Exact Selmer dimensions and signed averages in two elliptic families

Brice Pouly

Prepared with AI assistance

hexagon:2611.01209v1

Abstract

For every integer $k \geq 1$ we determine the full 2-Selmer dimension of $y^2 = x^3 + 4(18k+1)x^2 + 4x$. One global square class can contribute two dimensions. Subtracting the 2-rank of an imaginary quadratic class group leaves an explicit integer-valued law, independent of two Gaussian prime counts on the full parameter line and of four at triangular parameters $k = n(n+1)/2$. Correcting Selmer parity at $2, 3, 7$ recovers the Liouville sign of $4k(9k+1)$ with limiting frequencies of 97.5370% and 95.1742%, respectively, to four decimal places. On the full parameter line, raw and corrected parities are logarithmically balanced on each fixed progression, even at a prescribed value of the integer difference of positive density. In a second elliptic family, two relative Selmer parities from cyclic cubic trace-zero twists yield opposite root-number-weighted means of order $(\log X)^{-4/3}$. We compute the first two coefficients on every fixed stratum of positive even valuations. Each prescribed prime $11, 19, 23 \bmod 24$ delays the leading exponent by one, and each prime $1 \bmod 24$ reverses its sign. The second coefficients have the same sign for the two cubic parities.

1 Two ways to detect the parity of a product

For a positive integer t let $\lambda(t) = (-1)^{\Omega(t)}$, with prime factors counted with multiplicity. We study the sign $C(n) = \lambda(D(n))$ of

$$(L_1, L_2, L_3, L_4) = (n, n+1, 3n+1, 3n+2), \qquad D(n) = \prod_{i=1}^{4} L_i = 4k(9k+1), \qquad k = \frac{n(n+1)}{2}. \tag{1}$$

There is a natural reason to look for this parity in an isogeny descent. The squarefree part of $k(9k+1)$ determines C. It also determines which odd primes ramify in a quadratic field. On a curve with two colliding rational roots, the local isogeny index can see one further kind of prime, where the collision is split and has even positive depth. The discrepancy is sparse enough that it can be counted exactly in the limit.

The first question is how sampling changes what the Selmer dimension can detect. Theorem 2.1 is pointwise in every positive k. On the triangular subfamily, Theorem 3.1 separates the integer difference from the quadratic genus rank and four Gaussian prime counts. Proposition 3.3 gives the corresponding law for all k: the dyadic factor persists, but the multiplicity of each odd root changes. These comparisons follow from arithmetic formulas for the two dimensions.

For linear parameters, the target sign is $\lambda(k)\lambda(9k+1)$. Tao's binary theorem gives balance on each fixed progression after local corrections. Theorem 4.3 also computes the joint signs on the full parameter line. On the triangular subfamily, the same correction still gives a high agreement probability. Here the four-site conditional signed mean in Corollary 4.2 remains an arithmetic input. *Elliptic root numbers and a signed squarefree variance for a quartic Liouville correlation* [21] studies this four-site sign through root numbers and signed variances. The first family here uses the parity of the full Selmer dimension of E'.

There is a second way to use the same four signs. Multiplication by the root number of another curve removes three Liouville factors, leaving a one-variable multiplicative function and local signs. A relative Selmer parity, constructed from a cyclic cubic extension and evaluated at the remaining factor, then changes the first surviving term of the mean. These are Selmer groups of trace-zero abelian surfaces, rather than the usual cubic twists of elliptic curves with $j = 0$. The two choices in Theorem 7.1 give opposite signs on the full interval. On the subset where 11 is the only prime $p > 3$ of positive even valuation in D, the first term has order $(\log X)^{-7/3}$. Replacing 11 by 73 restores the exponent $4/3$ and reverses the sign. The local calculation in Theorem 8.1 explains both effects. Theorem 8.3 also gives the next coefficient, at order $(\log X)^{-5/3}$ on the full population. Throughout, Selmer dimensions include rational torsion, logarithms are natural, and the root number is that of the global functional equation.

Relation to earlier work. Normal fluctuations of isogeny Selmer dimensions and Tamagawa ratios occur in several previously studied families. Xiong and Zaharescu [27, Theorem 1] treat quadratic twists with a small isogeny Selmer group and a normally distributed dual dimension. Klagsbrun and Lemke Oliver study Tamagawa ratios and quadratic twists with partial two-torsion [12, 13]. Morgan and Paterson obtain local formulas for full 2-Selmer dimensions on a density-one set after quadratic extension [19, Corollaries 6.4 and 6.6]. Chan and Verzobio prove normal laws for Selmer ratios in families with an isogeny, ordered by naive height [4, Theorems 1.1 and 5.9]. These precedents already include families over $\mathbb{Q}$. Our calculation concerns a non-isotrivial family, every positive parameter, and the integer difference that Gaussian normalization would erase.

There is also a close geometric comparison. Translating the source curve in (2) by $X = x - (18k + 2)$ gives

$$y^2 = X(X + 2)(X + 18k + 2).$$

This lies in the family considered by Yu [28, Theorem 1.1] and in the fixed-root example of Koymans and Smith [16, Example 7.3], with their parameters $b = -2$ and $u = -18k - 2$. Those results concern different averaged objects and parameter populations. Our exact formula is for the isogenous *target* E', whose full Selmer dimension is not determined by the source's full Selmer dimension. Moreover, the target's generic two-torsion field on the triangular subfamily is $\mathbb{Q}(n)(\sqrt{D(n)})$. The four simple roots of D prevent this extension from becoming constant after a number-field extension, so the constant-torsion hypothesis of [16, Theorem 1.3] does not apply to this target. The Gaussian coordinates fit the general prime-count theory of El-Baz, Loughran and Sofos [6]; the additional information here is the explicit discrete correction and its independence from those coordinates.

For the signed averages, the trace-zero construction is due to Mazur, Rubin and Silverberg [18], and the local parity comparison to Klagsbrun, Mazur and Rubin [14, 15]. Helfgott's root-number factorization [11] and the Selberg–Delange method [9, 2] provide the analytic framework. The family-specific calculations identify the surviving character components, the exact even-valuation strata on which coefficients vanish, and the signs and non-vanishing of the first two coefficients. An exponent supplied by Selberg–Delange is not by itself the contribution: the local cancellations determine which exponent actually occurs.

Route through the argument. Sections 2–4 give the exact dimension, compare the two parameter laws, and identify the signed conclusions available under each sampling. Sections 5 and 6 construct the cubic parities and the root-number identity. Section 7 states the uniform analytic inputs and explains how square-divisibility conditions pass through the expansion. Section 8 uses the zeros of those local factors to compute the two coefficients on each fixed stratum. The estimates there concern ordinary intervals; the conditional symmetry in Proposition 5.2 concerns logarithmic laws.

2 An exact descent for every positive parameter

In this section $k \geq 1$ is an arbitrary integer. The descent does not require it to be triangular. Put $b = 9k + 1$ and write $kb = dt^2$ with $d > 0$ squarefree. The factors k, b are coprime. Consider the curves

$$E:\ y^2 = x(x-A)(x-B), \quad A = 18k,\ B = 2b = A+2, \qquad E':\ y^2 = x^3+4(18k+1)x^2+4x. \tag{2}$$

The isogeny $\phi : E \to E'$ with kernel generated by $(0,0)$ has dual $\hat{\phi}$. At $p \geq 5$ dividing kb, call the collision *split* if $(-2/p) = 1$ on $p \mid k$, or $(2/p) = 1$ on $p \mid b$. Let

$$J = \#\{p \geq 5 :\ p \text{ split},\ v_p(kb) > 0 \text{ even}\}, \qquad N = \omega_{\geq 5}(d) + J,$$

where $\omega_{\geq 5}$ counts distinct prime divisors at least five. The two sets of primes counted by N are disjoint.

At 2, put $h = v_2(kb) \geq 1$ and $w = kb/2^h \bmod 8$. Define $t_2 = 1$ if w occurs in the indicated cell, and $t_2 = 2$ otherwise.

	$h = 1$	$h \geq 2$
k even	$3, 5$	$1, 3$
k odd	$1, 3$	$1, 7$

One global class must also pass all the local tests. Define $e = 1$ if k is even, every prime counted by N satisfies $(-2/p) = 1$, and either $(h, w) = (1, 3)$ or $h \geq 3$ and $w \in \{1, 3\}$. In every other case put $e = 0$.

Theorem 2.1. *For every integer $k \geq 1$, the full dimension is*

$$s_k := \dim_{\mathbb{F}_2} \mathrm{Sel}_2(E'/\mathbb{Q}) = N + t_2 - 1 + 2e. \tag{3}$$

For $k = 1$, $kb = 10$, $w = 5$, and the formula gives $s_1 = 2$. The non-triangular parameter $k = 2$ gives $kb = 38$, $N = 1$, $t_2 = 1$, $e = 1$, and hence $s_2 = 3$. At the triangular parameter $n = 75$, one has $k = 2850 = 2 \cdot 3 \cdot 5^2 \cdot 19$ and $b = 25651 = 113 \cdot 227$. The prime 5 has even valuation and is nonsplit. The three primes $19, 113, 227$ are each 1 or 3 mod 8, and $h = 1,\ w = 3$. Thus $N = 3,\ t_2 = 1,\ e = 1$, giving $s_{2850} = 5$. The two extra dimensions come from one global Kummer class and its contribution to the dual Selmer dimension. Passing from s_k to a Mordell–Weil rank would also require the dimension of the 2-torsion in the Tate–Shafarevich group of E'.

2.1 Odd places and the real place

For $y^2 = x(x^2+ax+b_0)$, the Kummer map for the isogeny *arriving* at this curve is represented by x, taking the classes $1, b_0$ at infinity and $(0,0)$, respectively [23, Ch. X, §4]. On E', $b_0 = 4$. At an odd prime, a nonzero valuation $v_p(x)$ makes the valuation of the cubic $3v_p(x)$ if negative and $v_p(x)$ if positive. Thus the global candidates for $\mathrm{Sel}_\phi(E)$ are $1, -1, 2, -2$.

Normalize $t_v = \log_2 |E(\mathbb{Q}_v)/\hat{\phi}E'(\mathbb{Q}_v)| - 1$. At an odd $p \mid AB$ write $v = v_p(AB)$. On E the discriminant is $16A^2B^2(A-B)^2$. On E' the standard two-isogeny equation has discriminant of valuation v. Both c_4's are units. The types are I_{2v} and I_v. At $p \mid A$ the double point on E has tangent square class $-B = -2 \bmod p$, and at $p \mid B$ it has tangent square class $-A = 2 \bmod p$. The isogeny preserves splitting of the reduction torus. Its differential and that of its dual are units at odd p, since their product is 2 on minimal differentials. Hence $2^{t_p} = c_p(E)/c_p(E')$. Split I_j has $c_p = j$, and nonsplit I_j has $c_p = \gcd(2, j)$. This gives $t_p = 1$ at a split collision or when v is odd, and $t_p = 0$ otherwise. Good odd primes have $t_p = 0$. In particular $t_3 = 1$, since $3 \mid A$ and -2 is a square modulo 3.

The local image on E' at an odd prime always lies in the unit square classes. By local duality the two Kummer images are orthogonal complements in $\mathbb{Q}_p^\times/\mathbb{Q}_p^{\times 2}$. Thus the image on E' is trivial for $t_p = 1$ and is the unit subgroup for $t_p = 0$. The place 3 leaves $1, -2$ among the four candidates. Testing -2 at the other odd places gives precisely the prime condition in e. At the real place the abscissae on E are nonnegative, giving $t_\infty = -1$, and both signs occur on E'. In total $\sum_v t_v = N + t_2$, once t_2 has the value in the table.

2.2 The three dyadic generators

Order the two roots as $H = 2^{h+1}U$ and $M = 2V$, with U, V odd. The local image on E is

$$E(\mathbb{Q}_2)/\hat{\phi}E'(\mathbb{Q}_2) = \langle [H], [M], [T] \rangle, \qquad T = 3 + 4k \bmod 8. \tag{4}$$

Indeed the two roots supply H, M. For odd x the cubic is $x - H - M \bmod 8$, and $1 + H + M \equiv T \bmod 8$. An odd abscissa occurs exactly in class T. A negative valuation is even, and $1 - H/x, 1 - M/x$ are then 1 mod 8, so the class of x is trivial. For $1 < v_2(x) < h + 1$ the cubic has odd valuation. At $v_2(x) = 1$ or $h+1$, translation by the torsion point with matching valuation replaces x by $M(x - H)/(x - M)$ or $H(x - M)/(x - H)$, of valuation at most zero. If $v_2(x) > h + 1$, translation by $(0, 0)$ replaces x by HM/x, again of valuation at most zero. These translations multiply the Kummer class by known generators.

The classes M, T are independent. Since $HM = 36kb$, its odd unit is w mod 8 up to an odd square. For even h, testing HM shows that the plane $\langle M, T \rangle$ contains H exactly when $w \in \{1, T\}$. For odd h the test is $U \in \{1, T\}$ mod 8. When $h \geq 3$, $V = 1$ mod 8 for even k, and $V = 7 = T$ mod 8 for odd k. When $h = 1$, the difference of the two roots gives $V = 1 + 2U$ mod 8 for even k and $V = 2U - 1$ mod 8 for odd k. Substitution gives the displayed table, since t_2 is one less than the dimension in (4).

For any $x = 2^j v$ with v odd, the Hilbert symbol satisfies $(-2, x)_2 = 1$ exactly for $v = 1, 3$ mod 8. For odd k, $T = 7$ excludes -2. For even k, $T = 3$, $V = 1 + 9k$, and $h = 2$ gives $V = 5$. If $h = 1$, the two tests on U, V allow precisely $U = 1,\ V = 3$, or $w = 3$. If $h \geq 3$, $V = 1$ and $w = 1, 3$ is necessary and sufficient. This establishes $\dim \mathrm{Sel}_\phi(E) = e$.

2.3 The passage from isogeny Selmer groups to full dimension

Cassels's isogeny formula [3, Theorem 1.1, p. 184, (1.26)], applied to $\hat{\phi}$, gives $\dim \mathrm{Sel}_{\hat{\phi}}(E') - \dim \mathrm{Sel}_\phi(E) = \sum_v t_v$. The global ratio of rational kernel sizes is one. In the Selmer sequence

$$0 \longrightarrow E[\phi](\mathbb{Q})/\hat{\phi}(E'[2](\mathbb{Q})) \longrightarrow \mathrm{Sel}_{\hat{\phi}}(E') \longrightarrow \mathrm{Sel}_2(E') \longrightarrow \mathrm{Sel}_\phi(E)$$

the term on the left has dimension one. Indeed the two remaining roots of E' require $\sqrt{kb}$, and $(3k)^2 < kb < (3k + 1)^2$. To identify the last image, use Fisher's notation

$$S_1 = \mathrm{Sel}_\phi(E), \qquad S_2 = \mathrm{im}\big(\mathrm{Sel}_2(E') \longrightarrow \mathrm{Sel}_\phi(E)\big).$$

His descent pairing [7, Theorem 2.1, $m = 1$] is alternating on S_1, with radical S_2. Since $\dim S_1 = e \leq 1$, the pairing vanishes and $S_2 = S_1$. This argument does not assume the finiteness of the Tate–Shafarevich group. Thus the last arrow is onto and

$$s_k = (N + t_2 + e) - 1 + e = N + t_2 - 1 + 2e.$$

3 Four prime counts and an integer difference

First specialize to $k = k(n) = n(n + 1)/2$ and write $s(n) = s_{k(n)}$. This specialization factors $4k(9k + 1)$ into the four linear forms in (1); it is not needed for Theorem 2.1. Let

$r = \dim_{\mathbb{F}_2} \mathrm{Cl}(\mathbb{Q}(\sqrt{-d}))[2]$, and let Δ be the fundamental discriminant. Gauss's genus formula gives $r = \omega(|\Delta|) - 1$, since the field is imaginary [24, Theorem 2.1]. For $i = 1, \dots, 4$, define

$$R_i = \#\{p \geq 5 : v_p(L_i) \text{ is odd}\}, \qquad \mathcal{R} = \{n : k \bmod 16 \in \{0, 7, 8, 11, 12, 15\}\}.$$

The roots of the four L_i are distinct at every $p \geq 5$. Consequently

$$r = \sum_{i=1}^{4} R_i + \mathbf{1}_{2|\Delta} + \mathbf{1}_{3|d} - 1, \tag{5}$$

$$s - r = J + t_2 - \mathbf{1}_{2|\Delta} - \mathbf{1}_{3|d} + 2e. \tag{6}$$

These are integer identities. The sector $\mathcal{R}$ consists of twelve classes modulo 32 and has density 3/8. We give its law separately because the signed comparison in the next section uses this population.

Theorem 3.1. *Take n uniformly from $[X, 2X)$, either on all integers or conditionally on $\mathcal{R}$, and set $\ell = \log\log X$. Then*

$$\left(\left(\frac{R_i - \ell}{\sqrt{\ell}} \right)_{1 \leq i \leq 4}, \, s - r \right) \implies (G_1, G_2, G_3, G_4, D_*).$$

The G_i are independent standard normal variables and are independent of D_*. Write $\delta_p = \nu_p/[p(p+1)]$, with $\nu_p = 4, 2, 0, 2$ for $p = 1, 3, 5, 7 \bmod 8$, respectively. The probability generating function of D_* is

$$\mathbb{E} z^{D_*} = (a + b_1 z + c_2 z^2) \frac{1 + z^{-1}}{2} \prod_{p \geq 5} (1 - \delta_p + \delta_p z). \tag{7}$$

Here $(a, b_1, c_2) = (5/12, 1/2, 1/12)$ on all integers and $(5/18, 1/2, 2/9)$ on $\mathcal{R}$. The generating functions of $s - r$ converge locally uniformly on $\mathbb{C} \setminus \{0\}$. In particular the pair $((r - 4\ell)/(2\sqrt{\ell}), (s - 4\ell)/(2\sqrt{\ell}))$ converges to (G, G), with $G = (G_1 + \cdots + G_4)/2$.

For example, the factor $(1 + z^{-1})/2$ vanishes at $z = -1$. Thus s and r have equal parity on half the parameters in each population. Their integer difference retains a bias: on $\mathcal{R}$ its limiting expectation is $4/9 + \sum_{p \geq 5} \delta_p$. This distinction is lost upon Gaussian normalization.

3.1 The exceptional square class

The example $n = 75$ makes it necessary to retain e in (3). The following upper bound permits its removal when computing the law.

Lemma 3.2. *The number of $X \leq n < 2X$ with $e = 1$ is $O(X/(\log X)^2)$.*

Proof. For each $p \equiv 5, 7 \bmod 8$, a valuation of exactly one at any of the four roots contributes to N and has $(-2/p) = -1$. Exclude the corresponding $\rho(p) = 4(p-1)$ residue classes modulo p^2. For a squarefree product d' of such primes, the count of simultaneous memberships in the excluded sets is $Xw(d') + O(\rho(d'))$, where $w(p) = \rho(p)/p^2$ and both functions are multiplicative. Write $g(p) = w(p)/(1 - w(p))$ and extend to squarefree products. With sums restricted to these products, set $G(z) = \sum_{d' \leq z} g(d')$.

For clarity, the Selberg upper-bound weights supported on $d' \leq z$ are

$$\lambda_{d'} = \frac{\mu(d')}{G(z)} \prod_{p | d'} (1 - w(p))^{-1} \sum_{\substack{l \leq z/d' \\ (l, d') = 1}} g(l).$$

One has $\lambda_1 = 1$ and $|\lambda_{d'}| \le 1$, by including all products ul, $u \mid d'$, in $G(z)$. Squaring the sum of the weights over simultaneous excluded sets majorizes the indicator of survival. The quadratic main term is $X/G(z)$. Indeed

$$\sum_{d',e'} \lambda_{d'}\lambda_{e'} w([d',e']) = \sum_{l \le z} \frac{1}{g(l)} \left(\sum_{l \mid d'} \lambda_{d'} w(d') \right)^2,$$

and the inner sum is $\mu(l)g(l)/G(z)$ by Möbius inversion. The error is $O_\epsilon(z^{4+\epsilon})$ because $\rho([d',e']) \le 4^{\omega([d',e'])}[d',e']$.

To bound $G(z)$ below, take the full product over the permitted $p \le z^\alpha$ for small fixed $\alpha > 0$. Its mass is $\prod(1+g(p)) \asymp_\alpha (\log z)^2$ by Mertens' theorem in the two fixed classes. Under the weights proportional to $g(l)$, the expected $\log l$ is $\sum w(p)\log p \ll \alpha \log z$. At least half the mass lies at $l \le z$ for a sufficiently small α. Now take $z = X^{1/16}$. □

3.2 The discrete law

For $p \ge 5$ there are two roots coming from k, split when $(-2/p) = 1$, and two from $b = (3n+1)(3n+2)/2$, split when $(2/p) = 1$. This gives ν_p. At a simple root, positive even valuation has Haar probability $\sum_{j \ge 1}(1-1/p)p^{-2j} = 1/[p(p+1)]$. Thus J contributes one Bernoulli variable of mean δ_p per prime. Primes $p > Y$ contributing to J require $p^2 \mid L_i$ for some i. The probability is

$$\ll \sum_{p > Y} p^{-2} + X^{-1/2}, \tag{8}$$

by counting the four linear congruences for $p \ll \sqrt{X}$. Conditioning on $\mathcal{R}$ costs a bounded factor.

At 3, the roots of k are $0, -1$, and b is a unit. On either root, the unconditional probability of odd valuation is $(1-1/3)\sum_{j \ge 0} 3^{-2j-1} = 1/4$. This gives the independent factor $(1+z^{-1})/2$. At 2, each $k \bmod 2^j$ has exactly two preimages $n \bmod 2^{j+1}$: $n^2 + n - 2k$ has two simple roots modulo 2. The map $k \mapsto k(9k+1)$ has odd derivative and sends each parity class bijectively onto $2\mathbb{Z}_2$ (Hensel lifting). Hence, conditionally on the parity of k, h has mass 2^{-h} for $h \ge 1$, and w is uniform on the four units at each depth. Two is unramified in $\mathbb{Q}(\sqrt{-d})$ precisely when h is even and $w = 3, 7 \bmod 8$. Checking $k \bmod 16$ shows that $\mathcal{R}$ is $h \ge 3$ together with $h = 2,\ w = 3, 7$. Insert these conditions into the table for t_2. The masses of $h = 1,\ h = 2,\ h \ge 3$ odd, and $h \ge 4$ even are $1/2, 1/4, 1/6, 1/12$, respectively. They give the coefficients in (7) for $t_2 - \mathbf{1}_{2|\Delta}$.

The Chinese remainder theorem makes any finite collection of these local observations independent under the limiting residue law, with the stated change at 2 on $\mathcal{R}$. First truncate prime sizes and valuation depths. Then use (8), the geometric tails at the fixed primes and Lemma 3.2 in (6). This proves convergence of the integer coordinate. For exponential moments, let $B_2(m) = \#\{p : p^2 \mid m\}$. For every $a_0 \ge 1$,

$$a_0^{B_2(m)} = \sum_{v^2 \mid m} \mu^2(v)(a_0 - 1)^{\omega(v)}.$$

Each L_i maps the interval injectively into $[1, 6X+2]$. Summing the positive expansion over this latter interval bounds the mean by $O(1)\prod_p(1 + (a_0-1)/p^2)$. Since $J \le \sum_i B_2(L_i)$, Hölder bounds every fixed exponential moment of J uniformly. The gap is bounded below by -1 and differs from J by a bounded integer. Truncating the gap itself now proves the locally uniform convergence and convergence of all its moments.

3.3 Why the four Gaussian variables are independent of the gap

We give the moment argument because conditioning on the integer difference is important here. It is the sieve form of the Erdős–Kac method [10, 6]. Fix real coefficients $\theta_1, \dots, \theta_4$ and a fixed residue class of n. For $p \geq 5$ outside its modulus the four events $p \mid L_i$ have probability $1/p$ each in the residue model and are mutually exclusive. Replace odd valuations by divisibility up to $y = X^{1/\log(3+\ell)}$. Repeated prime factors cost $O(1)$ in mean by (8) and the finitely many small primes. All larger prime divisors cost $O(\log(3+\ell))$ deterministically, since $L_i \leq 6X + 2$. These errors are $o(\sqrt{\ell})$ in probability. The finitely many primes of the fixed modulus have the same negligible normalized effect.

Take independent such categorical variables across primes $p \leq y$, with value θ_i at the i-th root and zero off the four roots. Their sum has mean $(\sum_i \theta_i)\ell + o(\sqrt{\ell})$, and variance $(\sum_i \theta_i^2)\ell + o(\ell)$. Indeed the off-diagonal covariances at one prime are $-1/p^2$. Bounded independent summands satisfy the normal limit in moments, by expanding centered moments and grouping equal primes. For each fixed order j the moments of the arithmetic sum and this model agree up to $O_{j,\theta,q}((4y)^j(1+\ell)^j/X)$ after centering. Expanding into ordered tuples of prime and site gives this bound directly from the error $O_q(1/X)$ in the density of each compatible residue class, also when primes repeat. It tends to zero. The Cramér–Wold criterion now gives the same four-dimensional normal law on every fixed residue class. Joint convergence with any finite residue observation follows. Approximate $s - r$ by the finite observations just used for (7) and then remove their tails. This proves the independence and finishes Theorem 3.1.

3.4 Sampling the whole family

The same curve admits two natural choices of parameter. Sampling n gives two roots each for $k(n)$ and $9k(n) + 1$ at every prime above three. Sampling k directly leaves one root for each factor. The change at 3 is already visible on a single progression at 2.

Let $r_k = \dim_{\mathbb{F}_2} \mathrm{Cl}(\mathbb{Q}(\sqrt{-d}))[2]$, with d from Section 2. For the two forms $A_1(k) = k$, $A_2(k) = 9k + 1$, put $U_i(k) = \#\{p \geq 5 : v_p(A_i(k)) \text{ is odd}\}$. Write $\mathcal{R}_{\mathrm{lin}}$ for the six classes of k mod 16 defining $\mathcal{R}$.

Proposition 3.3 (The integer difference on the full parameter line)**.** *Choose k uniformly in $[X, 2X)$, either without restriction or on $\mathcal{R}_{\mathrm{lin}}$. Put $\ell = \log\log X$. Then*

$$\left(\frac{U_1 - \ell}{\sqrt{\ell}}, \frac{U_2 - \ell}{\sqrt{\ell}}, s_k - r_k\right) \implies (G_1, G_2, D_{\mathrm{lin}}),$$

where the two standard normal variables and the integer D_{lin} are independent. With the dyadic polynomial $P_2(z) = a + b_1 z + c_2 z^2$ for the corresponding population in Theorem 3.1,

$$\mathbb{E}z^{D_{\mathrm{lin}}} = P_2(z)\frac{3 + z^{-1}}{4} \prod_{p \geq 5} \left(1 + \frac{\delta_p}{2}(z - 1)\right).$$

The generating functions converge locally uniformly for $z \neq 0$. The number of $k \in [X, 2X)$ with $e = 1$ is $O(X/\log X)$.

Proof. For any k, the genus formula and Theorem 2.1 give

$$s_k - r_k = t_2 - \mathbf{1}_{2|\Delta} - \mathbf{1}_{3|d} + J + 2e.$$

The dyadic distribution is the same in the two samplings: each k mod 2^j has two preimages among n mod 2^{j+1}. At 3, the second form is a unit, and the *single* root of the first gives $\mathbb{P}(3 \mid d) = (1 - 1/3)\sum_{j \geq 0} 3^{-2j-1} = 1/4$. At $p \geq 5$ the roots $0, -1/9$ are distinct. The number

of split roots is $\mathbf{1}_{\left(\frac{-2}{p}\right)=1}+\mathbf{1}_{\left(\frac{2}{p}\right)=1}=\nu_p/2$. Each root contributes $1/[p(p+1)]$ to the probability of positive even valuation.

For the exceptional class, the sieve in Lemma 3.2 now excludes $2(p-1)$ classes modulo p^2 at $p\equiv 5,7 \bmod 8$. Consequently $w(p)=2/p+O(p^{-2})$ at these primes. The same truncated-product argument gives $G(z)\gg \log z$, and the remainder is still $O_\epsilon(z^{4+\epsilon})$. Taking $z=X^{1/16}$ gives the asserted bound. A fixed progression of positive density costs at most a constant factor in this estimate.

A prime $p>Y$ counted by J satisfies $p^2\mid k$ or $p^2\mid 9k+1$. Counting the two progressions up to $p\le\sqrt{18X+1}$ gives an exceptional proportion $O(\sum_{p>Y}p^{-2}+X^{-1/2})$. The Chinese remainder theorem and the geometric tails at the fixed primes now give the stated generating function. The exponential-moment bound used for (7) applies to the two injective forms A_i, with values at most $18X+1$, and gives local uniform convergence.

For the normal coordinates, on any fixed residue class the two disjoint roots at a prime outside the modulus have masses $1/p, 1/p$. The covariance at that prime is $-1/p^2$. For any linear combination of the two counts, the moment comparison above applies with two sites in place of four. The contribution from repeated factors is bounded in mean and the discarded large primes cost $o(\sqrt{\ell})$. Thus the conditional limit on every fixed class is the same pair of independent normals. Approximating the integer coordinate by the finite residue conditions just used proves the joint assertion. □

For instance, fix $k\equiv 2 \bmod 16$. Then $h=1,\ w=3 \bmod 8$, so $t_2=\mathbf{1}_{2|\Delta}=1$. The identity in the proof shows that s_k and r_k have equal parity with limiting probability

$$\frac{1}{2}+\frac{1}{4}\prod_{p\ge 5}(1-\delta_p).$$

This lies between 0.7291961 and 0.7291963: the product differs from the one over $p>7$ bounded in Section 4.2 by the factor $27/28$. For triangular $k(n)$ in this same dyadic class, the answer is exactly $1/2$. The prime 3 accounts for the distinction: its parity bit is fair for $k(n)$ and has success probability $1/4$ for k. In both samplings the global term $2e$ drops out of this parity test.

4 What corrected Selmer parity predicts

4.1 Agreement on the triangular subfamily

Take $k=n(n+1)/2$. For a fixed $y\ge 3$, let $J_{\le y}$ count the primes contributing to J up to y, and define

$$S_y(n)=(-1)^{s(n)+1-t_2+\mathbf{1}_{2|d}+\mathbf{1}_{3|d}+J_{\le y}},\qquad M_y(n)=(-1)^{J-J_{\le y}}.$$

The sign of $D=4dt^2$ is $(-1)^{\omega(d)}$. Comparing with (3) gives the pointwise relation $S_y=CM_y$. In particular e disappears from this parity comparison.

Proposition 4.1. *In natural density, on all integers and on $\mathcal{R}$, the agreement probability is*

$$\mathbb{P}(S_y=C)=\frac{1+m_y}{2},\qquad m_y=\prod_{p>\max(3,y)}(1-2\delta_p). \tag{9}$$

Uncorrected Selmer parity $(-1)^s$ agrees with C with probability $1/2$ on both populations. The same conclusions hold on any windows of diverging harmonic mass, with weights $1/n$.

Proof. The product in (9) follows from the independent prime indicators and the square-divisor tail in (8). For the raw parity the exponent of $C(-1)^s$ modulo two is $t_2-1+\mathbf{1}_{2|d}+\mathbf{1}_{3|d}+J$. Its ternary bit is fair and independent of the other displayed local observations. Every bounded function used here is approximable in mean by fixed periodic functions, after prime and depth truncation. The dyadic estimates for the absolute errors also control harmonic windows. Indeed partition the interval into full dyadic intervals and at most two end intervals, each of harmonic mass $O(1)$. The finitely many small intervals disappear upon normalization, and the uniform tail estimate on the others persists. □

There is a useful distinction between agreement and the mean of C on $\mathcal{R}$. For $I=(U,V]$, let $\langle f\rangle_{I,\mathcal{R}}$ denote $\sum_{n\in I\cap\mathcal{R}} f(n)/n$ divided by $\sum_{n\in I\cap\mathcal{R}} 1/n$. An extraction is a sequence with $\log(V/U)\to\infty$ on which the finitely many means in the claim converge. Write $c_{\mathcal{R}}, s_{\mathcal{R},7}$ for the extracted conditional means of C, S_7 and c for the unconditional mean of C. In this notation

$$s_{\mathcal{R},7} = m_7 c_{\mathcal{R}} + \mathrm{Cov}_{\mathcal{R}}(C, M_7).$$

The covariance is taken in the same extracted law. The large agreement in (9) measures m_7; the signed dependence enters the other term.

Corollary 4.2. *Set* $B = 196567331/402653184$. *Along every extraction,*

$$|c_{\mathcal{R}} - s_{\mathcal{R},7}| \le B(1-m_7) < 0.0471173, \tag{10}$$

$$|c| \le \tfrac{3}{8}|s_{\mathcal{R},7}| + \tfrac{3}{8}B(1-m_7) + \tfrac{5}{8}B. \tag{11}$$

Thus $|s_{\mathcal{R},7}| \le 0.02813$ *on each extraction would suffice for a uniform bound* $|c|<1/3$.

Proof. Theorem 1.1 of *An explicit bound below one half for a quartic Liouville correlation* [20] bounds the absolute normalized harmonic mean of C on every fixed progression by B in the limit superior. Summing over residue classes bounds the signed average against any fixed nonnegative periodic function by B times its mean. Apply this to periodic approximations of $\mathbf{1}_{\mathcal{R}}(1-M_7)$ and then take their mean limit. Since $C-S_7 = C(1-M_7)$, this proves the first bound. The complement of $\mathcal{R}$ is periodic of density $5/8$, giving the second.

For the numbers, multiply the rational factors through $Y=10^7$ with outward rounding. Since $2\delta_p \le 8/(p(p+1))$, the remaining product lies between $1-8/(Y+1)$ and 1. This gives

$$0.90348390 < m_7 < 0.90348463, \qquad \tfrac{8}{9} - \tfrac{5}{3}B - B(1-m_7) > 0.0281378.$$

Inserting this in (11) proves the strict inequality under the stated sufficient condition. The bound for $s_{\mathcal{R},7}$ itself is open. □

4.2 A logarithmic balance on the full parameter line

With k as the sampling variable, the sign of $4k(9k+1)$ is $C_{\mathrm{lin}}(k) = \lambda(k)\lambda(9k+1)$. The two factors are affine in k. This lets us compute a signed distribution even after observing the integer difference $s_k - r_k$.

Write $\mathcal{J}(k)$ for the set of split primes of positive even valuation counted by J. Given a set $\mathcal{P}$ of primes at least five, define the corrected parity by

$$S_{\mathcal{P}}(k) = (-1)^{s_k+1-t_2+\mathbf{1}_{2|d}+\mathbf{1}_{3|d}+|\mathcal{J}(k)\cap\mathcal{P}|}.$$

As in Proposition 4.1, the pointwise identity is $S_{\mathcal{P}} = C_{\mathrm{lin}} M_{\mathcal{P}}$, where $M_{\mathcal{P}} = (-1)^{|\mathcal{J}(k)\setminus\mathcal{P}|}$. Write $\langle f\rangle_{x,\omega} = (\log\omega)^{-1}\sum_{x/\omega<k\le x} f(k)/k$.

Theorem 4.3 (Logarithmic balance after local corrections)**.** *Let $\mathcal{A}$ be a fixed arithmetic progression, let $1 \le \omega = \omega(x) \le x$, and suppose $\omega \to \infty$. Then*

$$\sup_{\mathcal{P}} \sup_{\|\Phi\|_\infty \le 1} \left| \langle \mathbf{1}_{\mathcal{A}}\, \Phi(s_k - r_k)\, S_{\mathcal{P}}(k) \rangle_{x,\omega} \right| \longrightarrow 0.$$

Here $\Phi : \mathbb{Z} \to \mathbb{C}$ and the set $\mathcal{P}$ are common to all parameters in a given interval. The same balance holds for the raw parity $(-1)^{s_k}$. In particular, on any fixed value of the difference having positive density in $\mathcal{A}$, even and odd full Selmer dimensions have equal logarithmic frequency.

For a fixed $\mathcal{P}$, put $m(\mathcal{P}) = \prod_{p \ge 5,\, p \notin \mathcal{P}} (1 - \delta_p)$. On the full line, and also conditionally on $\mathcal{R}_{\mathrm{lin}}$, the limiting logarithmic frequency of $C_{\mathrm{lin}} = \sigma,\ S_{\mathcal{P}} = \tau$ is

$$\frac{1 + \sigma\tau\, m(\mathcal{P})}{4} \qquad (\sigma, \tau \in \{-1, 1\}).$$

The agreement event alone has ordinary density $(1 + m(\mathcal{P}))/2$ on each of these populations.

Proof. On a progression $k = qj + a$, $1 \le a \le q$, the two factors have arguments $qj + a$ and $9qj + 9a + 1$; their determinant is $q \ne 0$. Tao's logarithmically averaged two-point theorem [25, Theorem 1.2, (1.3)] applies. Replacing $1/(qj + a)$ by $1/(qj)$ costs $O_q(1)$ before normalization. After omitting at most two initial terms take $v = \lfloor (x-a)/q \rfloor$, $u = \max(1, \lfloor (x/\omega - a)/q \rfloor)$. Then $\log(v/u) = \log\omega + O_q(1)$ and $v/u \to \infty$; use the theorem on $(u, v]$. Thus $\langle \mathbf{1}_{k = a \bmod q} C_{\mathrm{lin}} \rangle_{x,\omega} \to 0$. For fixed q, summing the absolute values over its residue classes gives convergence uniformly against all q-periodic weights of modulus at most one.

We explain why one fixed modulus suffices for the approximation regardless of $\mathcal{P}$ and Φ. Truncate the prime factors of $\mathcal{J}(k)$ at Y, and then all valuations used at $2, 3$ and at primes up to Y at depth H, keeping the three dyadic unit bits. Off the discarded classes and the set $e = 1$, the integer $s_k - r_k$ and the sign $M_{\mathcal{P}}$ are functions of residues modulo a single $q(Y, H)$. The modulus may also include the period of $\mathcal{A}$. On a dyadic interval $[X, 2X)$, the limiting upper proportion discarded is bounded by

$$O\left(\sum_{p > Y} p^{-2} \right) + O_Y(2^{-H}).$$

Indeed Proposition 3.3 bounds the large-square tail and the exceptional class, while the two linear forms have at most two roots at each odd prime power. The bound also holds up to an absolute constant on the harmonic windows in the statement. Partition into full dyadic intervals and two end pieces of bounded harmonic mass, then discard any fixed initial range before taking $x \to \infty$. The bounded periodic approximation to $\mathbf{1}_{\mathcal{A}} \Phi(s_k - r_k) M_{\mathcal{P}}$ therefore has an error tending to zero as Y and then $H = H(Y)$ are chosen large. This proves the uniform balance. For the raw parity multiply C_{lin} by $(-1)^{t_2 - 1 + \mathbf{1}_{2|d} + \mathbf{1}_{3|d} + J}$; the same truncations apply.

Finally $C_{\mathrm{lin}} = S_{\mathcal{P}}$ for the choice of all primes, so both C_{lin} and $S_{\mathcal{P}}$ have zero logarithmic mean. Their product is $M_{\mathcal{P}}$. The independent local indicators in Proposition 3.3 give it ordinary mean $m(\mathcal{P})$, also under the dyadic condition. Periodic approximation gives the same harmonic mean. Expanding the two sign indicators gives the four stated frequencies. □

Correcting through 7 gives agreement between 0.9753698 and 0.9753701 on the full parameter line, compared with 0.9517419 to 0.9517424 on the triangular subfamily. The first interval follows by multiplying $1 - \delta_p$ for $7 < p \le 10^7$: the omitted product is between $1 - 4/(10^7 + 1)$ and one. On the triangular subfamily the first sign has four affine arguments in n. The binary theorem used on the full parameter line supplies no estimate for the signed input to Corollary 4.2.

5 Cubic Selmer signs at the four linear forms

Return to the triangular parameters and the four affine forms in n. The first family compares one dimension with the squarefree kernel of D. We now attach a cyclic cubic field to *one integer* m. Its ramified primes will turn Selmer parity into a completely multiplicative function of m. We can then evaluate different parities either at the four sites of (1) or together at $m = 3n + 2$.

5.1 Fixing the local conditions of the cubic character

Let $E_0 : y^2 = x(x-1)(x+3)$ and $E_i = E_0^{(d_i)}$ for $(d_1, d_2, d_3, d_4) = (1, -1, 2, 5)$. The primes of bad reduction of these curves, together with 3, lie in $S = \{2, 3, 5\}$. For each $p \equiv 1 \bmod 3$, take the primitive cubic character ψ_p sending the least positive primitive root g_p to ζ_3. Use the following four auxiliary primes D_0. The last column gives the exponents of ζ_3 in $\psi_q(2), \psi_q(3), \psi_q(5)$.

q	g_q	$\#E_0(\mathbb{F}_q)$	exponents at $2, 3, 5$
241	7	224	$(1, 2, 0)$
601	7	640	$(0, 1, 1)$
1801	11	1792	$(0, 1, 2)$
2161	23	2144	$(2, 1, 1)$

Begin with ψ_{2161} and multiply by ψ_p for every $p \equiv 1 \bmod 3$ outside $S \cup D_0$ with $v_p(m)$ odd. Multiply further by powers of $\psi_{241}, \psi_{601}, \psi_{1801}$ so the character takes value one at $2, 3, 5$. The first three rows have determinant one over $\mathbb{F}_3$, so these three powers are unique modulo three. The conductor always contains 2161. The resulting character cuts out a nontrivial cyclic cubic field F_m in which $2, 3, 5$ split.

Let $b_i = \dim_{\mathbb{F}_3} \mathrm{Sel}_3(E_i/\mathbb{Q})$. The trace-zero construction of Mazur, Rubin and Silverberg [18, Theorems 5.5 and 5.8] gives the abelian surface and relative Selmer parity

$$A_{i,m} = \ker\Big(\mathrm{Res}_{F_m/\mathbb{Q}}(E_i) \xrightarrow{\mathrm{Tr}} E_i\Big), \qquad Z_i(m) = (-1)^{\dim_{\mathbb{F}_3} \mathrm{Sel}_{1-\zeta_3}(A_{i,m}/\mathbb{Q}) - b_i}.$$

The trace kernel is connected, as is seen over a splitting field by summing three copies of E_i. The order-three Galois action gives the indicated isogeny. Reversing the generator changes $1 - \zeta_3$ by a unit.

Lemma 5.1. *Each Z_i is a completely multiplicative sign. At $p \notin S \cup D_0$ with $p \equiv 1 \bmod 3$, write F_p for a Frobenius matrix on $E_0[3]$. Then $Z_i(p) = -1$ exactly when $(d_i/p)F_p$ is nonidentity unipotent. At every remaining prime $Z_i(p) = 1$.*

Proof. Apply the global parity comparison and its identification with the isogeny Selmer group in [14, Theorem 4.11, Corollary 4.12 and Proposition 5.9] to each $E_i[3]$, comparing the cubic character with the trivial one. The local conditions agree at S and at the real place. At good primes away from 3, an unramified character preserves the unramified condition. A ramified character contributes the local width $\dim E_i(\mathbb{F}_p)[3] \bmod 2$. At the four auxiliary primes every d_i is a square and the displayed cardinalities are prime to 3, so these widths vanish. For the remaining ramified primes the Weil pairing gives $\det F_p = 1$. A matrix of determinant one in characteristic three has a one-dimensional fixed space precisely when it is nonidentity unipotent, agreeing with the width criterion in [14, Lemma 4.3]. The width of the twist tests the matrix $(d_i/p)F_p$. Ramification at these primes is selected by $v_p(m) \bmod 2$, which proves the multiplicative formula. □

5.2 Separating the four diagonal parities

The action on the four lines of $E_0[3]$ has division polynomial $3x^4+8x^3-18x^2-9$. Its reduction modulo 5 is irreducible. Modulo 13 it is $3(x+5)(x^3+2x^2-3x+2)$ with irreducible cubic. The projective group therefore contains a 4-cycle and a 3-cycle. Its order is divisible by 12, and the index-two subgroup A_4 contains no 4-cycle. Hence the projective group is S_4. A lift of a double transposition has determinant one (the determinant character on $\mathrm{PGL}_2(\mathbb{F}_3)$ is the sign), and square $-I$: a determinant-one matrix with square I in odd characteristic is scalar. The full image is thus $\mathrm{GL}_2(\mathbb{F}_3)$. Its commutator subgroup is $\mathrm{SL}_2(\mathbb{F}_3)$, since conjugation by a diagonal sign matrix inverts elementary transvections and these generate SL_2. The maximal abelian subfield of $\mathbb{Q}(E_0[3])$ is $\mathbb{Q}(\sqrt{-3})$.

Proposition 5.2. *On every limiting logarithmic law with* $\log(V/U)\to\infty$*, the four variables* $Z_i(L_i(n))$ *are independent fair signs conditionally on the four* $\lambda(L_j)$*, the profinite residue of* n*, and all twelve products* $Z_i(L_j)Z_i(L_i)$ *with* $j\neq i$*.*

Proof. Test a finite residue class and a Walsh character $\prod_{ij} Z_i(L_j)^{e_{ij}}\prod_j \lambda(L_j)^{a_j}$. If some row has odd sum, let $I\neq\varnothing$ be the set of such rows. The product of the four site functions at a prime is a constant sign times $\prod_{i\in I} Z_i(p)$. Fix a Dirichlet character ξ, and take Q divisible by its conductor and 120. When $|I|$ is odd, the class 1 mod Q makes an odd number of the symbols (d_i/p), $i\in I$, positive. When $|I|$ is even, the square class $\prod_{i\in I} d_i$ is nontrivial and independent of -3; choose a reduced class modulo Q splitting $\mathbb{Q}(\sqrt{-3})$ with this product of symbols negative. Again an odd number are positive. Such a class exists by independence of these quadratic characters.

The division field meets $\mathbb{Q}(\zeta_Q)$ exactly in $\mathbb{Q}(\sqrt{-3})$. In the same chosen class, Chebotarev therefore realizes both the identity and a nonidentity unipotent on $E_0[3]$ on positive-density sets of primes. They give opposite signs to the site product. For any ξ, the sum of $1-\Re(g(p)\overline{\xi(p)})$ weighted by $1/p$ is thus $\gg_\xi \log\log x$, where g is the product of site functions. This excludes weak pretence to every Dirichlet character. On a fixed class $n=qv+b$, complete multiplicativity permits replacing the first two arguments by $3n, 3(n+1)$ at the cost of two constant signs. All four arguments now have slope $3q$ and distinct shifts. The harmonic weight in v differs from the rescaled weight in n by a summable error. The vanishing assertion of Tao and Teräväinen [26, Theorem 1.2 and Appendix A, Theorem A.1], applied with dilation parameter one, therefore annihilates this mode. Passing from cylinder tests to the profinite residue uses bounded approximation. Only modes invariant under independent sign reversals of the four rows remain. Those reversals preserve the within-row products and act freely and transitively on the sixteen diagonal assignments, proving the proposition. □

The same-site products $Z_+=Z_1Z_3$ and $Z_-=Z_2Z_3$ retain different information when sampled at $m=3n+2$. Write $\chi=\chi_{-3}$ and take $\eta=(2/\cdot)$ for Z_+, $\eta=(-2/\cdot)$ for Z_-, with both η's primitive of conductor 8. Put $h_\pm=\chi\lambda Z_\pm$. For $p\notin S\cup D_0$, fix its reduced residue class modulo any fixed multiple of 120. Average the remaining Frobenius data by Chebotarev in the compositum with that cyclotomic field. In determinant -1 the value $Z_\pm(p)$ is one. In determinant one, the uniform projective element of A_4 is a 3-cycle with probability 8/12. On these elements the two possible lifts have a unique choice of sign giving a nonidentity unipotent. Thus the product of the two twist marks equals $\eta(p)$. On the other four projective elements both marks are $+1$. The conditional mean is

$$\mathbb{E}(Z_\pm(p)\mid \text{cyclotomic data})=\tfrac{2}{3}-\tfrac{1}{3}\chi(p)+\tfrac{1}{3}\bigl(\eta(p)+\eta(p)\chi(p)\bigr). \tag{12}$$

As $\lambda(p)=-1$, the only primitive Dirichlet characters ψ for which $h_\pm\overline{\psi}$ has nonzero prime mean, and those means z_ψ, are

ψ	1	η	$\eta\chi$	χ
z_ψ	$1/3$	$-1/3$	$-1/3$	$-2/3$

Here the principal primitive character has conductor one. The finitely many exceptional Euler factors will be kept with their actual values.

6 How an elliptic root number removes three signs

Write $t = 3n,\ m = t + 2$ and let $W(n)$ be the global root number of

$$G_n : \quad y^2 = x(x - A)(x - B), \qquad A = t + 3, \quad B = 3(t + 1).$$

In this section A, B denote the roots of G_n. Helfgott's factorization isolates a Liouville contribution and local corrections supported at square divisors [11, Proposition 6.5, Theorem 6.6 and (6.19)]. For this model the three collision divisors are linear, and we compute their angular signs at the two small places.

6.1 The pointwise identity

Put $g_0 = 3,\ g_1 = -2,\ g_3 = 6$. For $z > 0$ write $z_{(p)} = z/p^{v_p(z)}$. On odd units put $\epsilon_4 = (-1/\cdot)$ and $\epsilon_{-8} = (-2/\cdot)$. Define

$$D_2(n) = W_2(G_n)(-1)^{v_2(t+1)}\epsilon_4(t_{(2)})\epsilon_{-8}((t+1)_{(2)})\epsilon_{-8}((t+3)_{(2)}), \tag{13}$$

$$D_3(n) = (-1)^{v_3(n+1)}\chi(n_{(3)}(n+1)_{(3)}). \tag{14}$$

Values at the isolated zero of a linear form in a local integral may be assigned arbitrarily.

Lemma 6.1. *For every $n \geq 1$,*

$$C(n)W(n) = -\lambda(m)D_2(n)D_3(n) \prod_{\substack{i\in\{0,1,3\},\, p>3 \\ v_p(t+i)>0 \text{ even}}} [-(g_i/p)]. \tag{15}$$

Moreover $\int_{\mathbb{Z}_3} D_3(n)\, dn = -1/3$, for Haar probability measure.

Proof. The discriminant is $16A^2B^2(A - B)^2$. For $p > 3$ the three collisions occur on $t, t + 1, t + 3 = 0$, with tangent square classes $3, -2, 6$, respectively. The nodal local root number is $-(g_i/p)$. For a prime of positive valuation v in $t + i$, multiplying this by $(-1)^v$ gives $(g_i/p)^v$ and an extra factor $-(g_i/p)$ precisely when v is even. At 3, untwisting by 3 leaves roots $0, n + 1, 3n + 1$, with good or multiplicative reduction according to n mod 3. Its underlying Weil representation has unramified semisimplification. Twisting by the ramified quadratic character kills the inertia invariants, and the product of the two quadratic Gauss factors has sign $(-1/3) = -1$. This also covers the multiplicative case since the monodromy correction on inertia invariants is then trivial.

For the signs introduced by reciprocity, write $t + i = 2^{a_i}3^{b_i}u_i$, $(u_i, 6) = 1$. On u_i the three Jacobi symbols are $\epsilon_4\chi,\ \epsilon_{-8},\ \epsilon_{-8}\chi$, in this order. Moving the factors 3^{b_i} into the binary units and the factors 2^{a_i} into the ternary units uses $\epsilon_4(3) = -1,\ \epsilon_{-8}(3) = 1,\ \chi(2) = -1$. Including $\lambda(2^{a_i}3^{b_i})$, the resulting exponent modulo two is $a_1 + b_3$, since $b_1 = 0$. As $b_3 = 1 + v_3(n + 1)$, its constant sign cancels $W_3 = -1$. The real sign contributes the minus in (15). Finally $\lambda(t)\lambda(t + 3) = \lambda(n)\lambda(n + 1)$. At 3, $n = 1$ mod 3 contributes $-1/3$ to the integral. In each other class, at any fixed depth of the vanishing form the angular component has mean-zero character and all remaining factors are constant. □

6.2 The dyadic integrals

Let $j = v_2(m)$, $m_0 = m2^{-j}$. For $\eta = (\pm 2/\cdot)$ and $e' = 0, 1$ set

$$A_{e'}(\eta) = \int_{\mathbb{Z}_2} D_2(n)\eta(m_0)(-1)^{e'j}\, dn.$$

Lemma 6.2. *In Haar probability on $\mathbb{Z}_2$, $\int_{\mathbb{Z}_2} D_2 = 0$. The four ramified integrals are*

η	$A_0(\eta)$	$A_1(\eta)$
$(2/\cdot)$	$-1/8$	$1/8$
$(-2/\cdot)$	$1/8$	$-1/8$

Proof. First remove the three balls $16 \mid n,\ 64 \mid n+1,\ 64 \mid 3n+1$. If v is the valuation of the corresponding form, then $v_2(j(G_n)) = 6 - 2v$ on the first ball and $10 - 2v$ on the other two, away from their centres. Indeed $c_4 = 16(A^2 - AB + B^2)$ and the discriminant was given above. Moreover

$$-c_6 = -1728(2n+1)(n-1)(5n+1)$$

has constant square class $3, 6, -2$ on the respective balls, by reducing the three unit parts modulo 8. For potentially multiplicative reduction the splitting character is determined by $-c_6$ and determines the local root number [22, Theorem 2(ii)]. Thus W_2 is constant on each ball. On each shell exactly one of the three angular characters in (13) is evaluated on the free deep unit, with zero integral. The remaining signs, j, and $\eta(m_0)$ are constant on that shell. The omitted balls have zero contribution to all five integrals.

On the complement we give a finite determination *on classes*, so that checking representatives suffices for the final sums. Put $a' = -(A+B)$, $b' = AB$. For the 2-isogeny with kernel $(0,0)$, the local formula of the Dokchitsers [5, Theorem 5.7] gives $W_2 = (-1)^{s_2-1}(a', -b')_2$, where s_2 is the dimension of the Kummer image on the target. Its second Hilbert factor is trivial since $a'^2 - 4b' = (A-B)^2$. For each of the eight local classes $d_0 = 1, 3, 5, 7, 2, 6, 10, 14$, the corresponding homogeneous space tests primitive dyadic pairs u, v in

$$z^2 = d_0^3 u^4 + 12(2n+1)d_0^2 u^2 v^2 + 36n^2 d_0 v^4. \tag{16}$$

When v is even, u is odd. An even d_0 gives valuation three, and an odd d_0 gives unit d_0 mod 8, forcing the trivial class. Otherwise scale to $v = 1$. At precision 2^N a residue of even valuation at most $N-3$ with odd part 1 mod 8 certifies solubility for every lift of n; absence of any square residue excludes the class. Count certain and possible classes and keep only powers of two between the counts. The Hilbert symbol and the unit characters are evaluated only when their three unit bits are determined.

At $N = 13$ the counts of negative, unresolved and positive D_2 are $(1632, 64, 608)$ on $\mathcal{R}$ outside the balls and $(2048, 0, 3072)$ off $\mathcal{R}$. The four lifts of the 64 unresolved classes to $N = 15$ give 128 of each sign. Consequently the integrals on $\mathcal{R}$ and its complement are $-1/8, +1/8$ before conditional normalization. Grouping the same representatives by $j = v_2(3n+2)$ gives

Shell	$j=0$	$j=1$	$j=2$	$j \geq 3$
Signed sum of D_2 (denominator 32768)	0	0	4096	-4096

The deep balls lie in $j = 0, 1$ and integrate to zero separately. Each of the last two columns has 4096 representatives. Consequently $D_2 = 1$ throughout $j = 2$ and $D_2 = -1$ throughout $j \geq 3$. These shell values will determine the coefficient after the leading term. For the ramified table, multiply each resolved class modulo 2^{15} by its two weights if $j \leq 12$. For $j > 12$ the angular integral is zero: at fixed j, the residue of m_0 is fixed to precision at most 4, and both primitive characters change sign on adding 4 mod 8. The resulting numerator matrix is $4096\left(\begin{smallmatrix} -1 & 1 \\ 1 & -1 \end{smallmatrix}\right)$ with denominator 32768. The complete residue enumeration for (16) is in the sources. Direct evaluation of W_2 by PARI on the 29696 ordinary representatives gives the same matrix. □

7 The first surviving term of the weighted mean

Let $K(n)$ be the product of the primes $p > 3$ for which $v_p(D(n))$ is positive and even. The empty product is one. At most one of the four sites can be divisible by any such prime. For numbers $0 \le \theta_p \le 1$, fixed in advance for $p > 3$, form the weight

$$P_{\boldsymbol{\theta}}(n) = \prod_{p \mid K(n)} \theta_p.$$

Thus all $\theta_p = 1$ keeps the full population, all $\theta_p = 0$ gives $\mathbf{1}_{K=1}$, and $\theta_p = \mathbf{1}_{p \le y}$ restricts every prime of K to $p \le y$. For a bounded function P define the ordinary interval means

$$B_{\pm}(P; X) = \frac{1}{X} \sum_{X \le n < 2X} P(n) C(n) W(n) Z_{\pm}(3n+2).$$

The variable of summation is n, with each origin counted once.

For a fixed choice of $Z_{\pm}$, abbreviate $h = h_{\pm}$, let $\psi_e = \eta\chi^e$, $f_e = h\psi_e$ for $e = 0, 1$, and interpret χ^0 as the character of conductor one. For $p > 3$ let $u_p = \#\{i \in \{0,1,3\} : (g_i/p) = 1\}$. Define

$$c_e = \frac{1}{\Gamma(-1/3)} \prod_p \frac{(1-1/p)^{-1/3}}{1 - f_e(p)/p}, \qquad \kappa_{p,e} = \frac{1 - f_e(p)/p}{p^2 - 1}. \tag{17}$$

The product for c_e uses increasing prime order. Its local factors are positive. Chebotarev in the fixed field and partial summation give a nonzero limit for their product. Thus $c_e < 0$.

Theorem 7.1. *For each $\boldsymbol{\theta}$ as above,*

$$B_{\pm}(P_{\boldsymbol{\theta}}; X) = \mp b_{\pm}(\boldsymbol{\theta}) (\log X)^{-4/3} + O((\log X)^{-5/3}), \tag{18}$$

$$b_{\pm}(\boldsymbol{\theta}) = -\frac{1}{8} \sum_{e=0}^{1} c_e \prod_{p>3} [1 - 4\kappa_{p,e} + (4 - 2u_p)\theta_p \kappa_{p,e}] > 0. \tag{19}$$

The error constant can be chosen independently of $\boldsymbol{\theta}$. Omitting $Z_{\pm}$ gives $O_A((\log X)^{-A})$ for each fixed $A > 0$, uniformly in these weights.

At $\theta_p = 0$ the factor is $1 - 4\kappa_{p,e}$, and at $\theta_p = 1$ it is $1 - 2u_p\kappa_{p,e}$. Both are positive because $u_p \le 3$ and $\kappa_{p,e} \le 1/[p(p-1)]$. Absolute convergence at $p > 3$ then gives strict positivity of the product. The constants for the two choices of parity use their respective f_e.

7.1 Prime means in a varying progression

We give the analytic argument through arbitrary fixed order because prescribing exact valuations can remove several leading coefficients. By (15) and $\chi(3n+2) = -1$, the remaining factor $-\lambda(m)Z_{\pm}(m)$ is $h(m)$. Outside a fixed finite set, its prime values form a class function in the fixed compositum of the 3-division field with the quadratic fields defining the twists. The Frobenian framework is described in [17, §2]. The additional issue here is uniformity while the local conditions introduce moduli that grow as powers of $\log X$.

We first specify the analytic input. The number field in this statement is fixed; only the finite-order character and its conductor vary.

Lemma 7.2 (Hecke–Siegel–Walfisz in a fixed field)**.** *Let F be a fixed number field. For every $A, B > 0$, uniformly over finite-order Hecke characters ξ whose finite conductor $\mathfrak{f}_\xi$ satisfies $\mathrm{N}\mathfrak{f}_\xi \le (\log t)^B$,*

$$\sum_{\substack{\mathfrak{p} \text{ prime} \\ \mathrm{N}\mathfrak{p} \le t}} \xi(\mathfrak{p}) \log \mathrm{N}\mathfrak{p} = \mathbf{1}_{\xi=1}\, t + O_{F,A,B}\big(t(\log t)^{-A}\big).$$

The character is zero on primes dividing its conductor. The same conclusion holds for a character induced to a modulus $\mathfrak{q}$ with $\mathrm{N}\mathfrak{q} \le (\log t)^B$. The constants may be ineffective.

Proof. This is the fixed-field, finite-order form of the Hecke–Siegel–Walfisz theorem of Goldstein [8]. A formulation with all the required uniformity is also obtained from [1, Theorem 1.3 and (4.11)], taking both representations to be the trivial character of $\mathrm{GL}_1(\mathbb{A}_F)$. For completeness, the four hypotheses of that theorem have the following interpretation here. The short-interval bound for ideal von Mangoldt coefficients follows from

$$\sum_{t<\mathrm{N}\mathfrak{a}\le t+u}\Lambda_F(\mathfrak{a}) \ll_F u\log t + t^{1/2}(\log t)^2 \ll_F u\log t \qquad (u\ge t^{3/4}).$$

Indeed prime ideals of residue degree one give at most $[F:\mathbb{Q}]$ terms above each rational prime, and the remaining prime powers contribute the second term. The Hecke L-functions are holomorphic except for the pole of the principal character; the required zero-free region and Siegel bound are [1, Corollary 2.5 and Theorem 2.12]. Their formula (4.11) therefore gives a uniform estimate for the sum over prime ideal powers. Removing powers costs $O_F(t^{1/2}(\log t)^2)$. Inducing a character only removes primes dividing $\mathfrak{q}$, whose total logarithmic weight is $O(\log \mathrm{N}\mathfrak{q})$. □

Put $\Psi=\{1,\eta,\eta\chi,\chi\}$, with prime means z_ψ as in (12), and set

$$f_\psi = h\overline{\psi}, \qquad \sum_{m\ge1}\frac{f_\psi(m)}{m^s}=\zeta(s)^{z_\psi}G_\psi(s) \quad (\Re s>1), \qquad Q_{\psi,q}(s)=\prod_{p\mid q}(1-f_\psi(p)p^{-s}).$$

All four characters in Ψ are real. We retain conjugation to make the progression convention explicit.

Lemma 7.3 (Uniform progression expansion)**.** *For every $A,B>0$, there is an integer $J=J(A,B)$ such that, uniformly for $q\le(\log x)^B$ and $(a,q)=1$,*

$$\sum_{\substack{m\le x\\ m\equiv a \bmod q}} h(m) = \frac{x}{\varphi(q)}\sum_{\substack{\psi\in\Psi\\ \mathrm{cond}(\psi)\mid q}}\psi(a)\sum_{j=0}^{J}\frac{\gamma_{\psi,j}(q)}{\Gamma(z_\psi-j)}(\log x)^{z_\psi-1-j}+O_{A,B}\big(x(\log x)^{-A}\big),$$

$$\gamma_{\psi,j}(q)=[w^j]\left\{\frac{(w\zeta(1+w))^{z_\psi}G_\psi(1+w)Q_{\psi,q}(1+w)}{1+w}\right\}.$$

Here $[w^j]$ denotes the coefficient determined by the derivatives from the right at $w=0$. The same conclusion can be used after extracting $m=dv$ with $d\le(\log x)^B$: on a reduced residue class of the residual variable v, complete multiplicativity supplies $h(d)$, and the upper endpoint becomes x/d.

Proof. Character orthogonality reduces the claim to twists $h\overline{\psi}$ by characters modulo q. If the primitive character inducing ψ is outside Ψ, the prime mean is zero. Uniformly for these twists and every $A>0$,

$$\sum_{p\le t}h(p)\overline{\psi(p)}\log p \ll_{A,B} t(\log t)^{-A} \qquad (\exp(\sqrt{\log x})\le t\le x). \tag{20}$$

To obtain this estimate, use Brauer induction to express the fixed Frobenius class function as a finite complex linear combination of induced linear characters. Each term, twisted by $\overline{\psi}$, is a finite-order Hecke character over one of finitely many fixed fields. Its conductor divides a fixed ideal times $q\mathcal{O}_F$, so its norm is $O_F(q^{[F:\mathbb{Q}]})$. Since $q\le(\log t)^{2B}$, Lemma 7.2 applies with a fixed enlarged exponent. The trivial-character terms add to the prime mean, which is zero in the present case. Prime powers and residue degrees at least two cost $O(t^{1/2}(\log t)^{O(1)})$; the fixed ramified primes and the primes dividing q have smaller cost.

The same arbitrary logarithmic saving holds for the sums over integers. Put $y=\exp(\sqrt{\log x})$ and write every non-y-smooth integer uniquely as $m=\ell p$, where p is its largest

prime factor. The inner sum in p runs from $\max(y, P^+(\ell))$ to x/ℓ, with the equality case $p = P^+(\ell)$ allowed. Complete multiplicativity and partial summation in (20) bound it by $O_A((x/\ell)(\log y)^{-A})$, after increasing the prime-sum precision. The sum of $1/\ell$ costs $O(\log x)$. Rankin's bound at $1 - 1/\log y$ gives

$$\#\{m \le x : P^+(m) \le y\} \ll x \exp(-\sqrt{\log x})(\log x)^{O(1)},$$

since $\prod_{p\le y}(1 - p^{-1+1/\log y})^{-1} \ll (\log y)^{O(1)}$. Endpoint primes cost $O(x/y)$. This treats all zero-mean twists, including imprimitive ones.

For each of the four fixed primitive twists, Chebotarev gives $\sum_{p\le x}(f_\psi(p) - z_\psi)\log p \ll_A x(\log x)^{-A}$ for every fixed A. Since $|f_\psi| \le 1$, the conditions of [2, Theorem 1.2] hold with $r = 1$ and any $1/2 < \sigma < 1$. The Selberg–Delange expansion therefore applies to every fixed order, including the negative means $-1/3$ and $-2/3$; its coefficients are given by [2, (1.10)].

For an induced character modulo q, removing the Euler factors at $p \mid q$ multiplies the Dirichlet series by $Q_{\psi,q}$. Finite Möbius inversion expands this product into terms indexed by divisors of $\mathrm{rad}(q)$, all at most q. Apply the fixed expansion at the resulting endpoints x/d. The number of terms and their derivatives have a polylogarithmic bound, and $\log(x/d) = \log x + O_B(\log\log x)$. Taking the fixed expansion and J to sufficiently high order gives the claimed uniform remainder. Orthogonality supplies $\psi(a)/\varphi(q)$. The last assertion follows directly from $h(dv) = h(d)h(v)$. □

More generally, an additional local factor $U(s)$ gives the coefficient of $x(\log x)^{z_\psi - 1 - j}$

$$\frac{[w^j]\{(w\zeta(1+w))^{z_\psi} G_\psi(1+w)U(1+w)/(1+w)\}}{\Gamma(z_\psi - j)}. \tag{21}$$

We use right derivatives for the fixed factor G_ψ; no holomorphic extension of that factor to a fixed disc is needed. The local factors U introduced below are holomorphic on a common disc, which is the separate property used to control their derivatives. The leading fixed constant is the increasing-prime product $G_\psi(1) = \prod_p (1 - 1/p)^{z_\psi}/(1 - f_\psi(p)/p)$.

7.2 Passing the expansion through the square conditions

The preceding lemma handles finitely many congruences. We now justify retaining its expansion after all positive even valuations have been considered. This requires an absolute bound on the expanded sum, as well as convergence of the coefficient factors.

Lemma 7.4 (Uniform transfer through the local conditions)**.** *For every fixed $A > 0$, the averages $B_\pm(P_{\boldsymbol{\theta}}; X)$ admit, up to an error $O_A((\log X)^{-A})$, an expansion to sufficiently high finite order in the four scales*

$$(\log X)^{z_\psi - 1 - j}, \qquad \psi \in \Psi, \quad j \ge 0.$$

The error and each coefficient of fixed order are bounded uniformly for $0 \le \theta_p \le 1$. The coefficients are obtained from (21) by imposing the local conditions and multiplying by the interval factor $6^s - 3^s$. The additional local factors and their derivatives are limits of the finite truncations on a common disc about $s = 1$.

The same assertion holds for $P = \mathbf{1}_{K=a}$, for any fixed squarefree a prime to 6, with constants allowed to depend on a. No uniformity as a grows with X is asserted.

Proof. At $p > 3$, the multiplier of $P_{\boldsymbol{\theta}}$ times the last product in (15) is one unless a site has positive even valuation. It is θ_p on the m-site and $-(g_i/p)\theta_p$ on the other three. Open each event using

$$\mathbf{1}_{v_p(L)>0 \text{ even}} = \sum_{a\ge 2} (-1)^a \mathbf{1}_{p^a \mid L}.$$

The coefficient at a selected prime has modulus at most two, and at most one site can be selected. Retain assignments whose product of prime powers is at most $H = (\log X)^{B_1}$. If a discarded assignment occurs, at least one L_i has a powerful divisor $v > H^{1/4}$, meaning that every prime exponent of v is at least two. Injectivity of L_i and $\sum_{v \text{ powerful}} v^{-3/4} < \infty$ give probability $O(H^{-1/16})$.

The total absolute expansion at n is bounded by $\prod_i T(L_i)$, where T is multiplicative and $T(p^a) = 1 + 2\max(a-1, 0)$. Every fixed power T^b has uniformly bounded mean on each L_i: its convolution with μ has nonnegative coefficients, supported at exponents at least two at every occurring prime, with summable reciprocal weights. Sum first over all integers up to $6X + 2$, then use Hölder and Cauchy–Schwarz. The mean absolute contribution of the discarded assignments is $O(H^{-1/32})$, uniformly in $\boldsymbol{\theta}$.

At $2, 3$, discard valuations of the four forms above $C_0 \log\log X$. Their interval probability is $O((\log X)^{-cC_0} + X^{-1})$, with $c > 0$ absolute. This truncation is made in the bounded original product, before expanding the other primes. The retained local functions have periods $2^{O(1+C_0 \log\log X)}$ and $3^{O(1+C_0 \log\log X)}$. Choose at least three unit bits beyond the deepest retained $v_2(m)$, so both characters of conductor 8 remain visible after extraction. Split off the powers assigned to the m-site. At other sites the roots are distinct from $m = 0$, and the remaining congruences prescribe unit classes. At 3, m is always a unit. The moduli, divisors and number of terms are polylogarithmic. Lemma 7.3 applies at the endpoints $3X + 2$ and $6X + 2$; the shifts by two are harmless in the bounded original sum.

We next control the restored coefficients. On $|s-1| \le 1/8$, an assignment p^a at the moving site costs at most $p^{-7a/8}$; at a unit site it costs at most $2/\varphi(p^a)$, including Euler removal. With the bounded coefficients and four possible sites, the absolute sums are dominated by

$$\sum_{d \text{ powerful}} C_1^{\omega(d)} d^{-7/8} \prod_{p \mid d} (1 - 1/p)^{-1} < \infty.$$

The corresponding series at exponent $3/4$ also converges, so the tail $d > H$ is $O(H^{-1/8})$ on the disc. At two, a shell $v_2(m) = j$ contributes $2^{-j\Re s}$ times a bounded unit average. The number of unit classes cancels their totient denominator, and the deep tail is geometric. Near each other root, m is a unit or has bounded valuation, while the deleted class has mass $O(2^{-C_0 \log\log X})$. The same argument at three uses only units. Cauchy's inequalities control every fixed number of derivatives on a smaller disc.

The choices are made in this order: fix the desired precision A; choose B_1, C_0 so the absolute truncation errors and the coefficient tails are smaller than that precision; then choose the order and remainder strength in Lemma 7.3 to absorb the polylogarithmic number of retained terms. Summing the coefficients before estimating them uses the uniform disc bound just proved. Thus the intermediate coefficients remain bounded independently of the truncation and of $\boldsymbol{\theta}$. This proves the expansion and permits multiplication of the local factors. The change of endpoints from the prefix sum to the interval contributes $6^s - 3^s$, whose value at one is three.

Finally $\mathbf{1}_{K=a}$ is a finite difference of product weights: allow or forbid each of the finitely many primes dividing a, and forbid all other primes greater than three. Apply the preceding conclusion to these finitely many terms. Its constants may depend on a, as stated. □

7.3 The coefficient on the full interval

At a prime $p > 3$ let $f = h\overline{\psi}$ be any of the four fixed twists; $f(p) \in \{1, -1\}$. On $v_p(m) = 2j > 0$ extraction of the exact valuation contributes $(1 - f(p)p^{-s})p^{-2js}$. At each of the other three roots, the unit mass of its even valuations after Euler removal is $(1 - f(p)p^{-s})/(p^2 - 1)$, by summing the differences of $1/\varphi(p^j)$ for exact depths. Thus at $s = 1$ all four masses are $(1 - f(p)/p)/(p^2 - 1)$. The sum of the three signs on the unit branches is $3 - 2u_p$. This yields

the local factor in (19). Its analytic version, needed below, is

$$1+(1-f(p)p^{-s})\left\{\frac{\theta_p-1}{p^{2s}-1}+\frac{(3-2u_p)\theta_p-3}{p^2-1}\right\}. \tag{22}$$

For the principal mode $z=1/3$, $h(2)=1$. Its normalized mass at 2 is Haar, and its coefficient of $(\log X)^{-2/3}$ vanishes by Lemma 6.2. Its term of the next order and the first possible term from χ are $O((\log X)^{-5/3})$.

For $\psi_e=\eta\chi^e$, $f_e(2)=0$ but $h(2)=1$. In the shell $m=2^jm_0$ the factor $2^{-j}/\varphi(2^{N-j})$ for a fixed unit class gives twice its Haar mass at $s=1$. Its η-factor is $\eta(m_0)$. The χ^e-factor is $(-1)^{e(j+1)}$, since $m\equiv 2 \bmod 3$. This leaves the dyadic integral $2(-1)^eA_e(\eta)$. At 3, $h(3)=0$; reduced classes have full mass divided between the two units modulo 3. The class $m=2 \bmod 3$ has mass $1/2$ and the conditional D_3-integral in the variable n is $-1/3$. Including the interval-length factor 3, the common ternary and interval multiplier is $-1/2$. Thus the two coefficients are weighted by $-A_0(\eta)$ and $+A_1(\eta)$, respectively. Lemma 6.2 gives $+1/8,+1/8$ for Z_+, and $-1/8,-1/8$ for Z_-. Combining with the negative c_e proves (18).

When the cubic mark is omitted, use λ in the conditional progressions arising from (15). The ordinary Siegel–Walfisz estimate for Möbius gives arbitrary logarithmic saving in the reduced classes of polylogarithmic modulus. It gives the same estimate for λ via $\lambda(m)=\sum_{d^2|m}\mu(m/d^2)$, truncating d at a sufficiently large power of $\log X$. The preceding truncation in the four sites then gives the asserted bound for $P_{\boldsymbol{\theta}}CW$.

8 Prescribing every even-valuation prime

Fix a squarefree $a\ge 1$ with $(a,6)=1$. The condition $K=a$ forces four competing branches at each $p\mid a$: the even valuation can occur on the remaining Liouville site m or at one of the three multiplicative collisions. These branches have equal mass at $s=1$. Their signs can sum to zero. To determine the first term one must then compare how their masses vary with s.

Define

$$k(a)=\#\{p\mid a: p\equiv 11,19,23 \bmod 24\},\qquad r_1(a)=\#\{p\mid a: p\equiv 1 \bmod 24\}.$$

The local count from (17) is as follows.

$p \bmod 24$	1	5	7	11	13	17	19	23
u_p	3	1	0	2	1	1	2	2
$4-2u_p$	-2	2	4	0	2	2	0	0

The constant involves the four functions $f=h,h\chi,f_0,f_1$ for a given parity. Write $\kappa_p(f)=(1-f(p)/p)/(p^2-1)$ and $H_f(\mathbf{1}_{K=a})=\prod_{p>3}t_{p,f}(a)$, where

$$t_{p,f}(a)=\begin{cases}1-4\kappa_p(f), & p\nmid a,\\ |4-2u_p|\kappa_p(f), & p\mid a,\ u_p\neq 2,\\ (1-f(p)/p)\,2p^2\log(p)/(p^2-1)^2, & p\mid a,\ u_p=2.\end{cases} \tag{23}$$

The third line is the magnitude of the derivative at a prime where the four signed contributions cancel.

Theorem 8.1. *For either parity and every fixed squarefree a prime to 6,*

$$B_\pm(\mathbf{1}_{K=a};X)=\mp(-1)^{r_1(a)}b_{\pm,a}(\log X)^{-4/3-k(a)}+O_a((\log X)^{-5/3-k(a)}),\qquad b_{\pm,a}>0. \tag{24}$$

For $(4/3)_k=\prod_{j=1}^k(j+1/3)$, with $(4/3)_0=1$, the coefficient is

$$b_{\pm,a}=-\frac{(4/3)_{k(a)}}{8}\sum_{e=0}^{1}c_eH_{f_e}(\mathbf{1}_{K=a}). \tag{25}$$

Proof. To force a prime $p \mid a$, take the difference of allowing and forbidding it in (22). Forbidding every $p \nmid 6a$ at the same time gives the exact condition $K = a$. Equivalently take a finite difference at each of the finitely many $p \mid a$, starting with the local weight identically zero above 3. All expansions and arbitrary-order remainders proved in Section 7 apply, with constants allowed to depend on a.

At each of the four modes and for $f = h\overline{\psi}$ the forced factor is

$$V_{p,f}(s) = (1 - f(p)p^{-s}) \left(\frac{1}{p^{2s} - 1} + \frac{3 - 2u_p}{p^2 - 1} \right). \tag{26}$$

At $s = 1$ it equals $(4 - 2u_p)(1 - f(p)/p)/(p^2 - 1)$. Its sign reverses for $p \equiv 1 \bmod 24$. For $u_p = 2$ all four signed masses cancel at 1, and

$$V'_{p,f}(1) = -\left(1 - \frac{f(p)}{p}\right) \frac{2p^2 \log p}{(p^2 - 1)^2} < 0. \tag{27}$$

The derivative comes from the $m = 0 \bmod p$ branch. On the other three branches m is a unit, and their relative unit masses are constant in s. In particular this zero is simple for either sign of $f(p)$.

There are k such zeros in every fixed mode, and the principal mode has the additional zero at 2 from Lemma 6.2. Its contribution starts at most at order $(\log X)^{-5/3-k}$, as does that of the χ mode. For either mode with $z = -1/3$, the first possible coefficient in (21) is at $j = k$. Its gamma multiplier relative to the coefficient at order zero is

$$\frac{\Gamma(-1/3)}{\Gamma(-1/3 - k)} = (-1)^k (4/3)_k.$$

Its sign cancels the k negative derivatives in (27). The remaining reversals contribute exactly $(-1)^{r_1(a)}$, independently of the mode. The two dyadic coefficients therefore have the same sign for the chosen parity, so the first possible coefficient is nonzero. Evaluating the forbidden factors at 1, the other forced factors at 1, and the k first derivatives gives (25). This proves the stated error for fixed a. □

For example, $K = 11 \cdot 19$ gives exponent $10/3$ with the same sign as $K = 1$ for a given parity. On $K = 11 \cdot 73$ the exponent is $7/3$ and both signs reverse. Each is a positive-density population. More precisely,

$$d(a) = \prod_{\substack{p>3 \\ p \nmid a}} \left(1 - \frac{4}{p(p+1)}\right) \prod_{p \mid a} \frac{4}{p(p+1)} > 0 \tag{28}$$

by the four disjoint simple roots and the tail bound (8). Primes in any fixed one of the three delaying residue classes can be chosen in arbitrarily large finite numbers. Hence every exponent $4/3 + k$, $k \geq 0$, occurs.

Corollary 8.2. *Let $\sigma = +1$ for the mark Z_+ and $\sigma = -1$ for Z_-. Given $K = a$ and $Z_\sigma(3n + 2) = \varepsilon$, with $\varepsilon = \pm 1$, the ordinary-interval probability of agreement is*

$$\mathbb{P}_X(C = W \mid K = a,\ Z_\sigma = \varepsilon) = \frac{1}{2} - \frac{\varepsilon\sigma(-1)^{r_1(a)} b_{\sigma,a}}{2d(a)} (\log X)^{-4/3-k(a)} + o_a((\log X)^{-4/3-k(a)}).$$

Here Z_σ in the conditioning is always evaluated at $3n + 2$.

Proof. On every fixed progression in m, Z_σ has ordinary mean zero. Indeed (12) gives prime means $2/3, -1/3, 1/3, 1/3$ against $1, \chi, \eta, \eta\chi$; each is strictly less than one. The progressions argument in Section 7 gives a vanishing average on unit classes, also for the

remaining zero-mean modes. On a nonunit class extract the common divisor with its fixed modulus. The resulting variable ranges over a unit class of the reduced modulus, and complete multiplicativity supplies a constant factor. Approximate $\mathbf{1}_{K=a}$ by fixed residue conditions in density using (8). Each mark value then has joint density $d(a)/2$ with $K = a$. The numerator for the conditional mean of CW uses $\mathbf{1}_{Z_\sigma=\varepsilon} = (1 + \varepsilon Z_\sigma)/2$. The term with the mark is given by (24), and the term without it has arbitrary logarithmic saving by the same finite differences of Theorem 7.1. Finally $\mathbf{1}_{C=W} = (1 + CW)/2$, as required. □

8.1 The next coefficient has the same sign for both parities

The two terms of order $(\log X)^{-4/3-k}$ came from characters ramified at 2. The next exponent receives two contributions: the derivative at the zero of the principal mode and the value of the mode χ. We can compare them because the two associated Euler products agree at every $p \equiv 1 \bmod 3$.

For a fixed parity let $f^{(0)} = h,\ f^{(1)} = h\chi$, and let $G_0 = \prod_p \frac{(1-1/p)^{1/3}}{1-h(p)/p} > 0$ be the principal Euler constant. For an exact stratum use $H_f(P)$ from (23). For a product weight take $H_f(P_{\boldsymbol\theta}) = \prod_{p>3}(1 - [4 - (4 - 2u_p)\theta_p]\kappa_p(f))$. All these products are positive.

Theorem 8.3. *For $P = P_{\boldsymbol\theta}$ set $k = r = 0$. For $P = \mathbf{1}_{K=a}$ set $k = k(a)$, $r = r_1(a)$. With $b_\pm(P)$ from (19) or (25), respectively, one has, for $L = \log X$,*

$$B_\pm(P; X) = \mp(-1)^r b_\pm(P) L^{-4/3-k} - (-1)^r d_\pm(P) L^{-5/3-k} + O_P(L^{-7/3-k}), \tag{29}$$

where

$$d_\pm(P) = \frac{G_0}{8|\Gamma(-2/3-k)|}\left(\frac{\pi}{\sqrt{3}} H_{f^{(1)}}(P) - (\log 2) H_{f^{(0)}}(P)\right) > 0. \tag{30}$$

For $P_{\boldsymbol\theta}$ the error is uniform in $\boldsymbol\theta$.

Proof. First take $\psi = \chi^b$, $b = 0, 1$, and put $x = (-1)^b 2^{-s}$. On extracting $2^j \mid m$ the unit character at 3 supplies $(-1)^{b(j+1)}$, as in Section 7. Keep its constant $(-1)^b$ with the ternary factor. If $I_j = \int_{v_2(m)=j} D_2(n)\, dn$, the dyadic factor in the analytic coefficient is

$$U_b(s) = 2(1 - x) \sum_{j \ge 0} I_j (2x)^j.$$

Indeed normalized unit mass at two gives twice Haar mass per class, with Euler removal $1 - f^{(b)}(2)2^{-s} = 1 - x$. The shell table in Lemma 6.2 gives $I_0 = I_1 = 0,\ I_2 = 1/8,\ I_j = -2^{-j-1}$ for $j \ge 3$. Hence

$$U_b(s) = (1 - x)\left(x^2 - \frac{x^3}{1 - x}\right) = x^2(1 - 2x), \qquad U_0(1) = 0, \quad U_0'(1) = \tfrac{1}{4}\log 2, \quad U_1(1) = \tfrac{1}{2}.$$

At three the analytic multiplier is the constant $-(-1)^b/6$: the outer variable m is a unit, one of its two residue classes is used, and the conditional D_3-integral is $-1/3$. The interval contributes $6^1 - 3^1 = 3$.

Both modes carry the common leading factor $(-1)^{r+k} H_{f^{(b)}}(P)$ at degree k from primes $p > 3$, by (26) and (27). At $L^{-5/3-k}$ the principal mode enters at degree $k + 1$ and χ at degree k. Their gamma denominators are both $\Gamma(-2/3 - k)$. Before this denominator and $(-1)^{r+k}$, the two contributions are

$$-\tfrac{1}{8}(\log 2) G_0 H_{f^{(0)}}(P), \qquad +\tfrac{1}{4} G_\chi(1) H_{f^{(1)}}(P).$$

Terms from the two ramified modes after their leading coefficients have order $O_P(L^{-7/3-k})$. All higher terms of the unramified modes are smaller. The arbitrary-order control in Section 7 therefore gives the claimed error.

For $p \equiv 2 \bmod 3$ one has $h(p) = 1$, including $p = 2$, and at 3 both functions vanish. Euler factors give, for $\Re s > 1$,

$$\frac{\sum f^{(1)}(m)m^{-s}}{\sum f^{(0)}(m)m^{-s}} = \prod_{p\equiv 2\ (3)} \frac{1-p^{-s}}{1+p^{-s}} = \frac{L(s,\chi)}{\zeta(s)(1-3^{-s})}.$$

The two prime means differ by -1, so $G_\chi(1)/G_0 = \frac{3}{2}L(1,\chi) = \pi/(2\sqrt{3})$. Here $L(1,\chi) = \int_0^1 (1+t+t^2)^{-1}dt$. Together with the sign of $\Gamma(-2/3-k)$ this proves the formula for the second coefficient.

It remains to exclude cancellation in its parentheses. At $p \equiv 1 \bmod 3$ both local factors agree. At a forced prime $p \equiv 2 \bmod 3$, the ratio of the positive local factors for $f^{(1)}$ and $f^{(0)}$ is $(p+1)/(p-1) > 1$, also when the factor is a derivative. At the other inert primes the ratio is

$$R_p = \frac{1 - A/[p(p-1)]}{1 - A/[p(p+1)]} \quad \text{with } 0 \le A \le 4.$$

Indeed $u_p = 1$ or 2 for these primes, and forbidding a prime gives $A = 4$. For $p \ge 5$,

$$1 - R_p = \frac{2A}{p(p^2-1)(1 - A/[p(p+1)])} \le \frac{10}{p(p^2-1)}.$$

The sum of $1/[n(n^2-1)]$ over all integers $n \ge 5$ telescopes to $1/40$. Multiplying the positive ratios gives $H_{f^{(1)}}(P)/H_{f^{(0)}}(P) \ge 3/4$. Since $3\pi/(4\sqrt{3}) > \log 2$, the difference in (30) is strictly positive. □

The arithmetic distinction is visible in (15). The root number cancels three Liouville factors before the asymptotic expansion starts. Removing that weight restores the four-site signed mean c of Section 4; estimating the conditional signed mean in (11) remains necessary for that application.

Computations and AI assistance

The accompanying Zenodo source archive, also embedded in the original Zenodo PDF, contains the local algorithms and their execution records. It compares both isogeny descents on 203 triangular parameters and on 513 choices of k; the sets overlap. The linear-parameter checks compare the genus and Selmer identities for $1 \le k \le 4096$, enumerate residue classes in both samplings, and bound the tail of the agreement product. The second-family computations compare dyadic root numbers from two rational kernels at 239 parameters. PARI/GP evaluates the local signs on 29 696 residue representatives and checks the global root-number identity at 1 012 parameters. Separate finite calculations check the five integrals in Lemma 6.2, the cubic characters, the four prime means in (12) and the local factors on exact strata.

OpenAI ChatGPT and Codex assisted with mathematical development, literature checks, exposition, programming and finite computations. The limiting laws and signed asymptotics are established by the arguments above.

References

[1] A. Akbary and P.-J. Wong, *Fourier coefficients of automorphic L-functions over primes in ray classes*, J. Ramanujan Math. Soc. 39 (2024), no. 3, 265–280. Author's published copy.

[2] R. de la Bretèche and G. Tenenbaum, *Remarks on the Selberg–Delange method*, Acta Arith. 200 (2021), 349–369. 10.4064/aa201024-26-5.

[3] J. W. S. Cassels, *Arithmetic on curves of genus 1. VIII. On conjectures of Birch and Swinnerton-Dyer*, J. Reine Angew. Math. 217 (1965), 180–199. 10.1515/crll.1965.217.180.

[4] S. Chan and M. Verzobio, *Selmer groups of families of elliptic curves with an ℓ-isogeny*, arXiv:2508.21406 (2025). 10.48550/arXiv.2508.21406.

[5] T. Dokchitser and V. Dokchitser, *Root numbers and parity of ranks of elliptic curves*, J. Reine Angew. Math. 658 (2011), 39–64. 10.1515/crelle.2011.060.

[6] D. El-Baz, D. Loughran and E. Sofos, *Multivariate normal distribution for integral points on varieties*, Trans. Amer. Math. Soc. 375 (2022), 3089–3128. 10.1090/tran/8545.

[7] T. Fisher, *Higher descents on an elliptic curve with a rational 2-torsion point*, Math. Comp. 86 (2017), 2493–2518. 10.1090/mcom/3163.

[8] L. J. Goldstein, *A generalization of the Siegel–Walfisz theorem*, Trans. Amer. Math. Soc. 149 (1970), 417–429. 10.1090/S0002-9947-1970-0274416-6.

[9] A. Granville and D. Koukoulopoulos, *Beyond the LSD method for the partial sums of multiplicative functions*, Ramanujan J. 49 (2019), 287–319. 10.1007/s11139-018-0119-3.

[10] A. Granville and K. Soundararajan, *Sieving and the Erdős–Kac theorem*, in *Equidistribution in Number Theory, An Introduction*, NATO Science Series II 237, Springer (2007), 15–27. 10.1007/978-1-4020-5404-4_2.

[11] H. A. Helfgott, *On the behaviour of root numbers in families of elliptic curves*, arXiv:math/0408141 (2009). 10.48550/arXiv.math/0408141.

[12] Z. Klagsbrun and R. J. Lemke Oliver, *The distribution of the Tamagawa ratio in the family of elliptic curves with a two-torsion point*, Res. Math. Sci. 1 (2014), article 15. 10.1186/s40687-014-0015-4.

[13] Z. Klagsbrun and R. J. Lemke Oliver, *The distribution of 2-Selmer ranks of quadratic twists of elliptic curves with partial two-torsion*, Mathematika 62 (2016), 67–78. 10.1112/S0025579315000121.

[14] Z. Klagsbrun, B. Mazur and K. Rubin, *Disparity in Selmer ranks of quadratic twists of elliptic curves*, Ann. of Math. 178 (2013), 287–320. 10.4007/annals.2013.178.1.5.

[15] Z. Klagsbrun, B. Mazur and K. Rubin, *Erratum to "Disparity in Selmer ranks of quadratic twists of elliptic curves"*, Ann. of Math. 198 (2023), 879–880. 10.4007/annals.2023.198.2.9.

[16] P. Koymans and A. Smith, *Tamagawa ratios and unbounded Selmer moments*, arXiv:2606.31649 (2026). 10.48550/arXiv.2606.31649.

[17] D. Loughran and L. Matthiesen, *Frobenian multiplicative functions and rational points in fibrations*, J. Eur. Math. Soc. 26 (2024), 4779–4830. 10.4171/JEMS/1374.

[18] B. Mazur, K. Rubin and A. Silverberg, *Twisting commutative algebraic groups*, J. Algebra 314 (2007), 419–438. 10.1016/j.jalgebra.2007.02.052.

[19] A. Morgan and R. Paterson, *On 2-Selmer groups of twists after quadratic extension*, J. Lond. Math. Soc. 105 (2022), 1110–1166. 10.1112/jlms.12533.

[20] B. Pouly, *An explicit bound below one half for a quartic Liouville correlation*, Zenodo (2026). 10.5281/zenodo.24443184.

[21] B. Pouly, *Elliptic root numbers and a signed squarefree variance for a quartic Liouville correlation*, Zenodo (2026). 10.5281/zenodo.24578917.

[22] D. E. Rohrlich, *Galois theory, elliptic curves, and root numbers*, Compositio Math. 100 (1996), 311–349. NUMDAM, CM_1996___100_3_311_0.

[23] J. H. Silverman, *The Arithmetic of Elliptic Curves*, 2nd ed., Graduate Texts in Mathematics 106, Springer (2009). 10.1007/978-0-387-09494-6.

[24] P. Stevenhagen, *Rédei reciprocity, governing fields and negative Pell*, Math. Proc. Cambridge Philos. Soc. 172 (2022), 627–654. 10.1017/S0305004121000335.

[25] T. Tao, *The logarithmically averaged Chowla and Elliott conjectures for two-point correlations*, Forum Math. Pi 4 (2016), e8, 1–36. 10.1017/fmp.2016.6.

[26] T. Tao and J. Teräväinen, *The structure of logarithmically averaged correlations of multiplicative functions, with applications to the Chowla and Elliott conjectures*, Duke Math. J. 168 (2019), 1977–2027. 10.1215/00127094-2019-0002.

[27] M. Xiong and A. Zaharescu, *Distribution of Selmer groups of quadratic twists of a family of elliptic curves*, Adv. Math. 219 (2008), 523–553. 10.1016/j.aim.2008.05.005.

[28] G. Yu, *Average size of 2-Selmer groups of elliptic curves, II*, Acta Arith. 117 (2005), 1–33. 10.4064/aa117-1-1.