hexagon:2611.00672v1

Pointwise Correlations of Quadratic Form Values and Splitting Semigroups

Roy van Rijn

Abstract

Under the stated analytic inputs, we establish the fixed-shift pair asymptotics predicted by Smilansky for the ordinary values of x^2+dy^2, $d=1,2,3,4,7$. In particular, the number of $n \le X$ for which both n and $n+1$ are sums of two squares is asymptotic to $X/(2\log X)$. For primitive positive definite binary forms of fixed discriminants $\Delta_i = f_i^2 D$ sharing one negative fundamental discriminant D, we count values admitting coprime coordinates. Replacing either form within its respective genus preserves the leading binary count, with a relative logarithmic saving through an independent classical route. Under the recent analytic inputs we give an explicit mixed leading term with relative $o(1)$, including different fixed conductors in that one field, and no class-number divisor. The central mechanism is a shared comparison between sparse indicators and arithmetic models retaining every surviving congruence character. Following Walsh's proposed indicator extension, we supply the models, local factors and a uniform reflected-target argument, using the recent zero-free and weighted-dilation theorems of OpenAI. Applications to fixed subgroup restrictions and prime-coset self-sums give target-uniform asymptotics with logarithmic savings and exact eventual congruence criteria. They include prime support 1 (mod 12) on targets 2 (mod 12) and prime support 3 (mod 4) on even targets. All counts concern numerical values without coordinate or ideal multiplicities.

1 Introduction

How many positive integers $n \le X$ have the property that both n and $n+1$ are sums of two squares? Write $b_d(n)$ for the indicator of the positive values of x^2+dy^2, with $x,y \in \mathbb{Z}$, and put $b_d(n)=0$ for $n \le 0$. Coordinates are unrestricted and may vanish; each numerical value is counted once. Under the analytic inputs stated below, we obtain

$$\#\{1 \le n \le X : b_1(n)=b_1(n+1)=1\} \sim \frac{X}{2\log X}. \tag{1}$$

This is the shift-one Connors–Keating prediction. It belongs to Smilansky's five-form family of pair-correlation predictions.

Corollary 1.1. *Under the inputs of Section 3, for $d \in \{1,2,3,4,7\}$ and fixed $h \neq 0$,*

$$\sum_{1\le n\le X} b_d(n)b_d(n+h) = \mathscr{T}_{d,h}\frac{X}{\log X} + o_{d,h}(X/\log X),$$

with precisely the coefficients predicted in [10, Section 6.3, equations (6.2)–(6.4)]. In particular,

$$\sum_{n\le X} b_2(n)b_2(n+1) \sim \frac{1}{\sqrt{2}}\frac{X}{\log X}. \tag{2}$$

Section 10 evaluates the coefficients and exceptional-prime conditions.

1.1 A second headline: genus governs primitive value counts

Fix one negative fundamental discriminant D and two fixed discriminants $\Delta_i = f_i^2 D$. Let Q_i, Q_i' be primitive positive definite integral binary forms of discriminant Δ_i, in the same genus for each $i = 1, 2$, in the one field $\mathbb{Q}(\sqrt{D})$. If P_Q is the set of positive values of Q at coprime coordinates and $R_{Q_1,Q_2}(N) = \#\{(a,b) \in P_{Q_1} \times P_{Q_2} : a+b = N\}$, then the classical route gives

$$R_{Q_1,Q_2}(N) = R_{Q_1',Q_2'}(N)\left(1 + O((\log N)^{-1/5})\right) \tag{3}$$

on sufficiently large common locally admissible targets. The slots may have different genera and conductors. Under the recent analytic inputs the same count has an explicit leading term of order $NJ_S(N)/\log N$, with relative error $o(1)$. No class-number divisor occurs: a value can belong to every class in its compatible genus. The common-field theorem and its principal specialization are in Section 12.

The counting conventions and strengths are different:

Result family	What is counted	Conclusion
CK/Smilansky	Ordinary values; unrestricted coordinates	Fixed-shift asymptotic with the stated little-o error
Subgroup/coset sums	Integers with prescribed prime support	Target-uniform asymptotic with a logarithmic saving
Primitive-form sums	Values admitting coprime coordinates in two fixed orders of one field	Explicit leading term with relative $o(1)$
Genus replacement	Comparison of two primitive-form value counts	Relative logarithmic saving through the classical route

The first three asymptotic families use the imports in Section 3; genus replacement does not. A *primitive form* $[a,b,c]$ has $\gcd(a,b,c) = 1$. A *primitive representation* $Q(x,y)$ has $\gcd(x,y) = 1$. *Coprime summands* in $a+b = N$ means $\gcd(a,b) = 1$, a separate condition imposed only when stated. For example, $4 = 2^2 + 0^2$ is an ordinary sum of two squares, but any representation of four has both coordinates even and is not primitive. Consequently the primitive-form theorem does not extend the ordinary CK statement to arbitrary forms. Its class-error saving also does not supply a logarithmic remainder for the full conductor-prefix leading term.

1.2 The common mechanism, in the CK example

The central result is the indicator-to-model comparison of Theorem 2.1. Multiplication by suitable primes preserves an indicator and its arithmetic model. One-variable means center their difference, and a positive sparse sieve controls exceptional prime-factor records. A shared graph reduction then leaves two frequency estimates: one for fixed shifts and one for reflected targets. Their assembly replaces the indicators in the binary count by the models. The model correlation can then be evaluated from local arithmetic.

Here is the model for b_1. At outer scale X, put $L = \log X$, $y = \exp(L^{2/5})$, $\mathcal{P}_y = \{p : p \le y\}$, and $V_y = \prod_{p \le y}(1 - 1/p)$. Write every n uniquely as $n = rv$, where $r = n_{\mathcal{P}_y}$ contains its full small-prime powers and v has no prime factor at most y. The rough ordinary-value series is

$$F_y(s) = \prod_{\substack{p>y \\ p\equiv 1 \ (4)}} (1-p^{-s})^{-1} \prod_{\substack{p>y \\ p\equiv 3 \ (4)}} (1-p^{-2s})^{-1} = (s-1)^{-1/2} H_y(s).$$

For a sufficiently deep fixed K, its smooth density expansion is

$$M(v) = \sum_{j=0}^{K} \frac{H_y^{(j)}(1)}{j!} \frac{(\log v)^{-1/2-j}}{\Gamma(1/2-j)}, \qquad \mathcal{D}(rv) = b_1(r)\big(1 + \chi_{-4}(v)\big)\frac{M(v)}{V_y}.$$

These formulas apply on the polynomial rough scales used in the proof; the smooth bounded extension is specified in Section 9 and the quadratic application. The leading term of M is $H_y(1)/\sqrt{\pi \log v}$. Dividing by V_y, the density of rough integers, gives the conditional density on the rough variable. Every allowed rough value has $\chi_{-4}(v) = 1$. The correction $1 + \chi_{-4}(v)$ retains this bias; average density alone would miss a leading contribution. For a split prime $p \le y$, multiplication by p changes only r, preserves $b_1(r)$, and leaves both v and $\mathcal{D}$ unchanged. Replacing $b_1(n)b_1(n+h)$ by $\mathcal{D}(n)\mathcal{D}(n+h)$ with an error $o(X/\log X)$ is the analytic task. The remaining model product gives the coupled local factors and the coefficient $1/2$ in (1).

1.3 Semigroups and the reflected-target obstacle

The same mechanism reaches a different geometry. Let $C_{12} = \{n \ge 1 : p \mid n \Rightarrow p \equiv 1 \pmod{12}\}$, including one. Every sufficiently large $N \equiv 2 \pmod{12}$ is a sum of two elements of C_{12}, and the ordered value-pair count $R_{12}(N) = \#\{(a,b) \in C_{12}^2 : a + b = N\}$ satisfies

$$R_{12}(N) = A_{12}J_{12}(N)\frac{N}{(\log N)^{3/2}}\left(1 + O((\log N)^{-\delta})\right), \qquad \delta > 0. \tag{4}$$

The positive factors are explicit in Theorem 11.2. For $C_4^- = \{n \ge 1 : p \mid n \Rightarrow p \equiv 3 \pmod 4\}$, every sufficiently large even target likewise has an explicit asymptotic of order $N/\log N$, with a positive logarithmic saving. The subgroup and coset theorems allow fixed arithmetic data and retain all surviving characters, including the complex branches forced by the prime class 3 (mod 10). Numerical saving exponents are collected in Section 8.

One-variable density does not control the particular line $a + b = N$, and averages can conceal exceptional targets. The nonconstant linear parts of n and $N - n$ are dependent, so the independent-linear-form correlation theorems of Loughran–Matthiesen do not directly apply [5, Theorems 1.9 and 3.10]. The reflected frequency argument must retain cancellation and normalization when the target enters the rational denominators.

Iwaniec–Pomykała's quartic-norm work is a direct antecedent [3]. For $\mathbb{Q}(\zeta_{12})$, its splitting conditions select the primes $1 \pmod{12}$. Their Theorem 2 gives conditional pointwise order bounds for coprime sums; hypothesis (P4) is uniform in growing shifts. Their Theorem 1 treats fixed differences unconditionally. They count numerical values once. Our unrestricted leading term gives a coprime asymptotic by Möbius inversion, with the extra target factor in Corollary 11.3.

1.4 Attribution and dependencies

Walsh's multiplicative-invariance framework treats positive integer divisor orders and fixed nonzero shifts [9]. His Section 1.2.4, pp. 5–6, proposes sums-of-two-squares indicator correlations. Here we supply their complete character models, arithmetic hypothesis checks and local factors, and the uniform reflected-target argument.

The recent inputs are the weighted-dilation root-moment theorem [7, Section 3] and the all-conductor zero-free half-plane $\Re s > 11/12$ [8, Part I, Theorem 3.1]; Part II is not used. Walsh notes a possible classical recovery of his fixed-shift theorem [9, p. 2], but his printed Lemma 5.1 and Section 6.4 retain that zero-free input. Our polynomial-logarithmic fixed-shift moduli still occur at Mellin heights up to $x(\log x)^C$. Section 6.4 proves the classical intermediate-height estimate and states the remaining high-frequency obligation. Both recent inputs remain in our fixed-shift result; such a replacement would also leave the reflected growing-conductor problem.

Appendix A proves classical quadratic order bounds by Kloosterman completion and positive restoration. With the class-error sieve these give genus comparison independently of the recent inputs. All data are fixed; thresholds are not numerical. The companion paper [15] gives exact classifications and finite certificates for four primitive quadratic norm sumsets. The deposited window paper concerns the placement of summands.

1.5 Notation

Symbol	Role
X, $L = \log X$, h	Outer scale, its logarithm, and the fixed shift
$\mathcal{P}_y$, $n_{\mathcal{P}_y}$, V_y	Small primes, complete small-prime part, and rough-integer density
$f_i, \mathcal{D}_i, \xi_i$	Indicators, arithmetic models, and centered errors
D, D_i, U_i	Auxiliary prime products and their dyadic lengths
H_i, $\mathscr{A}$, $\mathcal{E}$	Residue subgroups, a prime coset, and forbidden prime sets
B_0, S	Finite exceptional-prime sets in the ordinary-indicator and primitive-form applications
$\Delta_i = f_i^2 D$	Order discriminants sharing one fixed fundamental discriminant D; here f_i denotes a conductor
P_Q, Λ_Q, G	Primitive numerical values, their local shadow, and the applicable class group

The modulus and the arithmetic data stay fixed. Definitions in each application specify its groups and exceptional primes; h is reserved for the additive shift throughout.

Contents

1 Introduction . . . 1
2 Replacing indicators by arithmetic models . . . 4
3 Analytic inputs and the comparison framework . . . 6
4 The shared graph comparison . . . 9
5 Frequency tools for both branches . . . 17
6 Fixed shifts and wide major arcs . . . 20
7 Reflected targets and exact rational grids . . . 24
8 Reconstruction and quantitative comparison . . . 27
9 Fractional rough models and one-variable means . . . 30
10 Ordinary quadratic indicators and fixed-shift pairs . . . 32
11 Splitting semigroups and prime cosets . . . 36
12 Primitive forms and invariance within a genus . . . 49
13 Consequences and boundaries . . . 58
A An independent classical bound for quadratic cores . . . 58

2 Replacing indicators by arithmetic models

We isolate the prime invariance, one-variable means and sparse sieve bounds needed to replace an indicator by its arithmetic model.

Fix an even integer $L_0 \geq 2$ and a reduced class $c \in (\mathbb{Z}/L_0\mathbb{Z})^\times$. At the outer scale X, write

$$L = \log X, \qquad y = \exp(L^{2/5}), \qquad \mathcal{P}_y = \{p : p \leq y\}, \qquad E = \exp(\sqrt{L}).$$

For a set of primes $\mathcal{P}$, let $n_{\mathcal{P}}$ be the part of n supported on $\mathcal{P}$, including its full valuations. A rough integer has no prime factor in $\mathcal{P}_y$. There are two families of lines in the theorem:

$$\mathcal{L}_X^{\mathrm{sum}} = \{(u,v) \in \mathbb{Z}_{>0}^2 : Au + Bv = X\}, \qquad t(u,v) = Au, \qquad J_X = [0, X],$$
$$\mathcal{L}_X^{\mathrm{shift}} = \{(u, u+h) : u \in \mathbb{Z},\ X \leq u \leq 2X\}, \qquad t(u,v) = u, \qquad J_X = [X, 2X].$$

In the first family X is an integer target and A, B are fixed coprime positive integers supported on primes dividing L_0. In the second, $h \neq 0$ is a fixed integer and X tends to infinity; both coordinates are then positive. Set $k = X$ in the sum family and $k = h$ in the shift family. Dependence on the fixed data, including h in the second family, is permitted.

Let $0 \leq f_i \leq 1$, $i = 1, 2$, be fixed functions on the positive integers. For every sufficiently large fixed integer K, let $\mathcal{D}_{i,X,K}$ be an arithmetic model, and write $\xi_{i,X,K} = f_i - \mathcal{D}_{i,X,K}$. We

often suppress X, K. The entire model, its cutoff y, and its small-prime records stay tied to this outer scale during all subsequent dilations.

Choose a positive comparison scale $\mathcal{S}_X$ satisfying

$$XL^{-\kappa_0} \le \mathcal{S}_X \le XL^{\kappa_0} \tag{5}$$

for a fixed $\kappa_0 \ge 0$. Arithmetic factors may be included in this scale. The hypotheses are as follows.

(C1) Bounded models. For every fixed K, sufficiently large X gives $0 \le \mathcal{D}_{i,X,K}(n) \le 1$ on $1 \le n \le X^3$. The bound one is independent of the model depth; its threshold may depend on K.

(C2) A common prime class. For every prime $p \le y$ with $p \equiv c \pmod{L_0}$,

$$f_i(pn) = f_i(n), \qquad \mathcal{D}_{i,X,K}(pn) = \mathcal{D}_{i,X,K}(n)$$

whenever the arguments are in the range under consideration. This includes repeated multiplication by p. A residue filter is part of f_i and must satisfy the same equality.

(C3) Centered one-variable rough means. For every fixed $B_0, \mathcal{D}_0 > 0$, there is a depth K_0 such that every fixed $K \ge K_0$ has the following property. Uniformly for $\mathcal{P}_y$-supported integers $r \le E$, $X^{1/2} \le Y \le X^2$, bounded periodic functions $a(v)$ of period at most L^{B_0}, and smooth functions W supported in a fixed compact subinterval of $(0, \infty)$,

$$\left| \sum_{v \text{ rough}} \xi_{i,X,K}(rv) a(v) W(rv/Y) \right| \ll_{B_0,\mathcal{D}_0,K} \frac{Y}{r} L^{-\mathcal{D}_0}. \tag{6}$$

Here $|a| \le 1$, and every fixed logarithmic derivative of W used in the estimate is at most a fixed power of L^{B_0}. The hypothesis permits Mellin factors with polynomial-logarithmic frequency, by including them in W. Its constants may depend on the chosen derivative orders. Every surviving character branch, including a negative or complex pole parameter, must be retained when verifying this estimate.

(C4) A positive sparse record envelope. Take disjoint prime banks in the two ranges

$$[e^{L^{1/8}}, e^{L^{1/4}}], \qquad [e^{L^{1/3}}, e^{L^{2/5}}],$$

each bank consisting of the primes in an interval which are $c \pmod{L_0}$, after deleting primes dividing k. Assume each bank has reciprocal mass between $1/2$ and 2, and that the number in either range is at most $(\log L)/(16\varphi(L_0))$. Write $V_g = \sum_{p \in \mathcal{P}_g} p^{-1}$ for its reciprocal mass. Let $\mathcal{B}$ be their union. Define the independent joint record law $\mu_{\mathcal{B}} = \prod_{p \in \mathcal{B}} \mu_p$ by

$$\mu_p(0,0) = 1 - 2/p, \qquad \mu_p(e,0) = \mu_p(0,e) = (1 - 1/p)p^{-e} \quad (e \ge 1), \tag{7}$$

and zero mass on two positive valuations. For every fixed $B_0 > 0$, nonnegative record function $F(r,s) \le L^{B_0}$, interval $J \subset J_X$, and $w_i \in \{f_i, \mathcal{D}_{i,X,K}\}$, require

$$\sum_{\substack{(u,v) \in \mathcal{L}_X \\ t(u,v) \in J}} F(u_{\mathcal{B}}, v_{\mathcal{B}}) w_1(u) w_2(v) \ll_{B_0,K} \frac{|J|}{X} \mathcal{S}_X \mathbb{E}_{\mu_{\mathcal{B}}} F + X^{1/2} L^{C_{B_0}} + XL^{C_{B_0}} \exp(-c_0 L^{1/10}). \tag{8}$$

The positive constant c_0 and exponent C_{B_0} are independent of the final fixed model depth. Constants and thresholds may depend on that depth. Uniformity includes the bank intervals, the exact valuations in the record, and the interval endpoints.

For a bank g, let $\omega_g(n)$ count its distinct prime divisors, and define

$$I_g(u,v) = \mathbf{1}_{\omega_g(u)=\omega_g(v)=1}, \qquad \Omega_T(n) = \prod_{g\in T} V_g^{-1}\mathbf{1}_{\omega_g(n)=1}.$$

Here T is a set of bank indices. We also write $\Omega_{\mathcal{P}}$ when $\mathcal{P}$ denotes the union of the active banks. The joint valuation law above is the ordinary CRT law on either line: every bank prime avoids k and is a unit at the fixed coefficients.

Condition (C3) is a statement about one variable on a progression; it allows arbitrary retained small-prime records, but no arbitrary weight in the rough variable. Condition (C4) is only an upper bound with nonnegative weights. In the applications it follows from CRT root counts and an upper sieve, with a power-smaller endpoint error. Neither condition assumes centered binary cancellation or a major-arc energy estimate. In particular, choosing a model with the right average density alone does not verify (C3).

Theorem 2.1 (Arithmetic model comparison)**.** *Assume the zero-free and weighted-dilation inputs of Section 3, and conditions* (C1)–(C4)*. Put* $\delta = 1/(8192\varphi(L_0))$*. For every sufficiently deep fixed* K*, uniformly for intervals* $J \subset J_X$*,*

$$\sum_{\substack{(u,v)\in\mathcal{L}_X\\ t(u,v)\in J}} \big(f_1(u)f_2(v) - \mathcal{D}_{1,X,K}(u)\mathcal{D}_{2,X,K}(v)\big) \ll_K \mathcal{S}_X L^{-\delta}. \tag{9}$$

For the sum family the estimate is uniform in the integer target and its prime valuations. For the shift family the integer h *is fixed, and the constant may depend on it. The two families have different major-arc geometries; no uniformity in the intervening range of growing sublinear shifts is asserted.*

The proof consists of the shared graph reduction (Section 4), the fixed-shift and reflected estimates (Sections 6–7), and reconstruction (Section 8). Subgroup pairs satisfy (C2) in the identity class; a self-coset uses a class in that coset. Arbitrary prescriptions may have no common class. The final asymptotic saving must also pay for evaluating the model.

3 Analytic inputs and the comparison framework

We use Walsh's comparison architecture [9, Sections 2–7], the zero-free input [8, Theorem 3.1], and the weighted-dilation transfer [7, Section 3]. The operator normalizations and parameter order are stated explicitly.

3.1 The zero-free input

Proposition 3.1 (Imported zero-free theorem)**.** *Every Dirichlet L-function, including* ζ*, is nonzero in*

$$\Re s > 11/12.$$

The pole at $s = 1$ *of a principal character is allowed. Both conductor and height are unrestricted.*

This is [8, Theorem 3.1, printed p. 10], restricted to Dirichlet characters and the open half-plane. We also use the usual functional equation, Euler product and holomorphic logarithm on zero-free rectangles.

3.2 Prime groups and the two operators

For this subsection only, write

$$\Lambda = \log x, \qquad T = \log \Lambda.$$

Fix constants

$$0 < \alpha < \beta < \gamma < \delta < 0.47, \quad c_0, C_0, v_-, v_+ > 0, \quad 0 < \vartheta < 1, \quad \ell \geq 1.$$

There are disjoint collections $\mathcal{G}_s, \mathcal{G}_b$ of pairwise disjoint finite prime sets P_g, with

$$c_0 T \leq |\mathcal{G}_s|, |\mathcal{G}_b| \leq C_0 T.$$

They satisfy

$$\begin{aligned} e^{\Lambda^\alpha} \leq p \leq e^{\Lambda^\beta} \quad (g \in \mathcal{G}_s,\ p \in P_g), \\ e^{\Lambda^\gamma} \leq p \leq e^{\Lambda^\delta} \quad (g \in \mathcal{G}_b,\ p \in P_g), \\ v_- \leq V_g := \sum_{p \in P_g} \frac{1}{p} \leq v_+, \qquad \mu_g(p) = \frac{1}{V_g p}. \end{aligned} \tag{10}$$

The sets need not be intervals. This freedom is needed when taking identity residue classes and deleting primes dividing the target.

Fix a sufficiently large integer J, and put $M = J + \ell$. A pattern ν specifies subsets $C_g \subseteq \{1, \dots, J\}$, with

$$C_g = \varnothing \quad (g \in \mathcal{G}_s), \qquad |C_g| \leq m \quad (g \in \mathcal{G}_b), \qquad t_g = \ell + |C_g|.$$

The slots in $\{1, \dots, J\} \setminus C_g$ are shared between the two endpoint lists. The remaining t_g slots at each endpoint are unshared. Form D from the shared labels and, in the C_g-slots, fresh independent draws with laws μ_g. Thus D has J prime factors from each group, counted with multiplicity. Write $D = D_s D_b$ according to group size.

Each pattern has a complex coefficient K_ν. It may depend on the ordered big-group source labels, target labels, and auxiliary labels used in D_b, and on no small-group labels. Assume

$$\sum_\nu \sup |K_\nu| \leq \Lambda^{C_1}, \tag{11}$$

where m, C_1 are fixed independently of J. All operators below average source and target slot permutations independently in each applicable group. These averages are orthogonal projections.

The *ideal operator* acts on

$$\mathcal{H}_{\mathrm{id}} = L^2\left(\prod_{g \in \mathcal{G}_b} P_g^M, \bigotimes_{g \in \mathcal{G}_b} \mu_g^{\otimes M} \right).$$

Numerical repetitions are permitted. Its row operation retains shared labels, samples the unshared target labels and auxiliary labels independently, and multiplies by

$$K_\nu D_b^{i\zeta} e(\Theta D_b), \qquad e(t) = \exp(2\pi i t).$$

Sum over patterns and compose with the permutation projections on both sides; denote the result by $\mathcal{T}(\Theta)$.

A *physical state* consists of an integer n and an ordered M-tuple in every group, with distinct labels within each group, all dividing n. Give every such state at n mass $\prod_g V_g^{-M}$, and use counting measure on the integer position. This defines $\mathcal{H}_{\mathrm{ph}}$. Put

$$\omega_g(n) = \#\{p \in P_g : p \mid n\}, \quad \omega(n) = \sum_g \omega_g(n), \quad s = |\mathcal{G}_s| + |\mathcal{G}_b|.$$

For a fixed integer k with $0 < |k| \le \Lambda^{C_2}$, the *physical operator* A moves from n to $n' = n + kD$. It sums over the physical target lists with the prescribed shared labels, using row normalization $\prod_g V_g^{-t_g}$; auxiliary labels retain their independent probability laws. Every auxiliary label is forbidden from both endpoint lists, although auxiliary labels may equal each other. Multiply the row by

$$K_\nu D^{\,i\zeta} \vartheta^{(\omega(n)-Ms)/2} \vartheta^{(\omega(n')-Ms)/2}, \tag{12}$$

sum over patterns and apply all endpoint permutation projections. The adjoint reverses the displacement and conjugates the coefficient with the endpoint roles interchanged. The joint endpoint normalization in each group is $V_g^{-M-t_g}$ in either direction.

Proposition 3.2 (Imported absolute bounds)**.** *Replace every elementary transition coefficient of A by its absolute value before summing patterns and averaging permutations. The resulting operator and its adjoint have row sums at most $\Lambda^{C_{\mathrm{abs}}}$. The exponent may depend on m, C_1 and the fixed group data, but is independent of J. The same assertion holds with integer position replaced by independent uniform residues at all group primes.*

This is [7, Lemma 3.4]. The independence from J is part of the assertion. To see its source, once the shared labels are fixed, the new target slots in a group cost at most

$$\sup_{z \ge 0} V_g^{-t_g} (z + t_g)^{t_g} \vartheta^{z/2}.$$

Here $t_g \le \ell + m$. This is a fixed constant per group and its product over $O(T)$ groups is a fixed power of Λ. Reversing the joint endpoint measure gives the column bound.

3.3 Root moments, residue replacement, and endpoints

Proposition 3.3 (Imported local transference)**.** *For every fixed $E > 0$, there is $E_{\mathrm{id}} > 0$, depending only on E and the fixed group data, such that the following holds. Suppose*

$$\sup_{\Theta \in \mathbb{R}} \|\mathcal{T}(\Theta)\|_{2\to 2} \le \Lambda^{-E_{\mathrm{id}}}. \tag{13}$$

After m, C_1 have been fixed, take a sufficiently large fixed J. Put $R = \lceil \Lambda^{0.52} \rceil$. For every fixed $\eta > 0$ and $I = [X, 2X)$, with $x^\eta \le X \le x^{1/\eta}$,

$$\frac{1}{|I \cap \mathbb{Z}|} \sum_{n \in I \cap \mathbb{Z}} \langle u_n, (AA^*)^R u_n \rangle \le \Lambda^{-2RE}. \tag{14}$$

Here u_n is one on every physical mark list at position n and zero elsewhere. The threshold may depend on all the fixed parameters, including J, C_2, η. No bound on ζ is required beyond (13).

This is [7, Theorem 3.5]. It imposes return of the integer position. The final mark list need not equal the initial list, and u_n is not normalized. Substituting a trace with identical initial and final mark lists would change the statement.

The arithmetic replacement inside that theorem has the following precise form, which also explains its scale limitation. For a fixed compatible path with $r = 2R$ edges, let h_i be

the integer offsets and $\mathcal{A}_i$ its active prime set at visit i. Set $\vartheta_i = \vartheta$ at internal visits and $\vartheta_0 = \vartheta_r = \sqrt{\vartheta}$. The position-dependent factor is

$$W(n) = \prod_{i=0}^{r} \left(\prod_{p \in \mathcal{A}_i} 1_{p \mid n+h_i} \right) \vartheta_i^{\#\{p \notin \mathcal{A}_i:\ p \text{ is a group prime},\ p \mid n+h_i\}}. \tag{15}$$

Proposition 3.4 (Imported residue replacement)**.** *For every fixed $B > 0$, the average of (15) on $I \cap \mathbb{Z}$ equals its expectation with independent uniform residues at all group primes, multiplied by $1 + O(\exp(-BrT))$. An incompatible collection of active congruences has zero weight in both models.*

This is [7, Lemma 3.6]. For clarity, its relative form retains all deterministic damping belonging to active primes. The product Q_{act} of active primes satisfies

$$\log Q_{\mathrm{act}} = O_J(rT\Lambda^{\delta}) = o(\Lambda).$$

After those factors are extracted, a Bonferroni expansion of degree $O_B(rT)$ controls the remaining damping. Its independent expectation is at least $\exp(-O(rT))$; its total interval counting error is $\exp(-\eta\Lambda + o(\Lambda))$. The inequality $0.52 + \delta < 1$ makes that error negligible even after division by the active congruence density. Together with Proposition 3.2, this permits summation over paths with absolute coefficients. These observations describe the imported replacement and do not assert replacement on intervals shorter than its stated range.

Corollary 3.5 (Imported endpoint consequence)**.** *Under Proposition 3.3, let $f \in \mathcal{H}_{\mathrm{ph}}$ be supported on positions $I = [X, 2X)$, constant on the chosen mark list at each position, with value f_n. If*

$$\sup_{n \in I} |f_n| \le \exp(O(\sqrt{\Lambda})), \qquad \|f\|, \|g\| \le X^{1/2}\Lambda^{C_3},$$

then

$$|\langle f, Ag \rangle| \ll X\Lambda^{-E+2C_3}. \tag{16}$$

Only f must be constant on mark lists.

This is [7, Corollary 3.11]. Its proof partitions I into intervals on the displacement scale $\exp(O_J(T\Lambda^{\delta}))$, uses spectral Hölder for AA^*, and bounds the positive matrix $(\langle u_n, (AA^*)^R u_m \rangle)$ by its trace. It makes no equal-fiber-mass assumption.

3.4 The application interfaces

Three steps remain: construct the fractional-density models retaining every surviving character; prove the progression-filter ideal comparison with its big-label coefficient bound (11) fixed before J; and transfer to the growing reflected target with uniform endpoint and localization bounds. The root-moment theorem supplies none of these steps on its own. Sections 4–8 prove the comparison; the applications then check its arithmetic hypotheses.

4 The shared graph comparison

This section proves the graph reduction common to the reflected and fixed-shift families of Theorem 2.1. Its output is (33)–(34): every error correlation is replaced by comparisons containing two independent large prime products. Sections 6 and 7 estimate those comparisons in their different frequency windows. No arithmetic form of the indicators is used here beyond (C1)–(C3).

The labels lie in one fixed reduced class. The character frame must therefore identify restrictions differing by a known unit phase.

Write $L = \log x$. There are $\asymp \log L$ small banks and big banks, with primes respectively in

$$[\exp(L^{1/8}), \exp(L^{1/4})], \qquad [\exp(L^{1/3}), \exp(L^{2/5})].$$

Every label is $c \bmod L_0$; primes dividing the nonzero displacement k are removed. We use $k = T = x$ for sums and fixed $k = h$ for shifts. All banks are disjoint. Their laws are $\mu_g(p) = (V_g p)^{-1}$, where $V_g = \sum_{p \in \mathcal{P}_g} p^{-1}$ is bounded above and below by positive fixed constants. Banks in this section may be any active subfamilies of the banks constructed from the class in (C2), provided that each size has at least a fixed positive multiple of $\log L$ banks. All constants may depend on these fixed bounds and on L_0.

4.1 The character frame for progression banks

The ideal-kernel construction in [7, Section 4] uses primitive characters of conductor at most $Q = L^{q_*}$. On labels $c \bmod L_0$, this list is redundant up to phases. For example, the principal character and χ_{-4} agree on primes $1 \bmod 4$ and are opposite on primes $3 \bmod 4$. The quotient below removes this redundancy while retaining its phase.

Let $\mathcal{X}_Q$ be the primitive characters of conductor at most Q, and define

$$\eta \sim \eta' \quad \Longleftrightarrow \quad \operatorname{cond}(\eta\overline{\eta'}) \mid L_0. \tag{17}$$

The conductor here belongs to the primitive character inducing the quotient on a common modulus. Choose one representative from each class and denote the resulting set by $\mathcal{R}_Q$. Thus $|\mathcal{R}_Q| \le Q^2$.

To verify the equivalence and its meaning, put $M_0 = \operatorname{lcm}(L_0, \operatorname{cond}\eta, \operatorname{cond}\eta')$. Reduction from $(\mathbb{Z}/M_0\mathbb{Z})^\times$ onto $(\mathbb{Z}/L_0\mathbb{Z})^\times$ is surjective. Its kernel consists of the units $1 \bmod L_0$. A character is trivial on that kernel exactly when it factors through the latter group. This is equivalent to its primitive conductor dividing L_0. Consequently, if $\eta \sim \eta'$ and their quotient is induced by $\vartheta \bmod L_0$, then on a product of r labels,

$$\eta(D) = \vartheta(c)^r \eta'(D).$$

Both comparison products have the same slot type and the same phase. Inequivalent characters have zero average on the kernel, and also on any fixed coset of it: translation multiplies their sum by a unit phase.

Fix a nonempty set Λ of at most m slots per big bank, where m is fixed before the eventual graph mark length. Let D_Λ be the product of independent labels in these slots. Partition its logarithm into cells of width $\tau > 0$ $I_j = [j\tau, (j+1)\tau)$, and write $\nu_j = \mathbb{P}(\log D_\Lambda \in I_j)$. For inequivalent $\eta, \eta' \in \mathcal{R}_Q$, and any prescribed $A > 0$, Siegel–Walfisz gives

$$\left| \mathbb{E}\big[\mathbf{1}_{\log D_\Lambda \in I_j} \eta(D_\Lambda)\overline{\eta'(D_\Lambda)}\big] \right| \ll_A L^{-A}. \tag{18}$$

Here and below the accuracy in Siegel–Walfisz is chosen after all fixed logarithmic exponents.

Indeed, freeze all but one label. Its allowable values form an interval inside its original bank interval, still restricted to $c \bmod L_0$. On a common modulus at most L_0Q^2, the sum of $\eta\overline{\eta'}$ over the appropriate reduced residue classes is zero. Siegel–Walfisz, summed over these classes, therefore gives

$$\sum_{\substack{p \le z \\ p \equiv c \ (L_0)}} \eta(p)\overline{\eta'(p)} \ll_K L_0 Q^2 z (\log z)^{-K}.$$

Partial summation bounds the reciprocal-prime sum on any subinterval of $[\exp(L^{1/3}), \exp(L^{2/5})]$ by $O_K(L_0Q^2(L^{1/3})^{1-K})$. This is an absolute bound; no lower bound for the subinterval's length

or harmonic mass is needed. Divide by V_g, then average the frozen labels. Deleting target divisors changes the answer by at most

$$\sum_{\substack{p\mid k\\ p\ge \exp(L^{1/3})}} \frac{1}{p} \le \frac{\log(2|k|)}{L^{1/3}} \exp(-L^{1/3}), \tag{19}$$

which is smaller than every fixed inverse power of L. This proves (18) uniformly for the actual banks.

On cells with $\nu_j \ge \xi$, define

$$K_\Lambda(D,X) = \sum_{\nu_j\ge\xi} \frac{\mathbf{1}_{\log D\in I_j}\mathbf{1}_{\log X\in I_j}}{\nu_j} \sum_{\eta\in\mathcal{R}_Q} \overline{\eta(D)}\eta(X). \tag{20}$$

Both product variables have the same law. The elementary bounds are

$$|K_\Lambda| \le Q^2\xi^{-1}, \qquad \sup_X \mathbb{E}_D|K_\Lambda(D,X)| \le Q^2, \qquad \#\{j:\nu_j>0\} \ll_m 1+(\log L)L^{2/5}/\tau. \tag{21}$$

The last inequality uses only the range of the logarithm.

On a retained cell, the kernel is the positive frame operator

$$f \longmapsto \sum_{\eta\in\mathcal{R}_Q} \eta\,\langle \eta, f\rangle$$

for the conditional product probability measure. Its Gram matrix has diagonal entries one, since all labels exceed L_0Q^2, and off-diagonal entries $O_A(\xi^{-1}L^{-A})$ by (18). Thus its norm is at most $1+O_A(Q^2\xi^{-1}L^{-A}) \le 2$. Distinct cells are orthogonal blocks, so

$$\|K_\Lambda\|_{2\to2} \le 2. \tag{22}$$

This also holds for tuple tests: conditional expectation from tuples to their product is a contraction. If ψ is a character of modulus $r \le Q$, its primitive inducing character is equivalent to one member of $\mathcal{R}_Q$. For r_Λ selected slots and $\psi = \vartheta\eta$ on these units, that summand gives $\vartheta(c)^{r_\Lambda}\eta(X) = \psi(X)$. The other summands give

$$\mathbb{E}_D K_\Lambda(D,X)\psi(D) = \psi(X) + O_A(Q^2\xi^{-1}L^{-A}) \tag{23}$$

at every retained X. This is where quotienting the frame is essential: the unquotiented frame can reproduce the same constant several times.

Lemma 4.1 (Progression ideal comparison)**.** *For every prescribed ideal accuracy E_{id} and fixed phase range $|\zeta| \le L^{C_4}$, the progression banks admit the signed comparison family of [7, Theorem 4.1]. Its ideal norm is at most $L^{-E_{\mathrm{id}}}$, uniformly in every real additive frequency. Its coefficients ignore small-bank labels and shared labels, free at most m slots per big bank, and have total supremum at most L^{C_1}. The parameters m, C_1 and all kernel powers are fixed before the eventual mark length.*

Proof. We give the comparison estimates used in the construction, including the only step where restricting the prime labels changes the proof.

For a product formed from $k_g \le m$ draws in each of at most $C_0\log L$ disjoint banks, unique factorization gives its exact law

$$n\mathbb{P}(D_\Lambda = n) = \prod_g \frac{k_g!}{V_g^{k_g}\prod_{p\in\mathcal{P}_g} a_p!}, \qquad \mathbb{P}(D_\Lambda = n) \le L^{K_0}/n, \tag{24}$$

where $K_0 = C_0(\log(m!) + m \log \max(1, v_-^{-1}))$ and $V_g \geq v_- > 0$. The factor C_0 multiplies both terms of this exponent. For a tuple test, Cauchy–Schwarz on each product fibre gives an integer coefficient sequence of squared norm at most L^{K_0}/U times the tuple squared norm on a product dyad $[U, 2U)$.

The elementary integer bilinear estimate [7, Lemma 4.2], or [9, Lemma 6.2], gives

$$\left| \sum_{u,v} a_u b_v e^{2\pi i \theta uv} \right|^2 \ll \|a\|_2^2 \|b\|_2^2 (1 + U/r)(V + r \log(2r)) \tag{25}$$

when the two supports lie in containing intervals of lengths $O(U), O(V)$, $(a, r) = 1$, and $|\theta - a/r| \leq r^{-2}$. For completeness, Cauchy–Schwarz and expansion leave $V + \sum_{1 \leq j \ll U} \min(V, \|\theta j\|^{-1})$. Divide the j's into blocks of diameter less than $r/2$. Within a block the rational parts are $1/r$ separated and their perturbations differ by at most $1/(2r)$; hence the sum is $O(V + r \log(2r))$ per block. This proves (25) and permits arbitrary subsets of the containing intervals.

Apply this estimate to two nonempty selected product sets A, B. On product dyads U, V, use Dirichlet approximation with denominator ceiling $\lfloor UV/S \rfloor$. When the resulting denominator exceeds Q, (24) gives the normalized bound

$$\ll L^{K_0} \big(o(L^{-A}) + Q^{-1} + L^{7/5}/S \big)^{1/2}. \tag{26}$$

There are at most $L^{7/5}$ product dyads per variable. The same bound applies to the kernel terms by moving each kernel onto its endpoint test using (22). The factor $(uv)^{i\zeta}$ separates and preserves test norms.

For denominators $r \leq Q$, all products are units modulo r. The multiplicative-character expansion of the rational phase has total absolute coefficient at most $\sqrt{\varphi(r)}$. Replace each primitive character by its representative in $\mathcal{R}_Q$ and apply (23). Grouping characters after inserting their unit phases cannot increase coefficient mass, so the reproduction error is $O_A(Q^{5/2} \xi^{-1} L^{-A})$. Write the remaining phase as $M(v) = \exp(i\zeta v) e^{2\pi i \beta \exp v}$. On matched cells its derivative is $O(L^{C_4} + S)$, and the two logarithmic products move by at most 2τ. By (21), this costs $O(Q^4 \tau (S + L^{C_4}))$. The union of discarded cells has probability at most $(N_A + N_B)\xi$, where N_A, N_B count the possible cells. Cauchy–Schwarz bounds its contribution by the square root of this probability times the two test norms. These estimates prove the two-product comparison, for arbitrary L^2 tuple tests, to any prescribed power L^{-A_0}.

There is no circular choice of precisions. One possible hierarchy is

$$\begin{aligned} B_0 &> A_0 + 2(7/5) + 10, \\ q_*, s_* &> 2(B_0 + K_0 + 10) + 7/5, \\ \tau_* &> 4q_* + \max(s_*, C_4) + B_0 + 10, \\ x_* &> 2B_0 + 7/5 + \tau_* + 10, \end{aligned}$$

with $Q = L^{q_*}, S = L^{s_*}, \tau = L^{-\tau_*}, \xi = L^{-x_*}$. Choose the accuracy in (18) afterwards. All these choices are independent of the eventual mark length.

Here is the signed cancellation which completes the ideal construction. Split the big banks into two blocks of comparable size and designate m probe slots per group. For nonempty probe subsets A, B in the two blocks, free their slots and give the comparison coefficient

$$-(-1)^{|A|+|B|} K_A(D_A, X_A) K_B(D_B, X_B). \tag{27}$$

Here D_A, D_B are auxiliary-label products and X_A, X_B are the unshared endpoint products. The coefficient has no small or shared labels. There are at most $2^{mC_0 \log L}$ comparisons; by (21) their total supremum is at most

$$1 + 2^{mC_0 \log L} Q^4 \xi^{-2} \leq L^{C_1}. \tag{28}$$

Expanding both kernels also separates each coefficient into a sum of products of four bounded factors: one on each auxiliary product and one on each unshared endpoint list. The number of terms and total scalar coefficient mass are fixed powers of L.

We now make the exceptional projection and the cancellation explicit. Write the two big-bank blocks as $\mathcal{B}_1, \mathcal{B}_2$, of sizes s_1, s_2, and their designated probe coordinates as $\mathcal{I}_1, \mathcal{I}_2$. For a coordinate i, let E_i average its independent label. These commuting operators are orthogonal projections. Let Q_{tail} average the last ℓ input or target coordinates in every group, and let Π be within-group symmetrization on all $M = J + \ell$ coordinates. The ideal multipliers do not use those last coordinates. For symmetric endpoint tests F, H, put $F' = Q_{\mathrm{tail}}F,\ H' = Q_{\mathrm{tail}}H$. Both averaging operations are contractions.

Let D_{sh} be the product of the first J coordinates, and write

$$\Phi_{\Theta,\zeta} = D_{\mathrm{sh}}^{i\zeta} e^{2\pi i \Theta D_{\mathrm{sh}}}, \qquad R(F', H') = \mathbb{E}\big[F'\overline{H'}\Phi_{\Theta,\zeta}\big].$$

Thus R is a bilinear pairing of norm at most one. Resolve the probes in the first block of F' and the second block of H':

$$F' = \sum_{S \subset \mathcal{I}_1} F_S, \qquad F_S = \prod_{i \in S}(I - E_i) \prod_{i \in \mathcal{I}_1 \setminus S} E_i F',$$
$$H' = \sum_{T \subset \mathcal{I}_2} H_T, \qquad H_T = \prod_{i \in T}(I - E_i) \prod_{i \in \mathcal{I}_2 \setminus T} E_i H'.$$

These are orthogonal decompositions. Define $P_1 = \prod_{i \in \mathcal{I}_1}(I - E_i)$, $P_2 = \prod_{i \in \mathcal{I}_2}(I - E_i)$, so that $F_{\mathrm{full}} = P_1 F'$, $H_{\mathrm{full}} = P_2 H'$.

In one group the original symmetric vector decomposes orthogonally over subsets of its M coordinates. Symmetry makes the squared norms equal among subsets of cardinality j. Requiring all m probes to occur and all last ℓ slots to be absent therefore retains exactly the fraction

$$\frac{\binom{M-m-\ell}{j-m}}{\binom{M}{j}} = \frac{(j)_m (M-j)_\ell}{(M)_{m+\ell}} \le 2(j/M)^m (1 - j/M)^\ell \le 2\ell^\ell m^{-\ell} =: \theta_m$$

when $M \ge (m+\ell)(m+\ell-1)$; impossible cardinalities give zero. For a vector which is not a tensor product across groups, first resolve by the tuple of cardinalities in the groups of a block. Independent within-group permutations make all subset tuples of those cardinalities have equal squared norms. The retained fraction is consequently the product of the displayed fractions. Summing over the cardinality tuples proves the operator estimates

$$\|P_1 Q_{\mathrm{tail}} \Pi\| \le \theta_m^{s_1/2}, \qquad \|P_2 Q_{\mathrm{tail}} \Pi\| \le \theta_m^{s_2/2}. \tag{29}$$

Averaging tails outside the block is a further contraction. We use $\ell = 1$, and each $s_i \ge c_0 \log L/3$.

For a component pair F_S, H_T, put $U = \mathcal{I}_1 \setminus S$, $V = \mathcal{I}_2 \setminus T$. A comparison freeing $A \subset \mathcal{I}_1$, $B \subset \mathcal{I}_2$ vanishes unless $A \subset U, B \subset V$. Indeed, if a freed input coordinate belongs to S, it is mean zero, absent from the target, and absent from the comparison multiplier: integrating it gives zero. The corresponding statement holds for a freed target coordinate belonging to T. This uses the orientation in (27): the A endpoint labels are from the target and the B endpoint labels are from the input.

If $A \subset U, B \subset V$, condition on all shared coordinates outside $A \cup B$, with product C. The input test is independent of the freed A labels and the target test is independent of the freed B labels. The two-product kernel comparison established above therefore replaces the unsigned comparison by $R(F_S, H_T)$, with error

$$L^{-A_0} \|F_S\|_2 \|H_T\|_2.$$

Its frequency is ΘC; uniformity for every real frequency is essential. Cauchy–Schwarz over the conditioned coordinates gives exactly the displayed unconditional norms.

For nonempty U, V, the total unsigned subset sign is

$$\sum_{\substack{\varnothing\neq A\subset U\\ \varnothing\neq B\subset V}} (-1)^{|A|+|B|} = \left(\sum_{\varnothing\neq A\subset U} (-1)^{|A|}\right)\left(\sum_{\varnothing\neq B\subset V} (-1)^{|B|}\right) = 1.$$

The initial minus sign in (27) cancels the raw pairing. If either missing set is empty, no comparison survives. All exceptional raw component pairs together have the exact identity

$$R_{\mathrm{exc}}(F,H) = R(P_1F', H') + R((I-P_1)F', P_2H').$$

In operator notation, if M_Φ multiplies by $\Phi_{\Theta,\zeta}$, their operator on the symmetric space is

$$T_{\mathrm{exc}} = \Pi Q_{\mathrm{tail}}\big[M_\Phi P_1 + P_2 M_\Phi (I-P_1)\big] Q_{\mathrm{tail}}\Pi.$$

Consequently

$$\|T_{\mathrm{exc}}\| \le \theta_m^{s_1/2} + \theta_m^{s_2/2}.$$

This identity sums the exceptional components before taking an absolute value. It is the reason their bound has no subset-count or comparison-coefficient factor.

There are at most 3^P admissible component/free-subset triples, where $P = m(s_1+s_2) \le mC_0 \log L$: each probe is indicated, missing but not freed, or freed. Thus the full signed ideal operator satisfies

$$\|T_{\mathrm{signed}} - T_{\mathrm{exc}}\| \le 3^P L^{-A_0}.$$

Choose m so that $\theta_m < 1$ and $-c_0 \log \theta_m/6 > E_{\mathrm{id}}+3$, and then choose $A_0 > E_{\mathrm{id}} + mC_0 \log 3 + 10$. The kernel precisions above and the exponent C_1 in (28) are chosen after these choices. Only afterwards require $J \ge m$ and $J + \ell \ge (m+\ell)(m+\ell-1)$; increasing this fixed J to meet the physical root-moment theorem changes no kernel parameter or coefficient exponent. The resulting ideal norm is at most $L^{-E_{\mathrm{id}}}$. This proves the asserted ideal norm and coefficient conditions. □

4.2 A large displacement in the physical graph

Lemma 4.2 (Large-target root moment and pairing)**.** *The root moment and endpoint pairing from [7, Theorem 3.5 and Corollary 3.11] hold for any nonzero integer k coprime to the bank primes, with the same logarithmic accuracies. In particular they hold for $k = T \asymp x$ after target-prime deletion. The first endpoint is required to be independent of its marks, exactly as in the imported pairing.*

Proof. Fix the labels, patterns, and permutations of a closed walk of length $2R$, where $R = \lceil L^{0.52} \rceil$. Its position offsets are ks_i. Return to the initial integer is equivalent to $s_{2R} = 0$, independent of the nonzero k. The root dependence consists of divisibility and damping conditions at the offsets ks_i.

The independent-residue replacement [7, Lemma 3.6] is uniform in the numerical values of these offsets. Indeed its active modulus has logarithm $O_M(R(\log L)L^{2/5}) = o(L)$. The Bonferroni truncation has the same modulus and term count, and each interval CRT error is $O(1)$, independently of the specified residues. Once the deterministic damping from active primes is factored out, the error is relative and can be summed with absolute path coefficients using the imported absolute row bound.

In the resulting independent residue space, replace the root residue z_p by $k^{-1}z_p \bmod p$ for every bank prime p. Every fixed-path divisibility and damping expectation is then exactly the expectation for $k = 1$. The remaining signed coefficients are unchanged. Thus the $k = 1$ root moment gives the identical bound for this k. Negative root intervals follow either from the same CRT count or from $n \mapsto -n,\ k \mapsto -k$.

It remains to avoid a large factor in the endpoint pairing. Let $D_{\max}$ bound the edge products, so that $\log D_{\max} = O_M((\log L)L^{2/5})$. Partition the first-endpoint interval I into blocks of length $\lceil |k| D_{\max} \rceil$, as in the source pairing, and refine each block by residue modulo $|k|$. The operator preserves these residues. Each refined block has $O(D_{\max}+1)$ integer positions. For $B = \mathcal{A}_k \mathcal{A}_k^*$, the matrix $(\langle u_n, B^R u_m \rangle)$ on one such block is positive semidefinite. Its largest eigenvalue is at most its trace. Here u_n is the unnormalized vector equal to one on all marks at n. Mark independence of the first endpoint f_n gives

$$\langle f_j, B^R f_j \rangle \le \Big(\sum_{n \in I_j} |f_n|^2 \Big) \sum_{n \in I_j} \langle u_n, B^R u_n \rangle.$$

Spectral Hölder and Cauchy–Schwarz give the source block estimate; its only dimension factor is now

$$\big(O(D_{\max}) \sup |f_n|^2\big)^{1/(2R)} \le \exp\left(O_M\left(\frac{(\log L)L^{2/5} + \sqrt{L}}{L^{0.52}} \right) \right) = O(1). \tag{30}$$

Second-endpoint neighborhoods have bounded overlap within a residue class and are disjoint between residue classes. Hölder over blocks uses only the sum of root traces on the whole original interval, bounded by the root theorem. No CRT estimate on an individual residue class modulo k is used. This proves the pairing uniformly in k. □

For bounded invariant endpoint functions, the list moments needed here do not create a mark-dependent logarithmic loss. For a fixed ordered distinct list with product P, divisibility gives $O(X/P)$ integers in an interval of length $O(X)$ contained in $[-CX, CX]$. For fixed mark length, $P = x^{o(1)}$. After summing the list weights,

$$\sum_{\text{lists}} \frac{1}{P} \prod_g V_g^{-M} \le 1. \tag{31}$$

Thus the endpoint r-moments are $O_r(X)$. Condition (C1) bounds every fixed-depth model without a depth-dependent logarithmic exponent. Thus this observation is valid before the final depth is chosen.

4.3 Shared labels and the two comparison geometries

The following consequence of the pairing applies to both geometries. For a vector $\mathbf{t} = (t_g)$ and a bounded function b on its ordered distinct lists, put

$$W_{\mathbf{t}}^b(n) = \prod_g \big(\mathbf{1}_{\omega_g(n) = t_g} V_g^{-t_g} \big) \sum_{\substack{\mathbf{p} \text{ a } \mathbf{t}\text{-list} \\ p_{g,j} | n}} b(\mathbf{p}). \tag{32}$$

Here $\omega_g(n)$ counts distinct bank-prime divisors of $|n|$. For fixed $t_g \le m+1$, these weights are bounded by a fixed power of L. In particular $W_{\mathbf{1}}^1 = \Omega_{\mathcal{P}}$.

Proposition 4.3 (Graph-to-arithmetic reduction)**.** *Assume the imported graph inputs of Section 3. Use the progression banks of this section, with primes dividing k deleted, and take $k = x$ for reflected sums or fixed $k = h \neq 0$ for shifts. Let α, β be functions of the integer endpoint, of modulus at most one, invariant under repeated multiplication by every bank prime whenever both arguments are in the endpoint ranges. For sums use the negative first half-line and positive second half-line, extending the first function by absolute value; for shifts use positive endpoints. The pullback of the first endpoint to physical states must be constant on its mark lists.*

Take $\ell = 1$, $M = J + 1$, and a smooth Ψ supported in a fixed compact subset of $(0, \infty)$. For every fixed accuracy $A > 0$, Lemmas 4.1 and 4.2 give fixed m, C and a sufficiently large fixed J for which

$$C_0 + \sum_{\nu} c_\nu C_\nu = O_A(L^{-A}), \qquad \#\{\nu\} + \sum_{\nu} |c_\nu| \le L^C, \tag{33}$$

where

$$C_0 = \sum_{w \neq 0} \frac{\Psi(|w|/x)}{|w|} \alpha(w)\beta(w+k) W_{\mathbf{1}}^1(w) W_{\mathbf{1}}^1(w+k)$$

and each comparison has the form

$$\begin{aligned} C_\nu = \mathbb{E}_{D_1, D_2} a_\nu(D_1) b_\nu(D_2) \sum_{w \neq 0} \frac{\Psi(|w|/(xD_1D_2))}{|w|} \alpha(w)\beta(w + kD_1D_2) \\ \times W_{\mathbf{t}}^{u_\nu}(w) W_{\mathbf{t}}^{v_\nu}(w + kD_1D_2). \end{aligned} \tag{34}$$

The following restrictions and normalizations are part of the conclusion.

1. *The physical pairing is restricted to $\omega_g(n) = \omega_g(n') = M$ in every bank. Its state mass is $\prod_g V_g^{-M}$, its transition row normalization is $\prod_g V_g^{-t_g}$, and hence its joint endpoint normalization in bank g is $V_g^{-M-t_g}$. After shared labels are removed, the exact counts are $\omega_g(w) = \omega_g(w + kD_1D_2) = t_g$, as imposed by (32), where $t_g = 1 + \#\{\text{freed slots in } g\} \le m + 1$. In every small bank $t_g = 1$.*

2. *Each comparison frees nonempty probe sets in two disjoint blocks of big banks. The products D_1, D_2 use these respective sets, with at most m draws per bank. All auxiliary draws are independent with laws μ_g, including between the two products; repetitions among auxiliary draws are allowed. No auxiliary-label exclusion from endpoint lists remains in the displayed comparison.*

3. *Before separation, a signed pattern coefficient may depend on the freed big-bank auxiliary labels and the unshared big-bank endpoint labels specified in (27). It uses no small-bank labels, shared labels, or distinguished last endpoint slots. After separation, a_ν depends only on D_1, b_ν only on D_2, and u_ν, v_ν only on the ordered unshared big-bank labels at their respective endpoints. Each of these four factors has modulus at most one. In particular, both list coefficients ignore every small-bank label.*

4. *The parameters m, C, the comparison family and its scalar coefficient bounds are fixed before the final mark length J and the final model depth. Increasing J supplies shared slots without changing them. For endpoints satisfying (C1), no logarithmic loss exponent depends on the model depth; the threshold in x may depend on that depth and on J.*

Proof. The signed coefficients are those of Lemma 4.1; their allowed label dependence and four-factor expansion were proved in (27)–(28). Each auxiliary factor depends only on its numerical product: (20) uses only that product's log cell and character. Thus the expectation over ordered draws pushes forward to the product law (24), including its repetition multiplicities. The first endpoint is mark independent, as required by the imported pairing, while all pattern-dependent mark factors remain in the operator. A pattern has $M - 1$ dilation slots per bank, some shared and some freed, and one distinguished unshared endpoint slot. Write its edge product as $D = D_R D_C$, with shared product D_R and auxiliary product D_C. For a source of either sign in $|n| \in [Z, 2Z)$, an originally nonzero edge weight satisfies $\Psi(|n|/(xD)) \neq 0$, hence $D \ll Z/x$. Since $|k| \ll x$, this implies $|n + kD| \ll Z$. Insert a second-endpoint cutoff on $[-C_\Psi Z, C_\Psi Z]$ before Mellin inversion. It equals one on every contributing edge. After inversion retain this same cutoff. An edge originally having zero Ψ still has zero full Mellin integral, so the insertion does not change the identity. No cutoff involving shared-label products has been introduced.

Mellin inversion of Ψ, with the harmonic weight $1/|n|$, now applies the pairing of Lemma 4.2 on each of $O(L)$ source dyads. Its tails are bounded by the absolute row/column estimate, its main part by the small ideal norm, and its endpoint norms by (31). Compressing the physical operator to exact-count vertices $\omega_g(n) = \omega_g(n') = M$ is legitimate: the corresponding projection depends only on the integer, preserves mark independence, and sets both damping factors equal to one.

For completeness, put $n = D_R w$, $n' = D_R w'$, so $w' = w + kD_C$, and put $t_g = 1 + \#\{\text{freed slots in } g\}$. Invariance and harmonic scaling give

$$\alpha(D_R w) = \alpha(w), \quad \beta(D_R w') = \beta(w'), \quad \frac{\Psi(|n|/(xD))}{|n|} = \frac{1}{D_R}\frac{\Psi(|w|/(xD_C))}{|w|}.$$

The graph normalization at a shared prime combines with $1/D_R$ to give its probability law μ_g. More explicitly, with $(D_R)_g$ its product in bank g,

$$\frac{V_g^{-M-t_g}}{(D_R)_g} = V_g^{-2t_g} \prod_{\text{shared slots in } g} \mu_g(p). \tag{35}$$

For fixed reduced endpoints and fixed auxiliary labels, the shared-label probability is

$$\prod_g \mathbf{1}_{\omega_g(w)=\omega_g(w')=t_g} + O(L^{C_M} e^{-L^{1/8}}). \tag{36}$$

Indeed there are only $O_M(\log L)$ draws, each atom is $O(e^{-L^{1/8}})$, and $ww'D_C$ has $O(L)$ prime divisors. If the reduced counts are t_g, distinct draws avoiding those divisors reconstruct the exact graph edge. If a reduced count is larger, a valid reconstruction must choose a shared prime already dividing that endpoint, an event of the displayed small probability. This proves (36), also for negative w.

Summing its absolute error costs only fixed logarithmic powers, by bounded endpoints and the harmonic weight. The restrictions excluding auxiliary labels from endpoint lists can then be removed: for an auxiliary prime $p \mid D_C$, $p \mid w$ if and only if $p \mid w + kD_C$; these positions have density $O(1/p)$ in the containing interval. Every auxiliary prime is at least $\exp(L^{1/3})$. An unshared prime common to the two endpoints divides kD_C; bank-prime deletion implies it divides D_C, so the same estimate covers it. Thus these errors remain smaller than every logarithmic power. Finally expand the separated coefficients from (27). Formula (35) yields exactly (32), while $D_C = D_1 D_2$. The raw pattern has $D_C = 1, \mathbf{t} = \mathbf{1}$. This proves (33)–(34) with their stated independence properties. The kernel parameters and coefficient exponents were fixed before J in Lemma 4.1. Endpoint boundedness and (31) pay the list norms without a depth-dependent logarithmic power. Thus the final model depth can be chosen later to meet the one-variable accuracy, changing only the threshold for boundedness in (C1). □

5 Frequency tools for both branches

The shared comparison has left two independent products of large prime labels. We first record its common Fourier normalization, then the prime estimates used by the two frequency branches.

5.1 The common Fourier reduction

Dyadically localize the products in (34) to $D_i \asymp U_i$. Their ranges give

$$U_i \geq \exp(L^{1/3}), \qquad U_i \leq \exp(O_m((\log L)L^{2/5})), \qquad Y = xU_1U_2.$$

After inserting endpoint dyadic cutoffs before Mellin inversion, the ordinary comparison has the form

$$\frac{1}{Y}\int_0^1 \widehat{F}(\theta)\widehat{G}(-\theta)m(k\theta)\,d\theta, \qquad m(\theta)=\mathbb{E}b_1(D_1)b_2(D_2)e^{2\pi i\theta D_1D_2}. \tag{37}$$

The endpoint reflection for sums changes only signs and fixed character factors. External Mellin transforms have fixed logarithmic costs. The centered endpoint has the following common form. Let ξ be either centered function in (C3), at a fixed depth which will be chosen after the required precision. For a comparison term, define

$$F(n)=\psi(n/Y)n^{i\sigma}\xi(n)W_{\mathbf{t}}^{b}(n), \quad Y=x^{1+o(1)}, \quad |\sigma|\le L^C, \tag{38}$$

where the support of ψ is a fixed compact positive interval. All its logarithmic derivatives are bounded by fixed powers of L. The coefficient b ignores a selected small-bank label, whose count is exactly one. This holds for every comparison above. Reflection changes only fixed character factors and additive signs, so we state the estimate at a positive endpoint.

The other endpoint satisfies $\sum_n |G(n)|^2 \ll YL^C$, by bounded invariant endpoints and the pointwise list bound. No multiplicativity of the list coefficients is required.

By (24) and (25), a rational approximation $|\theta-a/r|\le r^{-2}$ gives

$$|m(\theta)|^2 \ll L^{C_{\mathrm{prod}}}\left(r^{-1}+U_1^{-1}+\frac{\log(2r)}{U_2}+\frac{r\log(2r)}{U_1U_2}\right). \tag{39}$$

The two branches apply this bound at $k=h$ and $k=x$, respectively; their resulting physical major windows differ.

5.2 Sharp twisted prime sums

The sharp estimate uses the zero-free input at conductors polynomial in x. The reflected branch needs that range; the fixed-shift branch uses the smooth version below at polynomial-logarithmic conductors and polynomial height.

Lemma 5.1 (Prime factor estimate)**.** *Assume the fixed all-conductor zero-free half-plane stated in Section 3. Fix $\sigma_1\in(11/12,1)$ and $a,B,B_1>0$. For $q\le x^B$, every character $\chi \bmod q$, and $\exp(L^a)\le P\le x^{B_1}$, uniformly on subintervals $I\subset[P,2P]$,*

$$\sum_{p\in I}\chi(p)p^{it}=\delta_\chi\int_I\frac{u^{it}}{\log u}\,du+O((1+|t|)P^{\sigma_1}L^C). \tag{40}$$

Here $\delta_\chi=1$ for a principal character and zero otherwise, and the estimate holds for every real t. The exponent C and implied constant may depend on a,B,B_1,σ_1 and the fixed bank modulus L_0. For progression banks the main vanishes unless χ is induced from a character modulo L_0. Removing modulus divisors or prime divisors of a nonzero target $|k|\le x^B$ preserves the error bound.

Proof. A fixed distance to the right of the zero-free boundary, the local logarithmic-derivative formula and functional equation give

$$\left|\frac{L'}{L}(s,\chi)+\frac{\delta_\chi}{s-1}\right|\ll\log(q(|\Im s|+3)). \tag{41}$$

Imprimitive Euler factors contribute $O(\log q)$, since the real part is fixed and positive. Apply truncated Perron to $-L'/L(s,\chi)$, at $c=1+1/\log Z$, height Z^2, and a half-integer endpoint $Z\in[P/2,4P]$. The truncation error is $O(\log^2 Z)$: near Z, sum the bounds $O(Z\log Z/(Z^2|n-Z|))$; away from Z, use absolute convergence at c. Shift to σ_1. The principal pole gives $\delta_\chi Z$, the new vertical integral is $O(Z^{\sigma_1}\log^2(qZ))$, and horizontal edges cost $O(Z^c\log(qZ)/(Z^2\log Z))$.

Prime powers cost $O(Z^{1/2}\log^2 Z)$. Replacing a real endpoint by the adjacent half-integer changes its prime-weighted count by $O(\log Z)$. Therefore

$$\sum_{p\le Z}\chi(p)\log p=\delta_\chi Z+O(Z^{\sigma_1}L^C).$$

Here $Z\asymp P\le x^{B_1}$ and $q\le x^B$ give $\log(qZ)\ll_{B,B_1}L$, justifying the displayed logarithmic cost. Partial summation against $u^{it}/\log u$ proves (40) with its linear $1+|t|$ growth.

Expand the progression condition into the finitely many characters modulo L_0. A constituent is principal exactly when the original character is induced from L_0. There are at most $O(\log(q(2+|k|)))=O(L)$ deleted prime divisors; their absolute unweighted contribution is absorbed by the error. □

Every prime scale used below lies in one of the two bank ranges, so $P\le\exp(L^{2/5})\le x$ for large x. Thus $B_1=1$ suffices for both comparison geometries; in the reflected application the conductors are at most $xL^{O(1)}$, and in the fixed-shift application they are polynomial logarithmic.

In particular, for a nonexceptional character the prime factor saves P^{σ_1-1} against its trivial length. The all-height form above is convenient: smooth product-cutoff Mellin transforms have integrable $(1+|t|)$-weighted absolute value bounded by a fixed power of L, so there is no boundary-smoothing loss from a sharp bank.

5.3 Smooth prime sums at polynomial height

The high-frequency argument uses a smooth version of Lemma 5.1. Its height dependence is essential: the linear $1+|t|$ loss in the sharp-interval estimate is not adequate when $|t|$ can be of order x. Fix $\sigma_1\in(11/12,1)$. For w smooth and supported on $[1/2,2]$, with each required derivative bounded by a fixed power of L, the same zero-free input gives

$$\sum_p\chi(p)w(p/P)p^{it}=\delta_\chi\int_0^\infty\frac{w(u/P)u^{it}}{\log u}\,du+O(P^{\sigma_1}L^{C_w}) \tag{42}$$

uniformly for $q\le L^B$, characters χ mod q, $|t|\le x^C$, and $P\ge\exp(cL^{1/8})$. Here C_w depends on the fixed derivative bounds and on B,C. This is the smooth prime estimate used in [9, Lemma 5.1 and Section 6.4].

To obtain it directly from the present input, insert the von Mangoldt weight and the smooth factor $w(u/P)/\log u$, apply Mellin inversion, and move the line from $\Re s>1$ to $\Re s=\sigma_1$. Formula (41) bounds the regularized logarithmic derivative on this line by $O(\log(q(|\Im s|+3)))$. The Mellin transform is translated in height by t; its decay depends on the derivatives of w, rather than on $|t|$. Its weighted integral is consequently $O(L^{C_w})$ even for $|t|\le x^C$. The principal pole gives the integral in (42); prime powers cost $O(P^{1/2}L^{C_w})$, absorbed by the displayed error. Truncation followed by dominated convergence justifies the contour move. This uses the same fixed zero-free half-plane, without a second prime-distribution theorem.

Impose simultaneously the bank class $c\pmod{L_0}$ and a prime class modulo q. Incompatible classes contribute zero; otherwise character orthogonality at $\operatorname{lcm}(q,L_0)$ applies. For large x, every prime in the small bank exceeds $qL_0|h|$, so nonreduced prime classes and primes dividing h do not occur. Integration by parts in the principal integral, for every fixed j, shows that the normalized prime sum $P^{-1}\sum_p\cdots$ is bounded by

$$L^{C_{w,j}}(1+|t|)^{-j}+L^{C_w}\exp(-cL^{1/8}). \tag{43}$$

All character and congruence costs here are polynomial logarithmic.

6 Fixed shifts and wide major arcs

For fixed $k = h \neq 0$, the shared-label comparison (33)–(34) remains valid. The frequency estimate has a different scale. After localizing $D_i \asymp U_i$, put

$$U_* = U_1U_2, \qquad Y = xU_*, \qquad W = U_*L^{-B_0}.$$

The multiplier is now $m(h\theta)$. The integer bilinear bound (39) confines its significant part to $O_h(L^{2B_0})$ arcs about rationals a/q, with

$$q \le |h|L^{B_0}, \qquad |\theta - a/q| \le C_h/W.$$

In particular, their physical width is L^{B_0}/U_*, rather than L^{B_0}/Y. The target-grid proposition does not estimate these arcs: on them $|\beta|Y$ can be of order xL^{B_0}. We extract the already established Dirichlet-frequency argument of [9, Sections 6.2–6.4], with the arithmetic hypotheses (C1)–(C3) in place of its particular divisor model.

Proposition 6.1 (Fixed-shift comparison energy)**.** *Let F have the form (38), with a selected small bank whose label does not enter b, and with its exact prime count equal to one. Suppose*

$$U_i \ge \exp(L^{1/3}), \qquad U_i \le \exp(O((\log L)L^{2/5})), \qquad Y = xU_1U_2.$$

For every fixed $B_0, D > 0$, conditions (C1)–(C3) *and the zero-free input imply, at sufficiently deep fixed model depth,*

$$\int_{|\beta|\le C_h/W} |\widehat{F}(a/q+\beta)|^2\, d\beta \ll_{D,h} YL^{-D} \tag{44}$$

uniformly for $q \le |h|L^{B_0}$, where $\widehat{F}(\theta) = \sum_n F(n)e^{2\pi i\theta n}$. Consequently every fixed-shift comparison in (34) has an arbitrary prescribed logarithmic saving. No binary cancellation estimate is assumed.

6.1 Conversion to Dirichlet frequencies

We use the following existing two mean-square inequalities, in their normalizations from [9, Lemma 6.3 and (6.14)]. If a_n is supported on $n \asymp Y$, then

$$\int_{|\beta|\le C/W} \left|\sum_n a_n e^{2\pi i\beta n}\right|^2 d\beta \ll_C \frac{1}{Y}\int_{\mathbb{R}} (1+W|t|/Y)^{-2}\left|\sum_n a_n n^{it}\right|^2 dt. \tag{45}$$

Here is the averaging step in that inequality, following [13, Lemmas 2.1–2.2]. Put $\eta = W/(4C)$, $A(v) = \sum_{n\le v} a_n$, $P(t) = \sum_n a_n n^{it}$, and $J_\eta = \int_0^\infty |A(v+\eta) - A(v)|^2 dv$. The interval kernel has Fourier transform of modulus at least $c\eta$ on $|\beta| \le C/W$, so Plancherel bounds the left side by $O(\eta^{-2}J_\eta)$. For $2\eta \le \nu \le 3\eta$, split the increment at $v+\nu$, square, and average in ν. Changing v to $v+\eta$ in the second increment gives, if the coefficients are supported on $[aY, bY]$,

$$J_\eta \ll \eta^{-1}\int_\eta^{3\eta}\int_{aY/2}^{2bY} |A(v+\nu) - A(v)|^2 dv\, d\nu.$$

Substitute $\nu = \Delta v$. Since $v \asymp Y$, this is at most

$$\frac{CY^2}{\eta}\int_{c_1\eta/Y}^{c_2\eta/Y}\int_0^\infty |A(v(1+\Delta)) - A(v)|^2 \frac{dv}{v}\, d\Delta.$$

Logarithmic Plancherel evaluates the inner integral exactly as

$$\frac{1}{2\pi}\int_{\mathbb{R}}|P(t)|^2\left|\frac{(1+\Delta)^{it}-1}{t}\right|^2dt.$$

The kernel is $O(\min\{\eta^2/Y^2,|t|^{-2}\})$, and the Δ-interval has length $O(\eta/Y)$. Thus $J_\eta \ll Y\int\min\{\eta^2/Y^2,|t|^{-2}\}|P(t)|^2dt$. Multiplying by η^{-2} proves (45); no identification of fixed additive windows with fixed multiplicative windows is used. The other inequality, the ordinary Dirichlet-polynomial mean-value theorem, is

$$\int_I\left|\sum_{n\le CR}a_nn^{it}\right|^2dt\ll_C(|I|+R)\sum_n|a_n|^2, \tag{46}$$

uniformly in the location of the interval I. Neither inequality is a new distributional input.

Apply (45) to $a_n=F(n)e^{2\pi ian/q}$ and set

$$\mathcal{T}(t)=Y^{-1}\sum_nF(n)e^{2\pi ian/q}n^{it},\qquad T_0=(Y/W)L^{B_1}=xL^{B_0+B_1}.$$

The pointwise list bound and (C1) give $\sum_n|F(n)|^2\ll YL^C$. Thus on a dyadic interval $T\le|t|\le2T$, (46) gives

$$\int_{T\le|t|\le2T}|\mathcal{T}(t)|^2dt\ll L^C(1+T/Y).$$

Summing this with the weight in (45), for $|t|>T_0$, gives

$$\int_{|t|>T_0}(1+W|t|/Y)^{-2}|\mathcal{T}(t)|^2dt\ll L^C\big(L^{-2B_1}+W^{-1}L^{-B_1}\big). \tag{47}$$

The exponent B_1 is chosen after the required saving, so this tail is arbitrarily small. It remains to prove

$$\int_{|t|\le T_0}|\mathcal{T}(t)|^2dt\ll_D L^{-D}. \tag{48}$$

6.2 Low frequencies and recordwise centering

Choose a fixed exponent B_2, to be specified after the smooth prime estimates, and put $T_c=L^{B_2}$. In the low range write $n=rv$, where $r=n_{\mathcal{P}_y}$ is the complete smooth record and v is rough. Every list weight is a function of r alone and is bounded by L^C. The terms with $r>E$ have population $O(YL^C\exp(-cL^{1/10}))$, by Rankin's inequality with parameter $c/\log y$, as in (76); their contribution to $\mathcal{T}(t)$ is negligible uniformly in t. For each remaining record, $e^{2\pi iarv/q}$ is a bounded periodic test of period at most q. The phase $(rv)^{i(t+\sigma)}\psi(rv/Y)$ has fixed logarithmic derivatives bounded by fixed powers of L, throughout $|t|\le T_c$. Condition (C3) therefore gives any prescribed saving for that record. Summing uses only

$$\sum_{r\ \mathcal{P}_y\text{-smooth}}\frac{1}{r}=\prod_{p\le y}(1-1/p)^{-1}\ll L^C.$$

Choose the precision in (C3) after B_2 and all list costs. This gives

$$\int_{|t|\le T_c}|\mathcal{T}(t)|^2dt\ll_D L^{-D}. \tag{49}$$

The arithmetic progression test is on v, so this argument does not require r or the endpoint to be a unit at L_0.

6.3 High frequencies and a long cofactor

For this part remove selected-bank prime squares, whose population is $O(Y/P_{\min})$. Applied to their bounded coefficients, (46) bounds their normalized integrated energy on $[-T_0, T_0]$ by $L^C(1+T_0/Y)/P_{\min}$, hence by less than every logarithmic power. Do not impose the cutoff $n_{\mathcal{P}_y} \le E$ here: it would couple the selected prime to its cofactor.

The remaining terms have a unique selected-bank prime and can be written $n = pm$. Exact invariance (C2) gives $\xi(pm) = \xi(m)$. The distinguished list coefficient ignores this prime, so after removing its factor V_g^{-1} the remaining coefficient $\gamma(m)$ is independent of p and bounded by L^C. It contains the condition that m has no selected-bank prime, together with the other list weights.

Partition the selected bank into smooth dyadic pieces $p \asymp P$, including its two sharp outer endpoints. Replace each sharp endpoint by a smooth cutoff across a strip of length εP, where $\varepsilon = L^{-A_{\mathrm{sm}}}$. The normalized sum of prime labels in such a strip is at most $O(\varepsilon + P^{-1})$, simply by counting integers. The Dirichlet mean-value bound for the cofactor below and Minkowski show that this replacement changes the normalized integrated energy by at most

$$L^C(\varepsilon + P_{\min}^{-1})^2.$$

There are only polynomially many dyadic pieces in L. Choose A_{sm} large after the desired saving and these costs. The resulting smooth prime cutoffs have derivative bounds which are fixed powers of L.

Put $R = Y/P$, and insert a smooth cofactor cutoff $g(m/R)$ equal to one wherever the product cutoff can contribute. Separate the additive phase by the residue classes of p, m modulo q, and Mellin-invert the product cutoff. Apart from the errors just bounded, $\mathcal{T}(t)$ is a sum of at most polynomially many terms of the form

$$\frac{1}{2\pi V_g} \sum_{r,r' \bmod q} e^{2\pi i a r r'/q} \int_{\mathbb{R}} \widetilde{\psi}(\nu) Y^{-i\nu} P_r(t+\sigma+\nu) R_{r'}(t+\sigma+\nu)\, d\nu, \tag{50}$$

where harmless Mellin sign choices are fixed by this formula and

$$P_r(s) = P^{-1} \sum_{\substack{p \equiv r \ (q) \\ p \equiv c \ (L_0)}} w(p/P) p^{is},$$

$$R_{r'}(s) = R^{-1} \sum_{m \equiv r' \ (q)} g(m/R)\gamma(m) m^{is}.$$

There is no unit restriction on r'. For any interval of length at most $4T_0$, (46) gives, uniformly in its location,

$$\int_I |R_{r'}(t)|^2 dt \ll L^C(1 + T_0/R) \ll L^C. \tag{51}$$

The final inequality uses the decisive scale separation

$$\frac{T_0}{R} \ll \frac{P L^{B_0+B_1}}{U_1 U_2} \le \exp\bigl(CL^{1/4} - 2L^{1/3}\bigr) L^{B_0+B_1} = o(1).$$

It comes from the two independent big-prime products already present in the graph comparison, rather than a new energy bound.

The Mellin transform in (50) has Schwartz decay with polynomial-logarithmic seminorms. Choose a fixed exponent C_ν large enough that its tail $|\nu| > L^{C_\nu}$, in weighted L^1, is smaller than the required logarithmic power. Uniformity in the location of I in (51), and the trivial bound $|P_r| \ll 1$, justify removing this tail. Now choose B_2 larger than C_ν, the bound for $|\sigma|$, and the smooth-cutoff derivative costs. Throughout $T_c < |t| \le T_0$, the argument $s = t + \sigma + \nu$

then has $|s| \gg L^{B_2}$. In (43) fix, for example, the integration-by-parts order $j = 2$ first. After choosing the requested exponent A, enlarge B_2 beyond the resulting fixed derivative costs. This gives

$$\sup_{T_c<|t|\le T_0,\ |\nu|\le L^{C_\nu}} |P_r(t+\sigma+\nu)| \ll_A L^{-A}$$

for any requested A. The smooth prime estimate applies because $T_0 = xL^{B_0+B_1} \le x^2$ for large x. Choose this A after the fixed partition and Mellin costs. Combine the estimate with (51), then apply Minkowski to the Mellin integral and sum the dyadic and residue pieces. This proves

$$\int_{T_c<|t|\le T_0} |\mathcal{T}(t)|^2 dt \ll_D L^{-D}.$$

Together with (49) this is (48); together with (47) and (45) it proves (44).

The parameter order is explicit. First fix the graph comparison, its list bounds, and its external cutoff costs. Then fix the major arcs and the required energy saving, the tail cutoff B_1, the bank smoothing precision, and the Mellin truncation. Choose the high-frequency boundary B_2 after those choices. Finally choose the model depth through (C3) at the required low-frequency precision. No exponent in the graph or the endpoint moment bound depends on that final depth.

On the arcs, Cauchy–Schwarz, (44), and Parseval for the other endpoint bound the normalized Fourier comparison by every requested inverse power of L, after paying the polynomially many arcs. Off the arcs, (39) and the two endpoint norms do so. The normalization is $1/Y$, exactly as in (37), with $k = h$. The harmonic conversion $\Psi(t) = t\Phi(t)$ therefore gives the same centered banked replacement as for reflected sums. This completes the proof of Proposition 6.1 and supplies the fixed-shift branch used in Section 8.

6.4 Classical estimates and the height range

The polynomial-logarithmic modulus here is an advantage over the reflected target grid, but the height is a separate constraint. Write $Z = \exp(L^{1/16})$. The classical Dirichlet zero-free region, together with Siegel's theorem in its ineffective all-character form, gives (42) with error

$$O\Big(PL^{C_w} e^{-cL^{1/16}}\Big) \qquad (|t| \le 2Z, q \le L^B).$$

For exact classical statements see [14, Proposition 51, Theorem 62(ii) and Exercise 45(i)–(ii)]. Indeed, a rectangle of height $3Z^2$ has a zero-free strip of width 2δ, with $\delta = cL^{-1/16}$. For a possible exceptional character, choose $\varepsilon < 1/(32(B+1))$ in Siegel's theorem; its zero then has distance $\gg L^{-1/32}$ from one, larger than the strip width for sufficiently large x. On the line $\Re s = 1-\delta$, the regularized logarithmic derivative has polynomial-logarithmic size. The same translated smooth Mellin contour as above, truncated at imaginary part Z^2, therefore saves $P^{-\delta} \le e^{-cL^{1/16}}$; fixed sufficiently high derivatives control the horizontal and truncation terms. Prime powers are smaller. Constants and the threshold in this classical estimate may be ineffective.

The factor two absorbs the retained polynomial-logarithmic Mellin translations when the outer frequency has absolute value at most Z. This covers the high-prime argument only up to Z. Its required upper height is $T_0 = xL^{B_0+B_1}$. At that endpoint the classical strip has width of order L^{-1}, whereas $\log P \le CL^{1/4}$; the contour saving $\exp(-c\log P/L)$ does not tend to zero. Thus a direct Siegel–Walfisz substitution does not prove the remaining weighted prime–cofactor energy on $Z < |t| \le T_0$. Walsh's classical-recovery remark [9, p. 2] motivates a different high-frequency argument, but the proof here retains the zero-free input on that range. Such a replacement for fixed shifts would not remove the target-sized conductor issue for reflected sums.

7 Reflected targets and exact rational grids

This branch treats a displacement equal to the target. Unlike the fixed-shift windows in Section 6, its rational denominators contain that target. We prove the residue-energy estimate with those denominators retained.

7.1 The exact target-grid energy

Let $\mathcal{P}_y = \{p : p \le y\}$, $y = \exp(L^{2/5})$, and $E_0 = \exp(\sqrt{L})$, as in Section 2. The endpoint F is the centered function in (38).

Proposition 7.1 (Target-grid residue energy)**.** *Suppose that*

$$q \le xL^B, \qquad \frac{Y}{qP_{\max}} \gg_A L^A \quad \text{for every fixed } A, \tag{52}$$

where $P_{\max} \le \exp(CL^{1/4})$ is the selected small-bank scale. For every prescribed $D > 0$, a fixed sufficiently deep density model gives

$$E_q(F) := \sum_{a \bmod q} \left| \sum_{n \equiv a \ (q)} F(n) \right|^2 \ll_D \frac{Y^2}{q} L^{-D}. \tag{53}$$

The estimate is uniform in all valuations of q. It also holds with $F(n)$ multiplied by $e^{2\pi i\beta n}$ when $|\beta|Y$ is bounded by a fixed power of L.

Proof. Refine to $Q = \operatorname{lcm}(q, L_0)$. Cauchy–Schwarz on the refined classes gives $E_q(F) \le (Q/q)E_Q(F) \le L_0 E_Q(F)$. For every gcd profile $d \mid Q$, put $r = Q/d$. Multiplicative orthogonality gives the exact identity

$$E_Q(F) = \sum_{d \mid Q} \frac{1}{\varphi(r)} \sum_{\chi \bmod r} \left| \sum_m F(dm)\chi(m) \right|^2. \tag{54}$$

The character zeros impose $(m, r) = 1$, so this identity retains prime powers, including primes dividing L_0. We do not assume that the endpoint is prime to L_0. Call a character exceptional when its primitive inducing character has conductor dividing L_0. There are only a fixed number of such characters for each r. In particular $L_0 \mid r$ is not required.

We first identify the negligible sets in the norm actually needed. For a set $\mathcal{B}$ and $|G| \le L^C$, Cauchy–Schwarz in each residue class gives

$$E_Q(G\mathbf{1}_{\mathcal{B}}) \le (1 + CY/Q) \sum_{n \asymp Y} |G(n)|^2 \mathbf{1}_{\mathcal{B}}(n). \tag{55}$$

The factor is $O(Y/Q)$ here. The following populations are exponentially small relative to Y, up to fixed logarithmic factors:

$$\begin{array}{ll} n_{\mathcal{P}_y} > E_0 : & O(YL^C e^{-cL^{1/10}}), \\ p \mid n,\ p \mid Q,\ p > y : & O(YL/y), \\ p \mid n,\ p \mid Q,\ p \text{ in the selected bank} : & O(YL/P_{\min}), \\ p^2 \mid n,\ p \text{ in the selected bank} : & O(Y/P_{\min}). \end{array} \tag{56}$$

For the first bound, take $a = c/\log y$. Counting multiples of each possible smooth record gives

$$\#\{n \le CY : n_{\mathcal{P}_y} > E_0\} \le CYE_0^{-a} \prod_{p \le y} (1 - p^{a-1})^{-1} \ll YL^C e^{-cL^{1/10}}.$$

This estimate uses no property of the indicators. The other bounds follow by counting multiples, using $P_{\min} \geq \exp(cL^{1/8})$. Thus all these sets are negligible in (55).

For nonexceptional characters remove only selected primes dividing Q and selected-prime squares. Do not yet impose $n_{\mathcal{P}_y} \leq E_0$, as that would couple the extracted prime to its cofactor. The unique selected small-bank prime gives $n = dpm$. On a dyadic prime interval $p \asymp P$, exact dilation invariance and the independence of b from this slot make its remaining coefficient $\gamma_d(m)$ independent of p, with $|\gamma_d(m)| \leq L^C$. It includes the condition that m has no other selected-bank prime, and every other bank weight. Its containing length is $M \asymp Y/(dP)$. Mellin inversion of the product cutoff separates the character sum as an integral of

$$A_\chi(t)B_{d,\chi}(t), \qquad A_\chi(t) = \sum_{p\asymp P,\ p\in\mathcal{P}_g} \chi(p)p^{it}, \quad B_{d,\chi}(t) = \sum_{m\asymp M} \gamma_d(m)\chi(m)m^{it}. \tag{57}$$

The omitted Mellin scaling has modulus one. The transform has polynomial-logarithmic weighted L^1 norm, including weight $1+|t|$. Condition (52) gives

$$M/r \asymp Y/(QP) \gg 1,$$

independently of d. Elementary character orthogonality therefore yields

$$\frac{1}{\varphi(r)} \sum_{\chi \bmod r} |B_{d,\chi}(t)|^2 \leq (1 + CM/r) \sum_{m\asymp M} |\gamma_d(m)|^2 \ll \frac{M^2}{r} L^C. \tag{58}$$

Apply Lemma 5.1 to A_χ and Minkowski to the integral in (57). Its nonexceptional energy at profile d is

$$\ll \frac{Y^2}{Qd} L^C \exp(-cL^{1/8}). \tag{59}$$

The prime dyadic decomposition costs only a fixed power of L. Summing d is harmless because $\sum_{d|Q} 1/d \leq Q/\varphi(Q) \ll \log\log(3Q)$. There is no number-of-characters or unweighted divisor-count loss.

For exceptional characters use the further restriction $n_{\mathcal{P}_y} = A \leq E_0$, and remove Q-prime divisors above y. The errors were already controlled in energy. Now the gcd profile is determined by the smooth record: $\gcd(A, Q) = d$, $n = Av$, with v $\mathcal{P}_y$-rough. If χ^* is the primitive character inducing an exceptional $\chi \bmod r$, then on the retained profile

$$\chi(n/d) = \chi^*(A/d)\chi^*(v).$$

Indeed $\operatorname{cond}\chi^* \mid r$ and $(A/d, r) = 1$. One must use χ^* here: a literal character modulo L_0 could vanish on a retained bad-prime prefix. Every other list weight depends only on A, is bounded by L^C, and the remaining periodic test in v has fixed period L_0. The one-variable hypothesis must center this full periodic test; matching only a scalar principal mean would not suffice.

There is still the condition that v avoid primes $p \mid Q$ above y. Remove it before applying the fixed-period centering lemma, extending the same bounded periodic test to those new points. At each retained record, multiples of such p number at most $CY/(Ap)$ on the containing interval $0 < v \leq CY/A$. Consequently the change in the profile sum is at most

$$L^C Y \sum_{\substack{A\ \mathcal{P}_y\text{-smooth}\\ d|A}} \frac{1}{A} \sum_{\substack{p|Q\\ p>y}} \frac{1}{p} \leq \frac{Y}{d} L^C V_y^{-1} \frac{\log Q}{y \log y} \ll_D (Y/d)L^{-D}. \tag{60}$$

We used the exact harmonic identity

$$\sum_{\substack{A\ \mathcal{P}_y\text{-smooth}\\ d|A}} \frac{1}{A} = \frac{1}{dV_y} \tag{61}$$

when d is $\mathcal{P}_y$-smooth; otherwise this restricted profile is empty.

After this extension, hypothesis (C3) applies record by record at length Y/A. It permits arbitrary bounded fixed-period tests, arbitrary choices of retained records, and the stated smooth phases. Choose its logarithmic accuracy after the list and cutoff costs. Using (61) gives, for every prescribed D_0,

$$\left|\sum_m F_{\mathrm{good}}(dm)\chi(m)\right| \ll (Y/d)L^{-D_0}. \tag{62}$$

Restoring the avoidance condition costs (60). In (54) the total exceptional contribution is at most

$$\frac{Y^2}{Q}L^{-2D_0}\sum_{d|Q}\frac{1}{d}\frac{Q/d}{\varphi(Q/d)} \ll \frac{Y^2}{Q}L^{-2D_0+C}. \tag{63}$$

Choose D_0 large, add the negligible sets, and use the refinement inequality to prove (53). Multiplying the cutoff by $e^{2\pi i\beta Yu}$ changes only its fixed logarithmic derivative bounds when $|\beta|Y \le L^B$, proving the last assertion. □

This proof uses only boundedness, exact prime invariance, and the recordwise periodic means in (C1)–(C3). In particular no subgroup support, pole order, positivity of an individual character branch, or unit-support condition at the fixed modulus enters the energy estimate.

7.2 From target-grid energy to comparison cancellation

Specialize the common Fourier reduction of Section 5 to the reflected family $k = T = x$. Use the product scales U_i and $Y = xU_1U_2$ from that reduction. Choose an outer cutoff B_{out}, after the finite comparison family has been fixed. Dirichlet approximation with denominator ceiling $U_1U_2/L^{B_{\mathrm{out}}-1}$ makes the multiplier smaller than any requested inverse logarithmic power unless $T\theta$ lies within $L^{B_{\mathrm{out}}}/(U_1U_2)$ of a rational of denominator $r \le L^{B_{\mathrm{out}}}$. The corresponding physical frequencies are covered by grids of denominator $q = Tr$, with widths $O(L^{B_{\mathrm{out}}}/(TU_1U_2)) = O(L^{B_{\mathrm{out}}}/Y)$. Moreover

$$\frac{Y}{qP_{\max}} \gg \exp(cL^{1/3} - CL^{1/4})L^{-B_{\mathrm{out}}},$$

so (52) holds uniformly in every gcd profile.

Discrete additive Parseval gives the exact normalization

$$\sum_{j \bmod q}|\widehat{F}(j/q+\beta)|^2 = qE_q(Fe^{2\pi i\beta n}). \tag{64}$$

Integrating over width $L^{B_{\mathrm{out}}}/Y$ and applying Proposition 7.1 gives YL^{-D} for every desired D, after demanding a larger residue-energy saving. The polynomial-logarithmic family of r's is paid at the same step. In particular the T grid points do not create a factor T: they are already included in (64). Cauchy–Schwarz and Parseval on the other endpoint bound (37) on these grids by an arbitrary inverse power of L. Off the grids, use (39) and the endpoint norms. This proves that every comparison in (33) is negligible to the requested logarithmic accuracy.

Finally $\Psi(t) = t\Phi(t)$ makes

$$\frac{\Psi(|w|/x)}{|w|} = \frac{\Phi(|w|/x)}{x}.$$

Thus the harmonic raw correlation is $1/x$ times the ordinary banked correlation on the reflected line. The comparison at scale Y has exactly the corresponding factor $1/Y$. Paying the finite

comparison coefficients therefore proves the centered banked replacement required in the assembly.

Fixed bad-supported coefficients in $Aa + Bb = T$ cause no new analytic condition. Embed, for example, $\widetilde{F}(n) = \mathbf{1}_{A|n} F(n/A)$. Its exact residue energy is

$$E_q(\widetilde{F}) = E_{q/(q,A)}(F),$$

and its core length is divided by the fixed A. All energy ratios change only by fixed factors. Bank primes are units at A, B, so dilation invariance and shared-label removal persist. This embedding also avoids treating the physical summands as units modulo L_0 when they are not.

8 Reconstruction and quantitative comparison

The results above require a positive logarithmic saving. We now record the numerical choices and prove that every approximation fits them. For the fixed subgroup theorem one may take $\delta = 1/(8192\varphi(L_0))$. In particular $\delta_{12} = 1/32768$, the coprime C_{12} saving is $\delta_{12}/2 = 1/65536$, and the prime-3 $(\mathrm{mod}\ 4)$ self-sum saving is $\delta_- = 1/16384$. These choices concern the semigroup asymptotics; the ordinary quadratic-indicator statement has the remainder stated in Section 10. We assemble the two frequency estimates into Theorem 2.1 at the sparse scale in (C4).

Write $\phi = \varphi(L_0)$, $L = \log X$, $\delta_0 = 1/(8192\phi)$, and use (5).

8.1 An explicit discarded-bank exponent

Primes in the reduced bank class $c \bmod L_0$ have reciprocal mass $(1/(8\phi) + o(1)) \log L$ in the small range and $(1/(15\phi) + o(1)) \log L$ in the big range. Choose

$$b_X = \left\lfloor \frac{\log L}{32\phi} \right\rfloor$$

disjoint banks of each size, greedily giving each reciprocal mass $V_g = 1 + o(1)$. The overshoot per bank is at most the largest reciprocal prime. Delete bank primes dividing the displacement. For a reflected target their total reciprocal mass is $O(L/(P_{\min} \log P_{\min}))$; for a fixed shift the deletion is eventually empty. Thus the total mass error is smaller than every fixed inverse logarithmic power, uniformly in the reflected target.

The independent law (7) gives

$$\mathbb{E} I_g = \prod_{p \in \mathcal{P}_g} (1 - 2/p) \sum_{\substack{p,q \in \mathcal{P}_g \\ p \neq q}} \frac{1}{(p-2)(q-2)} = V_g^2 e^{-2V_g} + O\left(\sum_{p \in \mathcal{P}_g} p^{-2} \right) \longrightarrow e^{-2} > 1/8. \quad (65)$$

Banks are independent in this law. Chernoff's inequality bounds the probability of fewer than $b_X/16$ successful banks in either range by

$$2e^{-b_X/64} \ll L^{-1/(4096\phi)} = L^{-2\delta_0}, \quad (66)$$

since $b_X \geq \log L/(64\phi)$ eventually. Apply (C4) to this bad-bank indicator and use $|f_i - \mathcal{D}_i| \leq f_i + \mathcal{D}_i$. The discarded centered correlation is $O(\mathcal{S}_X L^{-2\delta_0})$, plus its power-smaller and exponential endpoint errors.

The reconstruction is exact. For b banks, the Boolean function asserting at least k successes has expansion $\sum_T c_T \prod_{g \in T} I_g$, where $c_T = 0$ for $|T| < k$ and $|c_T| = \binom{|T|-1}{k-1}$ otherwise. Binomial inversion in the number of successful banks proves this identity, and its total coefficient mass

is at most 3^b. Multiply the two size-range identities. For every active set T, exact prime invariance gives

$$\xi_i(x)w_j(y)\prod_{g\in T} I_g(x,y) = \left(\prod_{g\in T} V_g^2\right)\xi_i(x)w_j(y)\Omega_T(x)\Omega_T(y), \qquad w_j \in \{f_j, \mathcal{D}_j\}. \tag{67}$$

These weights introduce no representation multiplicity. The full reconstruction coefficient mass is at most

$$3^{2b_X}\prod_g \max(1, V_g^2) \le L$$

for large X. Indeed the second product is $1+o(1)$, and $2b_X \log 3 \le (\log 3)\log L/(16\phi)$. Prime squares in active banks have population at most a fixed logarithmic factor times $X\sum_{p\in\mathcal{B}} p^{-2}$, which is exponentially small. This estimate follows from bounded endpoints and counting multiples, independently of a pair asymptotic.

8.2 Moving cutoffs and reflected target localization

Put $\epsilon = L^{-\delta_0}$. In the reflected case choose $0 \le \Phi_\epsilon \le 1$, equal to one on $[2\epsilon, 1-2\epsilon]$, supported in $[\epsilon, 1-\epsilon]$, with transitions of width ϵ. The same derivative estimates apply when additional interval endpoints are smoothed at that width. Set $\Psi_\epsilon(t) = t\Phi_\epsilon(t)$ and $F_\epsilon(u) = \Psi_\epsilon(e^u)$. Differentiation gives

$$\|F_\epsilon\|_1 \ll 1, \qquad \|F_\epsilon^{(j)}\|_1 \ll_j \epsilon^{1-j} \quad (j \ge 1).$$

At the lower transition logarithmic differentiation cancels its small scale. Fourier integration by parts yields

$$\int |\widehat{F}_\epsilon(\zeta)|(1+|\zeta|)^r\, d\zeta \ll_r \epsilon^{-r-1}, \qquad \int_{|\zeta|>L^3} |\widehat{F}_\epsilon(\zeta)|\, d\zeta \ll_j L^{-(3-\delta_0)(j-1)}. \tag{68}$$

The external phase range stays at L^3; its tail derivative order is chosen last.

On a contributing reflected edge $n<0$, $n' = n + XD > 0$, put $t = |n|/(XD)$. Then

$$\frac{n'}{|n|} = \frac{1-t}{t} \in \left[\frac{\epsilon}{1-\epsilon}, \frac{1-\epsilon}{\epsilon}\right]. \tag{69}$$

For a source dyad $|n| \asymp Y$, insert a target cutoff on $[c\epsilon Y, C\epsilon^{-1}Y]$, equal to one on that range, *before* Mellin inversion, and retain it in each pairing. Removed edges originally had zero full Mellin integral. Shared-label division preserves the ratio. There are $O(\log L)$ target dyads Y', each with uniform rescaled logarithmic derivatives. Their exact-list endpoint moments are $O(Y)$, $O(Y')$, by counting multiples of a fixed list product. This introduces no exponent depending on the eventual mark length or density depth. Pairing costs at most $\sqrt{Y'/Y} \ll \epsilon^{-1/2}$. The Mellin norm and dyadic count together cost at most $L^{5\delta_0/2}$; reserve $L^{8\delta_0}$ for all localization.

For comparison scale $Y \asymp XD_1D_2$, use the actual target dyad $c\epsilon Y \le Y' \le C\epsilon^{-1}Y$. The reflected major window has width $L^{B_{\mathrm{out}}}/Y$, so $|\beta|Y' \ll L^{B_{\mathrm{out}}+\delta_0}$. Proposition 7.1 applies at this enlarged logarithmic window exponent. Its integrated rescaling costs at most ϵ^{-2}, within the reserve. At the smallest dyad,

$$\frac{Y'}{qP} \gg e^{cL^{1/3}-CL^{1/4}} L^{-B_{\mathrm{out}}-\delta_0} \gg_A L^A. \tag{70}$$

Removing a gcd divisor scales length and modulus equally, so this bound holds in every target profile. For a fixed shift the two endpoints are comparable on the original support; the cutoff and pairing cost is no larger. Its wider major windows are handled by Proposition 6.1, not by extending (53) beyond its proved range.

8.3 Precision order and completion

Let $Q_0 = \lceil \kappa_0 \rceil + 30$. Request normalized graph-pairing error $L^{-(Q_0+3)}$ per source dyad. There are at most L^2 such dyads. Fix root and ideal accuracies for this request, then the probe complexity and eventual mark length, keeping the phase range L^3. Only after this family is fixed let C_{cmp} dominate its coefficient and dyadic costs, and C_w its endpoint pointwise exponent. These exponents do not depend on the density depth. If m_* bounds free draws per group, the dyadic product laws obey

$$\nu_i(n) \le L^{K_*}/n, \qquad K_* = \frac{\log(m_*!) + m_* \log 2}{16\phi}.$$

Their mass is at most one, so their squared norm at length T_i is at most L^{K_*}/T_i. In (39) take $C_{\mathrm{prod}} = 2K_* + 2$, and choose

$$\begin{aligned} A_{\min} &> C_{\mathrm{cmp}} + 2C_w + \kappa_0 + 50 + 8\delta_0, \\ B_{\mathrm{out}} &> 2A_{\min} + C_{\mathrm{prod}} + 4. \end{aligned} \tag{71}$$

Dirichlet approximation with denominator ceiling $T_1T_2/L^{B_{\mathrm{out}}-1}$ gives multiplier $O(L^{-A_{\min}})$ off the outer grids. Their denominators are Xr in the reflected case and $|h|r$ for a fixed shift, with $r \le L^{B_{\mathrm{out}}}$. The middle terms in the multiplier bound are superlogarithmically small because $T_i \ge \exp(L^{1/3})$.

Request integrated major-energy saving $T_E > 2(C_{\mathrm{cmp}} + C_w + \kappa_0 + 50 + 8\delta_0)$. In the reflected case enlarge the residue-energy accuracy to pay both arc widths and the number of outer denominators. Exact discrete Parseval accounts for all target-grid points. In the fixed-shift case use the low/high Dirichlet-frequency proof at that accuracy, choosing its smoothing and low-frequency threshold after the fixed comparison costs. The cofactor-length margin pays its mean-value bound. Both major contributions are then at most $L^{-T_E/2+C_w+C_{\mathrm{cmp}}+8\delta_0}$. The minor bound has the same ample margin.

Choose the model depth last, using (C3) at the largest period, smoothness and precision exponents now required. Its uniform-in-depth pointwise and sieve exponents in (C1) and (C4) make this order noncircular. Choose the external Mellin tail derivative order last as well. Nothing changes the earlier ideal family or phase range.

Completion of Theorem 2.1. After localization, source-dyad summation and reconstruction, the graph and comparison errors are

$$O(XL^{-Q_0+8\delta_0}) = o(\mathcal{S}_X L^{-\delta_0}).$$

The rare-bank term is $O(\mathcal{S}_X L^{-2\delta_0})$. All record, square and CRT tails are exponential or power-smaller than $\mathcal{S}_X L^{-\delta_0}$. For each ordered endpoint pair use

$$f_1(x)f_2(y) - \mathcal{D}_1(x)\mathcal{D}_2(y) = (f_1(x) - \mathcal{D}_1(x))f_2(y) + \mathcal{D}_1(x)(f_2(y) - \mathcal{D}_2(y)).$$

The banked comparison bounds each centered term to the required precision. The argument is symmetric in the endpoints; for a fixed shift, interchange them and replace h by $-h$.

To restrict to an arbitrary interval, smooth its endpoints at width ϵX; in the reflected case first remove the two extreme strips. At most four strips of total length $O(\epsilon X)$ are lost. Condition (C4), applied to each nonnegative pair weight, bounds their total contribution by $O(\epsilon \mathcal{S}_X + X^{1/2}L^C + XL^C e^{-cL^{1/10}})$. For an interval shorter than this width use that upper bound directly. Otherwise total variation remains $O(1)$, and a transition at t_0 has logarithmic derivative norm at most $O_j(t_0^j \epsilon^{1-j})$. The Mellin and localization budgets are therefore uniform in both endpoints. Since $\epsilon = L^{-\delta_0}$, every error has the asserted size. This proves the comparison theorem in both regimes. □

9 Fractional rough models and one-variable means

We develop the rough Hankel model and its one-variable centering estimate before applying the comparison theorem. Fix an even L_0, put $G = (\mathbb{Z}/L_0\mathbb{Z})^\times$, and take a proper subgroup $H \le G$. Write

$$\mathcal{P}_H = \{p : p \nmid L_0,\ p \bmod L_0 \in H\}, \qquad C_H = \langle \mathcal{P}_H \rangle, \qquad f = \mathbf{1}_{C_H},$$

and set $\rho = |H|/|G|$, $k_H = [G : H] = 1/\rho$. The outer scale is X; later reflected applications take $X = N$. Put

$$L = \log X, \quad y = \exp(L^{2/5}), \quad \mathcal{P}_y = \{p : p \le y\}, \quad E = \exp(\sqrt{L}), \quad V_y = \prod_{p \le y}(1 - 1/p).$$

Let $n_{\mathcal{P}_y}$ denote the $\mathcal{P}_y$-supported part of n, and $n_{>y} = n/n_{\mathcal{P}_y}$. Define

$$E_+(y) = \prod_{\substack{p \le y \\ p \in \mathcal{P}_H}} (1 - 1/p), \quad E_-(y) = \prod_{\substack{p \le y \\ p \nmid L_0,\ p \notin \mathcal{P}_H}} (1 - 1/p), \quad \alpha_0 = \varphi(L_0)/L_0.$$

Then $V_y = \alpha_0 E_+(y) E_-(y)$ for large X.

9.1 Fractional poles and centered rough means

Character orthogonality gives, for $\Re s > 1$,

$$F_H(s) = \prod_{p \in \mathcal{P}_H} (1 - p^{-s})^{-1} = G_H(s) \prod_{\eta \in H^\perp} L(s, \eta)^\rho. \tag{72}$$

At a forbidden good prime of coset order $r > 1$, the correction is $(1 - p^{-rs})^{1/r}$; thus G_H converges absolutely for $\Re s > 1/2$. The literal principal character carries the bad Euler factors. The positive marginal constant is

$$c_H = \frac{1}{\Gamma(\rho)} \lim_{s \downarrow 1} (s - 1)^\rho F_H(s);$$

its Euler expression is (89). Every $\eta \in H^\perp$ equals one on the support of f; every such twist therefore has the same pole branch.

Define the rough series and its regular part by

$$F_R(s) = F_H(s) \prod_{\substack{p \le y \\ p \in \mathcal{P}_H}} (1 - p^{-s}) = (s - 1)^{-\rho} H_R(s).$$

Here $H_R(1) = \Gamma(\rho) c_H E_+(y)$. Fixed-order Cauchy estimates on a disk of radius $c/\log y$ give

$$|H_R^{(j)}(1)|/j! \ll_j H_R(1)(C \log y)^j. \tag{73}$$

For a fixed depth K, define on polynomial scales in X

$$M_{R,K}(t) = \sum_{j=0}^{K} \frac{H_R^{(j)}(1)}{j!} \frac{(\log t)^{\rho - j - 1}}{\Gamma(\rho - j)}. \tag{74}$$

For $X^a \le t \le X^C$, with fixed $0 < a < C$, this is $c_H E_+(y)(\log t)^{\rho-1}(1 + O_K(L^{-3/5})) > 0$. For definiteness fix a smooth real function η_0, with $0 \le \eta_0 \le 1$, zero on $(-\infty, 1/4]$ and one on $[1/3, \infty)$. Set $M_{R,K}(t) = 0$ for $t \le X^{1/4}$, and multiply the displayed expression by $\eta_0(\log t/L)$ above that point. On the transition range it has the same positive upper bound. All rough scales used in (C3) are at least $X^{1/2}/E > X^{1/3}$, so the cutoff never changes their means. Only the polynomial bulk enters principal means.

Lemma 9.1. *For fixed $B, A_0 > 0$, sufficiently large fixed K has the following property. If b is bounded and periodic of period at most L^B, and W is a smooth dyadic weight at a scale $X^{2/5} \le T \le X^2$, with fixed derivatives bounded by fixed powers of L, then*

$$\sum_n W(n)b(n)f(n)\mathbf{1}_{(n,\prod_{p\le y}p)=1} - \sum_n W(n)b(n)\mathbf{1}_{(n,\prod_{p\le y}p)=1}\left(\sum_{\eta\in H^\perp}\eta(n)\right)\frac{M_{R,K}(n)}{V_y} \ll TL^{-A_0}.$$

The same holds with a factor n^{iu}, $|u| \le L^B$.

Proof. Let $q \le L^B$ be the period of b. Refine it to $Q = \operatorname{lcm}(q, L_0)$ and expand on unit residues by characters. All primes dividing Q are at most y. For $\psi \bmod Q$, the twisted rough series is

$$\prod_{\eta\in H^\perp} L(s,\psi\eta)^\rho\, G_\psi(s) \prod_{\substack{p\le y\\ p\in\mathcal{P}_H}} (1-\psi(p)p^{-s}).$$

If ψ is inflated from $H^\perp$, this equals F_R; otherwise all its character factors are nonprincipal. Use smooth Mellin inversion at height $\exp(L^{1/10})$, moving to $1 - c/\log y$. The classical zero-free region is wider there at these heights and polylogarithmic conductors. Siegel's theorem places a possible exceptional real zero outside the strip for large X, after choosing its fixed exponent depending on B. The finite Euler product costs $L^{O(1)}$; the displacement gives $\exp(-cL^{3/5})$, dominating polynomial L-function growth at that height. Smooth tails are smaller than every fixed logarithmic power. The threshold can be ineffective.

Hankel integration gives (74) for each principal branch. Its relative Taylor error is $O_K(L^{-3(K+1)/5})$, by (73). Choose K after the derivative and character-count costs. For the unweighted rough population the series $\zeta(s)\prod_{p\le y}(1-p^{-s})$ has principal density V_y, and only the trivial character contributes. The projection $\sum_{\eta\in H^\perp}\eta$ therefore reproduces every principal branch of the first sum with its exact coefficient. The factor n^{iu} changes only the derivative budget and choice of K. □

The contour argument uses the fractional pole and the absolute convergence and fixed-order derivative bounds of its Euler correction. It therefore also applies when that correction is inverted and remains absolutely convergent for $\Re s > 1/2$. The quadratic-indicator application uses precisely this variant, with exponent $1/2$ and its two principal character branches.

For $U \subset H$, set

$$\mathcal{D}_U(n) = k_H\mathbf{1}_{n \bmod L_0\in U} f(n_{\mathcal{P}_y})\frac{M_{R,K}(n_{>y})}{V_y}, \qquad \xi_U(n) = f(n)\mathbf{1}_{n \bmod L_0\in U} - \mathcal{D}_U(n). \tag{75}$$

Then $\mathcal{D}_U \ge 0$, $\mathcal{D}_U \ll L^{-3(1-\rho)/5}$ and $|\xi_U| \ll 1$. Conditioning on an exact record $n_{\mathcal{P}_y} = r \le E$ leaves residual state U/r, so Lemma 9.1 centers its principal mean. Fixed divisor-growth record coefficients cost only powers of L, since $\sum_{r\ \mathcal{P}_y\text{-supported}} d_j(r)/r \ll L^{O_j(1)}$. Rankin's bound with exponent $c/\log y$ gives

$$\#\{n \le T : n_{\mathcal{P}_y} > E\} \ll TL^{O(1)}e^{-cL^{1/10}} \quad (X^a \le T \le X^C). \tag{76}$$

Indeed the Euler moment is $L^{O(1)}$, whereas $E^{-c/\log y}$ gives the exponential saving. The same argument permits fixed divisor weights. For $H = G$, use the exact model $\mathcal{D}_U = \mathbf{1}_{n \bmod L_0\in U}$ and $\xi_U = 0$.

10 Ordinary quadratic indicators and fixed-shift pairs

The same comparison also treats ordinary quadratic norm indicators. In this application inert primes may occur to an even exponent; they are not excluded altogether. The rough Dirichlet series still has two principal character branches, and their projection retains the conditions at ramified primes. This gives a common statement for the five pair conjectures considered by Smilansky [10, Sections 6.1 and 6.3].

10.1 A quadratic character and finitely many valuation rules

Let χ be a primitive nonprincipal quadratic Dirichlet character of conductor q. Fix a finite set of primes B_0 containing every prime divisor of q, and masks $\alpha_p : \mathbb{Z}_{\ge 0} \to \{0,1\}$, $p \in B_0$, with $\alpha_p(0) = 1$. Define the multiplicative indicator

$$f(n) = \prod_{p \in B_0} \alpha_p(v_p(n)) \prod_{\substack{p \notin B_0 \\ \chi(p) = -1}} \mathbf{1}_{2 \mid v_p(n)} \quad (n \ge 1), \qquad f(n) = 0 \quad (n \le 0).$$

Thus split primes outside B_0 have no valuation restriction. The finite masks are fixed; no periodicity in their exponent is required. Write $F_p(s) = \sum_{e \ge 0} \alpha_p(e) p^{-es}$ for $p \in B_0$.

The bad-prime factor must retain a joint residue condition. Let $x = (x_p)_{p \in B_0} \in \prod_{p \in B_0} \mathbb{Z}_p$, and put $e_p = v_p(x_p)$ when all these valuations are finite. For each $\ell \mid q$, define a unit modulo $\ell^{v_\ell(q)}$ by

$$u_\ell(x) = \ell^{-e_\ell} x_\ell \prod_{\substack{p \in B_0 \\ p \ne \ell}} p^{-e_p} \pmod{\ell^{v_\ell(q)}}.$$

These residues define $u_{B_0}(x) \in (\mathbb{Z}/q\mathbb{Z})^\times$ by CRT. At an ordinary positive integer this is exactly $n/n_{B_0} \pmod q$. Set

$$\vartheta_{B_0}(x) = \prod_{p \in B_0} \alpha_p(e_p) \mathbf{1}_{\chi(u_{B_0}(x)) = 1},$$

$$W_{B_0}(h) = \int_{\prod_{p \in B_0} \mathbb{Z}_p} \vartheta_{B_0}(x) \vartheta_{B_0}(x+h)\, dx. \tag{77}$$

Normalized product Haar measure is used. Define the integrand to be zero on the null sets with an infinite valuation. In general W_{B_0} is coupled: it is not a product of separate primewise tests.

Theorem 10.1. *Assume the analytic inputs of Section 3. For the fixed data above and every fixed $h \ne 0$,*

$$\sum_{1 \le n \le X} f(n) f(n+h) = C_{\chi,\boldsymbol{\alpha}}(h) \frac{X}{\log X} + o_{\chi,\boldsymbol{\alpha},h}\left(\frac{X}{\log X}\right), \tag{78}$$

where

$$C_{\chi,\boldsymbol{\alpha}}(h) = \frac{4L(1,\chi)}{\pi} W_{B_0}(h) \prod_{p \in B_0} \frac{1 - \chi(p)/p}{1 - 1/p} \cdot \prod_{\substack{p \notin B_0,\ \chi(p) = -1 \\ p \mid h}} \frac{1 - p^{-v_p(h)-1}}{1 - p^{-1}}. \tag{79}$$

If $W_{B_0}(h) = 0$, the pair count is identically zero. Otherwise the constant is positive. The theorem concerns the specified Euler indicator; it does not identify that indicator with the values of an arbitrary principal quadratic form.

Proof. We separate the arithmetic model, its positive upper envelope, and its correlation constant. Only the final assembly invokes the binary comparison theorem.

Stage 1: model construction and rough means. Choose a fixed even L_0 divisible by q and all primes in B_0, and use the prime class one modulo L_0. Put $L = \log X$, $y = e^{L^{2/5}}$, $\mathcal{P}_y = \{p : p \le y\}$, $E = e^{\sqrt{L}}$, and $V_y = \prod_{p\le y}(1 - 1/p)$.

The exact Euler identity for $F(s) = \sum f(n)n^{-s}$ is

$$F(s)^2 = \zeta(s)L(s,\chi) \prod_{\substack{p\notin B_0\\ \chi(p)=-1}} (1-p^{-2s})^{-1} \prod_{p\in B_0} F_p(s)^2(1-p^{-s})(1-\chi(p)p^{-s}). \tag{80}$$

The first correction converges absolutely for $\Re s > 1/2$. Every finite factor is analytic near one and positive there on the real axis. With this branch convention, the one-variable constant c_f satisfies

$$c_f^2 = \frac{L(1,\chi)}{\pi} \prod_{\substack{p\notin B_0\\ \chi(p)=-1}} (1-p^{-2})^{-1} \prod_{p\in B_0} F_p(1)^2(1-p^{-1})(1-\chi(p)p^{-1}). \tag{81}$$

Here $F_p(1)$ denotes the local factor evaluated at $s = 1$, not at its series variable. The factor $\pi = \Gamma(1/2)^2$ is the Selberg–Delange normalization.

Remove all factors at primes in $\mathcal{P}_y$, and write the resulting rough series as $F_R(s) = (s-1)^{-1/2}H_R(s)$. Its leading coefficient is

$$H_R(1)/\Gamma(1/2) = c_f \prod_{p\le y} F_{f,p}(1)^{-1},$$

where $F_{f,p}$ is the actual local factor of f. The Cauchy bound $H_R^{(j)}(1)/j! \ll_j H_R(1)(C\log y)^j$ permits the fixed-depth density

$$M_K(t) = \sum_{j=0}^{K} \frac{H_R^{(j)}(1)}{j!} \frac{(\log t)^{-1/2-j}}{\Gamma(1/2-j)}.$$

Use the cutoff $\eta_0(\log t/L)$ from Section 9, setting the model to zero for $t \le X^{1/4}$. It equals one on the rough range in (C3), where $Y/r \ge X^{1/2}/E > X^{1/3}$. For $n = av$, $a = n_{\mathcal{P}_y}$, define

$$\mathcal{D}(n) = f(a)(1+\chi(v))M_K(v)/V_y. \tag{82}$$

For every fixed depth, $\mathcal{D} \ll L^{-3/10}$ on the relevant polynomial ranges, so (C1) holds eventually with a depth-independent bound one. Multiplication by a split bank prime changes only a, preserving $f(a)$ and v exactly. This proves (C2), including repeated prime factors.

For (C3), inflate a test character η to a modulus containing q. The square of its rough Dirichlet series has factors $L(s,\eta)L(s,\eta\chi)$, an absolutely convergent prime-square correction, and the removed small-prime factors. All primes dividing the polylogarithmic test modulus lie in $\mathcal{P}_y$. Only $\eta = 1$ and $\eta = \chi$ give a principal branch, and their rough coefficients agree pointwise. The factor $1 + \chi(v)$ in (82) reproduces both. The proof of Lemma 9.1 now applies with exponent $1/2$: the contour is $1 - c/\log y$, its height is $e^{L^{1/10}}$, and the Taylor remainder has relative size $O_K(L^{-3(K+1)/5})$. The inverted inert-prime correction in (80) has the same absolute convergence and derivative bounds as the correction in that lemma. The finite masked factors have been removed completely from the rough variable. Thus they introduce no further rough branches. Conditioning on $a \le E$ leaves only its fixed multiplier $f(a)$, outside the periodic rough sum. Choosing K after the desired logarithmic accuracy proves (C3).

Stage 2: the prime-square sieve envelope. For (C4), keep exact bank records $r, s \le E$; larger ones have contribution $XL^{O(1)}e^{-cL^{1/10}}$ by (76). The retained records are coprime because bank primes dividing h have been deleted. Their divisibilities give one progression

with $|J|/(rs)+O(1)$ terms. Exactness contributes two excluded bank roots outside rs, and one at a prime of rs. Together with $1/(rs)$, these are exactly the record masses (7).

At an inert prime outside B_0, a necessary condition is $v_p(n)\neq 1$. It excludes $p-1$ classes modulo p^2. For $p\nmid h$, the two endpoints exclude $2(p-1)$ classes; bank-record slopes are units at these primes. Drop the finitely many primes dividing $2h\prod_{p\in B_0}p$ for an upper bound. The nonbank sieve dimension is one, with a convergent $1+O(p^{-2})$ correction to the ordinary two-root product. A Selberg upper sieve with prime-square atoms, level $z=X^{1/100}$, and $w=z^{1/16}$, has endpoint error $O(z^{4+o(1)})$ per retained record. This follows by summing the CRT root errors: a least-common-multiple atom has at most $2^{\omega(d)}d$ roots and modulus d^2. Summing the $E^2=X^{o(1)}$ records remains power smaller than $X^{1/2}$.

For ff, $\mathcal{D}f$, and $\mathcal{D}\mathcal{D}$, the nonbank sieve denominators have orders respectively $\log w$, $(\log y\log w)^{1/2}$, and $\log y$. The bound $\mathcal{D}\ll L^{-3/10}$ supplies the missing factor on each model side. Explicitly, since $\log y=L^{2/5}$ and $\log w\asymp L$, their factors are

$$(\log w)^{-1}\ll L^{-1},\qquad \frac{L^{-3/10}}{(\log y\log w)^{1/2}}\ll L^{-1},\qquad \frac{L^{-3/5}}{\log y}=L^{-1}.$$

All three cases therefore give the common scale $\mathcal{S}_X=X/L$, with the exact bank masses. This proves (C4), uniformly in the interval J. The exponents in the exceptional-record errors do not depend on the final fixed depth.

Stage 3: model correlation and the arithmetic constant. We calculate the model main term separately from the analytic replacement. Let $I_y=\{p\le y: p\notin B_0,\ \chi(p)=-1\}$ and put

$$\Lambda_y=\frac{2c_f}{\sqrt{L}}\prod_{p\in B_0}\frac{F_p(1)^{-1}}{1-p^{-1}}\prod_{p\in I_y}(1+p^{-1}).$$

Outside the negligible records $n_{\mathcal{P}_y}>E$, the rough cofactor has $\log v=L+O(\sqrt{L})$. If the small inert exponents are even, its character is $\chi(n/n_{B_0})$. Hence

$$\mathcal{D}(n)=\Lambda_y\vartheta_{B_0}(n)\prod_{p\in I_y}\mathbf{1}_{2|v_p(n)}(1+O_K(L^{-1/2})). \tag{83}$$

At an inert prime the local pair density is

$$\delta_p(h)=\frac{p}{p+1}(1-p^{-v_p(h)-1}). \tag{84}$$

To see this, put $e=v_p(h)$ and partition $v_p(n)$ below, equal to, and above e. Below e, both endpoints have the same valuation, and the permitted even terms form a geometric sum. If e is odd there are no further terms. If e is even, the stratum with both valuations equal to e has mass $p^{-e}(p-2)/p$; the two possible higher even tails have combined mass $2p^{-e}/(p(p+1))$. The two parity cases give (84), also for $p=2$.

We justify the product of these densities at the growing cutoff. Extract the inert square parts d^2,e^2, with d,e supported on I_y, and retain $d,e\le E$. Truncate each bad-prime power at E. The omitted integer count is $O_{B_0,h}(X/E+\sqrt{X})$, and the omitted Haar mass is $O_{B_0,h}(E^{-1})$. Each retained joint record and residual unit state gives an empty stratum or finitely many progressions modulo $X^{o(1)}$. On them the two residual forms must be units at the primes in I_y. Empty root cells are discarded; the remaining sieve dimension is at most one. There are at most $O_{B_0,q}(E^2(\log E)^{2|B_0|})$ retained cells: two square prefixes, two bounded valuation lists, and a fixed number of residual unit choices. Their moduli are at most $O_{B_0,q}(E^{4+2|B_0|})$. Use upper and lower fundamental-lemma weights of level $D_0=X^{1/100}$. Their parameter $\log D_0/\log y\asymp L^{3/5}$

gives relative error smaller than every log power. Since the root count at a squarefree sieve divisor is at most $2^{\omega(d)}$, a conservative total endpoint bound is

$$E^{O_{B_0,q}(1)}D_0^2(\log D_0)^{O(1)} = X^{1/50+o(1)}.$$

This is negligible compared with X/L^A for every fixed A. Their main masses are summable: compatibility requires $(d,e)^2 \mid h$, and their progression densities are $\ll_h 1/\operatorname{lcm}(d^2,e^2)$, with geometric bad-prime tails. Summing the disjoint valuation cells proves

$$\sum_{X\le n\le 2X} \mathcal{D}(n)\mathcal{D}(n+h) = \Lambda_y^2 X W_{B_0}(h)\prod_{p\in I_y}\delta_p(h) + o_h(X/L).$$

For $p\nmid h$, $(1+p^{-1})^2\delta_p(h) = 1-p^{-2}$; for $p\mid h$ the additional ratio is $(1-p^{-v_p(h)-1})/(1-p^{-1})$. Substitution of (81) cancels every $F_p(1)$ and the inert Euler product, leaving (79). Thus the model correlation has the asserted constant on each dyadic interval.

Assembly. Stages 1 and 2 establish (C1)–(C4) at scale X/L. The comparison theorem now gives

$$\sum_{X\le n\le 2X}\{f(n)f(n+h) - \mathcal{D}(n)\mathcal{D}(n+h)\} = o_h(X/L). \tag{85}$$

Combining this with Stage 3 proves the dyadic asymptotic for the ordinary indicator. Summing dyadically gives (78); the upper bound from (C4) controls the part below εX, after which ε tends to zero.

Finally, if an actual positive pair has $f(n)f(n+h) = 1$, then both residual characters equal one. Its valuations and residual units are preserved on a sufficiently deep joint B_0-adic cylinder. That cylinder has positive Haar measure, so $W_{B_0}(h) > 0$. This proves the exact zero assertion. □

10.2 The five quadratic forms

Let $b_d(n)$ indicate the positive integers of the form x^2+dy^2, with $x,y\in\mathbb{Z}$. Coordinates may vanish, and no coprimality condition is imposed. Put $b_d(n) = 0$ for $n\le 0$. Smilansky's prime-factor criterion [10, Theorem 10] identifies these ordinary indicators with the preceding Euler indicators:

d	χ	B_0	nontrivial mask in B_0
1	χ_{-4}	$\{2\}$	none
2	χ_{-8}	$\{2\}$	none
3	χ_{-3}	$\{3\}$	none
4	χ_{-4}	$\{2\}$	$\alpha_2(e) = \mathbf{1}_{e\neq 1}$
7	χ_{-7}	$\{2,7\}$	$\alpha_2(e) = \mathbf{1}_{e\neq 1}$

The corresponding order discriminants are $-4,-8,-12,-16,-28$. In particular the third, fourth and fifth forms are conductor-two orders. For $d=3$, the ordinary value set nevertheless has exactly the inert-even rule for χ_{-3}, including the prime two; no extra dyadic mask is imposed.

Proof of Corollary 1.1. Write $Q_d = \{p : \chi(p) = -1\}$, except that no prime in B_0 is included; for the displayed five forms this exception changes none of the inert sets. These are respectively

$$\{p\equiv 3\ (4)\},\quad \{p\equiv 5,7\ (8)\},\quad \{p\equiv 2\ (3)\},\quad \{p\equiv 3\ (4)\},\quad \{p\equiv 3,5,6\ (7)\}.$$

Let $R_1 = R_2 = R_4 = \{2\}$, $R_3 = \{3\}$, and $R_7 = \{2,7\}$. The bad-prime integrand factors in these cases. For $d=7$, this uses $\chi_{-7}(2) = 1$, which removes the cross-valuation phase from its

ramified unit test. Put $a_p = v_p(|h|)$. The local densities are

$$W_{1,2}(h) = \begin{cases} 1/4, & a_2 = 0, \\ 1/2 - 3/2^{a_2+2}, & a_2 \geq 1, \end{cases} \qquad W_{2,2}(h) = \begin{cases} 1/4, & a_2 = 0, 1, \\ 1/2 - 3/2^{a_2+1}, & a_2 \geq 2, \end{cases}$$

$$W_{3,3}(h) = \tfrac{1}{2}(1 - 2/3^{a_3+1}), \qquad W_{7,7}(h) = \tfrac{1}{2}(1 - 4/7^{a_7+1}),$$

$$W_{4,2}(h) = \begin{cases} 1/8, & a_2 = 0, \\ 0, & a_2 = 1, \\ 5/16, & a_2 = 2, \\ 3/8 - 3/2^{a_2+2}, & a_2 \geq 3, \end{cases} \qquad W_{7,2}(h) = \begin{cases} 1/2, & a_2 = 0, 1, \\ 3/4, & a_2 \geq 2. \end{cases}$$

For completeness, the odd ramified unit predicate is a square modulo $p = 3$ or 7. At $a_p = 0$, there are $(p-3)/4$ nonzero square pairs modulo p, and exactly one zero neighborhood contributes $1/(2p)$. The result is $(p-1)/(4p)$. For $a_p \geq 1$, the recurrence $W(a_p) = (p-1)/(2p) + W(a_p - 1)/p$ gives the displayed formulas. At two the ramified unit predicates are respectively one modulo four and one or three modulo eight. Their initial pair masses are $1/4$, and, for the latter, $1/4$ also when $a_2 = 1$. The former has mass $1/8$ at $a_2 = 1$; thereafter its recurrence is $W(a) = 1/4 + W(a-1)/2$. The latter has mass $1/8$ at $a_2 = 2$ and obeys that recurrence for $a \geq 3$. Removing valuation one gives the $d = 4$ row: for $a \geq 2$ its mass is $1/4 + W_{1,2}(h/4)/4$. The $d = 7$ dyadic row counts only the two valuation masks, giving its stated masses.

Using primitive characters, the prefactors are

$$c_d = \frac{4L(1,\chi)}{\pi} \prod_{p \mid \mathrm{cond}(\chi)} (1 - p^{-1})^{-1}, \qquad (c_1, c_2, c_3, c_4, c_7) = (2, 2\sqrt{2}, 2/\sqrt{3}, 2, 2\sqrt{7}/3).$$

For $d = 7$ the extra bad prime two contributes ratio one in (79). Consequently

$$\mathscr{T}_{d,h} = c_d \prod_{p \in R_d} W_{d,p}(h) \prod_{\substack{p \in Q_d \\ p \mid h}} \frac{1 - p^{-a_p - 1}}{1 - p^{-1}},$$

as required. The primitive values are $L(1, \chi_{-3}) = \pi/(3\sqrt{3})$ and $L(1, \chi_{-7}) = \pi/\sqrt{7}$; Smilansky's series in his equation (6.3) omits the Euler factor at two. At $h = 1$ the five constants are

$$(\mathscr{T}_{1,1}, \mathscr{T}_{2,1}, \mathscr{T}_{3,1}, \mathscr{T}_{4,1}, \mathscr{T}_{7,1}) = \left(\frac{1}{2}, \frac{1}{\sqrt{2}}, \frac{1}{3\sqrt{3}}, \frac{1}{4}, \frac{1}{2\sqrt{7}} \right).$$

Only $d = 4$ has a zero coefficient for nonzero shifts: $v_2(h) = 1$. Its pair count is exactly zero because its values are zero or one modulo four. All other displayed coefficients are positive. □

11 Splitting semigroups and prime cosets

11.1 Subgroup sums and congruence admissibility

Throughout, $L_0 \geq 2$ is even, $G = (\mathbb{Z}/L_0\mathbb{Z})^\times$, and $H_1, H_2 \leq G$ are fixed. Define

$$\mathcal{P}_i = \{p \nmid L_0 : p \bmod L_0 \in H_i\}, \quad C_i = \langle \mathcal{P}_i \rangle, \quad \rho_i = \frac{|H_i|}{\varphi(L_0)}, \quad \kappa = 2 - \rho_1 - \rho_2.$$

An empty product is one. Fix $U \subset H_1$, $V \subset H_2$, and coprime positive L_0-supported integers A, B. Write

$$R^{A,B}_{U,V}(N; I) = \#\{Ax + By = N : x \in C_1,\ y \in C_2,\ x \bmod L_0 \in U,\ y \bmod L_0 \in V,\ Ax/N \in I\}.$$

Omit I when $I=[0,1]$. Put

$$c_i=\frac{1}{\Gamma(\rho_i)}\lim_{s\downarrow 1}(s-1)^{\rho_i}\prod_{p\in\mathcal{P}_i}(1-p^{-s})^{-1}, \tag{86}$$

$$\Gamma_N=\frac{L_0}{|H_1||H_2|}\#\{(u,v)\in U\times V: Au+Bv\equiv N \pmod{L_0}\}, \tag{87}$$

$$\mathcal{E}=\{p\nmid L_0: p\notin\mathcal{P}_1,\ p\notin\mathcal{P}_2\},$$

$$C_{\mathcal{E}}=\prod_{p\in\mathcal{E}}\left(1-\frac{1}{(p-1)^2}\right),\qquad J_{\mathcal{E}}(N)=\prod_{\substack{p\mid N\\ p\in\mathcal{E}}}\frac{p-1}{p-2}. \tag{88}$$

The infinite product in (88) converges absolutely and is positive, since $2\mid L_0$. The constants c_i are positive. For an explicit convergent expression, let $r_i(p)>1$ denote the order of pH_i in G/H_i at a forbidden good prime. Then

$$c_i=\frac{(\varphi(L_0)/L_0)^{\rho_i}}{\Gamma(\rho_i)}\prod_{\substack{\eta\in H_i^{\perp}\\ \eta\neq 1}}L(1,\eta)^{\rho_i}\prod_{\substack{p\nmid L_0\\ p\notin\mathcal{P}_i}}(1-p^{-r_i(p)})^{1/r_i(p)}. \tag{89}$$

Characters here are literal modulo L_0. Continuation along the positive real axis from $s>1$ fixes the powers; conjugate factors are paired. The Euler identity proving this formula is given in Section 11.2.

Theorem 11.1. *Under the analytic inputs of Section 3, with all these data fixed, there is an explicit $\delta>0$ depending only on L_0 such that, uniformly in N and every interval $I\subset[0,1]$,*

$$R_{U,V}^{A,B}(N;I)=c_1c_2\Gamma_N C_{\mathcal{E}}|I|\frac{NJ_{\mathcal{E}}(N)}{AB(\log N)^{\kappa}}+O\left(\frac{NJ_{\mathcal{E}}(N)}{AB(\log N)^{\kappa+\delta}}\right).$$

If $\Gamma_N=0$, the count is identically zero. For $\Gamma_N>0$ and $I=[0,1]$, the error is a relative $O((\log N)^{-\delta})$. More generally it is relative $o(1)$ when $|I|(\log N)^{\delta}\to\infty$.

The implicit constant and threshold may depend on all fixed data. An admissible numerical choice of δ is recorded in Section 8, together with the localization budgets. No numerical threshold, uniformity in a growing modulus, or power-short interval assertion is included. For noncoprime fixed coefficients, first divide by their gcd when it divides N. The even-modulus convention makes the prime two a bad prime throughout the analytic theorem. Passing from an odd modulus to its even multiple requires retaining this convention; freely admitting the prime two is a separate valuation-profile operation, not a change of notation.

Theorem 11.2. *Assume the analytic inputs of Section 3. Let $\mathcal{E}_{12}=\{p\nmid 6: p\not\equiv 1 \pmod{12}\}$, and set*

$$A_{12}=\frac{4\pi}{\Gamma(1/4)^2}\sqrt{\frac{\log(2+\sqrt{3})}{6}}\prod_{p\in\mathcal{E}_{12}}\left(1-\frac{2}{p(p-1)}\right),\qquad J_{12}(N)=\prod_{\substack{p\mid N\\ p\in\mathcal{E}_{12}}}\frac{p-1}{p-2}.$$

There is an explicit $\delta_{12}>0$ such that, for all sufficiently large $N\equiv 2 \pmod{12}$,

$$R_{12}(N)=A_{12}J_{12}(N)\frac{N}{(\log N)^{3/2}}\left(1+O((\log N)^{-\delta_{12}})\right).$$

All other residue classes have zero count.

Deduction from Theorem 11.1. Take $L_0 = 12$, $H_1 = H_2 = U = V = \{1\}$, $A = B = 1$. Then $\rho_i = 1/4$ and $\Gamma_N = 12$ on the indicated class. The four characters give the exact identity

$$F_{12}(s)^4 = \zeta(s)L(s,\chi_{-4})L(s,\chi_{-3})L(s,\chi_{12})\,(1-2^{-2s})(1-3^{-2s})\prod_{p\in\mathcal{E}_{12}}(1-p^{-2s})^2.$$

Here the three nonprincipal characters displayed are primitive. Use $L(1,\chi_{-4}) = \pi/4$, $L(1,\chi_{-3}) = \pi/(3\sqrt{3})$, and

$$L(1,\chi_{12}) = \int_0^1 \frac{1-x^4}{1+x^6}\,dx = \left[\frac{1}{2\sqrt{3}}\log\frac{x^2+\sqrt{3}x+1}{x^2-\sqrt{3}x+1}\right]_0^1 = \frac{\log(2+\sqrt{3})}{\sqrt{3}}.$$

Consequently

$$c_{12} = \frac{1}{\Gamma(1/4)}\left(\frac{\pi^2\log(2+\sqrt{3})}{54}\prod_{p\in\mathcal{E}_{12}}(1-p^{-2})^2\right)^{1/4}.$$

In $12c_{12}^2 C_{\mathcal{E}_{12}}$, combine the factors by $(1-p^{-2})(1-(p-1)^{-2}) = 1-2/(p(p-1))$. This is exactly A_{12}. The real density integrates to one; there is no extra beta-integral factor. □

The uniformity in the target also permits a coprimality restriction. Only the primes allowed in the summands can divide their gcd. Removing them changes the leading term by the following finite product.

Corollary 11.3. *Assume the analytic inputs of Section 3. Put*

$$R_{\mathrm{cop}}(N) = \#\{(a,b)\in C_{12}^2 : a+b=N,\ (a,b)=1\}, \qquad P_{12}(N) = \prod_{\substack{p\mid N\\ p\equiv 1\ (12)}}(1-p^{-1}).$$

Uniformly over sufficiently large $N \equiv 2 \pmod{12}$,

$$R_{\mathrm{cop}}(N) = A_{12}J_{12}(N)P_{12}(N)\frac{N}{(\log N)^{3/2}}\left(1+O((\log N)^{-\delta_{12}/2})\right).$$

Every other residue class has zero count. In particular, every sufficiently large admissible target has coprime summands.

Proof. Möbius inversion gives the exact identity

$$R_{\mathrm{cop}}(N) = \sum_{\substack{d\mid N\\ d\in C_{12}}}\mu(d)R_{12}(N/d). \tag{90}$$

Indeed, division by any common divisor preserves both prime supports, and every common divisor lies in C_{12}. Moreover $d \equiv 1 \pmod{12}$, so $N/d \equiv 2 \pmod{12}$, and $J_{12}(N/d) = J_{12}(N)$.

Write $L = \log N$, $\ell = \log\log(3N)$, $P = P_{12}(N)$, and $\delta = \delta_{12}$. We use the elementary divisor estimates

$$S_0 = \sum_{\substack{d\mid N\\ d\in C_{12}}}\frac{|\mu(d)|}{d} = \prod_{\substack{p\mid N\\ p\equiv 1\ (12)}}(1+p^{-1}) \ll \ell, \qquad P^{-1}\ll\ell, \tag{91}$$

$$S_1 = \sum_{\substack{d\mid N\\ d\in C_{12}}}\frac{|\mu(d)|\log d}{d} = S_0\sum_{\substack{p\mid N\\ p\equiv 1\ (12)}}\frac{\log p}{p+1} \ll \ell^3. \tag{92}$$

For (91), split the prime products at L. Mertens' product bound controls primes at most L by $O(\log L)$; the remaining factors have bounded product, since $\sum_{p\mid N,\ p>L} p^{-1} \le 1/\log L$. For (92), the primes at most L contribute at most $\sum_{n\le L}(\log n)/n \ll (\log L)^2$, and the others contribute at most $L^{-1}\sum_{p\mid N}\log p \le 1$.

Split (90) at $d = \sqrt{N}$. The bound $R_{12}(M) \le M$ shows that its large-divisor part is $O(\sqrt{N}\,\tau(N))$, where τ is the divisor function. For $d \le \sqrt{N}$, Theorem 11.2 applies uniformly, and

$$(L - \log d)^{-3/2} = L^{-3/2} + O(L^{-5/2}\log d).$$

Consequently, using $\sum_{d\mid N,\ d\in C_{12}} \mu(d)/d = P$ and extending the main sum back to all divisors,

$$\begin{aligned} R_{\mathrm{cop}}(N) = A_{12}J_{12}(N)\frac{N}{L^{3/2}}\Big(P + O(L^{-\delta}S_0 + L^{-1}S_1 + N^{-1/2}\tau(N))\Big) \\ + O(\sqrt{N}\,\tau(N)). \end{aligned}$$

After division by the main term, the error is

$$O\Big(L^{-\delta}\ell^2 + L^{-1}\ell^4 + N^{-1/2}\tau(N)L^{3/2}\ell\Big).$$

Here $J_{12}(N) \ge 1$ and $\tau(N) = N^{o(1)}$. The displayed error is $O(L^{-\delta/2})$, proving the claim. □

Remark 11.4 (Comparison with the quartic-norm target weight). The weight in [3, Theorem 2] is $W_{12}(N) = \prod_{p\mid N}(1 - \varepsilon_p/p)$, with $\varepsilon_p = 1$ on allowed primes and -1 otherwise. On the admissible class,

$$\frac{J_{12}(N)P_{12}(N)}{W_{12}(N)} = \frac{2}{3}\prod_{\substack{p\mid N\\ p\equiv 5,7,11 \pmod{12}}}\frac{p(p-1)}{(p-2)(p+1)} \asymp 1$$

uniformly in N, since the factors are $1 + O(p^{-2})$. The factor $2/3$ comes from $2 \mid N$; the prime three does not divide an admissible target. Thus the coprime asymptotic has the same order as the earlier conditional coprime bounds.

We next clarify the congruence criterion without analytic inputs. Write $C_H(U) = \{n \in C_H : n \bmod L_0 \in U\}$.

Lemma 11.5. *A residue $r \bmod M$ is represented by $C_H(U)$ if and only if its reduction modulo $\gcd(M, L_0)$ is that of some $h \in U$, and every forbidden prime $p \mid M$, $p \nmid L_0$, satisfies $p \nmid r$.*

Proof. Necessity follows from prime support. Choose such an h, and for each allowed $p^k \| M$ put $e_p = \min(v_p(r), k)$. Set $D = \prod p^{e_p}$, so $D \in C_H$ and $(D, L_0) = 1$. We seek a prime $\ell \equiv hD^{-1} \pmod{L_0}$. Where $e_p < k$, impose

$$\ell \equiv (r/p^{e_p})(D/p^{e_p})^{-1} \pmod{p^{k-e_p}}.$$

Where $e_p = k$, impose $\ell \equiv 1 \pmod p$. At primes of M not dividing D, impose the unit condition $\ell \equiv rD^{-1} \pmod{p^k}$. These congruences are compatible with the L_0-condition precisely by the stated hypotheses. They give a reduced CRT class. Dirichlet's theorem supplies a prime ℓ outside the finite exceptional set; then $D\ell$ has the required residue and state. □

Theorem 11.6. *The condition $\Gamma_N > 0$ is equivalent to solvability of $Ax + By \equiv N \pmod M$ with actual $x \in C_{H_1}(U), y \in C_{H_2}(V)$, for every $M \ge 1$. Under the analytic inputs of Section 3, for all sufficiently large N, it is also equivalent to positive additive representability.*

Proof. Taking $M = L_0$ proves necessity. For sufficiency choose a root counted by Γ_N and work modulo $K = \operatorname{lcm}(M, L_0)$. At a prime dividing L_0, at least one coefficient is a unit; adjust its variable to lift the equation while preserving the fixed unit states. At a prime outside L_0, both coefficients are units and $p \geq 3$. Of the p solutions modulo p, requiring a unit on each forbidden side removes at most two. A surviving pair lifts to every higher power. CRT and Lemma 11.5 realize the required two residues by actual semigroup values. Finally Theorem 11.1 is positive for all sufficiently large targets with $\Gamma_N > 0$. □

Remark 11.7 (A joint congruence obstruction). Take $L_0 = 12$, $H = \{1, 11\} = \ker \chi_{12}$. Then $H + H = \{0, 2, 10\} \pmod{12}$, so $N \equiv 4 \pmod{12}$ is never a sum. Nevertheless H projects onto the full unit groups modulo four and three. The separate closures are $\mathbb{Z}_2^\times$ and $\mathbb{Z}_3^\times$; at other primes they are $\mathbb{Z}_p$ or $\mathbb{Z}_p^\times$, according as that prime is allowed or forbidden. Two units sum to every element of $\mathbb{Z}_p$ for odd p, and to every even element for $p = 2$. Thus these targets pass every separate p-adic test. The missing condition is the simultaneous residue state modulo twelve.

11.2 Subgroup applications

Use the separate models of Section 9 for the two subgroup restrictions, with outer scale $X = N$. Thus $L = \log N$, $y = \exp(L^{2/5})$, and $E = \exp(\sqrt{L})$. We evaluate their correlation and verify the positive record envelope.

11.2.1 The model correlation

Use separate constructions for the two subgroups with common $\mathcal{P}_y$.

Proposition 11.8. *For a smooth $0 \leq \Phi \leq 1$ supported away from zero and one,*

$$\sum_{Ax+By=N} \Phi(Ax/N)\mathcal{D}_U(x)\mathcal{D}_V(y)$$
$$= c_1c_2\Gamma_N C_{\mathcal{E}} J_{\mathcal{E}}(N)\frac{N}{ABL^\kappa}\left(\int_0^1 \Phi(t)\,dt + O(L^{-1/2})\right) + O(N^{1/2}).$$

The statement is uniform in target valuations. Endpoint strips of relative length ϵ have upper bound $O(\epsilon N J_{\mathcal{E}}(N)L^{-\kappa} + N^{1/2})$. For every interval $I \subset [0,1]$, the unsmoothed model correlation has the same leading coefficient multiplied by $|I|$, with absolute error $O(NJ_{\mathcal{E}}(N)L^{-\kappa-\delta_0})$, where $\delta_0 = 1/(8192\varphi(L_0))$.

Proof. Discard records larger than E using (76). The rough logarithms on the rest are $L + O(\sqrt{L})$, giving

$$\frac{M_{R,K}(n_{>y})}{V_y} = \frac{c_H L^{\rho-1}}{\alpha_0 E_-(y)}(1 + O_K(L^{-1/2})).$$

Count the remaining small-prime sieve set in the fixed residues. The fundamental lemma applies at level N^{ϵ_0}, for small fixed $\epsilon_0 > 0$, with $u = \log(N^{\epsilon_0})/\log y \asymp L^{3/5}$. Its relative error is $e^{-cu\log u}$, and CRT gives endpoint error $O(\nu(d))$ at each squarefree sieving modulus d. Summing those errors costs $N^{\epsilon_0}L^{O(1)}$. This is the standard fixed-dimensional fundamental-lemma application [2, Theorem 11.7]; target-root mergers decrease its dimension and preserve uniformity.

The positive integer line is $(x, y) = (x_0 + Bt, y_0 - At)$, with parameter length $N/(AB)$. Modulo L_0, this parameter bijects onto all roots: at each prime at least one coefficient is a unit. If R roots have the prescribed states, their raw probability is R/L_0. Multiplying the normalized marginal projections gives

$$\frac{R}{L_0}\frac{L_0c_1}{|H_1|}\frac{L_0c_2}{|H_2|} = c_1c_2\Gamma_N.$$

A prime forbidden on only one side has relative factor one. A prime forbidden on both sides has two excluded roots, merging exactly when $p \mid N$, and relative factor

$$\frac{1-(2-\mathbf{1}_{p|N})/p}{(1-1/p)^2}.$$

Their product up to y is $C_{\mathcal{E}}(y)J_{\mathcal{E},y}(N)$. The ordinary tail is $O(1/y)$, and the target-prime tail is $O(L/(y\log y))$. Partial summation inserts Φ, costing its total variation. The same upper sieve on an endpoint strip gives its length times the sparse density, with a power-smaller endpoint error. To obtain the uniform interval assertion, put $\epsilon = L^{-\delta_0}$ and remove strips of this width near zero and one. On the rest the rough logarithms are $L + O(\sqrt{L} + \delta_0 \log L)$, giving the same relative $O(L^{-1/2})$ replacement. The fundamental-lemma count in any remaining interval has the same CRT endpoint error; partial summation costs the total variation, bounded independently of the interval. Restoring the strips costs $O(\epsilon)$ of the sparse scale. All other tails are superlogarithmic or power-smaller, proving the claim. □

11.2.2 Verification of the comparison hypotheses

The indicators lie between zero and one, and the model bound after (75) gives $\mathcal{D}_U \le 1$ eventually for each fixed depth. Thus (C1) holds, with a threshold allowed to depend on that depth. The full-unit-group exact model also lies between zero and one.

Choose banks of primes $p \equiv 1 \pmod{L_0}$, with small banks in $[\exp(L^{1/8}), \exp(L^{1/4})]$ and big banks in $[\exp(L^{1/3}), \exp(L^{2/5})]$. Delete primes dividing N. Every bank prime is a unit at A, B, belongs to both allowed sets, and preserves the fixed states. Hence

$$f_U(pn) = f_U(n), \quad \mathcal{D}_U(pn) = \mathcal{D}_U(n), \quad \xi_U(pn) = \xi_U(n) \tag{93}$$

exactly, including higher powers. The deletion changes reciprocal mass by at most $O(L/(P_{\min}\log P_{\min}))$, which is smaller than every fixed inverse logarithmic power. This verifies (C2).

For (C3), fix an $\mathcal{P}_y$-smooth record $r \le E$. If $f(r) = 0$, both the indicator and model on that record vanish. Otherwise $r \bmod L_0 \in H$, and insert the residual state U/r into the periodic test of Lemma 9.1. Its smooth scale is Y/r, so the error is $(Y/r)L^{-D}$, as required. An additional periodic test of period at most L^B increases the period by only a fixed factor. For every required accuracy, the Hankel remainder works for every sufficiently large fixed depth; constants and thresholds, rather than earlier graph exponents, may depend on that depth. This proves (C3) with all surviving character branches retained.

It remains to check the positive envelope (C4). Write $\mu_{\mathcal{B}}$ for the law (7).

Lemma 11.9. *For a nonnegative record function $F(r,s) \le L^B$, and either indicator or model on each side, its weighted correlation on the positive line restricted by $Ax \in J \subset [0,N]$ is at most*

$$C\frac{|J|J_{\mathcal{E}}(N)}{ABL^{\kappa}}\mathbb{E}_{\mu_{\mathcal{B}}}F + NL^{O(1)}e^{-cL^{1/10}} + O(N^{1/2}).$$

The constant may depend on the fixed coefficients.

Proof. Discard records $r, s > E$ using (76) and bounded weights. Their total contribution is $NL^{O(1)}e^{-cL^{1/10}}$. There are at most E^2 retained record pairs. A shared bank prime makes a stratum empty, since every bank prime avoids N. For coprime r, s, their divisibilities select one progression of length $|J|/(ABrs)$.

At $p^e \| r$, write $x = r\ell$. Exact valuation excludes $\ell \equiv 0 \pmod{p}$, while y is automatically a unit because $p \nmid N$. The same argument applies to s. At an unused bank prime there are two distinct excluded roots. Thus the complete bank factor, including the progression density, is

$$\frac{1}{rs}\prod_{\substack{p\in\mathcal{B}\\ p\nmid rs}}(1-2/p)\prod_{p\mid rs}(1-1/p) = \mu_{\mathcal{B}}(r,s).$$

All forbidden-prime slopes remain units. Apply the upper Selberg sieve with $z = N^{1/10}$ and primes up to $w = z^{1/16}$, which contains every bank for sufficiently large N. The nonbank sieve denominators are bounded below, up to fixed positive constants, by

weights	sieve denominator
f_1, f_2	$(\log w)^{\kappa}/J_{\mathcal{E}}(N)$,
$\mathcal{D}_1, f_2$	$(\log y)^{1-\rho_1}(\log w)^{1-\rho_2}/J_{\mathcal{E}}(N)$,
$\mathcal{D}_1, \mathcal{D}_2$	$(\log y)^{\kappa}/J_{\mathcal{E}}(N)$.

The middle row also applies with the sides interchanged. A model excludes its forbidden primes up to y; an indicator excludes them up to w. A prime forbidden on only one side enters only that side's denominator. Target-root mergers cost at most the displayed factor $J_{\mathcal{E}}(N)$. Each model bound contributes $L^{-3(1-\rho_i)/5}$, so all three rows give the same sparse factor $L^{-\kappa}J_{\mathcal{E}}(N)$.

Consequently every retained stratum, for $w_i \in \{f_i, \mathcal{D}_i\}$, has the single uniform estimate

$$\sum_{\substack{Ax+By=N,\ Ax\in J\\ x_{\mathcal{B}}=r,\ y_{\mathcal{B}}=s}} w_1(x)w_2(y) \le C_K \frac{|J|}{AB}\mu_{\mathcal{B}}(r,s)\frac{J_{\mathcal{E}}(N)}{L^{\kappa}} + O_K(z^{4+\epsilon}),$$

where $\epsilon > 0$ is arbitrarily small and fixed. The remainder follows from the squared-weight CRT root counts. Multiplying by $F(r,s) \le L^B$ and summing at most E^2 retained records leaves

$$L^B E^2 z^{4+\epsilon} = O(N^{1/2})$$

after fixing, for example, $\epsilon = 1/2$. Summing the main terms gives $\mathbb{E}_{\mu_{\mathcal{B}}}F$. The endpoint exponent and sieve level are independent of the final model depth; constants and thresholds may depend on it. □

Proof of Theorem 11.1. The preceding checks verify (C1)–(C4), with $\mathcal{S}_N = NJ_{\mathcal{E}}(N)/(ABL^{\kappa})$. This scale is between NL^{-C} and NL^{C} for a fixed C, by $J_{\mathcal{E}} \ge 1$ and Mertens' upper product bound. Theorem 2.1 replaces both indicators by their models with error $O(\mathcal{S}_N L^{-\delta_0})$, uniformly in the interval. Proposition 11.8 evaluates that model count with exactly the coefficient in the theorem. If $\Gamma_N = 0$, the residue restrictions give zero identically. Otherwise its finitely many positive values have a fixed positive minimum, yielding the relative assertion. □

11.3 Self-sums with prime support in a fixed coset

The comparison theorem also covers a prime coset which contains no identity class. The additional arithmetic issue is that its character twists can have negative or genuinely complex singular exponents. Those branches enter the model used for centering, even though they are smaller than its positive leading part.

Fix an even integer m, put $G = (\mathbb{Z}/m\mathbb{Z})^{\times}$, and let $K \le G$, $a \in G$. Write

$$\mathscr{A} = aK, \qquad H = \langle K, a\rangle, \qquad r = [G:K], \qquad d = \operatorname{ord}_{G/K}(aK), \qquad \rho = 1/r.$$

Define

$$\mathcal{P}_{\mathscr{A}} = \{p : p \nmid m,\ p \bmod m \in \mathscr{A}\}, \qquad C_{\mathscr{A}} = \langle \mathcal{P}_{\mathscr{A}}\rangle, \qquad f = \mathbf{1}_{C_{\mathscr{A}}}.$$

Thus $1 \in C_{\mathscr{A}}$. The generated residue group is H, whereas the prime density is ρ; these are

different data. Put

$$c_{\mathscr{A}} = \frac{1}{\Gamma(\rho)} \lim_{s\downarrow 1}(s-1)^{\rho} \prod_{p\in\mathcal{P}_{\mathscr{A}}} (1-p^{-s})^{-1}, \tag{94}$$

$$\Gamma_{\mathscr{A}}(N) = \frac{m}{|H|^2}\#\{(u,v)\in H^2 : u+v\equiv N \pmod m\},$$

$$\mathcal{E}_{\mathscr{A}} = \{p : p\nmid m,\ p \bmod m \notin \mathscr{A}\},$$

$$C_{\mathcal{E}_{\mathscr{A}}} = \prod_{p\in\mathcal{E}_{\mathscr{A}}} \left(1-\frac{1}{(p-1)^2}\right), \qquad J_{\mathscr{A}}(N) = \prod_{\substack{p\mid N\\ p\in\mathcal{E}_{\mathscr{A}}}} \frac{p-1}{p-2}.$$

The constant $c_{\mathscr{A}}$ is positive; an explicit expression is given below. For $d>1$, set

$$g_{\mathscr{A}} = \rho\big(1-\cos(2\pi/d)\big), \qquad \delta_{\mathscr{A}} = \min\left\{\frac{1}{8192\varphi(m)}, \frac{3g_{\mathscr{A}}}{10}\right\}.$$

For $d=1$, use $g_{\mathscr{A}}=+\infty$ and $\delta_{\mathscr{A}} = 1/(8192\varphi(m))$.

Theorem 11.10. *Under the analytic inputs of Section 3, uniformly in all sufficiently large targets with $\Gamma_{\mathscr{A}}(N)>0$,*

$$\#\{(x,y)\in C_{\mathscr{A}}^2 : x+y=N\} = c_{\mathscr{A}}^2\Gamma_{\mathscr{A}}(N)C_{\mathcal{E}_{\mathscr{A}}}J_{\mathscr{A}}(N)\frac{N}{(\log N)^{2-2\rho}}\left(1+O((\log N)^{-\delta_{\mathscr{A}}})\right).$$

If $\Gamma_{\mathscr{A}}(N)=0$, the count is identically zero. The condition $\Gamma_{\mathscr{A}}(N)>0$ is equivalent to solvability in every finite quotient. Consequently it characterizes eventual representability. The modulus and coset remain fixed, and the error is uniform in all target prime valuations.

The exponent incorporates both discarded banks and the gap between the leading and secondary character branches. In particular a fixed positive exponent works for every fixed coset. For $d\ge 2$, $1-\cos(2\pi/d) = 2\sin^2(\pi/d) \ge 8/d^2$, so one may always replace $\delta_{\mathscr{A}}$ by the simpler, weaker $1/(8192\varphi(m)^3)$. There is no uniformity in growing m. The new assertion concerns unfiltered self-sums. A filter on the whole numerical value need not be preserved by multiplication by a prime in a nonidentity coset.

11.3.1 Every fixed-character branch

For a character ψ, define $F_\psi(s) = \sum_n \psi(n)f(n)n^{-s}$, with all characters inflated to a common modulus containing m. Fourier expansion of the prime coset gives

$$F_\psi(s) = G_\psi(s)\prod_{\chi\in K^\perp} L(s,\psi\chi)^{\rho\overline{\chi(a)}}. \tag{95}$$

The correction is uniformly absolutely convergent for $\Re s > 1/2$. Indeed the coefficient of p^{-s} in the logarithm is

$$\rho\psi(p)\sum_{\chi\in K^\perp}\overline{\chi(a)}\chi(p) = \psi(p)\mathbf{1}_{p \bmod m\in\mathscr{A}};$$

the remaining logarithmic terms are bounded by a fixed multiple of $\sum_p p^{-2\Re s}$. Literal Euler factors include all conductor deletions. Analytic logarithms continued from the Euler half-plane specify all complex powers.

A zeta factor occurs precisely for a fixed character $\psi\in K^\perp$; its exponent is

$$\alpha_\psi = \rho\psi(a). \tag{96}$$

All other twists have no singular branch at one. The exponents ρ occur exactly for $\psi \in H^\perp$. The remaining exponents are ρ times nontrivial d-th roots of unity, each with the same multiplicity. In particular, $\rho - \Re\alpha_\psi \ge g_{\mathscr{A}}$ outside $H^\perp$.

For the untwisted series, an explicit positive correction is

$$\log G_{1,p}(s) = \sum_{j\ge 2} \frac{\mathbf{1}_{p \bmod m \in \mathscr{A}} - \mathbf{1}_{(p \bmod m)^j \in \mathscr{A}}}{jp^{js}} \qquad (p \nmid m).$$

Thus

$$c_{\mathscr{A}} = \frac{(\varphi(m)/m)^\rho G_1(1)}{\Gamma(\rho)} \prod_{\substack{\chi \in K^\perp \\ \chi \ne 1}} L(1,\chi)^{\rho\overline{\chi(a)}}. \tag{97}$$

Here $G_1(1)$ is an absolutely convergent positive Euler product. Conjugate characters are paired. Continuation along $s > 1$ makes the expression positive and unambiguous. The ordinary one-variable asymptotic, and its complex twisted versions, also follow from [12, Theorem 1.2 and (1.9)–(1.10)]. That fixed-function theorem is not used in place of the moving-cutoff estimate which follows.

11.3.2 The arithmetic model and its one-variable means

If $K = G$, then f is the indicator of integers coprime to m, and we use the exact model $\mathcal{D} = f$. Counting the finite residue classes on the line proves Theorem 11.10 with $O_m(1)$ error. In the model construction below assume $K \ne G$, so $\rho < 1$ and the model upper bound tends to zero.

Keep an outer scale X, with $L = \log X$, and let

$$y = e^{L^{2/5}}, \quad \mathcal{P}_y = \{p : p \le y\}, \quad V_y = \prod_{p\le y}(1 - 1/p), \quad E = e^{\sqrt{L}}.$$

For $n = Av$, where $A = n_{\mathcal{P}_y}$, the integer v is $\mathcal{P}_y$-rough. The outer scale and cutoff are held fixed under multiplication by bank primes. For $\psi \in K^\perp$, write

$$F_{\psi,R}(s) = F_\psi(s) \prod_{\substack{p\le y \\ p\in\mathcal{P}_{\mathscr{A}}}} (1 - \psi(p)p^{-s}) = (s-1)^{-\alpha_\psi} H_{\psi,R}(s).$$

Since $\psi(p) = \psi(a)$ on the allowed prime set, fixed-character Mertens and a disk of radius $c/\log y$ give

$$|H_{\psi,R}(1)| \asymp (\log y)^{-\Re\alpha_\psi}, \qquad \frac{|H^{(j)}_{\psi,R}(1)|}{j!} \ll_j |H_{\psi,R}(1)|(C\log y)^j. \tag{98}$$

For the derivative bound, the logarithm of each finite Euler product relative to its value at one is $O(|s-1|\sum_{p\le y}(\log p)/p) = O(1)$ on that disk; the fixed regular factor is analytic and nonzero there.

Choose a sufficiently large fixed depth J, and put

$$M_\psi(t) = \sum_{j=0}^{J} \frac{H^{(j)}_{\psi,R}(1)}{j!} \frac{(\log t)^{\alpha_\psi - j - 1}}{\Gamma(\alpha_\psi - j)}. \tag{99}$$

Write $M_0 = M_1$ for the positive principal branch. On any fixed power range, $M_0(t) > 0$, and

$$|M_\psi(t)| \ll L^{-3g_{\mathscr{A}}/5} M_0(t) \qquad (\psi \notin H^\perp). \tag{100}$$

In fact its individual exponent is $3(\rho - \Re\alpha_\psi)/5$: the factors L and $\log y$ in (98) and (99) give exactly this ratio. Conjugate branches are conjugate functions.

Every $\eta \in H^\perp$ is one on the support of f, so $F_{\psi\eta,R} = F_{\psi,R}$ and $M_{\psi\eta} = M_\psi$ exactly. Define

$$\mathcal{D}_J(n) = \frac{f(A)}{V_y} \sum_{\psi \in K^\perp} \overline{\psi(v)} M_\psi(v), \qquad \xi_J(n) = f(n) - \mathcal{D}_J(n). \tag{101}$$

Grouping characters modulo $H^\perp$ factors the bracket by $|G/H|\mathbf{1}_{v \bmod m \in H}$. Its leading term is positive and all other terms satisfy (100). Therefore $\mathcal{D}_J$ is real and nonnegative for large X, and

$$\mathcal{D}_{0,J}(n) = |G/H|\mathbf{1}_{v \bmod m \in H} \frac{f(A)M_0(v)}{V_y}, \qquad |\mathcal{D}_J - \mathcal{D}_{0,J}| \ll L^{-3g_{\mathscr{A}}/5}\mathcal{D}_{0,J}. \tag{102}$$

For $d = 1$, the difference is zero.

For a precise extension, use the cutoff from Section 9, immediately after (74), with the present outer scale X: set every $M_\psi(t) = 0$ for $t \le X^{1/4}$, and multiply (99) by $\eta_0(\log t/L)$ above that point. It equals the displayed formula for $t \ge X^{1/3}$. The common real cutoff preserves conjugacy, $M_{\psi\eta} = M_\psi$, positivity of the combined model, and (100) on the transition range. Condition (C3) only uses $t \gg X^{1/2}/E > X^{1/3}$. Consequently, on $1 \le n \le X^3$,

$$0 \le \mathcal{D}_J(n) \ll L^{-3(1-\rho)/5}, \qquad |\xi_J(n)| \ll 1.$$

These exponents are independent of the fixed depth. For an allowed $p \le y$, multiplication changes A to pA and leaves v unchanged. Hence, exactly,

$$f(pn) = f(n), \qquad \mathcal{D}_J(pn) = \mathcal{D}_J(n), \qquad \xi_J(pn) = \xi_J(n). \tag{103}$$

The characters in (101) are evaluated at the rough part v, rather than at the whole integer n.

Lemma 11.11. *For every fixed $B, A_0 > 0$, a sufficiently large fixed depth J gives the recordwise estimate (6) for ξ_J. It holds uniformly for every $\mathcal{P}_y$-supported record $r_0 \le E$, every $Y \in [X^{1/2}, X^2]$, bounded periodic tests of period at most L^B, and the smooth weights permitted there. In particular it includes every fixed-character contribution.*

Proof. First consider an $\mathcal{P}_y$-rough variable at scale $T \in [X^{2/5}, X^2]$. Expand a periodic test in characters after refining its period to $Q = \operatorname{lcm}(q, m)$. All primes dividing Q lie below y. For a character $\vartheta \bmod Q$, (95) holds literally, with every small allowed Euler factor removed. A principal constituent occurs only if ϑ is inflated from $K^\perp$. In that case the rough series is precisely $F_{\vartheta,R}$ defined above; it does not acquire a new conductor-dependent regular part.

Complex powers require a bound on the analytic logarithm, rather than merely on the modulus of an L-function. For a principal character put $L^\#(s) = (s-1)L(s)/(s+1)$, and for a nonprincipal character put $L^\# = L$. The zero-free input makes $L^\#$ analytic and nonzero on $\Re s > 11/12$. Its logarithm is fixed by the Euler logarithm on $\Re s = 2$, together with the displayed pole-removing factor. On a fixed disk centered at $2 + it$, contained in that half-plane, convexity gives $\Re \log L^\# \ll \log(Q(2 + |t|))$, while the center logarithm is bounded by the same quantity. Borel–Carathéodory on a smaller fixed disk therefore gives

$$|\log L^\#(\sigma + it)| \ll \log(Q(2 + |t|)) \qquad (23/24 \le \sigma \le 2). \tag{104}$$

The disks can have fixed radii because their left boundary remains strictly to the right of $11/12$. The consistent center logarithm is uniformly controlled by its Euler series. Thus every fixed complex power in (95) has polynomial growth in conductor and height; its imaginary exponent causes no uncontrolled argument factor.

Use smooth Mellin inversion to height $\exp(L^{1/10})$, and move the line to $1 - c/\log y$. This lies inside the preceding strip for large X. The finite Euler factors cost only $L^{O(1)}$, and (104) costs $\exp(O(L^{1/10}))$. The factor T^{s-1} supplies $\exp(-cL^{3/5})$, which dominates both. Sufficiently

many fixed derivatives of the smooth weight dispose of the Mellin tails. A twist without a principal constituent therefore has mean $O_A(TL^{-A})$ for every fixed A.

For $\vartheta \in K^\perp$, Hankel integration around one gives exactly the density M_ϑ in (99). The derivative bound (98) gives relative Taylor remainder $O_J(L^{-3(J+1)/5})$ for that branch. All branch magnitudes are at most a fixed multiple of the principal magnitude. This argument applies without change to negative and complex α_ϑ; the reciprocal Gamma factor is the Hankel integral of $(s-1)^{-\alpha_\vartheta+j}$.

On the model side the unrestricted rough population has series $\zeta(s)\prod_{p\le y}(1-p^{-s})$, principal density V_y, and no other principal branch. Its twist by $\vartheta\overline{\psi}$ is principal exactly when $\vartheta=\psi$ on the units modulo Q. Hence $V_y^{-1}\sum_{\psi\in K^\perp}\overline{\psi}(v)M_\psi(v)$ matches each of the original character means with its exact coefficient. This proves arbitrary logarithmic centering after choosing J beyond the periodic-character and smooth-derivative costs. A factor v^{iu}, $|u|\le L^B$, only enlarges those fixed costs.

Finally, for a fixed record $r_0 \le E$, multiplicativity gives $f(r_0v)=f(r_0)f(v)$, and (101) has the same factor $f(r_0)$. The remaining scale is $T=Y/r_0$, within the preceding power range for large X. If the record is forbidden both terms vanish. Otherwise the just-proved estimate is $O((Y/r_0)L^{-A_0})$, uniformly in the record and test. This is the required recordwise assertion, with the depth chosen after the requested accuracy. □

11.3.3 Common banks and the independent sieve check

Choose one residue $a_0 \in aK$ and take all small and big banks from primes $p \equiv a_0 \pmod m$, deleting $p \mid N$. Their fixed density is $1/\varphi(m)$, irrespective of ρ. Equation (103) verifies condition (C2) of Theorem 2.1; the bounds above verify (C1), and Lemma 11.11 verifies (C3).

Apply Lemma 4.1 with $L_0=m$ and $c=a_0$. Its quotient character frame already retains the common unit phase on products with the same slot type. Thus these banks require no new graph construction or mark-length loss.

We verify the remaining envelope assumption independently of the comparison. Put $\kappa_{\mathscr{A}}=2-2\rho$ and $\Sigma_N=NJ_{\mathscr{A}}(N)L^{-\kappa_{\mathscr{A}}}$. This satisfies (5), for example with $\kappa_0=3$, since $0\le\kappa_{\mathscr{A}}<2$, $J_{\mathscr{A}}(N)\ge 1$, and the prime-product bound $J_{\mathscr{A}}(N)\ll(\log\log(3N))^2$ is uniform in N. At a bank prime $p\nmid N$, the exact two-coordinate valuation law is (7). For coprime bank records $r_1,r_2\le E$, their divisibilities select one interval progression of length $|\mathcal{I}|/(r_1r_2)$ on $x+y=N$. Exactness excludes two roots at every unused bank prime, and one remaining root at a prime dividing r_1r_2. Together with $1/(r_1r_2)$, these factors are exactly its probability under the independent bank law. Shared bank-prime divisibility is impossible.

Use the upper Selberg sieve at $z=N^{1/10}$, with primes up to $w=z^{1/16}$. This contains all banks. Every good forbidden prime has two roots, merging precisely when it divides N. Fixed-coset prime Mertens gives the following lower bounds for the nonbank sieve denominator, up to fixed positive constants:

weights	sieve denominator
f,f	$(\log w)^{2(1-\rho)}/J_{\mathscr{A}}(N)$,
$\mathcal{D}_J,f$	$(\log y)^{1-\rho}(\log w)^{1-\rho}/J_{\mathscr{A}}(N)$,
$\mathcal{D}_J,\mathcal{D}_J$	$(\log y)^{2(1-\rho)}/J_{\mathscr{A}}(N)$.

For a model factor use its nonnegative domination by $2D_{0,J}$, its bound $L^{-3(1-\rho)/5}$, and its exclusion of every forbidden prime at most y. These bounds turn each row into exactly $L^{-\kappa_{\mathscr{A}}}J_{\mathscr{A}}(N)$. Fixed residue conditions can be discarded for this upper estimate. Its CRT endpoint error per retained record pair is $O(z^{4+\epsilon})$; summing at most E^2 pairs gives $N^{0.45+o(1)}$. The exponent is independent of the model depth.

Rankin's bound with parameter $c/\log y$ gives

$$\#\{n\le N : n_{\mathcal{P}_y}>E\}\ll NL^{O(1)}e^{-cL^{1/10}}.$$

It applies to every actual/model pair because the weights are bounded. Multiplying the record-stratum estimate by an arbitrary nonnegative record function bounded by L^B, and summing, proves (8), including its interval-length factor and its errors. This is condition (C4). Its fixed logarithmic error exponents do not depend on the eventual density depth.

All four hypotheses of Theorem 2.1 have now been checked by invariance, one-variable estimates, or a positive upper sieve. No additive cancellation or residue-energy conclusion has been included as an arithmetic assumption.

11.4 The main term and the exact congruence criterion

Lemma 11.12. *A residue $t \bmod M$ is represented by $C_{\mathscr{A}}$ if and only if its reduction modulo $\gcd(M,m)$ is that of some $h \in H$, and $p \nmid t$ at each forbidden prime $p \mid M$, $p \nmid m$.*

Proof. Necessity follows from prime support. For sufficiency extract an allowed-prime prefix D with valuations $\min(v_p(t), v_p(M))$ at allowed primes dividing M. It belongs to $C_{\mathscr{A}}$, and its residue belongs to H. The desired unit tail has state $hD^{-1} \in H$.

The positive monoid generated by $\mathscr{A}$ in the finite group is H, so choose a nonempty word $s_1 \cdots s_k = hD^{-1}$, with $s_i \in \mathscr{A}$. An identity state also has a nonempty word, by taking a positive power of any element of $\mathscr{A}$. Realize all but the last letter by primes in their reduced classes modulo m, avoiding M. For the last letter, impose its class modulo m, and the remaining unit-tail congruences modulo the prime powers of M. At an allowed prime with prefix valuation $e < v_p(M)$, this is the quotient of t/p^e by the other unit factors modulo $p^{v_p(M)-e}$. At a saturated prefix no further congruence is needed, beyond avoiding that prime. All other good primes require a unit tail by the hypothesis. At primes dividing m, compatibility is precisely the prescribed reduction of h. The resulting CRT class is reduced, and Dirichlet's theorem supplies a prime in it. Multiplying all chosen primes and D realizes the required residue. A single prime would not suffice in general: the tail state need not itself lie in $\mathscr{A}$. □

Proof of Theorem 11.10. The full-group case was handled by its exact model. Suppose $K \neq G$. Theorem 2.1, applied using the checks above, replaces both copies of f by $\mathcal{D}_J$, at the sparse scale $NJ_{\mathscr{A}}(N)L^{-\kappa_{\mathscr{A}}}$, with relative error $O(L^{-1/(8192\varphi(m))})$. Its interval version and the upper sieve handle the endpoint strips.

On the retained records and power bulk, (102) replaces $\mathcal{D}_J\mathcal{D}_J$ by $\mathcal{D}_{0,J}\mathcal{D}_{0,J}$ with relative error $O(L^{-3g_{\mathscr{A}}/5})$. Put

$$\alpha_0 = \varphi(m)/m, \qquad E_-(y) = \prod_{\substack{p \le y \\ p \nmid m,\ p \notin \mathcal{P}_{\mathscr{A}}}} (1 - p^{-1}).$$

Since $\log v = L + O(\sqrt{L} + \log L)$, the leading model is

$$|G/H|\mathbf{1}_{n \bmod m \in H}\, f(n_{\mathcal{P}_y})\lambda_X, \qquad \lambda_X = \frac{c_{\mathscr{A}} L^{\rho-1}}{\alpha_0 E_-(y)},$$

with relative error $O(L^{-1/2})$. Here $n_{\mathcal{P}_y} \bmod m \in H$, so the rough and whole-value state tests are equivalent for this leading model.

Let R_N be the number of pairs counted in $\Gamma_{\mathscr{A}}(N)$. The raw fixed-residue probability on the integer line is R_N/m. Its two projections and marginal normalization give

$$\frac{R_N}{m}\frac{|G/H|^2}{\alpha_0^2} = \frac{mR_N}{|H|^2} = \Gamma_{\mathscr{A}}(N).$$

This calculation uses H for the residue support and ρ for the logarithmic density; it never identifies ρ with $|H|/|G|$. At a good forbidden prime the normalized pair factor is

$$\frac{1 - (2 - \mathbf{1}_{p|N})/p}{(1 - 1/p)^2}.$$

The fundamental sieve lemma at a fixed power level, as in [2, Theorem 11.7], therefore evaluates the leading model as

$$c_{\mathscr{A}}^2 \Gamma_{\mathscr{A}}(N) C_{\mathcal{E}_{\mathscr{A}}} J_{\mathscr{A}}(N) \frac{N}{L^{\kappa_{\mathscr{A}}}}.$$

Its sieve error is smaller than every fixed logarithmic power; CRT endpoints are $N^{0.45+o(1)}$. The omitted ordinary Euler tail is $O(1/y)$, and its target-prime tail is $O(L/(y \log y))$, uniformly in all target valuations. The density replacement and secondary-branch errors are bounded by the asserted $L^{-\delta_{\mathscr{A}}}$. This proves the asymptotic.

For the local assertion, testing modulo m proves necessity of $\Gamma_{\mathscr{A}}(N) > 0$. Conversely fix such a root. At primes dividing m, lift the equation while preserving its unit states. At every other prime, both coefficients are units and $p \geq 3$. The equation modulo p has p pairs; requiring a unit on each forbidden coordinate removes at most two. A surviving pair lifts, and CRT combines the choices. Lemma 11.12 realizes each coordinate by an actual element of $C_{\mathscr{A}}$. This proves solvability in every finite quotient. Finally the positive asymptotic proves eventual global representability. □

11.4.1 Two cosets which test the branches

For $m = 4$, $K = \{1\}$, $a = 3$, the generated group is $H = G$, while $\rho = 1/2$. The exponents are $1/2, -1/2$. Let C_4^- contain the positive integers supported on primes $3 \pmod 4$. With $P_3 = \prod_{p \equiv 3(4)} (1 - p^{-2})^{-1}$, its marginal is $c_{\mathscr{A}}^2 = 2P_3/\pi^2$. Indeed

$$F(s)^2 = (1 - 2^{-s}) \frac{\zeta(s)}{L(s, \chi_{-4})} F(2s).$$

Here $\Gamma_{\mathscr{A}}(N) = 2$ on every even target. Since $\prod_{p \text{ odd}} (1 - p^{-2})^{-1} = \pi^2/8$,

$$2c_{\mathscr{A}}^2 C_{\mathcal{E}_{\mathscr{A}}} = \frac{1}{2} \prod_{p \equiv 1(4)} \left(1 - \frac{2}{p(p-1)}\right) =: A_-.$$

Thus Theorem 11.10 gives

$$\#\{x + y = N : x, y \in C_4^-\} = A_- \prod_{\substack{p \mid N \\ p \equiv 1(4)}} \frac{p-1}{p-2} \frac{N}{\log N} \left(1 + O((\log N)^{-1/16384})\right)$$

for every sufficiently large even N. Odd targets have zero count. The negative branch is needed for centering, although it does not change this leading constant.

For a genuinely complex example, take $m = 10$, $K = \{1\}$, $a = 3$, and let $C_{3,10}$ be supported on primes $3 \pmod{10}$. Choose the literal quartic character χ mod 10 with $\chi(3) = i$. Then $H = G$, $\rho = 1/4$, $d = 4$, and the exponents of $1, \chi, \chi^2, \overline{\chi}$ are

$$\frac{1}{4}, \quad \frac{i}{4}, \quad -\frac{1}{4}, \quad -\frac{i}{4}.$$

The rough model is explicitly real:

$$\mathcal{D}_J(n) = \frac{f(A)}{V_y} \left\{ M_1(v) + 2\Re\big(\overline{\chi(v)} M_\chi(v)\big) + \chi(v)^2 M_{\chi^2}(v) \right\}.$$

The complex terms have relative size $O(L^{-3/20})$, and the negative branch has relative size $O(L^{-3/10})$. They cannot be omitted from an arbitrarily accurate centered model. The leading constant is given by (94); equivalently,

$$c_{3,10} = \frac{(2/5)^{1/4} G_1(1)}{\Gamma(1/4)} L(1, \chi)^{-i/4} L(1, \chi^2)^{-1/4} L(1, \overline{\chi})^{i/4},$$

with the positive determination and convergent correction specified in (97). The forbidden good residues are $1, 7, 9 \pmod{10}$. For even N,

$$\Gamma_{\mathscr{A}}(N) = \begin{cases} 5/2, & 5 \mid N, \\ 15/8, & 5 \nmid N. \end{cases}$$

This follows by counting pairs of nonzero residues modulo five: there are four at a zero target and three otherwise. Consequently every even target is congruence-admissible, and

$$\#\{x + y = N : x, y \in C_{3,10}\} = c_{3,10}^2 \Gamma_{\mathscr{A}}(N) C_{\mathcal{E}_{\mathscr{A}}} J_{\mathscr{A}}(N) \frac{N}{(\log N)^{3/2}} \left(1 + O((\log N)^{-1/32768})\right).$$

This includes all sufficiently large even targets, uniformly in their prime valuations. The example tests complex character contributions without changing the comparison theorem or its inputs.

The common-bank condition is substantive. Two different prime prescriptions can be disjoint and supply no common allowed prime class. Theorem 11.10 makes no assertion about arbitrary pairs of prime cosets.

12 Primitive forms and invariance within a genus

Fix a negative fundamental discriminant D, put $K = \mathbb{Q}(\sqrt{D})$, and let

$$\mathcal{O}_i = \mathbb{Z} + f_i \mathcal{O}_K, \qquad \Delta_i = f_i^2 D, \qquad G_i = \mathrm{Pic}(\mathcal{O}_i) \quad (i = 1, 2).$$

Let Q_i be a primitive positive definite integral binary form of discriminant Δ_i. For any such form Q, define

$$P_Q = \{Q(x, y) : x, y \in \mathbb{Z},\ \gcd(x, y) = 1\},$$
$$\mathscr{V}_{p,Q} = Q(\mathbb{Z}_p^2 \setminus p\mathbb{Z}_p^2), \qquad \Lambda_Q = \{n \geq 1 : n \in \mathscr{V}_{p,Q} \text{ for every } p\}.$$

Only positive values are included. Coordinates may vanish, and each numerical value is counted once. Write

$$R_{T_1,T_2}(N) = \#\{(a, b) \in T_1 \times T_2 : a + b = N\}, \qquad R_{Q_1,Q_2} = R_{P_{Q_1},P_{Q_2}}.$$

Two forms are in the same genus when they are equivalent over every $\mathbb{Z}_p$ and over $\mathbb{R}$. These equivalences preserve primitive vectors and therefore the primitive local value sets.

Put $\chi = \chi_D$ and

$$S = \{p : p \mid 2\Delta_1\Delta_2\}, \qquad J_S(N) = \prod_{\substack{p \mid N \\ p \notin S,\ \chi(p) = -1}} \frac{p - 1}{p - 2}.$$

A target is locally admissible for Q_1, Q_2 if $N \in \mathscr{V}_{p,Q_1} + \mathscr{V}_{p,Q_2}$ for every prime p. For a finite abelian group G, put

$$\eta_*(G) = \begin{cases} +\infty, & 2G = 0, \\ \min\limits_{\substack{\ell \mid\mid 2G| \\ \ell \text{ prime}}} \dfrac{\ell - 1}{2\ell}, & 2G \neq 0. \end{cases}$$

The mixed local value densities and universal factor are

$$d_{p;1,2}(N) = \int_{\mathbb{Z}_p} \mathbf{1}_{\mathscr{V}_{p,Q_1}}(a) \mathbf{1}_{\mathscr{V}_{p,Q_2}}(N - a)\, da \quad (p \in S),$$

$$\mathfrak{b}_{D,S} = \frac{4L(1, \chi)}{\pi} \prod_{p \in S} \frac{1 - \chi(p)/p}{1 - 1/p} \prod_{\substack{p \notin S \\ \chi(p) = -1}} \left(1 - \frac{2}{p(p - 1)}\right). \tag{105}$$

Haar measure gives $\mathbb{Z}_p$ mass one.

Theorem 12.1 (Common-field primitive form values)**.** *For the fixed data above, the following assertions hold.*

1. Classical order bounds and genus comparison. *Uniformly on sufficiently large locally admissible targets,*

$$R_{Q_1,Q_2}(N) \asymp R_{\Lambda_{Q_1},\Lambda_{Q_2}}(N) \asymp \frac{NJ_S(N)}{\log N}.$$

If Q_i' belongs to the genus of Q_i, then, for every fixed $0 < \eta < \min_i \eta_*(G_i)$, uniformly on their common sufficiently large locally admissible targets,

$$R_{Q_1,Q_2}(N) = R_{Q_1',Q_2'}(N)\left(1 + O((\log N)^{-\eta})\right).$$

If both doubled class groups vanish, the counts agree for every target. Otherwise $\eta = 1/5$ is always permissible. These assertions use Appendix A and the class-error sieve, independently of the recent zero-free and weighted-dilation inputs.

2. Explicit leading term. *Under the analytic inputs of Section 3, uniformly through sufficiently large locally admissible targets,*

$$R_{Q_1,Q_2}(N) \sim \mathfrak{b}_{D,S}\left(\prod_{p\in S} d_{p;1,2}(N)\right) J_S(N)\frac{N}{\log N}. \tag{106}$$

Each bad-prime density is a computable rational periodic function. The leading coefficient depends only on the respective genera, with no class-number divisor.

Nonadmissible targets have zero count. Constants and thresholds may depend on the fixed forms and, in the comparison, on η. The two conductors may differ; the quadratic field remains fixed.

A saturated numerical value can occur in every class of its compatible genus; the classes do not partition the values. Taking $f_1 = f_2 = f$ gives the one-discriminant case, with $\Delta = f^2 D$ and $J_\Delta = J_S$. Write $\mathcal{F} = (Q_1, Q_2)$ for the fixed pair in the proof.

12.1 Proper ideals and the genus map

Fix either order temporarily, write it as $\mathcal{O}$, and put $G = \mathrm{Pic}(\mathcal{O})$. The set S may contain additional primes besides those dividing $2\,\mathrm{disc}\,\mathcal{O}$. We use proper invertible ideals of the order. If I is an invertible fractional ideal, write $I_p = i_p\mathcal{O}_p$ and let $n(I) > 0$ be its covolume norm. Local determinants give

$$v_p(n(I)) = v_p(\mathrm{N}(i_p)), \qquad n(IJ) = n(I)n(J), \qquad I\overline{I} = n(I)\mathcal{O}. \tag{107}$$

These statements hold at conductor primes as well. Call an integral ideal primitive if it is not contained in $p\mathcal{O}$ for any rational prime p. Finitely many prescribed locally principal lattices, equal to $\mathcal{O}_p$ elsewhere, glue to an invertible fractional ideal: place them between $m\mathcal{O}_p$ and $m^{-1}\mathcal{O}_p$, glue the residue submodules by the Chinese remainder theorem, and glue their inverses in the same way.

The following principal-genus statement is classical for arbitrary orders; compare [4, Appendix A, Proposition A.1]. We give the norm argument because retaining conductor-local ideal types is necessary here.

Lemma 12.2. *Put $U_p = \mathrm{N}(\mathcal{O}_p^\times)$. The homomorphism*

$$\Theta : G \longrightarrow \prod_{p\in S} \mathbb{Z}_p^\times/U_p, \qquad \Theta_p([I]) = n(I)/\mathrm{N}(i_p) \pmod{U_p}$$

is well defined and has kernel $2G$.

Proof. The quotient in Θ_p is a unit by (107). A change of local generator changes it by an order-unit norm. Multiplication of I by a principal ideal changes numerator and denominator by the same positive element norm. Every U_p contains the rational unit squares, so the target has exponent two and $2G$ is in the kernel. Outside S every unit is a norm from $\mathcal{O}_p^\times$.

Conversely, if $\Theta([I]) = 1$, choose units $u_p \in \mathcal{O}_p^\times$ such that $\mathrm{N}(i_p u_p) = n(I)$ at every finite prime. Positivity handles the infinite place. The cyclic Hasse norm theorem [6, Chapter VIII, Theorem 3.1] supplies $\alpha \in K^\times$ with $\mathrm{N}(\alpha) = n(I)$. Local Hilbert 90 writes $i_p u_p/\alpha = z_p/\overline{z}_p$. For a norm-one element $c \neq -1$ one may take $z = 1 + c$; for $c = -1$, take $z = \sqrt{D}$. At almost all primes a unit solution exists: in an unramified field normalize its valuation by a rational uniformizer, and in a split algebra use $z = (c_1, 1)$. The lattices $z_p\mathcal{O}_p$ therefore glue to an invertible fractional ideal J. They give $I/(\alpha) = J/\overline{J}$, hence $[I] = 2[J]$ by (107).

□

12.2 The target ideal class and its local shadow

For either form $Q(x, y) = ax^2 + bxy + cy^2$, use its own order, discriminant and class group, and take

$$\beta = (b + \sqrt{\Delta})/2, \qquad A = \mathbb{Z}a + \mathbb{Z}\beta \subset \mathcal{O}, \qquad c_Q = -[A] \in G.$$

Here $\mathcal{O} = \mathbb{Z}[\beta]$, $n(A) = a$, and $N(ax + \beta y)/a = Q(x, y)$. The ideal is proper invertible:

$$A\overline{A} = a(a, \beta, \overline{\beta}, c)\mathcal{O} = a\mathcal{O},$$

since $\beta + \overline{\beta} = b$ and $\gcd(a, b, c) = 1$. This construction allows a to share primes with the conductor. It also fixes the class sign in our convention. For the classical normalized ideal-form description compare [4, Lemma 10.2].

Lemma 12.3. *For every positive integer n,*

$$n \in P_Q \iff \text{some primitive integral invertible ideal } I \text{ has } n(I) = n,\ [I] = c_Q,$$
$$n \in \Lambda_Q \iff \text{some primitive integral invertible ideal } I \text{ has } n(I) = n,\ [I] \in c_Q + 2G.$$

Proof. An element $\alpha \in A$ gives $I = (\alpha)A^{-1}$, with $n(I) = N(\alpha)/a$ and class $-[A]$. It is integral because $(\alpha) \subset A$. Multiplication by the invertible ideal A shows $I \subset p\mathcal{O}$ if and only if $\alpha \in pA$. This is exactly the coordinate-primitivity condition. Conversely, if $[I] = -[A]$, then $IA = (\alpha) \subset A$, and the same identities give a primitive representation of $n(I)$.

Write $A_p = z_p\mathcal{O}_p$. The norm identities in (107) give the units

$$\lambda_p = N(z_p)/a \in \mathbb{Z}_p^\times, \qquad \mathscr{V}_{p,Q} = \lambda_p \mathscr{V}_p,$$

where $\mathscr{V}_p = \{N(w) : w \in \mathcal{O}_p \setminus p\mathcal{O}_p\}$ is the primitive norm-value set for $\mathcal{O}_p$. The change from an integral basis of A to the local generator is in $\mathrm{GL}_2(\mathbb{Z}_p)$. Moreover $\Theta_p(c_Q) = \lambda_p \pmod{U_p}$, with Θ and U_p from Lemma 12.2.

If $n \in \Lambda_Q$, choose primitive local elements $\alpha_p = z_p w_p \in A_p$ with $N(\alpha_p) = an$. The ideals $w_p\mathcal{O}_p$ are primitive integral invertible, and equal $\mathcal{O}_p$ outside finitely many primes. Glue them as in Lemma 12.2. Their norm valuations are $v_p(n)$, so the resulting ideal I has norm n, and

$$\Theta_p([I]) = n/N(w_p) = \lambda_p \pmod{U_p}.$$

The kernel of Θ is $2G$, proving the coset assertion. Conversely, for such an I, write $I_p = i_p\mathcal{O}_p$. Equality of the Θ-classes supplies a unit u_p with $N(u_p) = n/(\lambda_p N(i_p))$. Then $\alpha_p = z_p i_p u_p$ is primitive in A_p and has norm an. Outside S the order-unit norm group is all of $\mathbb{Z}_p^\times$, so the same conclusion holds there. This proves the local-shadow equivalence at conductor primes as well. □

The coset condition is existential. A conductor-divisible value need not determine a unique genus. For example, at discriminant -32, x^2+8y^2 and $3x^2+2xy+3y^2$ have primitive odd values $1 \pmod 8$ and $3 \pmod 8$, respectively, and hence different genera. Both nevertheless primitively represent 24, at $(4,1)$ and $(3,-1)$.

12.3 Orientation saturation and maximal subgroups

For either form Q, work in its own group G, with target class $c_Q = -[A]$, and write $h_{\mathcal{O}} = |G|$ and $B = 2G$. For every inverse pair $\{g,-g\} \subset G$ with $2g \neq 0$, let Ω_g be the rational split primes outside S whose two prime ideal classes are $g,-g$, and put

$$m_g = \operatorname{ord}(2g) - 1, \qquad k_g(n) = \#\{p \in \Omega_g : v_p(n) = 1\}.$$

These prime families are disjoint. Enlarging S deletes only finitely many primes. The fixed-discriminant prime-class estimate [1, Lemma 2, equation (5.2)] gives

$$\sum_{\substack{p \le y \\ p \in \Omega_g}} \frac{1}{p} = \frac{1}{h_{\mathcal{O}}} \log\log y + O_{\mathcal{O},S}(1). \tag{108}$$

The coefficient is $1/h_{\mathcal{O}}$ for a nonambiguous inverse pair. The source allows nonfundamental negative form discriminants. For a fixed order its prime-counting error is $O_{\mathcal{O}}(x\exp(-c\sqrt{\log x}))$; partial summation gives the bounded remainder in (108).

Call the pair $\{g,-g\}$ saturated for n if $k_g(n) \ge m_g$, and set

$$B_n = \langle 2g : \{g,-g\} \text{ is saturated for } n \rangle \le B.$$

If $n \in \Lambda_Q$ and $B_n = B$, then $n \in P_Q$. To see this, start with a primitive ideal from Lemma 12.3, whose class belongs to $c_Q + B$. At a good split prime, primitivity puts the entire prime power in one conjugate prime ideal. From each saturated family reserve m_g exponent-one factors, initially orienting each as g. This changes the ideal class only by an element of B, so its class still belongs to $c_Q + B$. Reversing j reserved orientations changes its class by $-2jg$, for $0 \le j \le m_g$; these changes fill $\langle 2g \rangle$. Families have disjoint prime support, so their choices are independent and their sums fill $B_n = B$. They therefore cancel its difference from c_Q. The resulting ideal is primitive, has class c_Q, and has the same numerical norm and every conductor-local ideal unchanged. Hence $n \in P_Q$ by Lemma 12.3.

Suppose $B \neq 0$, and let $\mathcal{H}$ be its finite set of maximal proper subgroups. For $M \in \mathcal{H}$ define

$$\ell_M = [B : M], \qquad \Omega_M = \bigcup_{\{g,-g\}:\, 2g \notin M} \Omega_g,$$

$$K_M = \sum_{\{g,-g\}:\, 2g \notin M} (m_g - 1), \qquad k_M(n) = \#\{p \in \Omega_M : v_p(n) = 1\}.$$

Each inverse pair is counted once. Since B is finite abelian, B/M is cyclic of prime order ℓ_M. If $n \in \Lambda_Q \setminus P_Q$, then B_n is proper and lies in some $M \in \mathcal{H}$. Every pair with $2g \notin M$ is consequently deficient, giving

$$\Lambda_Q \setminus P_Q \subset \bigcup_{M \in \mathcal{H}} \{n : k_M(n) \le K_M\}. \tag{109}$$

If $2G = 0$, Lemma 12.3 already gives $P_Q = \Lambda_Q$.

The doubling map $G \to B$ has equal fibres. Precisely $h_{\mathcal{O}}(\ell_M - 1)/\ell_M$ classes satisfy $2g \notin M$. None is ambiguous, since $2g \neq 0$, so they form $h_{\mathcal{O}}(\ell_M - 1)/(2\ell_M)$ inverse pairs. Summing (108) over these disjoint families proves

$$\sum_{\substack{p \le y \\ p \in \Omega_M}} \frac{1}{p} = d_M \log\log y + O_{\mathcal{O},S}(1), \qquad d_M = \frac{\ell_M - 1}{2\ell_M} \ge \frac{1}{4}. \tag{110}$$

Every prime divisor of $|B|$ occurs as the index of a maximal subgroup: in the corresponding primary component, reduce a cyclic factor modulo that prime. Thus $\min_{M\in\mathcal{H}} d_M = \eta_*(G)$. The proper-subgroup mechanism is classical in the exceptional-value argument of Bourgain–Fuchs [1, Theorem 1, alternative (4.2), and Section 2, part (5)]. Here its use in (109) makes an entire positive-density collection deficient on the binary line. For $G \simeq \mathbb{Z}/8\mathbb{Z}$ or $\mathbb{Z}/2\mathbb{Z}\times\mathbb{Z}/4\mathbb{Z}$ the threshold is $1/4$, while for $G \simeq (\mathbb{Z}/3\mathbb{Z})^2$ it is $1/3$.

Proposition 12.4 (The class error)**.** *In the common-field setting above, for every fixed $0 < \eta < \min_i \eta_*(G_i)$, uniformly for $N \geq 3$,*

$$0 \leq R_{\Lambda_{Q_1},\Lambda_{Q_2}}(N) - R_{Q_1,Q_2}(N) \ll_{\mathcal{F},\eta} \frac{NJ_S(N)}{(\log N)^{1+\eta}}. \tag{111}$$

If $2G_i = 0$, then $P_{Q_i} = \Lambda_{Q_i}$, and that summand contributes no class error. If both doubled groups vanish the displayed difference is zero. In every case the exponent $\eta = 1/5$ is permissible. Only this saving exponent is independent of the orders; the implied constant remains dependent on the fixed forms. The proposition is classical and uses neither recent analytic input.

12.4 A sieve for the class error

We prove Proposition 12.4 directly on the line $a + (N - a) = N$. First charge a bad pair to its first summand $a \in \Lambda_{Q_1} \setminus P_{Q_1}$, and use the maximal-subgroup families in G_1. If $2G_1 = 0$, there is no such error. Otherwise fix $M \in \mathcal{H}$ in this group. At every good inert prime p, both values must be units, so the forbidden residues of a modulo p are $0, N$, with $\nu_N(p) = 2 - \mathbf{1}_{p|N}$ roots. At a marked prime $p \in \Omega_M$, forbid exactly $v_p(a) = 1$: this removes the $p - 1$ residues $p, 2p, \ldots, (p-1)p$ modulo p^2. Its forbidden density is

$$f_*(p) = \frac{p-1}{p^2}.$$

Higher prime powers remain permitted, including when $p \mid N$.

Here are the needed uniform sieve details. Given any mark set $\mathcal{M} \subset \Omega_M \cap [1, y]$, use modulus p at inert primes and p^2 at marked primes. All sieve indices below are squarefree and supported on the chosen inert and marked primes at most y. For such an index d, let $\ell(d)$ be the product of these moduli, $\nu(d)$ the product of the root counts, and $f(d) = \nu(d)/\ell(d)$. Exact interval counting by the Chinese remainder theorem gives $Xf(d) + O(\nu(d))$ on an interval of length X. In all cases $\nu(d) \leq d$ and $0 < f(p) \leq 2/p < 1$.

For $d \leq z$ the Selberg quadratic majorant has main term $X/G(z)$, where

$$w(d) = \prod_{p|d} \frac{f(p)}{1 - f(p)}, \qquad G(z) = \sum_{d\leq z} \mu^2(d) w(d).$$

To see the normalization, set $x_d = \lambda_d f(d)$ and $y_u = \sum_{u|d} x_d$. Then

$$\sum_{d,e\leq z} \lambda_d\lambda_e f([d,e]) = \sum_{u\leq z} \frac{y_u^2}{w(u)}, \qquad \lambda_1 = \sum_{u\leq z} \mu(u) y_u.$$

The minimum with $\lambda_1 = 1$ is $1/G(z)$, attained at $y_u = \mu(u)w(u)/G(z)$. Finite inversion gives

$$\lambda_d = \mu(d) \prod_{p|d} (1 - f(p))^{-1} \frac{G_d(z/d)}{G(z)}, \qquad G_d(v) = \sum_{\substack{e\leq v \\ (d,e)=1}} \mu^2(e) w(e).$$

Uniformly in the mark set and target, $|\lambda_d| \ll_\epsilon d^\epsilon$. Since $\nu([d,e]) \leq de$, the total interval error is $O_\epsilon(z^{4+\epsilon})$.

Take $y = z^{1/16}$. The unrestricted Euler mass for the same prime indices is

$$E(y) = \prod_{\substack{p \le y \\ p \notin S,\, \chi(p) = -1}} (1 - \nu_N(p)/p)^{-1} \prod_{p \in \mathcal{M}} (1 - f_*(p))^{-1}.$$

In its normalized positive expansion, a prime is included with probability $f(p)$. Therefore $\mathbb{E} \log d \le 2 \sum_{p \le y} (\log p)/p \le 4 \log y$ for sufficiently large y. Markov's inequality retains at least three quarters of the mass on $d \le z$, giving $G(z) \ge 3E(y)/4$. Fixed quadratic-character prime number theory gives

$$\prod_{\substack{p \le y \\ p \notin S,\, \chi(p) = -1}} (1 - \nu_N(p)/p)^{-1} = \frac{\displaystyle\prod_{p \le y,\, p \notin S,\, \chi(p) = -1} (1 - 2/p)^{-1}}{\displaystyle\prod_{p \le y,\, p \mid N,\, p \notin S,\, \chi(p) = -1} (p-1)/(p-2)} \gg_{\mathcal{F}} \frac{\log y}{J_S(N)}.$$

Consequently, the number surviving the marked sieve is at most

$$\ll_{\mathcal{F}} \frac{X J_S(N)}{\log y} \prod_{p \in \mathcal{M}} (1 - f_*(p)) + O_\epsilon(z^{4+\epsilon}). \tag{112}$$

This also holds without marked primes. Only root counts and exact interval errors have been used, so the estimate is uniform in target valuations and root positions.

Fix $M \in \mathcal{H}$. Independently mark each prime of $\Omega_M \cap [1, y]$ with probability $t \in (0, 1)$. An integer with $k_M(a) \le K_M$ survives with probability at least $(1-t)^{K_M}$: its exponent-one factors below y are at most its total such factors. This lower bound is a positive constant depending on the fixed order and t. On taking the finite expectation in (112), the marked factor becomes

$$\prod_{\substack{p \le y \\ p \in \Omega_M}} (1 - t f_*(p)) \ll_{\mathcal{F}, t} (\log y)^{-t d_M},$$

by (110), since $f_*(p) = p^{-1} + O(p^{-2})$. Take $z = N^{1/10}$, $y = z^{1/16}$ and $\epsilon = 1/2$. The number of pairs in $\Lambda_{Q_1} \times \Lambda_{Q_2}$ whose first summand has $k_M(a) \le K_M$ is at most

$$\ll_{\mathcal{F}, t} \frac{N J_S(N)}{(\log N)^{1 + t d_M}} + N^{9/20}.$$

Given $\eta < \eta_*(G_1)$, choose $t \in (0, 1)$ with $\eta < t\eta_*(G_1)$. Sum over this group's finitely many maximal subgroups. Repeat the argument for the second summand, using only its own group G_2 and prime families. Their Mertens densities are unchanged by deleting the finitely many extra primes in S. The common good inert sieve is the same in both arguments, while only one summand is marked at a time. Overlap between the two orders' split-prime families is therefore irrelevant. Taking η below both thresholds and adding the two bounds proves (111); the power error is absorbed since $J_S(N) \ge 1$. Increasing the constant covers all $N \ge 3$. For the uniform choice $\eta = 1/5$, one can take $t = 9/10$, because $t d_M \ge 9/40 > 1/5$. This argument uses only the existing fixed-order prime-class estimate, CRT and the finite upper sieve, independently of the recent zero-free and graph theorems. The use of $t < 1$ explains the strict range $\eta < \min_i \eta_*(G_i)$; no endpoint or optimality assertion is made.

12.5 Two local filters on a common split core

The remaining calculation concerns the two local shadows. It is shared by all pairs above and by the principal specialization below.

First the local sets have a finite description, even when a split conductor prime permits unbounded valuations. Put

$$e_p = 1 + 2\max_i v_p(\Delta_i) \qquad (p \in S).$$

At a primitive vector v, the identity

$$\operatorname{adj}(\operatorname{Hess} Q_i)\nabla Q_i(v) = -\Delta_i v$$

shows that some partial derivative has valuation $s \le v_p(\Delta_i)$. For any target congruent to $Q_i(v)$ modulo p^{e_p}, strong one-variable Hensel lifting in that coordinate applies because $e_p > 2s$. The lifted vector stays primitive. Thus each $\mathscr{V}_{p,Q_i}$ is a union of full cosets modulo p^{e_p}, found by enumerating primitive vectors in that quotient, and

$$d_{p;1,2}(N) = \frac{1}{p^{e_p}} \# \left\{ a \bmod p^{e_p} : \begin{array}{l} a \in \mathscr{V}_{p,Q_1} \bmod p^{e_p}, \\ N - a \in \mathscr{V}_{p,Q_2} \bmod p^{e_p} \end{array} \right\}. \tag{113}$$

In particular these densities are rational periodic functions. Outside S, the primitive value set of either form is $\mathbb{Z}_p$ at split primes and $\mathbb{Z}_p^\times$ at inert primes. These primes are odd, so every target is a sum of two permitted local values. Local admissibility is therefore the finite criterion $\prod_{p\in S} d_{p;1,2}(N) > 0$, and this product has a fixed positive minimum on admissible targets.

The local genus biases must still be retained. Set

$$L = 8 \prod_{\substack{p\in S \\ p \text{ odd}}} p, \qquad \alpha = \varphi(L)/L, \quad H = \ker(\chi \bmod L), \quad \mathcal{C} = \langle p : p \notin S,\ \chi(p) = 1 \rangle.$$

For each S-supported r, put

$$A_{r,i} = \{u \in (\mathbb{Z}/L\mathbb{Z})^\times : ru \in \mathscr{V}_{p,Q_i} \text{ for all } p \in S\}.$$

These sets are well defined and are unions of unit-square states: equal units modulo an odd p, or modulo eight at two, differ by a local unit square; multiplication of a primitive vector by a rational unit scales its value by that square.

Crucially $A_{r,i} \subset H$, although individual local genus characters can be nonprincipal. To verify this, use an ideal A_i as above, with norm a_i and local multipliers $\lambda_{i,p}$. At every prime outside S, an ideal norm has even valuation at an inert prime and has no field-norm obstruction at a split prime. Positivity and Hilbert reciprocity therefore give

$$\prod_{p\in S} (\lambda_{i,p}, D)_p = \prod_{p\in S} (a_i, D)_p = 1.$$

Here the numerator in $\lambda_{i,p} = N(z_{i,p})/a_i$ is a field norm. If $u \in A_{r,i}$, then $(ru, D)_p = (\lambda_{i,p}, D)_p$ at every bad prime. The positive S-supported number r has product of bad symbols one. For an integer unit representative u, reciprocity identifies $\prod_{p\in S}(u, D)_p$ with $\chi(u)$. Multiplication of the displayed local relations gives $\chi(u) = 1$. This does not identify $A_{r,i}$ with the principal-form filter.

Consequently the exact disjoint decomposition is

$$\Lambda_{Q_i} = \coprod_{r\in\langle S\rangle} r\{x \in \mathcal{C} : x \bmod L \in A_{r,i}\}. \tag{114}$$

Let $c_{\mathcal{C}}$ denote the Gamma-normalized marginal constant of this core. Its Euler series satisfies

$$\left(\sum_{n\in\mathcal{C}} n^{-s}\right)^2 = \zeta(s)L(s,\chi) \prod_{p\in S} (1 - p^{-s})(1 - \chi(p)p^{-s}) \prod_{\substack{p\notin S \\ \chi(p)=-1}} (1 - p^{-2s}),$$

so

$$c_{\mathcal{C}}^2 = \frac{L(1,\chi)}{\pi} \prod_{p\in S}(1-p^{-1})(1-\chi(p)p^{-1}) \prod_{\substack{p\notin S\\ \chi(p)=-1}} (1-p^{-2}). \tag{115}$$

Also put

$$C_{\mathrm{in}} = \prod_{\substack{p\notin S\\ \chi(p)=-1}} \left(1-\frac{1}{(p-1)^2}\right).$$

12.6 Finite profiles, the classical bounds, and genus comparison

For fixed prefixes $r, s \in \langle S\rangle$, put $g = (r, s)$. When $g \mid N$, write $A = r/g$, $B = s/g$, $M = N/g$, and set

$$\Gamma_{r,s;1,2}(N) = \frac{L}{|H|^2}\#\{(u,v) \in A_{r,1} \times A_{s,2} : Au + Bv = M \bmod L\};$$

otherwise set it to zero. The original two filters are retained after dividing by g.

A finite collection of profiles covers every locally admissible target. Put $Z = \prod_{p\in S} p^{e_p+1}$. Given local summands $a + b = N$, replace them by $a + p^{e_p}t, b - p^{e_p}t$. They remain in their full value cosets. A valuation below e_p stays fixed; otherwise at most one residue t mod p^2 makes that summand divisible by p^{e_p+2}. Since $p^2 > 2$, both forbidden residues can be avoided. The summands can therefore be chosen with valuations at most $e_p + 1$. Their valuations determine $r, s \mid Z$, and CRT assembles their unit quotients into a root counted by $\Gamma_{r,s;1,2}(N)$. Conversely, any such root lifts to an exact linear solution at each bad prime: one of A, B is a unit, and fixing the other unit residue preserves the two unit-square states. The local definitions of $A_{r,1}, A_{s,2}$ therefore remain valid. Outside S there is already no additive obstruction. Hence

$$N \text{ locally admissible} \quad\Longleftrightarrow\quad \Gamma_{r,s;1,2}(N) > 0 \text{ for some } r, s \mid Z.$$

Proof of Theorem 12.1, part 1. Lemma A.1 allows two distinct filters $U, V_0 \subset H$. In a positive fixed profile it supplies $\gg_{\mathcal{F}} MJ_S(M)/\log M$ core pairs on $Ax + By = M$. Multiplication by r, s is injective, and $J_S(M) = J_S(N)$ because g is supported on S. Only finitely many profiles and root patterns occur above, so the minimum of the positive lower constants and maximum of their thresholds give

$$R_{\Lambda_{Q_1},\Lambda_{Q_2}}(N) \gg_{\mathcal{F}} NJ_S(N)/\log N$$

uniformly on sufficiently large admissible targets. Sieving the original line only at good inert primes gives the matching upper bound by the unmarked case of (112). Proposition 12.4 transfers the lower bound to R_{Q_1,Q_2}.

If Q_i' is in the genus of Q_i for each i, their local value sets coincide exactly, as do the two mixed local shadows. Subtract their two class-error estimates and divide by the classical positive lower bound. This proves part 1 of Theorem 12.1. If both doubled groups vanish the individual global and local value sets are equal, proving the exact assertion for every target. Nonadmissible targets have zero count. This argument uses the older quadratic filter estimate, not Theorem 11.1. □

12.7 Prefix summation and Haar normalization

Under the inputs of Section 3, apply Theorem 11.1 with the coprime fixed L-supported coefficients A, B in each profile. Write $R_{r,s;1,2}(N)$ for its contribution in (114). Since $J_S(N/g) = J_S(N)$, its contribution divided by $NJ_S(N)/\log N$ is

$$c_{\mathcal{C}}^2 C_{\mathrm{in}} \Gamma_{r,s;1,2}(N)\frac{g}{rs} + o(1). \tag{116}$$

The original filters $A_{r,1}, A_{s,2}$ remain in this formula.

To sum the fixed-profile asymptotics, parameterize $rx + sy = N$. Its positive interval has length at most $Ng/(rs)$. At any good inert prime the two coefficients in this parameterization are units, and its two forbidden roots merge exactly when the prime divides N. Discarding both residue filters and applying the unmarked sieve (112) with $z = N^{1/10}$ gives

$$R_{r,s;1,2}(N) \ll_{\mathcal{F}} \frac{NgJ_S(N)}{rs \log N} + O(N^{9/20}), \tag{117}$$

uniformly in r, s. The interval error also covers intervals shorter than the sieve level. Only $O_{\mathcal{F}}((1 + \log N)^{2|S|})$ prefix pairs can occur, while

$$\sum_{r,s \in \langle S \rangle} \frac{(r,s)}{rs} = \prod_{p \in S} \sum_{i,j \geq 0} p^{-\max(i,j)} = \prod_{p \in S} \frac{1 + p^{-1}}{(1 - p^{-1})^2} < \infty.$$

The total interval error is negligible relative to $NJ_S(N)/\log N$; the remaining tail is uniformly small as a fixed prefix bound tends to infinity. This justifies a finite-prefix limit followed by a summable-tail limit, without assuming a binary asymptotic uniform in growing coefficients.

Partition the local Haar line $a + b = N$ by the exact valuation prefixes r, s. At a prime with $v_p(r) = i$, $v_p(s) = j$, and $v_p(L) = k$, each compatible unit-pair root modulo p^k has measure $p^{-\max(i,j)-k}$. For $i \leq j$, prescribing the second summand modulo p^{j+k} fixes the first summand's required unit residue; the case $j \leq i$ is symmetric. Thus each global root has measure $g/(Lrs)$, and the full stratum has measure

$$\frac{g}{Lrs} \#\{(u,v) \in A_{r,1} \times A_{s,2} : Au + Bv = M \bmod L\} = \frac{\alpha^2}{4} \Gamma_{r,s;1,2}(N) \frac{g}{rs},$$

because $|H| = \varphi(L)/2$. Zero values have Haar measure zero. Summing the disjoint strata gives the single normalization identity

$$\sum_{r,s} \Gamma_{r,s;1,2}(N) \frac{g}{rs} = \frac{4}{\alpha^2} \prod_{p \in S} d_{p;1,2}(N). \tag{118}$$

Combining this identity with (116) and the summable tail proves

$$R_{\Lambda_{Q_1}, \Lambda_{Q_2}}(N) \sim \frac{4c_{\mathcal{C}}^2}{\alpha^2} C_{\text{in}} \left(\prod_{p \in S} d_{p;1,2}(N) \right) J_S(N) \frac{N}{\log N}.$$

The coefficient is explicit. In (115), divide the bad-prime factors by $\alpha^2 = \prod_{p \in S} (1 - p^{-1})^2$, and use

$$(1 - p^{-2}) \left(1 - \frac{1}{(p-1)^2} \right) = 1 - \frac{2}{p(p-1)}$$

at every good inert prime. The result is precisely $\mathfrak{b}_{D,S}$ in (105). The finite positive minimum of the local density product and (111) transfer the asymptotic to the actual values. This proves part 2 of Theorem 12.1.

12.8 Principal norms

Take $Q(x,y) = N_{K/\mathbb{Q}}(x + y\omega)$ in a basis $1, \omega$ of a fixed order $\mathcal{O}$ of discriminant $\Delta = f^2 D$. Set $P = P_Q$, $\Lambda = \Lambda_Q$, $R_T = R_{T,T}$, $S = \{p : p \mid 2\Delta\}$, $\chi = \chi_D$, and $J_{\mathcal{O}} = J_S$. Theorem 12.1 with both slots equal to Q gives the principal order bounds through part 1 and, under the analytic inputs, the leading asymptotic through part 2. Its coefficient is

$$\mathfrak{A}_{\mathcal{O}}(N) = \frac{4L(1,\chi)}{\pi} \prod_{p \in S} \frac{1 - \chi(p)/p}{1 - 1/p} \prod_{\substack{p \notin S \\ \chi(p) = -1}} \left(1 - \frac{2}{p(p-1)} \right) \prod_{p \in S} d_p(N). \tag{119}$$

Here $d_p = d_{p;1,2}$ for $Q_1 = Q_2 = Q$. The norm form has $A = \mathcal{O}$, hence $c_Q = 0$, so Lemma 12.3 becomes

$$n \in \Lambda \quad \Longleftrightarrow \quad \text{some primitive integral invertible } I \text{ has } n(I) = n,\ [I] \in 2\,\mathrm{Pic}(\mathcal{O}); \tag{120}$$

the condition $[I] = 0$ characterizes P. In particular, Proposition 12.4 gives $R_\Lambda - R_P \ll_{\mathcal{O}} NJ_{\mathcal{O}}/(\log N)^{6/5}$, and $P = \Lambda$ when the doubled class group vanishes.

Remark 12.5 (Gaussian normalization). For $\mathcal{O} = \mathbb{Z}[i]$ one has $P = C_4 \sqcup 2C_4$ and

$$d_2(N) = \begin{cases} 1/4, & N \bmod 8 \in \{2,3,6,7\}, \\ 1/8, & N \bmod 8 = 4, \\ 0, & N \bmod 8 \in \{0,1,5\}. \end{cases}$$

Consequently the coefficient of $J_{\mathcal{O}}(N)N/\log N$ is $\frac{1}{2}\prod_{p\equiv 3 \pmod 4}(1 - 2/(p(p-1)))$ in the first four classes and half that coefficient in the class $4 \bmod 8$. This also calibrates the value-density normalization: the two Gamma factors are already included in $c_{\mathcal{C}}^2$, and no additional beta-integral factor occurs.

13 Consequences and boundaries

Corollary 13.1. *Under the analytic inputs of Section 3, the set $C_{12} \cup \{0\}$ is an asymptotic basis of exact order thirteen: every sufficiently large integer is a sum of at most thirteen elements of C_{12}, and twelve do not suffice.*

Proof. For a large target choose $j \in \{0, \dots, 11\}$ with $N - j \equiv 2 \pmod{12}$. Theorem 11.2 writes $N - j$ as two core values; add j copies of one. This uses at most thirteen positive summands. Conversely, $25 + 60r$, $r \geq 0$, is $1 \pmod{12}$ and divisible by the forbidden prime five. At most twelve positive core summands would force exactly one summand on that residue class, which is impossible. Padding with zeros gives the equivalent fixed-order assertion. □

The field formulation follows from classical abelian class field theory. A subgroup condition on Frobenius is complete splitting in its fixed field. Finite conjunctions become intersections of inverse image subgroups in the compositum. Kronecker–Weber places this compositum in a cyclotomic field [6, Chapter I, Theorem 4.16], giving precisely the fixed subgroups in Theorem 11.1, after discarding the fixed bad-prime set. The reduction is classical; the pointwise indicator correlation is the analytic assertion used here.

Exact classifications and their finite certificates are treated in the companion *Exact Sumsets of Primitive Quadratic Norms* [15]. The separate article *Pointwise representation counts for Eisenstein restricted sums* concerns the placement of summands. Higher-fold sums and independent coupled systems can use classical linear-form correlations, with a different proof from this reflected binary problem. Growing-shift uniformity and higher indicator correlations are not covered by the statements here.

A An independent classical bound for quadratic cores

The quadratic divisor pairing has two factors, so completed Kloosterman sums and positive restoration give the required classical order bounds.

Fix an even modulus L_0, a nonprincipal quadratic literal character $\chi \bmod L_0$, and put

$$H = \ker \chi, \qquad \mathcal{C} = \langle p : p \nmid L_0,\ \chi(p) = 1 \rangle,$$
$$V(n) = \tau(n)\mathbf{1}_{\mathcal{C}}(n), \qquad J(N) = \prod_{\substack{p \mid N \\ p \nmid L_0,\ \chi(p) = -1}} \frac{p-1}{p-2}.$$

All arithmetic data in the following lemma are fixed.

Lemma A.1 (Classical quadratic core bound)**.** *Let A, B be coprime positive L_0-supported integers and $U, V_0 \subset H$ fixed residue subsets. For any positive A', B', define*

$$\Gamma_{A',B';U,V_0}(N) = \frac{L_0}{|H|^2}\#\{(a,b) \in U \times V_0 : A'a + B'b = N \bmod L_0\}.$$

Uniformly over all sufficiently large targets with $\Gamma_{A,B;U,V_0}(N) > 0$,

$$\#\{Ax + By = N : x, y \in \mathcal{C},\ x \bmod L_0 \in U,\ y \bmod L_0 \in V_0\} \asymp \frac{NJ(N)}{\log N}.$$

Zero finite density gives zero count identically. In particular the lemma applies to unions of unit-square residue states. Its proof uses only fixed-character means, fixed-progression prime Mertens, the classical Weil–Estermann bound, and finite sieve inequalities.

Proof plan and parameter order. Completed two-factor boxes give a shifted weighted family, uniform under the divisibility predicates needed by a finite lower sieve. The rough mass has lower-bound scale $NJ(N)/(\log N)^2$. Restoration supplies the missing logarithm through the positive product in (135); the valuation patterns $(1,1),(1,0),(0,1)$ recover its labels d, u, v, and the residual rough divisor weights are bounded. This makes restoration injective at the numerical-value level.

Choose the box exponent and core family first, then the saving η and external exponent θ from (127). Next choose α, ξ by (132), the Bonferroni depth, the roughness exponent, and finally the restoration range, in that order. All remain fixed as $N \to \infty$.

A.1 Two-factor boxes with a fixed period

For bounded functions ω_1, ω_2 of fixed period c, intervals I_1, I_2, and $(a,m) = (c,m) = 1$, let

$$S(a) = \sum_{\substack{r \in I_1,\ s \in I_2 \\ rs \equiv a\ (m)}} \omega_1(r)\omega_2(s).$$

Then, for every $\epsilon > 0$,

$$S(a) - \frac{1}{\varphi(m)} \sideset{}{^*}\sum_{b \bmod m} S(b) \ll_{c,\epsilon} m^{1/2+\epsilon}. \tag{121}$$

The estimate is uniform in the intervals, including intervals longer than cm, and permits subintervals of arbitrary lengths.

Here is a direct completion proof. Put $Q = cm$,

$$W_i(\nu) = \sum_{n \in I_i} e^{2\pi i - \nu n/Q}, \qquad G_i(t) = \sum_{r \bmod c} \omega_i(r) e^{2\pi i t r/c}.$$

CRT and finite Fourier inversion express $S(a)$ as

$$Q^{-2} \sum_{\nu,k \bmod Q} W_1(\nu)W_2(k)G_1(m^{-1}\nu)G_2(m^{-1}k)K_1(c^{-1}\nu, ac^{-1}k; m),$$

where $K_1(u,v;m) = \sum_{z \bmod m}^{*} e^{2\pi i(uz+v\bar{z})/m}$; inverses in G_i are taken modulo c. The terms $\nu = 0 \bmod m$ or $k = 0 \bmod m$ are independent of the unit a, since the resulting Ramanujan sum is unchanged by unit multiplication. For all other frequencies use

$$|K_1(u,v;m)| \le \tau(m) m^{1/2} (u,v,m)^{1/2}.$$

This is the ordinary Weil–Estermann bound stated and applied in [11, printed p. 44]. The primary source's Lemma 6 is its singleton-padded divisor-box consequence. In centered frequency ranges, $|W_i(\nu)| \ll \min(|I_i|, Q/|\nu|)$ and $|G_i| \le c$. The remaining double sum is bounded using

$$\sum_{\substack{0<|\nu|,|k|\le Q/2\\ m\nmid\nu,\ m\nmid k}} \frac{(\nu,k,m)^{1/2}}{|\nu k|} \le \sum_{d\mid m} d^{1/2} \left(\sum_{\substack{0<|\nu|\le Q/2\\ d\mid\nu}} \frac{1}{|\nu|} \right)^2 \ll_c \log^2(2m) \sum_{d\mid m} d^{-3/2}.$$

The Q^2 from the two geometric sums cancels the completion normalization. This proves a bound $O_c(m^{1/2}\tau(m)\log^2(2m))$ around an a-independent quantity. Averaging that quantity over unit a changes the error by at most a factor two and proves (121). The case $m = 1$ has discrepancy zero. Tensor partial summation also allows bounded-variation interval weights.

A.2 Shifted kernels, means, and fixed filters

For $t \in \mathcal{C}$ put $F_t(n) = V(tn)/V(t)$. At a split prime with $b = v_p(t)$,

$$F_t(p^e) = 1 + \frac{e}{b+1};$$

at a nonsplit prime it is zero for $e > 0$. With literal characters, its Dirichlet series is

$$\sum_{n\ge1} \frac{F_t(n)}{n^s} = L(s,\chi_0^{L_0})L(s,\chi) \sum_{j\ge1} \frac{e_t(j)}{j^s}.$$

The correction factors are $1 - p^{-2s}$ at inert primes and $1 - \frac{b}{b+1}p^{-s}$ at split primes dividing t. Deleting all Euler factors at primes dividing another integer D replaces its split local correction by $(1-p^{-s})^2$. Thus all these kernels satisfy, uniformly in valuations,

$$|e(j)| \le \tau(j)^C, \qquad \sum_j |e(j)| j^{-3/5} \ll_\epsilon (tD)^\epsilon. \tag{122}$$

Every kernel index is an L_0-unit with $\chi(j) = 1$.

We record explicitly the elementary mean used below. If ξ is a character of a fixed modulus c with the same prime support as L_0, then the twisted series has character factors $\xi, \xi\chi$. It has a principal factor only when the restriction of ξ is 1 or χ; in those cases $F_t\xi = F_t$. Otherwise both character sums have bounded partial sums, and Dirichlet hyperbola at $\sqrt{X}$ gives $O_c(\sqrt{X})$ for their convolution. In the principal cases, split the principal factor from the nonprincipal one by the same hyperbola argument to obtain its residue times X, with error $O_c(\sqrt{X})$. Summing kernels by (122), including their main tails, therefore gives

$$\sum_{n\le X} F_t(n)\mathbf{1}_{(n,D)=1}\xi(n) = \delta_{\xi\in\{1,\chi\}}\rho_{t,D}X + O_\epsilon((tDX)^\epsilon X^{3/5}). \tag{123}$$

Here the two displayed characters mean their inflations to c. The error and main remain valid with any fixed periodic unit filter, after finite character expansion. For later normalization write

$$M_{t,p} = (1-p^{-1})\sum_{e\ge0} F_t(p^e)p^{-e}, \qquad \rho_t = \frac{\varphi(L_0)}{L_0}\prod_{p\nmid L_0} M_{t,p}.$$

The product is taken in natural prime order; its convergence and positivity follow equally from the character factorization and $L(1,\chi) > 0$. A deleted split prime divides this mean by its original local series. The residue measure on the fixed unit group is uniform on H.

Proposition A.2 (Uniform shifted weighted family)**.** *There exist $\eta, \theta > 0$, with θ chosen sufficiently small after η, such that the following holds. Let $t_1, t_2, u, v \in \mathcal{C}$, each at most N^θ, with no restriction on their gcds. For fixed residue subsets $U, V_0 \subset H$,*

$$\sum_{\substack{Aux+Bvy=N \\ x,y\geq 1,\ x \bmod L_0\in U,\ y \bmod L_0\in V_0}} F_{t_1}(x)F_{t_2}(y)$$
$$= \frac{N}{ABuv}\,\Gamma_{Au,Bv;U,V_0}(N)\rho_{t_1}\rho_{t_2}\prod_{p\nmid L_0}\frac{D_p}{M_{t_1,p}M_{t_2,p}} + O(N^{1-\eta}),$$

uniformly in every target valuation. Here

$$D_p = \lim_{m_p\to\infty} p^{-m_p} \sum_{\substack{a,b \bmod p^{m_p} \\ Aua+Bvb=N\ (p^{m_p})}} F_{t_1}(p^{\min(v_p(a),m_p)})F_{t_2}(p^{\min(v_p(b),m_p)}).$$

All factors, including zero factors, are retained. The same formula holds after imposing any fixed finite combination of small divisibility and exact-valuation tests whose resulting shifts and coefficients stay in this family.

Proof. First suppose $(u, v) = (uv, N) = 1$. For $g = (x, N)$, write $x = gx'$, $y = gy'$, $M = N/g$; a nonzero term has $g \in \mathcal{C}$. Its exact prefactor is $F_{t_1}(g)F_{t_2}(g)$, its new shifts are $t_1 g, t_2 g$, and $(x', M) = 1$. Transport the residue states to $U/g, V_0/g$. Every such multiplication is a unit bijection, so the corresponding finite root count, with its transported target, is unchanged.

Choose a small fixed $\lambda > 0$, put $K = N^\lambda$, and truncate g and the two kernels at K. Large gcds cost $O(N^{1-\lambda+\epsilon})$. For either kernel, counting multiples and using (122) gives

$$N^{1+\epsilon}\sum_{j>K}\frac{|e(j)|}{j} + N^\epsilon\sum_{j\leq N}|e(j)| \ll N^{1-2\lambda/5+\epsilon} + N^{3/5+\epsilon}.$$

Discard also $Aux \leq N^{1-\lambda}$, at $O(N^{1-\lambda+\epsilon})$. All implicit valuation and shift factors here are divisor-bounded.

We suppress primes on x', y' in the remainder of this core argument. The fixed filter on the original first unit variable is

$$w(x) = \mathbf{1}_{x \bmod L_0\in U/g}\,\mathbf{1}_{B|M-Aux}\,\mathbf{1}_{(z,L_0)=1,\ z \bmod L_0\in v(V_0/g)}, \qquad z = (M - Aux)/B.$$

It has fixed period $c = BL_0$. Expand it on unit classes by characters modulo c, retaining every scalar phase of every kernel index. The filter ensures that $t = (M - Aux)/(Bvk)$ is an L_0-unit with $\chi(t) = 1$. Consequently the exact symmetric divisor pairing is

$$(\chi_0^{L_0} * \chi)(t) = 2\sum_{\substack{a|t \\ a<\sqrt{t}}}\chi(a) + \mathbf{1}_{t=\square}\chi(\sqrt{t}).$$

The square endpoint costs $O(N^{1/2+O(\lambda+\theta)+\epsilon})$.

For a second kernel $k \leq K$, the remaining congruence is

$$Aux \equiv M \pmod{vka}, \qquad Aux + Bvka^2 < M, \qquad (x, M) = 1.$$

The entire bad-supported factor B has already been imposed in w; all varying moduli are coprime to c. Also $(k, uM) = (a, uM) = 1$. Expand the first kernel j and its two character variables. A nonzero congruence requires $(uj, vka) = 1$; otherwise it would force a prime divisor of uM into an excluded place, or make a unit equal to a nonunit. The same zero test holds in the unit model.

Partition the two character variables and a into relative-width $\Delta = N^{-\zeta}$ boxes, with $\zeta = \lambda/100$. There are $O(N^{3\zeta+\epsilon})$ boxes. On a box with its endpoints fixed, apply (121) to the first two factors, at the full modulus $m = vka$. It is enough that $m \le N^{1/2+O(\lambda+\theta)}$; summing its square-root error over $a \le N^{1/2}$, then the two kernel masses, gives

$$O(N^{3/4+100\lambda+100\theta+\epsilon}). \tag{124}$$

This deliberately loose bound includes the fine boxes, the fixed character expansions, and the g-strata. In particular, no medium-variable dispersion argument occurs in this quadratic specialization.

We justify the curved boundary separately. A crossing box satisfies $|Aux + Bvka^2 - M| \ll \Delta N$. On the actual congruence write $M - Aux = Bvkat$. Then $|Bvka(t-a)| \ll \Delta N$, which has

$$O\left(\frac{\Delta N \log N}{vk} + \sqrt{\frac{N}{vk}}\right)$$

integer pairs (a,t), after discarding only the divisibility which determines x. The unique scalar decompositions and divisor-bounded first-kernel multiplicities cost N^{ϵ}. Summing the second kernel by (122) leaves

$$O_\epsilon\big(N^{\epsilon}(\Delta N \log N + \sqrt{N}K^{1/10})\big). \tag{125}$$

We next display the finite main sum produced by this replacement, before completing either the paired-divisor series or the target gcd sum. This also fixes the normalization of every local factor. For a retained g, put

$$C_g = F_{t_1}(g)F_{t_2}(g), \qquad \mathcal{A} = F_{t_1 g}, \quad \mathcal{B} = F_{t_2 g}, \quad H_{\mathcal{A},p} = \sum_{i\ge 0} \mathcal{A}(p^i)p^{-i},$$

and define

$$\rho_{\mathcal{A},M} = \rho_{\mathcal{A}} \prod_{\substack{p\mid M\\ p\nmid L_0}} H_{\mathcal{A},p}^{-1}, \qquad E_{\mathcal{A}}(q) = \prod_{p\mid q} H_{\mathcal{A},p}^{-1}.$$

Let μ_g be the average of the displayed filter w over H_c, the lift of H to $c = BL_0$. The contribution T_g of this stratum is $C_g\mathcal{P}_g$, up to the errors estimated above, where the *finite* sum is

$$\begin{aligned}\mathcal{P}_g = \frac{2\mu_g\rho_{\mathcal{A},M}}{Au} &\sum_{\substack{k\le K\\ (k,uM)=1}} e_{\mathcal{B}}(k)\\ &\times \sum_{\substack{a<\sqrt{M/(Bvk)}\\ (a,L_0uM)=1}} \frac{\chi(a)E_{\mathcal{A}}(vka)}{\varphi(vka)}(M - Bvka^2).\end{aligned} \tag{126}$$

At a zero-length endpoint the summand is zero, so the strict endpoint convention has no effect.

To check the step producing (126), the box replacement contributes $1/\varphi(q)$, $q = vka$, times the unit population of the original first sequence, with its fixed filter w. The fixed-character mean gives, uniformly in q,

$$\sum_{\substack{n\le X\\ (n,Mq)=1}} \mathcal{A}(n)w(n) = \mu_g\rho_{\mathcal{A},M}E_{\mathcal{A}}(q)X + O_\epsilon((t_1gMqX)^{\epsilon}X^{3/5}).$$

Indeed the character expansion of w has principal contributions only at $1, \chi$, whose coefficients sum to μ_g. The first-kernel indices retain their scalar character phases; recombining them yields exactly this original-sequence mean. Their main terms retain $1/j$, while the errors retain $j^{-3/5}$. The interval length is $X = (M - Bvka^2)/(Au)$, and divisor pairing supplies the

factor two. This proves the finite formula. Its mean errors sum to $O(N^{3/5+100\lambda+100\theta+\epsilon})$: use $1/\varphi(vka) \ll_\epsilon (vka)^{-1+\epsilon}$ and (122). The main boundary intervals keep the same reciprocal kernel sums. Combining these errors with (124), (125), the truncations, and the square endpoints leaves the following uniform total reserve:

$$\begin{aligned}\mathcal{R}_N \ll_\epsilon\ & N^{1-2\lambda/5+\epsilon} + N^{1-\lambda+\epsilon} + N^{1-\zeta+\epsilon} \\ & + N^{3/4+100\lambda+100\theta+\epsilon} + N^{3/5+100\lambda+100\theta+\epsilon} + N^{1/2+100\lambda+100\theta+\epsilon}.\end{aligned} \tag{127}$$

The sum of the errors over retained g, including their divisor-bounded C_g, obeys this reserve, after decreasing ϵ in the preceding estimates.

The finite-place coefficient can now be checked independently. The group H_c has $B|H|$ elements. Each allowed pair modulo L_0 has B pair lifts modulo BL_0, since Au is a unit at primes dividing B. Each surviving first residue in the definition of w admits exactly B second-residue lifts. Thus the number of such first residues is the original number of pairs modulo L_0, and

$$\Gamma_g := \Gamma_{Au,Bv;U/g,V_0/g}(M) = \frac{2BL_0}{\varphi(L_0)}\mu_g = \frac{2B}{\alpha_0}\mu_g, \qquad \alpha_0 = \varphi(L_0)/L_0. \tag{128}$$

The map $(x', y') \mapsto (gx', gy')$ bijects these root pairs with the original ones. Hence Γ_g, and therefore μ_g, is constant through the g-sum. Denote this common mean by μ.

We complete the two finite sums in (126) with explicit error bounds. Set $M_{\mathcal{A},p} = (1 - p^{-1})H_{\mathcal{A},p}$ and $W_p = M_{\mathcal{A},p}^{-1}$. Then

$$\frac{E_{\mathcal{A}}(vka)}{\varphi(vka)} = \frac{1}{vka}\prod_{p\mid vk} W_p \prod_{\substack{p\mid a\\ p\nmid vk}} W_p.$$

For fixed k put

$$b_k(a) = \chi(a)\mathbf{1}_{(a,L_0uM)=1}\prod_{\substack{p\mid a\\ p\nmid vk}} W_p, \qquad Y_k = \sqrt{M/(Bvk)}.$$

For every allowed shift, $1 - p^{-1} \le M_{\mathcal{A},p} \le (1 - p^{-1})^{-1}$. Thus $W_p = 1 + O(1/p)$, uniformly in all its valuations. The Dirichlet series of b_k is $L(s, \chi)$ times a correction absolutely convergent for $\Re s > 0$. At a prime outside $uMvkL_0$, its correction is explicitly $1 + (W_p - 1)\chi(p)p^{-s}$; at primes of uM it has the finite deletion $1 - \chi(p)p^{-s}$, and at primes of vk no new correction is required. Since the partial sums of the fixed nonprincipal χ are bounded, convolution with this correction gives

$$\sum_{a\le Y} b_k(a) \ll_\epsilon (t_1 g\, uMvkY)^\epsilon.$$

In particular $\sum b_k(a)/a$ converges, its tail after Y is $O(N^\epsilon/Y)$, and $\sum_{a\le Y} a\, b_k(a) = O(N^\epsilon Y)$. Consequently

$$\begin{aligned}&\sum_{a<Y_k} \frac{\chi(a)E_{\mathcal{A}}(vka)}{\varphi(vka)}(M - Bvka^2) \\ &\qquad = \frac{M}{vk}\prod_{p\mid vk} W_p \sum_{a\ge 1}\frac{b_k(a)}{a} + O_\epsilon\left(N^\epsilon\sqrt{\frac{M}{vk}}\right).\end{aligned}$$

The fixed B is absorbed in the error constant. Summing against $e_{\mathcal{B}}(k)$ costs at most $K^{1/10}$ in the square-root error, by (122). The completed main has an absolutely bounded correction

times $e_{\mathcal{B}}(k)/k$, up to N^{ϵ}. Its $k > K$ tail is therefore $O(N^{1-2\lambda/5+\epsilon})$. Both errors are included in (127).

Here is the Euler calculation of the completed sum, without an implicit normalization. Write $z = p^{-1}$, $q_p = 1 - z$,

$$A_i = \mathcal{A}(p^i), \qquad B_i = \mathcal{B}(p^i), \qquad P_{\mathcal{B},p}(z) = \sum_{i\ge 0} e_{\mathcal{B}}(p^i) z^i.$$

The kernel identity says

$$\frac{P_{\mathcal{B},p}(z)}{1-\chi(p)z} = (1-z)\sum_{i\ge 0} B_i z^i = M_{\mathcal{B},p}.$$

The double sum over k, a, after extracting its real factor $M/(uv)$, has the following local factors:

$$\begin{aligned}
\mathcal{L}_p &= 1 + W_p\left(\frac{P_{\mathcal{B},p}(z)}{1-\chi(p)z} - 1\right) && (p \nmid uvM)\\
&= \frac{M_{\mathcal{A},p} + M_{\mathcal{B},p} - 1}{M_{\mathcal{A},p}} && (p \nmid uvM),\\
\mathcal{L}_p &= 1 && (p \mid uM),\\
\mathcal{L}_p &= \frac{W_p P_{\mathcal{B},p}(z)}{1-\chi(p)z} = \frac{M_{\mathcal{B},p}}{M_{\mathcal{A},p}} && (p \mid v).
\end{aligned}$$

For $p \mid uM$, both local indices k, a are excluded. For $p \mid v$, the factor W_p is already present even at $v_p(k) = v_p(a) = 0$; this explains that row of the table. The first mean is

$$\rho_{\mathcal{A},M} = \alpha_0 \prod_{p\nmid L_0} M_{\mathcal{A},p} \prod_{\substack{p\mid M\\ p\nmid L_0}} H_{\mathcal{A},p}^{-1}.$$

At $p \mid M$ its combined factor is therefore $M_{\mathcal{A},p}/H_{\mathcal{A},p} = q_p$. Multiplying the other rows by this first mean gives the completed fixed-g expression

$$\begin{aligned}
C_g \frac{M}{Auv}\, 2\alpha_0\mu \prod_{\substack{p\mid M\\ p\nmid L_0}} q_p \prod_{p\mid u} M_{\mathcal{A},p}\\
\times \prod_{p\mid v} M_{\mathcal{B},p} \prod_{p\nmid L_0 uvM} (M_{\mathcal{A},p} + M_{\mathcal{B},p} - 1).
\end{aligned} \tag{129}$$

Intermediate products are taken in natural prime order. Their factorization into the fixed $L(1,\chi)$ and absolutely convergent corrections justifies every completion; the final normalized product in the proposition is absolutely convergent.

We may extend the divisor sum from $g \le K$ to all $g \mid N$. The means and local factors in (129), with their normalized shifts, are $N^{O(\epsilon)}$ uniformly, whereas $C_g M \ll_\epsilon N^{1+\epsilon}/g$. The added main is $O(N^{1-\lambda+O(\epsilon)})$. At $p^n \| N$, use now the original local coefficients $A_i = F_{t_1}(p^i)$, $B_i = F_{t_2}(p^i)$. The choices $v_p(g) = i < n$ in (129) give $q_p A_i B_i z^i$. The choice $i = n$ gives

$$z^n A_n B_n \big(M_{A^{[n]},p} + M_{B^{[n]},p} - 1\big), \qquad A_j^{[n]} = A_{n+j}/A_n, \quad B_j^{[n]} = B_{n+j}/B_n.$$

A zero $A_n B_n$ means this choice is absent and no division by that coefficient is made. Expanding the two shifted means produces exactly

$$\begin{aligned}
D_p = I_{p,n} = q_p \sum_{i<n} A_i B_i z^i + z^n(1-2z) A_n B_n\\
+ q_p \left(B_n \sum_{i>n} A_i z^i + A_n \sum_{j>n} B_j z^j \right).
\end{aligned}$$

At $p \mid u$ or $p \mid v$, the target is a unit and the respective factors remain $M_{A,p}$ or $M_{B,p}$, exactly the original local line integral. Finally (128) implies

$$\frac{2\alpha_0\mu}{Auv} = \frac{\Gamma_{Au,Bv;U,V_0}(N)\alpha_0^2}{ABuv}, \qquad \rho_{t_1}\rho_{t_2} = \alpha_0^2 \prod_{p\nmid L_0} M_{t_1,p}M_{t_2,p}.$$

Together with (129), these identities give precisely the proposition's main term in the unit-coefficient target core, including the bad coefficient B.

For completeness, the general coefficient reductions also have an explicit local check. At a good prime write $a = v_p(u), b = v_p(v), n = v_p(N)$, and $t = \min(a, b)$. If $n < t$, the density is zero. Otherwise dividing the congruence by p^t gives

$$D_p = p^t D_p^{\mathrm{red}}(a - t, b - t, n - t).$$

The factor p^t follows from counting the two sets of lifts and the one-dimensional line normalization p^{-m_p}. If the remaining valuations are $a > 0, b = 0$, solving for the unit-coefficient variable gives

$$D_p^{\mathrm{red}} = \begin{cases} B_n M_{A,p}, & n < a, \\ B_a I_{p,n-a}(A, B^{[a]}), & n \geq a, \end{cases} \qquad B_j^{[a]} = B_{a+j}/B_a.$$

There is the symmetric formula when $b > 0, a = 0$. A zero normalizing coefficient makes the relevant density zero. These formulas retain arbitrarily large coefficient and target valuations.

Globally remove $r = (u, v)$, requiring $r \mid N$, and divide the coefficients and target by r. For the coprime remaining coefficients put $d = (u, N)$, $e = (v, N)$. Then $e \mid x, d \mid y$; the exact substitutions are

$$x = ex', \quad y = dy', \quad (t_1, t_2) \mapsto (t_1 e, t_2 d), \quad (u, v, N) \mapsto (u/d, v/e, N/(de)),$$

with prefactor $F_{t_1}(e)F_{t_2}(d)$ and states $U/e, V_0/d$. The new coefficients are prime to the new target. All removed factors belong to $\mathcal{C}$. Their unit multipliers biject the finite root sets, with the transported target and states. The displayed local formulas show that the radial integrals transform with exactly the same prefactors and common-coefficient gains. Thus reexpressing the transformed main in the original variables gives the stated D_p-product, for every gcd profile.

All estimates in (127) are uniform in these profiles and in the shift valuations. They depend on divisor bounds and (122), not on an unbounded number of local cases. When the external family exponent is decreased as below, transformed targets are at least $N^{1-3\theta}$, transformed parameters lie in the proved core range, and the divisor-bounded prefactors preserve a smaller positive saving. Choose λ small, then θ much smaller. Equations (124) and (125), the kernel saving $2\lambda/5$, and the gcd and small-summand saving λ, leave a fixed positive core saving. Choose η below a hundredth of the boundary saving, and decrease the external family exponent again. The transformed targets are at least $N^{1-3\theta}$ and all transformed parameters stay inside the core family. Divisor prefactors preserve a smaller saving. Finally $b \mid x$ is the exact substitution $x = bx'$, which changes the shift to $t_1 b$, coefficient to ub, state to U/b, and introduces $F_{t_1}(b)$. The analogous substitution on the second side, and finite differences for exact valuations, prove the predicate assertion. □

A.3 A finite lower sieve and positive restoration

We now prove Lemma A.1. Fix an admissible target. Its finite density $\Gamma(N)$ belongs to a fixed finite positive set, and hence has a positive lower bound depending only on the fixed data. At good primes write $x = p^{-1}$, $q_p = 1 - x$, and use the unshifted coefficients $A_i = B_i = V(p^i)$. Let $I_{p,e}$ be the radial integral in the preceding proof, where $e = v_p(N)$. The weighted mass with both variables units is

$$c_e(p) = \begin{cases} 1 - 2/p, & e = 0, \\ 1 - 1/p, & e \geq 1. \end{cases}$$

At an inert prime $I_{p,e} = c_e(p)$. At a split prime define the bad-event probability $g_e(p) = 1 - c_e(p)/I_{p,e}$. The explicit geometric series gives, uniformly in e,

$$0 \le g_e(p) < 1, \qquad g_e(p) = 4/p + O(p^{-2}). \tag{130}$$

For example, at $e = 0$, $I_{p,0} = 2(1 - p^{-1})^{-1} - 1$. For $e \ge 1$, the $i = 0, 1$ terms give $I_{p,e} = 1 + 3/p + O(p^{-2})$; the remaining polynomial-geometric tail is uniformly $O(p^{-2})$. The formula at $e = 1$ has the same expansion directly. At every finite collection of primes $I_{p,e}$ is uniformly bounded, whereas $c_e(p) \ge 1/3$, so the good-event probabilities have a uniform positive lower bound.

Put $M_p = (1 - p^{-1}) \sum_i V(p^i) p^{-i}$ and

$$\mathcal{V}(z) = \prod_{\substack{p \le z \\ p \nmid L_0}} \frac{1 - 2/p}{M_p^2}, \qquad E_p = \frac{p-1}{p-2}, \qquad E_{p,e} = \begin{cases} 1, & e = 0, \\ E_p, & e \ge 1, \end{cases} \qquad a_p = \frac{p-1}{p(p-2)}.$$

At a split prime $M_p = (1 - p^{-1})^{-1}$, so the factor in $\mathcal{V}$ is $1 - 4/p + O(p^{-2})$. At an inert prime it is $1 + O(p^{-2})$. Fixed-character prime Mertens, and the split-prime density $1/2$, therefore give

$$\mathcal{V}(z) \asymp (\log z)^{-2}. \tag{131}$$

This is the dimension-two rough weighted count which must recover one power of $\log N$.

Choose small fixed $\alpha, \xi > 0$ after the weighted family. We use squarefree pairwise coprime triples d, u, v with

$$d, u, v \in \mathcal{C}, \quad d \mid N, \quad (uv, N) = 1, \quad duv \le N^{\alpha},$$

all primes at most $w = N^{\omega}$. At target $M = N/d$, prescribe $v_p(x) = 1$ for $p \mid u$, $v_p(y) = 1$ for $p \mid v$, and require that no other prime at most $z = N^{\gamma}$ divide xy. Count with weight $V(x)V(y)/(2^{\omega(u)}2^{\omega(v)})$ and with residue states $x \in U/d$, $y \in V_0/d$. Thus the lifted values dx, dy have the original states.

All these predicates are covered by Proposition A.2. Indeed exact valuation one is the difference of divisibility by p and by p^2, and a roughness predicate expands using

$$\mathbf{1}_{p \mid xy} = \mathbf{1}_{p \mid x} + \mathbf{1}_{p \mid y} - \mathbf{1}_{p \mid x,\ p \mid y}.$$

For a tested squarefree divisor $\ell \le N^{\xi}$, each resulting shift and coefficient is at most $N^{2\alpha+\xi}$. The expansion has $N^{o(1)}$ terms. It suffices to choose

$$2\alpha + \xi < (1 - \alpha)\theta, \qquad \alpha + \xi < \eta/4. \tag{132}$$

Forbidden shifts give zero terms; common-divisibility terms use the noncoprime part of the weighted family.

The finite density remains exactly $\Gamma(N)$ in every lower term. After writing $x = ux'$, $y = vy'$, its states are $U/(du), V_0/(dv)$ and its equation is $Aux' + Bvy' = N/d$. The map $(x', y') \mapsto (dux', dvy')$ bijects its finite roots with the original roots modulo L_0. This is why split primes with nontrivial residue states may be restored. Restricting restoration to identity-state primes would lose the exponent needed below.

Here is the exact positive local model after imposing all roughness conditions. Its main mass for a triple is

$$\mathcal{M}_{d,u,v} = C_N(z) \frac{N}{ABd} \mathcal{V}(z) \prod_{\substack{p \mid N/d \\ p \le z,\ p \nmid L_0}} E_p \prod_{p \mid uv} a_p. \tag{133}$$

The positive factor $C_N(z)$ includes $\Gamma(N)\rho_1^2$ and the Euler tail above z. It is bounded above and below by fixed positive constants, uniformly in the target and the triples. To verify this, outside

the target the tail factors are $1+O(p^{-2})$; at target primes they are $1+O(1/p)$ uniformly in valuation. At most $1/\gamma$ target primes exceed z. All cofactor primes are below $w<z$, so this tail is unchanged when dividing by d. At a unit target prime, a valuation-one directive on either side has weighted mass $2(1-p^{-1})/p$. Its normalization by two and division by $1-2/p$ give exactly a_p, proving (133).

We justify that a finite sieve recovers a fixed fraction of this fully rough main, rather than asserting an asymptotic sieve. Treat finitely many small primes exactly so that $g_e(p)\le 1/2$ at the others. Split the remaining primes into blocks $(z^{2^{-j-1}}, z^{2^{-j}}]$. Their sums of local bad probabilities are bounded by a fixed constant M_0, uniformly in targets and directives. Let U_j be the even Bonferroni polynomial through degree $2(s+j)$, and L_j the odd one through degree $2(s+j)+1$. If I_j is the indicator of no bad event in block j, then $U_j\ge I_j\ge L_j$, $U_j\ge 0$, and telescoping gives

$$\prod_j U_j-\sum_j (U_j-L_j)\prod_{i\ne j}U_i\le\prod_j I_j. \tag{134}$$

This avoids multiplying lower polynomials which might be negative. In the independent prime model,

$$\mathbb{E}(U_j-L_j)\le\frac{M_0^{2(s+j)+1}}{(2(s+j)+1)!},\qquad \mathbb{E}U_j\ge\prod_{p\text{ in block}}(1-g_e(p))\ge e^{-2M_0}.$$

Choose the fixed s so that the sum of these relative errors is at most $1/4$. Factoring $\prod_j \mathbb{E}U_j$ shows that (134) has mean at least half the fully rough probability. Its squarefree divisor support is at most

$$C_0 z^{\sum_{j\ge0}(2(s+j)+1)2^{-j}}=C_0 z^{4s+6},$$

with coefficients $O(\log\log N)$. Choose $\gamma<\xi/(4s+7)$. Then its support is at most N^{ξ}, so Proposition A.2 applies term by term. There are at most $N^{\alpha+o(1)}$ cofactor triples. The total error, including sieve predicates, is

$$O(N^{1-\eta+\alpha+\xi+o(1)}),$$

negligible by (132). Thus the summed actual weighted count is at least a fixed fraction of the sum of (133), apart from that error.

Temporarily drop the size cap $duv\le N^{\alpha}$. The positive cofactor sum is an Euler product. At a split prime not dividing N, its options are $1, a_p, a_p$. At a split prime with target valuation $e\ge1$, the options omitting or including it in d give

$$E_p+p^{-1}E_{p,e-1}\ge E_p+1/p=1+2a_p.$$

At inert target primes no restoration option is present; the factor left over is E_p. Target factors above z contribute $1+o(1)$, uniformly. Consequently (131) and prime Mertens give

$$\sum_{d,u,v}\mathcal{M}_{d,u,v}\gg NJ(N)(\log N)^{-2}\prod_{\substack{p\le w\\ \chi(p)=1}}(1+2a_p)\asymp\frac{NJ(N)}{\log N}. \tag{135}$$

The cap must be restored in the actual positive model. After normalizing its Euler product, each prime is included in d,u, or v with probability $O(1/p)$, uniformly in target valuation. Therefore $\mathbb{E}\log(duv)\le C_1\log w$. Choose $\omega<\min(\gamma,\alpha/(4C_1))$. Markov retains at least three quarters of the mass in $duv\le N^{\alpha}$. Applied to this model, rather than merely a comparison product, it proves that the finite sieve retains the order of magnitude in (135).

Finally lift each residual pair to (dx,dy). At every cofactor prime the numerical valuation patterns are, respectively, $(1,1)$ for d, $(1,0)$ for u, and $(0,1)$ for v; all other small-prime

valuations are zero. Thus the ordered numerical pair recovers the entire triple and residual pair uniquely. Writing $x = ux'$, $y = vy'$, exact valuation one gives

$$\frac{V(x)V(y)}{2^{\omega(u)}2^{\omega(v)}} = V(x')V(y').$$

All prime factors of x', y' exceed N^{γ}, so $V(x')V(y') \le 4^{1/\gamma}$. No factor $V(d)^2$ is inserted: d lifts values, not their divisor weights. Injectivity and this uniform bound turn (135) into the required unweighted lower bound.

For the upper bound, sieve the integer parameter of $Ax + By = N$ using only good inert primes. Both coefficients are units there. There are two forbidden roots, merging exactly when $p \mid N$. A dimension-one upper Selberg sieve at fixed power level therefore gives

$$\ll N \prod_{\substack{p \le N^{\gamma_1} \\ p \nmid L_0,\ \chi(p)=-1}} \left(1 - \frac{2 - \mathbf{1}_{p|N}}{p}\right) + O(N^{1/2}\log^C N) \ll \frac{NJ(N)}{\log N}.$$

For an explicit finite version, use the nonnegative even block polynomials above at divisor support at most $N^{1/2}$. Their model is at most a fixed multiple of the true product, by the same summable factorial block errors. CRT gives $N\nu(\ell)/\ell + O(\nu(\ell))$, $\nu(p) \le 2$; the total endpoint error is $O(N^{1/2}\log^C N)$. Discarding the residue filters only increases this upper bound. This completes the proof of Lemma A.1.

AI assistance and review. ChatGPT (OpenAI) assisted with mathematical arguments, literature checks, exposition, proof review and verification code.

References

[1] J. Bourgain and E. Fuchs, *On representation of integers by binary quadratic forms*, arXiv:1105.3698v2, 23 May 2011. https://arxiv.org/abs/1105.3698v2.

[2] H.-E. Richert, *Lectures on Sieve Methods*, notes by S. Srinivasan, Tata Institute of Fundamental Research, Bombay, 1976, Theorems 11.5–11.8, pp. 138–142. Online edition.

[3] H. Iwaniec and J. Pomykała, *Sums and differences of quartic norms*, Mathematika **40** (1993), 233–245. https://doi.org/10.1112/S0025579300007014.

[4] I. Khayutin, *Joint equidistribution of CM points*, Annals of Mathematics (2) **189** (2019), 145–276. https://doi.org/10.4007/annals.2019.189.1.4.

[5] D. Loughran and L. Matthiesen, *Frobenian multiplicative functions and rational points in fibrations*, Journal of the European Mathematical Society **26** (2024), 4779–4830. https://doi.org/10.4171/JEMS/1374.

[6] J. S. Milne, *Class Field Theory*, version 4.03, 6 August 2020. https://www.jmilne.org/math/CourseNotes/CFT.pdf.

[7] OpenAI, *Weighted dilation graphs, smooth shifted primes and totient fibers*, 24 September 2026, Section 3. Immutable primary PDF.

[8] OpenAI, *The Quasi-Riemann Hypothesis: A Zero-Free Half-Plane* $\Re(s) > 7/8$, 30 September 2026, Part I, Theorem 3.1. Immutable primary PDF.

[9] M. N. Walsh, *Shifted convolutions of divisor functions*, arXiv:2610.40439v2, 12 November 2026. https://arxiv.org/abs/2610.40439v2.

[10] Y. Smilansky, *Sums of two squares—pair correlation and distribution in short intervals*, International Journal of Number Theory **9** (2013), 1687–1711. https://doi.org/10.1142/S1793042113500516. https://arxiv.org/abs/1208.2419v2.

[11] D. R. Heath-Brown, *The divisor function $d_3(n)$ in arithmetic progressions*, Acta Arithmetica **47** (1986), 29–56. https://doi.org/10.4064/aa-47-1-29-56.

[12] R. de la Bretèche and G. Tenenbaum, *Remarks on the Selberg–Delange method*, Acta Arithmetica **200** (2021), 349–369. https://doi.org/10.4064/aa201024-26-5. Corrected author edition.

[13] K. Soundararajan, *The Liouville function in short intervals [after Matomäki and Radziwiłł]*, arXiv:1606.08021v1, 26 June 2016, Lemmas 2.1–2.2. https://arxiv.org/abs/1606.08021v1.

[14] T. Tao, *254A, Notes 2: Complex-analytic multiplicative number theory*, 9 December 2014, Proposition 51, Theorem 62(ii), Exercise 45(i)–(ii). https://terrytao.wordpress.com/2014/12/09/254a-notes-2-complex-analytic-multiplicative-number-theory/.

[15] R. van Rijn, *Exact Sumsets of Primitive Quadratic Norms*, version 1, Zenodo, 2026. https://doi.org/10.5281/zenodo.24371462.