Output

TC1 contributes $o(N)$, singular tests are routed to already handled classes, and HighTC finite-grammar residuals are eliminated structurally. Hence

$$R_{\text{GoodAWACK}}(N) = o(N).$$

12 Local/Main Assembly

12.1 Explicit H4 Local Algebra

The local/main layer is assembled by a concrete finite local model, not by an undefined projection convention. Let

$$Q = \prod_{p \le w} p, \qquad w = w(N) \to \infty, \qquad w = o(\log N),$$

and define

$$\Lambda_Q(a) = \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1}, \qquad \sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a).$$

The local projection of the original Goldbach convolution is

$$\mathrm{Loc}_Q\, R_\Lambda(N) = N\sigma_Q(N).$$

For a terminal tagged cell $(\mathcal{B}, \tau)$, LPI defines $\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N)$ by replacing the arithmetic coefficients inside that same tagged cell by their residue-class local densities modulo Q, preserving the parent B1 block, routing tag, dyadic weights, and local congruence data.

12.2 H4 Admission

H4 admits a terminal local/main expression only if it satisfies the tagged admission condition

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N).$$

Thus a local-looking expression does not enter the main term merely because it resembles a local density. It must be attached to a parent B1 block and a terminal routing tag, and its normalization must be the single Λ_Q-replacement above.

The admitted local sources are:

1. LongAP/Local main terms whose F3P terminal predicate already forces the long-variable coefficients into the local coefficient algebra, with D1 expanding them into the tagged local projection;
2. CKP zero-frequency terms from G8a/B1LD, after the zero mode is identified with the same tagged Λ_Q-replacement;
3. LocalDiag terms, admitted only when the diagonal specialization is a tagged local projection;
4. harmless local boundary contributions explicitly admitted by H4 and already bounded through C1.

12.3 Linearity and No Double Counting

The exact B1 decomposition and the exact F3 tagged partition give

$$R_\Lambda(N) = \sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N).$$

The operator Loc_Q is linear on this finite tagged partition. Therefore

$$\sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau\in\mathcal{T}(\mathcal{B})} \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_\Lambda(N) + o(N) = N\sigma_Q(N) + o(N).$$

No double counting occurs because each local term is indexed by its parent B1 block $\mathcal{B}$ and its routing tag τ. Different B1 blocks are different summands of the exact B1 identity, and different tags inside the same block are disjoint cells of the F3 partition. Algebraically identical LocalDiag-looking formulas from different tags are therefore distinct partition summands, not duplicates.

12.4 Finite Local Factors

The finite density $\sigma_Q(N)$ factors by the CRT. For $p \le w$,

$$\sigma_p(N) = \frac{1}{p}\left(\frac{p}{p-1}\right)^2 \#\{a \bmod p : (a,p)=1,\ (N-a,p)=1\}.$$

For even N, $\sigma_2(N) = 2$. For odd p,

$$\sigma_p(N) = \begin{cases} \dfrac{p}{p-1}, & p \mid N, \\ 1 - \dfrac{1}{(p-1)^2}, & p \nmid N. \end{cases}$$

Hence

$$\sigma_Q(N) \to 2C_2 \prod_{\substack{p\mid N\\ p>2}} \frac{p-1}{p-2} = \mathfrak{S}(N).$$

12.5 Weighted Assembly

Combining the LPI admission condition consumed by H4, tagged linearity, no double counting, and the finite local factor computation gives

$$M_{\mathrm{local}}(N) = N\sigma_Q(N) + o(N) = \mathfrak{S}(N)N + o(N).$$

Lemma I1 then combines the terminal estimates and H4:

$$R_\Lambda(N) = M_{\mathrm{local}}(N) + o(N) = \mathfrak{S}(N)N + o(N).$$

GEB verifies that all terminal $o(N)$ terms remain $o(N)$ after the full polylogarithmic decomposition and routing summation.

13 Prime-Power Removal and Final Proof

13.1 Statement

Assume the weighted asymptotic I1:

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$$

for sufficiently large even N. Then, after removing the negligible nontrivial prime-power contribution by G2, the genuine prime-pair weighted sum is positive. Consequently every sufficiently large even N is a sum of two primes.

13.2 Setup

The von Mangoldt function is supported on prime powers:

$$\Lambda(n) = \begin{cases} \log p, & n = p^k,\ k \geq 1, \\ 0, & \text{otherwise.} \end{cases}$$

Let $R_{pp}(N)$ denote the ordered weighted sum restricted to genuine prime pairs:

$$R_{pp}(N) = \sum_{\substack{p_1+p_2=N \\ p_1,p_2 \text{ prime}}} (\log p_1)(\log p_2).$$

The ordered-pair convention is the same convention used in B1, I1, G2, G1, and G0H.

13.3 Proof

Lemma G2 removes the nontrivial prime powers p^k, $k \geq 2$. There are $O(N^{1/2})$ such prime powers up to N, and once one coordinate in $n_1 + n_2 = N$ is selected the other coordinate is determined. Since $\Lambda(n) \leq \log N$, the total contribution from pairs in which at least one coordinate is a nontrivial prime power is

$$O(N^{1/2}(\log N)^2) = o(N).$$

Therefore I1 and G2 imply

$$R_{pp}(N) = \mathfrak{S}(N)N + o(N).$$

For even N, the Goldbach singular series satisfies

$$\mathfrak{S}(N) \geq 2C_2 > 0.$$

Hence $R_{pp}(N) > 0$ for all sufficiently large even N. Since every summand in $R_{pp}(N)$ is nonnegative and the summands on genuine prime pairs are positive, this positivity implies that at least one ordered pair of primes p_1, p_2 satisfies

$$p_1 + p_2 = N.$$

13.4 Output

Lemma G0H records the ordered-pair normalization and the final logical handoff

$$I1 + G2 \Longrightarrow G1 \Longrightarrow G0.$$

The statement proved is the sufficiently-large binary Goldbach theorem. Finite verification below the asymptotic threshold is outside the asymptotic claim of this manuscript.

A Routing Tables

This appendix records the routing table used by Lemma F3T. The terminal predicates themselves are the intrinsic predicates fixed in Lemma F3P; the branch estimates are invoked only after the routing tag has been assigned.

A.1 Terminal Classes

Class	Entry condition	Exit theorem
Edge	strict C1P saving predicate	C1P/C1A/C1
LocalDiag	forced local dependence or diagonal relation	H4
LongAP/Local	F3P long AP/local fibre whose long-variable coefficients lie in the local coefficient algebra	D1/H4
CKP	balanced bilinear Kloosterman-fraction structure	G8a/X10
GoodAWACK	Branch B affine/global-testing residual with controlled origin	E10L

A.2 Canonical Routing Order

On each tagged B1-origin cell, F3T reads the following order.

1. Empty or incompatible cells are zero.
2. Strict C1P saving predicates go to Edge.
3. Forced equality, proportionality, repeated factor, or local dependence goes to LocalDiag.
4. Ordinary divisor and quotient predicates are decided by F4.
5. Square-divisor obstructions are routed by F3.
6. Controlled CRT restrictions are absorbed only when E5 verifies clean content stability.
7. Remaining grouping alternatives are selected or eliminated from the finite B3 grouping set.
8. The terminal predicate assigns one of the five terminal classes.

Every nonterminal step strictly decreases $\mathfrak{M}^\sharp$. Therefore the process terminates.

A.3 Routing Table

Row	Source regime	Routing outcome	Reason no other terminal class receives the cell
1	empty support or incompatible congruences	Edge-zero	no analytic mass remains
2	boundary, short volume, large content, square-divisor tail, high Fourier tail, small conductor, or Type I saving	Edge	a strict C1P saving predicate is present

3	forced equality, proportionality, repeated factor, or local dependence	LocalDiag	independent CKP/GoodAWACK variables are absent
4	one long AP variable and F3P local coefficient algebra	LongAP/Local	the positive F3P predicate already excludes nonlocal long-variable coefficients; D1.2A expands the resulting local algebra
5	central balanced bilinear Kloosterman-fraction form	CKP	Edge/LocalDiag have already failed; the structure is bilinear, not GoodAWACK
6	nonlocal non-CKP Branch B affine residual with controlled origin	GoodAWACK	E10Y/E10X/E10M/E10K exclude untagged rank-dropping AFF residuals
7	ordinary divisor predicate with short quotient or saving	Edge	F4 supplies a C1P predicate
8	ordinary divisor predicate forcing local dependence	LocalDiag	quotienting identifies active forms
9	ordinary divisor predicate preserving balanced bilinear structure	CKP	F4 removes local and short alternatives
10	ordinary divisor predicate preserving controlled Branch B affine residual	GoodAWACK	quotient origin is tagged and E5-clean
11	large square divisor	Edge	square-divisor saving applies
12	small controlled square divisor or full-rank CRT restriction	nonterminal decrease	F3 continues with smaller $\mathfrak{M}^\sharp$
13	unresolved finite grouping alternative	nonterminal decrease or terminal row	B3 grouping set is finite

Rows 12–13 are not terminal rows. They are included to show that all nonterminal transitions are among the allowed F3 operations.

A.4 Consequence

For fixed J_0, every B1-origin atom reaches exactly one tagged terminal class after finitely many steps. This is the routing-exhaustion input used in the assembly theorem.

B Edge Admission

This appendix records the three-layer Edge interface: C1P defines the strict Edge predicates, C1A verifies that every active Edge input satisfies one of them, and C1 estimates the admitted terms.

B.1 Edge Predicates

The active C1P predicates are:

1. boundary or short-volume loss;
2. Type I saving;
3. large gcd or large content;
4. large square-divisor tail;
5. high Fourier tail;
6. small conductor;
7. incompatible or zero support.

The exact numerical exponents are chosen inside the PAR/GEB hierarchy.

B.2 Admission Table

Source	Edge predicate	Saving mechanism
F3 empty or incompatible cell	zero support	no mass
F3 boundary cell	boundary/short volume	volume loss
F3 square-divisor tail	large square divisor	square-divisor summation
F4 short quotient or divisor	Type I or short-volume saving	divisor/quotient loss
CKP high-frequency tail	high Fourier tail	integration/Fourier decay
CKP small-conductor layer	small conductor	conductor saving
CKP large g or content	large gcd/content	divisor-bound loss
BRS singular carrier slice residual	boundary or slice-floor Edge	X16BRS/X16C plus C1
LongAP/Local boundary error	boundary/short volume	C1 boundary estimate

B.3 Summability

C1A pairs each active Edge source with one of the C1P predicates above. C1 then gives either a power saving in N or a logarithmic saving large enough to dominate the $L^{O(1)}$ routing multiplicity. Hence the total Edge contribution is $o(N)$.

C CKP/X10 Smooth-Weight Matching

C.1 Target

The CKP branch proves

$$R_{\mathrm{CKP}}(N) = M_{\mathrm{CKP}}(N) + o(N).$$

The nonzero-frequency contribution is estimated by X10. The zero-frequency contribution is local and is passed to H4.

C.2 DFI Theorem Used

The X10 input is the Duke–Friedlander–Iwaniec bilinear estimate for Kloosterman fractions. In the smooth weighted form used here, one estimates

$$\sum_{\substack{m\sim M,\ q\sim Q\\(m,q)=1}} \alpha_m \beta_q F(m,q) e\left(\frac{r\overline{m}}{q}\right),$$

where F is supported on a dyadic box and has controlled derivatives up to the required fixed order. In the CKP application the derivative-control parameter is a fixed power of L.

C.3 CKP Substitution

After G1a–G3a, the central CKP nonzero-frequency layer has the form

$$\mathcal{O}_{g,h} = \sum_{\substack{a\sim A_g,\ q\sim Q_g\\(a,q)=1}} \alpha_g(a)\gamma_g(q) W_{g,h}(a,q) e\left(\frac{hN_g\overline{a}}{q}\right).$$

The substitution into X10 is

$$m = a, \qquad M = A_g, \qquad r = |h|N_g, \qquad Q = Q_g.$$

The coprimality condition is exactly $(a,q) = 1$. For $h < 0$, the same estimate is applied to the conjugate phase, so the positive external integer parameter in X10 is $r = |h|N_g$. The case $h = 0$ is excluded from X10 and routed to the local term.

C.4 Smooth Weight

CKPD differentiates the actual CKP weight, not a separated surrogate. The weight contains the Fourier-fibre dependence

$$z(a,q,y) = \frac{N_g - ay}{q}.$$

On the central support, CKPD proves

$$\partial_a^i \partial_q^j \widetilde{W}_{g,h}(a,q) \ll L^C A_g^{-i} Q_g^{-j} \qquad (0 \le i,j \le 2).$$

Thus the X10 smooth-weight hypotheses hold with polylogarithmic loss.

C.5 Excluded Ranges

The following ranges are not sent to X10:

Excluded range	Routing
$h = 0$	B1LD and H4
high Fourier frequency	Edge
small conductor	Edge
large g or large content	X10ER, then Edge or CKP auxiliary exclusion
boundary or short-volume layer	Edge

After these exclusions, the remaining central nonzero-frequency contribution is $o(N)$ by X10 and GEB.

D TC1, BRS/TTH, and GoodAWACK Finite Grammar

D.1 TC1 Route

The TC1 branch tests structural B1-origin coarea families whose cells have not already been routed away. It never invokes Liouville/AP orthogonality on an arbitrary shifted short interval.

The route is packaged as Theorem TNG-A:

every unrouted TC1 test is either near-global or routed away.

The closure barrier preventing rogue short-interval refinements is TTH-SC. It proves that every short subtest of a released structural coarea test is either non-structural and reaggregated, or structural and routed through TTD/ROC/BRS/X16BRS/X16C before X9L-GT is invoked.

On the near-global alternative:

$$H \geq X(\log X)^{-B_\kappa}.$$

At this length, X9L-GT supplies the needed Davenport/AP orthogonality with polylogarithmic modulus and AP complexity.

The bridge from a non-small TC1 macro-template to a measured family of Liouville tests is TGT-MF. It proves the finite Fourier/coarea transfer

$$\text{GT-U2} \Longrightarrow \text{GT-Test},$$

with a probability measure ν_κ and fixed lower bound depending only on the macro-template. Thus the route tests an averaged family, not a selected pointwise short interval.

D.2 Singular Tests

If a testing measure is singular, it is not sent to X9L-GT. The routed alternative of TNG-A, proved by TTD/ROC/BRS with X16BRS/X16C, shows that an unrouted singular B1-origin coarea test must enter one of the already controlled destinations:

$$\text{Edge}, \quad \text{LongAP/Local}, \quad \text{CKP}, \quad \text{LocalDiag}, \quad 0.$$

The carrier-slice estimates used in this step are supplied by X16BRS and X16C. X16C uses the Shiu AP divisor-average input and the fixed divisor-function second moment recorded in the bibliography.

D.3 Finite GoodAWACK Grammar

The HighTC part is closed structurally. E10Y proves completeness of the finite grammar for actual terminal GoodAWACK skeletons. The start states are B1/B3 grouped cells, and the allowed transitions are fixing/projection, controlled CRT restriction, fixed divisor quotient, F4-tagged variable quotient, local/diagonal/gcd dependence, CKP-balanced structure, Edge routing, full-rank affine regrouping, post-terminal slicing after terminal vectors are fixed, E5 auxiliary inheritance, and terminal labelling.

E10X proves a grammar invariant on the E10Y-certified grammar: every rank-dropping affine operation generated by these transitions carries an allowed origin tag. E10S and E10S-MECH are non-logical source-maintenance and reproducibility records for the listed Branch B source list; they are not premises of the closure. E10M proves that every rank-dropping affine occurrence in an actual terminal GoodAWACK skeleton is tagged by an allowed origin. Therefore no untagged rank-dropping AFF source exists. E10K gives AFF-origin completeness, and E10X removes the FreeAffineHighTC residual.

D.4 Formal Interface Examples

The formal $4AP$-like pattern $Y_i = x+ir$ is retained as a diagnostic interface example. It shows why broad affine regrouping language is unsafe if read without origin data. It is not a terminal obstruction because any actual terminal occurrence must either be full-rank safe, tagged and rerouted, or excluded by the no-untagged-AFF theorem.

D.5 GoodAWACK Output

Combining TC1 testing, singular rerouting, and finite-grammar HighTC closure gives

$$R_{\mathrm{GoodAWACK}}(N) = o(N).$$

E Explicit Local Projection Algebra

This appendix records the local algebra used by H4. The term canonical local projection" is only shorthand for the concrete construction below.

E.1 The Finite Local Model

Let

$$w = w(N) \to \infty, \qquad w = o(\log N), \qquad Q = \prod_{p \le w} p.$$

The local model of the von Mangoldt weight modulo Q is

$$\Lambda_Q(a) = \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1}.$$

It has average value one on $\mathbb{Z}/Q\mathbb{Z}$. Define

$$\sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a).$$

The local Goldbach model at modulus Q is

$$\mathrm{Loc}_Q\, R_\Lambda(N) = N\sigma_Q(N).$$

Endpoint and smooth-partition discrepancies are already Edge errors in C1 and are $o(N)$.

E.2 Tagged-Cell Projection

The exact B1 decomposition and the F3 tagged partition give a finite tagged identity

$$R_\Lambda(N) = \sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N),$$

where $\mathcal{B}$ is the parent B1 block and τ is the complete routing tag of the terminal cell.

For a tagged cell $(\mathcal{B}, \tau)$, define

$$\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N)$$

by replacing the arithmetic coefficients in that same tagged cell by their residue-class local densities modulo Q, while preserving:

1. the parent B1 block $\mathcal{B}$;
2. the routing tag τ;
3. the dyadic and smooth weights;
4. the local congruence restrictions already present in the cell.

This is a tag-preserving operation. It is not a branch-specific density and it does not identify cells that merely have the same displayed algebraic shape.

E.3 H4 Admission Condition

A terminal local/main expression is admitted into H4 only if it satisfies

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N). \tag{E-adm}$$

The active sources are exactly:

1. LongAP/Local main terms from D1;
2. CKP zero-frequency terms from G8a together with B1LD;
3. LocalDiag cells produced by the routing table;
4. explicitly admitted local boundary terms whose discrepancy is already covered by C1.

D1 verifies (E-adm) for LongAP/Local cells after excluding nonlocal coefficients. G8a verifies it for CKP $h = 0$ cells by identifying the zero Fourier mode with the same tagged local replacement. LocalDiag cells are admitted only when the diagonal specialization remains a tagged local projection; otherwise the routing table sends the cell to Edge, CKP, GoodAWACK, impossible, or a continuing routed case.

E.4 Linearity Over the B1/F3 Partition

The operator Loc_Q is linear on the finite tagged decomposition:

$$\mathrm{Loc}_Q\left(\sum_{\mathcal{B},\tau} c_{\mathcal{B}} R_{\mathcal{B},\tau}\right) = \sum_{\mathcal{B},\tau} c_{\mathcal{B}} \,\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}.$$

Since B1 is an exact finite identity and F3 is an exact tagged partition, the tagged local projections reconstruct the local model of the original Goldbach convolution:

$$\sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau\in\mathcal{T}(\mathcal{B})} \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) = N\sigma_Q(N) + o(N). \tag{E-reconstruct}$$

Equivalently, applying the local replacement after the tagged proof-tree partition gives the same local main term as applying $\Lambda(n) \mapsto \Lambda_Q(n \bmod Q)$ directly to the two original von Mangoldt factors in $R_\Lambda(N)$.

E.5 No Double Counting

Local terms are indexed by $(\mathcal{B}, \tau)$. If two terms have different parent B1 blocks, they are distinct summands of the exact B1 identity. If they have the same parent block but different tags, F3.15 gives disjoint tagged cells.

Thus two local-looking formulas may be algebraically identical and still be different summands; they are recombined only through the tagged linear sum in (E-reconstruct). Conversely, a local-looking expression without a parent B1 block and routing tag is not an H4 input. Hence H4 neither double-counts nor loses admitted local/main terms.

E.6 Local Factors and the Singular Series

By the definition of Λ_Q,

$$\sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1} \frac{Q}{\varphi(Q)} \mathbf{1}_{(N-a,Q)=1}.$$

Since Q is squarefree, the CRT gives

$$\sigma_Q(N) = \prod_{p\le w} \sigma_p(N),$$

where

$$\sigma_p(N) = \frac{1}{p}\left(\frac{p}{p-1}\right)^2 \#\{a \bmod p : (a,p) = 1,\ (N-a,p) = 1\}.$$

For $p = 2$ and even N, $\sigma_2(N) = 2$. For odd p, the two forbidden residues are 0 and N. Hence

$$\sigma_p(N) = \begin{cases} \dfrac{p}{p-1}, & p \mid N, \\[2mm] 1 - \dfrac{1}{(p-1)^2}, & p \nmid N. \end{cases}$$

Therefore

$$\sigma_Q(N) = 2 \prod_{\substack{3 \le p \le w \\ p \nmid N}} \left(1 - \frac{1}{(p-1)^2}\right) \prod_{\substack{3 \le p \le w \\ p \mid N}} \frac{p}{p-1}.$$

Letting $w \to \infty$,

$$\sigma_Q(N) \to 2C_2 \prod_{\substack{p \mid N \\ p > 2}} \frac{p-1}{p-2} = \mathfrak{S}(N),$$

with

$$C_2 = \prod_{p>2} \left(1 - \frac{1}{(p-1)^2}\right).$$

E.7 H4 Local Algebra Theorem

The admitted terminal local/main terms satisfy

$$M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).$$

Indeed, by the admission condition (E-adm), every active local/main term is a tagged local projection up to $o(N)$. By tagged linearity and no double counting, the sum of these projections is $N\sigma_Q(N) + o(N)$. By the CRT factor computation, $\sigma_Q(N) \to \mathfrak{S}(N)$. Thus the local/main assembly contributes exactly the Goldbach singular series main term.

F Parameter Hierarchy and Global Error Budget

F.1 Parameter Witness

The proof fixes parameters in a nonempty hierarchy. One explicit witness is

$$J_0 = 20, \qquad \eta = \frac{1}{40}, \qquad \theta = \frac{1}{4000}.$$

For this witness $J_0 \ge J_*(\eta)$. After J_0 is fixed, all later logarithmic exponents are chosen in the order recorded by GEB.

F.2 Multiplicity Principle

The B1/B3/F3/F4 process creates at most L^{C_0} terminal cells. Hence:

1. $O(NL^{-C_0-A})$ per cell sums to $O(NL^{-A})$;
2. $O(N^{1-\rho}L^{C})$ per cell sums to $o(N)$;
3. a normalized $o(1)$ estimate remains $o(1)$ after polylogarithmic testing-family losses, once the relevant logarithmic exponent is chosen large enough.

F.3 Global Loss Table

Source	Estimate	Global conclusion
B1 decomposition	exact identity	no error
F3/F4 routing	exact tagged partition	no error
Edge	C1 strict saving	$o(N)$
LongAP/Local	LPI local projection plus Edge errors	local $+o(N)$
CKP excluded ranges	Edge/local/auxiliary routing	$o(N)$ or local
CKP central nonzero range	X10 after CKPD	$o(N)$
GoodAWACK TC1 regular	X9L-GT after TTH	$o(N)$
GoodAWACK singular	BRS/ROC rerouting	handled by existing branches
X16/BRS carrier slice	Shiu/AP carrier estimate	Edge or $o(N)$
HighTC/grammar	E10Y/E10X/E10M/E10K	no residual
LPI/H4 local projection	CRT local model	$\mathfrak{S}(N)N + o(N)$

F.4 Consequence

After summing over the complete terminal partition,

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N).$$

No hidden terminal-family summation remains outside GEB.

The prime-power removal is not a terminal branch summation. It is performed after I1 by G2, which gives $R_\Lambda(N) - R_{pp}(N) = o(N)$.

G Dependency Ledger and Traceability

This appendix is auxiliary. It records how the manuscript layer maps to the proof-source layer. File paths here are not part of the proof prose.

G.1 Main Logical Chain

$$X1 + B1 + B3 + F3 + F4 + F3T + E5 \Longrightarrow \text{terminal partition.}$$

$$C1P + C1A + C1, \quad D1 + H4, \quad G8a + X10, \quad E10L + X9L + X16, \quad GEB \Longrightarrow I1.$$

$$I1 + G2 + G1 + G0H \Longrightarrow \text{sufficiently-large binary Goldbach.}$$

G.2 Source Map

Manuscript item	Active proof-source item
X1	`../External/x_1_heath_brown_identity_verification_ltx.md`
X9L-GT	`../External/x_9l_gt_avg_polylog_verification_ltx.md`
X10	`../External/x_10_verification_ltx.md`
X16	`../External/x_16_divisor_sum_brs_verification_ltx.md`

B1	`../Lemmas/b_1_ltx.md`
B3	`../Lemmas/b_3_ltx.md`
F3P/F3	`../Lemmas/f3_intrinsic_terminal_predicate_catalogue_ltx.mo` `../Lemmas/f_3_ltx.md`
F3T	`../Lemmas/f3_complete_routing_exhaustion_ltx.md`
F4	`../Lemmas/f_4_ltx.md`
E5	`../Lemmas/e_5_ltx.md`
C1P	`../Lemmas/c1_strict_edge_predicate_catalogue_ltx.md`
C1A	`../Lemmas/c1_edge_admission_ledger_ltx.md`
C1	`../Lemmas/c_1_ltx.md`
D1	`../Lemmas/d_1_ltx.md`
H4	`../Lemmas/h_4_ltx.md`
CKPD	`../Lemmas/ckp_x10_smooth_weight_derivative_appendix_ltx.mo`
G8a	`../Lemmas/g_8_a_ltx.md`
TC1/TNG	`../Lemmas/tc1_measured_fourier_transfer_ltx.md`, `../Lemmas/tc1_structural_coarea_closure_ltx.md`, `../Lemmas/tc1_near_global_chain_ltx.md`
BRS/ROC/TTH	`../Lemmas/b1_range_skeleton_roc_slice_ltx.md`, `../Lemmas/tc1_singular_origin_roc_ltx.md`, `../Lemmas/tc1_theta_1_3_ltx.md`
X16C/X16BRS	`../Lemmas/x16_core_shiu_ap_proof_ltx.md`, `../Lemmas/x16_brs_carrier_slice_ltx.md`
E10Y	`../Lemmas/e10y_goodawack_routing_grammar_completeness_ltx`
E10X	`../Lemmas/e10_master_source_exhaustion_closure_ltx.md`
E10M/E10K/E10L	`../Lemmas/e10m_no_untagged_rank_dropping_aff_ltx.md`, `../Lemmas/e10k_aff_oc_affine_regrouping_origin_completenes` `../Lemmas/e10l_e10_clean_branch_b_ltx.md`; `../Lemmas/e10m_source_exhaustion_verification_ltx.md` is a non-logical source-maintenance record
GEB	`../Lemmas/global_error_budget_ltx.md`
I1	`../Lemmas/i_1_ltx.md`
G2	`../Lemmas/g_2_ltx.md`
G1	`../Lemmas/g_1_ltx.md`
G0H	`../Lemmas/g0_final_handoff_verification_ltx.md`
Proof ledger	`../Proof tree and ledger/g_proof_tree_n_ldg.md`

G.3 Synchronization Rule

If an active source file listed above changes, then the corresponding manuscript section or appendix must be checked before any LaTeX regeneration. Any mathematical correction to a manuscript passage must be made first in the active proof-source file and only then propagated into this manuscript layer.

References

[1] Tenenbaum, Gerald. *Introduction to Analytic and Probabilistic Number Theory*. American Mathematical Society. vol. 163. 2015. Ch. II.5, Theorem 5 is used for fixed divisor-function second moments..

[2] Davenport, Harold. *On some infinite series involving arithmetical functions (II)*. The Quarterly Journal of Mathematics. vol. os-8. no. 1. pp. 313–320. 1937. doi: 10.1093/qmath/os-8.1.313. Classical source for the near-global Liouville exponential-sum input; the active formulation is recorded in X9L-GT..

[3] Heath-Brown, D. R.. *Prime numbers in short intervals and a generalized Vaughan identity*. Canadian Journal of Mathematics. vol. 34. no. 6. pp. 1365–1377. 1982. doi: 10.4153/CJM-1982-095-9. Source for the fixed-depth Heath–Brown identity used in X1/B1..

[4] Shiu, P.. *A Brun–Titchmarsh theorem for multiplicative functions*. Journal für die reine und angewandte Mathematik. vol. 313. pp. 161–170. 1980. doi: 10.1515/crll.1980.313.161. Source for the Shiu AP divisor-average input used in X16C..

[5] Duke, William and Friedlander, John B. and Iwaniec, Henryk. *Bilinear forms with Kloosterman fractions*. Inventiones Mathematicae. vol. 128. no. 1. pp. 23–43. 1997. doi: 10.1007/s002220050135. Active CKP/X10 input: Theorem 2 with the smooth-weight formulation..

[6] Deshouillers, Jean-Marc and Iwaniec, Henryk. *Kloosterman sums and Fourier coefficients of cusp forms*. Inventiones Mathematicae. vol. 70. no. 2. pp. 219–288. 1982. Background source for classical Kloosterman-sum technology; the active CKP/X10 theorem is Duke–Friedlander–Iwaniec..