

Certified Constants and an Exact CM Reproduction for Koblitz's Prime-Order Conjecture

A finite-level quotient certificate for 540.f2, a fifty-checkpoint table verification to one billion, and the norm-pair form of the $j = 1728$ prime-order event

math.st¹ **@math__street**¹

with GPT 5.6 Sol · Claude Fable 5

¹Independent research collective, math.st · July 2026

ABSTRACT

Koblitz conjectured in 1988 that a fixed non-CM elliptic curve $E/\mathbb{Q}$ without congruence obstruction has $\sim C_{E,1}x/(\log x)^2$ primes $p \leq x$ of good reduction for which $\#E(\mathbb{F}_p)$ is prime; Zywina later corrected the constant through the adelic Galois image, which can force $C_{E,1} = 0$. This paper reports a five-session research program on P5.1 that closes every algebraic, computational, and source-audit sub-goal within reach of exact finite computation, and states precisely what remains open. We reproduce Zywina's universal Euler product and Serre-curve constant to within $4 \cdot 10^{-8}$, measure prime-order densities for three non-CM curves through $x = 2^{17}$ (observed/predicted 1.043 for the Serre curve 1728.w1; zero events, as proved, for curves with rational 2- and 3-torsion), and certify a quotient constant $C_{E,3} = (5824/5913)C$ for the 3-torsion curve 540.f2 by exhaustively enumerating its 699,840-element mod-90 Galois image — a result conditional only on LMFDB's published adelic data. For the CM curve $y^2 = x^3 - x$ we implement Walsh's exact trace formula and reproduce all 50 checkpoints of Zywina's published table through $x = 10^9$ with zero mismatches (final observed/predicted 0.99943). On the theory side we prove an exact reduction: apart from the single exceptional prime $p = 17$, every CM prime-quotient event lies in the class $p \equiv 5 \pmod{8}$ and is precisely the simultaneous primality of the two Gaussian norms $N(a + bi)$ and $N(a + bi + 1)/8$ — a prime-pair pattern that makes the sieve parity obstruction explicit. An audit of the 2025–2026 primary literature confirms that the unconditional fixed-curve asymptotic remains open; we trace exactly where GRH-strength hypotheses enter the strongest partial results and isolate the two logically independent missing ingredients. Every computation is deterministic, adversarially self-verified, and reproducible from seeded scripts.

Keywords. Koblitz conjecture · elliptic curves · prime order · adelic Galois image · Euler products · CM curves · sieve parity

1 Introduction

Elliptic-curve cryptography wants curves whose group of rational points over a prime field has prime order, and Koblitz's original question [1] is the natural arithmetic-statistics version: fix one curve $E/\mathbb{Q}$ and ask how often its reductions $E(\mathbb{F}_p)$ have prime order as p ranges over primes of

good reduction. Writing $N_p = \#E(\mathbb{F}_p) = p + 1 - a_p$ with $|a_p| \leq 2p^{1/2}$ by Hasse, the count of prime N_p up to x is conjecturally governed by the same $x/(\log x)^2$ scale as twin primes, with a curve-specific constant. Koblitz built the constant as a product of local densities read from the mod- ℓ Galois images [1]; Zywina later showed that independence across primes ℓ can fail — division fields entangle — and replaced the naive product by a finite-level adelic density limit which corrects the constant and can make it vanish [2]. The corrected conjecture has survived every published numerical test, yet no fixed-curve instance has ever been proved, and the 2025 literature still describes it as open [3], [4].

This paper is the record of research problem P5.1: measure the conjecture honestly at reachable ranges, certify every constant used rather than fitting anything, reproduce the strongest published computation exactly, and determine — by proof where possible, by source audit where not — exactly what separates the finite evidence from the asymptotic theorem. The work spans five sessions, twelve closed sub-goals (SG-01 through SG-12), 85 passing tests, and three deterministic data-producing scripts whose outputs back every figure and table below.

Main findings. (1) **EMPIRICAL: all good primes $5 \leq p \leq 2^{17}$** The corrected predictors track the data at toy scale: the Serre curve 1728.w1 gives observed/predicted 1.04293 (683 prime orders vs. 654.886 predicted), the rational-torsion curves give exactly zero prime orders as proved, and a certified quotient constant for 540.f2 gives ratio 1.01600.

(2) **CONDITIONAL: LMFDB 540.f2 adelic data** Exact enumeration of the level-90 Galois image (699,840 matrices, 98,280 favorable) proves $\delta_{E,3}(90) = 91/648$ and hence $C_{E,3} = (5824/5913)C$ for $y^2 = x^3 + 3x - 11$.

(3) **EMPIRICAL: all 50 checkpoints, $2 \cdot 10^7 \leq x \leq 10^9$** An independent exact implementation — Cornacchia representations plus Walsh's trace theorem — reproduces every actual count and every rounded expected count in Zywina's published CM table with zero mismatches, in 148.7 seconds.

(4) **PROVED** For $E : y^2 = x^3 - x$, apart from the single event $(p, N_p/8) = (17, 2)$, every prime-quotient event has $p \equiv 5 \pmod{8}$ and equals the simultaneous primality of $N(a + bi)$ and $N(a + bi + 1)/8$ in $\mathbb{Z}[i]$. The missing unconditional theorem is therefore a prime-pair distribution statement, not a missing trace formula, density, or test.

1.1 Honest scope

Nothing in this paper proves the Koblitz asymptotic, and no amount of the computation reported here could: a larger finite cutoff cannot by itself turn evidence into an asymptotic theorem. **PROVED** What the computations do is eliminate every non-asymptotic explanation for a failure — wrong constant, wrong predictor, wrong point counts, circular validation — through independent cross-checks (§8), so that the surviving gap is exactly the open conjecture recorded as Q026 (§7). The strongest 2025 fixed-curve result remains conditional on two unproved hypotheses [4], and the strongest unconditional results are averages over families [3] or almost-prime statements [5], [6].

2 Setting and notation

Fix an elliptic curve $E/\mathbb{Q}$ with conductor N_E . For a prime $p \nmid N_E$ of good reduction write $N_p = \#E(\mathbb{F}_p) = p + 1 - a_p$, $|a_p| \leq 2p^{1/2}$. Koblitz's conjecture in Zywin's corrected form [1], [2] reads:

CONJECTURE For $E/\mathbb{Q}$ non-CM and $t \geq 1$ a fixed integer,

$$\#\{p \leq x : p \nmid N_E, N_p/t \text{ an integer and prime}\} \sim C_{E,t} \frac{x}{(\log x)^2}, \quad (1)$$

where $C_{E,t} \geq 0$ is defined from the adelic Galois image $\rho_E : \text{Gal}(\overline{\mathbb{Q}}/\mathbb{Q}) \rightarrow \text{GL}_2(\hat{\mathbb{Z}})$ as a divisibility-ordered limit of finite-level densities; the constant captures entanglement between division fields and can vanish [2]. The case $t = 1$ is the prime-order conjecture; Koblitz's original positive-constant statement additionally excludes curves $\mathbb{Q}$ -isogenous to curves with nontrivial rational torsion [1].

Two normalizations of the finite-cutoff prediction matter below. The **raw asymptotic predictor** is $C_{E,t}x/(\log x)^2$. The **refined predictor** of Zywin §2.4 replaces the asymptotic mass by a prime sum,

$$\text{Pred}_t(x) = C_{E,t} \sum_{p \leq x} \frac{1}{\log(p+1) - \log t}, \quad (2)$$

which converges to the same asymptotic but is far less biased at small x [2]. **EMPIRICAL: $p \leq 2^{17}$** The distinction is not cosmetic at toy range: for the Serre curve below, the raw predictor is off by 28.9% at $x = 2^{17}$ while the refined predictor is off by 4.3% (Figure 1).

Throughout, C denotes Zywin's **universal constant**, the value of $C_{E,1}$ for a curve with maximal image at every level (equation (2.3) of [2]), and all measurements use the fixed seed 51012026 and deterministic pipelines.

3 Corrected constants and their certificates

3.1 The universal Euler product

CITED For a non-CM curve whose mod- ℓ image is all of $\text{GL}_2(\mathbb{F}_\ell)$, the local factor at ℓ is

$$f_\ell = 1 - \frac{\ell^2 - \ell - 1}{(\ell - 1)^3(\ell + 1)}, \quad (3)$$

and the universal product is

$$C = \prod_{\ell} f_\ell = 0.505166168239435774... \quad (4)$$

(Zywin [2], Proposition 2.4 and equation (2.3)). Each factor compares the density of image matrices without eigenvalue 1 — the Frobenius classes forcing $\ell \mid N_p$ — against the density $1 - 1/\ell$ for a random integer.

EMPIRICAL: Euler factors $\ell \leq 10^6$ Our implementation truncates the product over primes $\ell \leq 10^6$ and returns 0.505166202477432, within $3.43 \cdot 10^{-8}$ of the published limit; a 60-digit recomputation during the audit differs from the float value by $9.0 \cdot 10^{-15}$ (§8).

3.2 The Serre curve 1728.w1

CITED The curve $E_0 : y^2 = x^3 + 6x - 2$ (LMFDB 1728.w1 [7]) is a Serre curve — its adelic image is as large as the Weil pairing and the Kronecker–Weber entanglement allow — and Zywinia computes its corrected constant in closed form:

$$C_{E_0,1} = \frac{10}{9}C \approx 0.561295742488261971 \quad (5)$$

([2], Section 5 and equation (5.1)). The factor $10/9$ is exactly the entanglement correction that the naive local product misses. **EMPIRICAL: Euler factors $t \leq 10^6$** The implemented constant is 0.561295780530480 , within $3.80 \cdot 10^{-8}$ of the published value.

3.3 The rational-torsion obstruction

Proposition 1 (vanishing for rational torsion). **PROVED** Let $E/\mathbb{Q}$ have a rational point of prime order $t \in \{2, 3\}$. Then for every good prime $p > t$ the reduction satisfies $t \mid N_p$, so N_p is composite whenever $N_p > t$; in particular prime orders occur for at most finitely many p and $C_{E,1} = 0$.

Proof. Reduction modulo a good prime p is injective on prime-to- p rational torsion, so the order- t point survives into $E(\mathbb{F}_p)$ and $t \mid N_p$ by Lagrange. Hasse gives $N_p \geq p + 1 - 2p^{1/2} = (p^{1/2} - 1)^2 > t$ for all $p \geq 7$, so beyond finitely many primes N_p is a multiple of t exceeding t , hence composite. The corrected constant, which is an average of densities of prime values, is then zero [2]. $\square$

The two test curves realizing this obstruction are $y^2 = x^3 + x - 2$ (LMFDB 112.b4, torsion order 2 [7]) and $y^2 = x^3 + 3x - 11$ (LMFDB 540.f2, torsion $\mathbb{Z}/3\mathbb{Z}$ generated by $(3, 5)$ [7]).

Remark 2 (both curves are non-CM). **PROVED** The point $(3, 5)$ has order 3 on 540.f2: the tangent slope at $(3, 5)$ is 3, and the duplication formula gives $2(3, 5) = (3, -5) = -(3, 5)$. Neither curve has CM: their j -invariants are $432/7$ and $6912/125$, and a CM j -invariant is an algebraic integer, so a rational CM j -invariant must be a rational integer [8].

3.4 A certified finite-level quotient constant for 540.f2

Since $C_{E,1} = 0$ for 540.f2, the refined conjecture moves to the quotient $t = 3$: how often is $N_p/3$ prime? Zywinia's Proposition 2.4 reduces $C_{E,t}$ for a non-CM curve with adelic level M to one exact finite-level density $\delta_{E,t}$ at modulus $t \prod_{\ell \nmid tM} \ell$, times the universal factors at all other primes [2].

CITED LMFDB records the adelic image of 540.f2 as level $M = 30$, index 16, label 30.16.0-30.b.1.4, with seven explicit generators modulo 30 [7]. With $t = 3$ and $M = 30$ the required modulus is $3(2 \cdot 3 \cdot 5) = 90$.

Proposition 3 (exact level-90 density). **CONDITIONAL: LMFDB 540.f2 level and generators correct** Let $G(30) \subset \mathrm{GL}_2(\mathbb{Z}/30\mathbb{Z})$ be the subgroup generated by LMFDB's seven generators, and let $G(90)$ be its full preimage set of entrywise lifts. Exhaustive computation gives $|G(30)| = 8640$, hence $|G(90)| = 8640 \cdot 3^4 = 699,840$; among these exactly 98,280 matrices A satisfy $\gcd(\det(I - A), 90) = 3$, the condition that $N_p/3$ be coprime to 90. Therefore

$$\delta_{E,3}(90) = \frac{98280}{699840} = \frac{91}{648}. \quad (6)$$

The condition $\det(I - A) \in 3(\mathbb{Z}/90\mathbb{Z})^\times$ is equivalent to $\gcd(\det(I - A), 90) = 3$ because $90 = 2 \cdot 3^2 \cdot 5$ and divisibility of $N_p/3$ by 2, 3, or 5 is read off modulo 90. Assembling the constant per Proposition 2.4: dividing $91/648$ by the random-integer densities $(1 - 1/2)(1 - 1/3)(1 - 1/5) = 4/15$ gives $455/864$; the universal factors at the entangled primes are $f_2 = 2/3$, $f_3 = 27/32$, $f_5 = 365/384$ with product $1095/2048$; hence

$$C_{E,3} = \frac{455/864}{1095/2048} C = \frac{5824}{5913} C \approx 0.497562652330215. \quad (7)$$

EMPIRICAL: Euler factors $\ell \leq 10^6$ The numerical value uses the truncated universal product.

PROVED The factor $5824/5913$ was derived by exact rational arithmetic and was **not** fitted to the measured quotient counts; an independent CRT-decomposed recount during the audit reproduced all three fractions (§8). We flag one boundary of trust explicitly: the certificate consumes LMFDB's published level-30 generators as input and does not re-derive them from division polynomials.

4 The CM curve $y^2 = x^3 - x$: exact trace and the norm-pair reduction

CM curves need a separate formulation [1], [2]. For $E : y^2 = x^3 - x$ ($j = 1728$), all CM endomorphisms are defined over $\mathbb{Q}(i)$, the torsion of $E(\mathbb{Q}(i))$ has order 8, and a rational prime $p \equiv 3 \pmod{4}$ is supersingular with $N_p = p + 1$. **CITED** The CM conjecture therefore studies $N_p/8$ on the **split** stratum $p \equiv 1 \pmod{4}$ only, with constant

$$C_{E_{\mathbb{Q}(i)},8} = \prod_{\ell \neq 2} \left(1 - \chi(\ell) \frac{\ell^2 - \ell - 1}{(\ell - \chi(\ell))(\ell - 1)^2} \right) \approx 1.067350894, \quad (8)$$

where $\chi(\ell) = (-1)^{(\ell-1)/2}$; the product is evaluated absolutely convergently after extracting $L(1, \chi)^{-1} = 4/\pi$ ([2], Lemma 7.1). **EMPIRICAL: Euler factors $\ell \leq 10^6$** Our accelerated implementation returns 1.067350966817026 , within $7.3 \cdot 10^{-8}$ of the published approximation.

4.1 Walsh's trace formula, specialized

CITED For $E_d : y^2 = x^3 + dx$ and a split prime $p = a^2 + b^2$ normalized by $a \equiv 1 \pmod{4}$, b even, Walsh determines $a_p \in \{\pm 2a, \pm 2b\}$ from the quartic-residue class of d [9]. **PROVED** For $d = -1$: since -1 is a fourth power modulo p exactly when $p \equiv 1 \pmod{8}$, Walsh's cases specialize to

$$a_p = \begin{cases} 2a & \text{if } p \equiv 1 \pmod{8} \\ -2a & \text{if } p \equiv 5 \pmod{8} \\ 0 & \text{if } p \equiv 3 \pmod{4} \end{cases} \quad (9)$$

so N_p is computable from one Cornacchia representation of p — an $O(\log p)$ computation replacing the $O(p^{1/4})$ -group-operation BSGS counter. This is what makes the $x = 10^9$ reproduction of §5.4 feasible in pure Python.

4.2 The norm identity and its consequences

Everything in this subsection is unconditional given Walsh's cited theorem; the complete derivations are in the repository's `THEORY_CLOSURE.md`.

Theorem 4 (norm form of the CM order). **PROVED** Let $p \equiv 1 \pmod{4}$ be prime, $p = a^2 + b^2$ with $a \equiv 1 \pmod{4}$. Put $\varepsilon_p = 1$ if $p \equiv 1 \pmod{8}$ and $\varepsilon_p = -1$ if $p \equiv 5 \pmod{8}$. Then

$$N_p = p + 1 - 2\varepsilon_p a = (a - \varepsilon_p)^2 + b^2 = N_{\mathbb{Z}[i]/\mathbb{Z}}(a + bi - \varepsilon_p). \quad (10)$$

Proof. Substitute $p = a^2 + b^2$ into $p + 1 - 2\varepsilon_p a$ and use $\varepsilon_p^2 = 1$: the result is $(a - \varepsilon_p)^2 + b^2$, which is the Gaussian norm of $a + bi - \varepsilon_p$. $\square$

Lemma 5 (divisibility by 8). **PROVED** In the setting of the theorem, $8 \mid N_p$. More precisely, $16 \mid N_p$ when $p \equiv 1 \pmod{8}$, while $N_p \equiv 8 \pmod{16}$ when $p \equiv 5 \pmod{8}$.

Proof. If $p \equiv 1 \pmod{8}$: a is odd so $a^2 \equiv 1 \pmod{8}$, forcing $b^2 \equiv 0 \pmod{8}$, hence $4 \mid b$; and $a \equiv 1 \pmod{4}$ gives $4 \mid a - 1$. Then $(a - 1)^2$ and b^2 are both $\equiv 0 \pmod{16}$. If $p \equiv 5 \pmod{8}$: now $b^2 \equiv 4 \pmod{8}$ forces $b \equiv 2 \pmod{4}$, and $a + 1 \equiv 2 \pmod{4}$; each of $(a + 1)^2$, b^2 is $\equiv 4 \pmod{16}$, so their sum is $\equiv 8 \pmod{16}$. $\square$

Proposition 6 (the class $p \equiv 1 \pmod{8}$ is finite). **PROVED** If $p \equiv 1 \pmod{8}$, then $N_p/8$ is even, and it is prime only for $p = 17$, where $N_p/8 = 2$.

Proof. By the lemma $16 \mid N_p$, so $N_p/8$ is even; if prime it must equal 2, i.e. $N_p = 16$. The equation $(a - 1)^2 + b^2 = 16$ with $4 \mid a - 1$ and $4 \mid b$ admits $(a, b) = (1, \pm 4)$, giving the prime $p = 17$, and $(a - 1, b) = (\pm 4, 0)$, giving $p = a^2 \in \{9, 25\}$, not prime. Hence $p = 17$ is the only event in this class. $\square$

Corollary 7 (the norm-pair reduction). **PROVED** Apart from the single event $(p, N_p/8) = (17, 2)$, every prime p with $N_p/8$ prime satisfies $p \equiv 5 \pmod{8}$ and

$$p = N(a + bi), \quad \frac{N_p}{8} = \frac{N(a + bi + 1)}{8}, \quad (11)$$

i.e. the CM Koblitz event is exactly the **simultaneous** primality of two explicit Gaussian norms. Writing $u = (a + 1)/2$ and $v = b/2$ (both odd), the second value is the integral quadratic-form value

$$\frac{N_p}{8} = \frac{u^2 + v^2}{2} = \left(\frac{u + v}{2}\right)^2 + \left(\frac{u - v}{2}\right)^2. \quad (12)$$

EMPIRICAL: all 74,416 split primes $5 \leq p \leq 2 \cdot 10^6$ An independent Session-5 check verified the norm identity, the divisibility by 8, and the uniqueness of the $(17, 2)$ event on every split prime in range; a persistent regression repeats the identity through 10^5 .

HEURISTIC The reduction explains the parity barrier: the event is a Gaussian-integer analogue of a prime-pair problem, precisely the shape on which classical sieves lose to parity. This explanatory analogy would be falsified if a sieve using only divisibility data proved the required prime lower bound without a parity-breaking input.

5 Empirical validation

All measurements use exact point counting — the shared Hasse-interval BSGS/twist counter with a declared exhaustive fallback for the non-CM sweeps, and the Cornacchia–Walsh trace for the CM reproduction — validated against exhaustive counts at small primes before each run.

EMPIRICAL: 67 validation cases, $p \leq 97$ The production counter agreed with exhaustive counting on all validation cases (66 via BSGS/twist, one via the declared fallback). PARI/GP and SageMath were

unavailable in the environment; the pure-Python counter substitutes for them and is itself cross-validated (§8).

5.1 Non-CM prime-order densities through $x = 2^{17}$

Table 1: **EMPIRICAL: all good primes $5 \leq p \leq 2^{17}$** The three-curve sweep at the largest cutoff. Data: `measure_density_x131072_l1000000_s51012026_20260624.csv`, seed 51012026, 11.2 s wall time.

Curve	Torsion	$C_{E,1}$	Good p	Prime N_p	Refined pred.	Obs./pred.
1728.w1 $y^2 = x^3 + 6x - 2$	trivial	$(10/9)C$	12,249	683	654.886	1.04293
112.b4 $y^2 = x^3 + x - 2$	$\mathbb{Z}/2\mathbb{Z}$	0	12,248	0	0	—
540.f2 $y^2 = x^3 + 3x - 11$	$\mathbb{Z}/3\mathbb{Z}$	0	12,248	0	0	—

Table 1 summarizes the largest cutoff. Three observations, in decreasing order of strength.

Zero is exact. **EMPIRICAL: $p \leq 2^{17}$** Both rational-torsion curves produced **no** prime orders across 12,248 good primes each — the obstruction of Proposition 1 realized in data. With zero events the heuristic 95% upper bound on the measured constant is $-\ln(0.05)/B(x) \approx 0.00257$ at $x = 2^{17}$ (where $B(x)$ is the refined baseline), consistent with $C_{E,1} = 0$ and two orders of magnitude below the Serre-curve constant. A single contrary reduction would have falsified the implementation, since reduction injects prime-to- p torsion.

The refined predictor matters. **EMPIRICAL: $p \leq 2^{17}$** For 1728.w1 the refined prediction is 654.886 against 683 observed (ratio 1.04293), while the raw asymptotic expression $C_{E_0,1}x/(\log x)^2$ predicts only 529.85 (ratio 1.28904). **Figure 1** shows both ratios across cutoffs $2^{10} \leq x \leq 2^{17}$: the refined ratio moves 0.8030, 1.0372, 1.1561, 1.0950, 1.0944, 1.0567, 1.0466, 1.0429 — into and then along the heuristic band — while the raw ratio stays above 1.28 throughout.

The interval is heuristic. **HEURISTIC** Treating the 683 events as approximately Poisson gives the 95% interval $[0.9647, 1.1211]$ for observed/predicted, which contains 1; cross-prime dependence or finite-cutoff bias could invalidate this interval, so we report it as a heuristic band, not a confidence statement.




Figure 1: **EMPIRICAL: $5 \leq p \leq 2^{17}$** Observed/predicted prime-order counts for the Serre curve 1728.w1 under the refined prime-sum predictor (blue) and the raw asymptotic predictor (orange), with the heuristic Poisson band around the refined ratio. Data: `measure_density_x131072_l1000000_s51012026_20260624.csv`.

5.2 Certified quotient cases and a negative control

Table 2: **EMPIRICAL: $5 \leq p \leq 2^{17}$, seed 51012026** Quotient cases at the largest cutoff. The certified constants are 0.497563 (540.f2) and 1.067351 (CM); both measured intervals contain their predictions. The deliberately invalid pooled model is a negative control. Data: `measure_corrected_cases_x131072_l1000000_s51012026_20260713.csv`.

Case	Eligible p	Events	Predicted	Obs./pred.	Measured [95%]	const.
540.f2, $t = 3$, all good p	12,248	661	650.592	1.01600	0.5055 [0.4670, 0.5441]	
$y^2 = x^3 - x$, $t = 8$, split p	6,094	765	779.701	0.98114	1.0472 [0.9730, 1.1214]	
pooled full- GL_2 con- trol, $t = 8$	all good p	987	243.653	4.05084	—	

EMPIRICAL: $p \leq 2^{17}$ The certified 540.f2 constant of §3.4 predicts 650.592 quotient-prime events; the sweep observed 661, ratio 1.015998, and the measured-constant interval $[0.466985, 0.544060]$ contains the certified prediction 0.497563. The ratio across cutoffs $2^{10} \leq x \leq 2^{17}$ is 1.0767, 1.1480, 0.9618, 1.1734, 1.0715, 1.0232, 0.9796, 1.0160 (Figure 2). Because the constant was derived by exact finite-group enumeration **before** comparison — never fitted — this agreement is a genuine test of Zywinia’s Proposition 2.4 on a curve with an entangled level-30 image.

EMPIRICAL: split primes $\leq 2^{17}$ The CM case observed 765 events among 6,094 split primes against 779.701 predicted, ratio 0.981145, rising steadily with cutoff from 0.8443 at $x = 2^{10}$ (Figure 2); the measured interval again contains the published constant.

EMPIRICAL: all good rational primes $\leq 2^{17}$ The **negative control** pools the CM curve’s split and supersingular strata and applies the inapplicable full- GL_2 quotient model: it predicts 243.653 events but 987 occur — ratio 4.05084. The split-prime convention is mathematically substantive, not presentational: using the wrong stratum model fails by a factor of four while both certified predictors sit within a few percent of 1.




Figure 2: **EMPIRICAL: $5 \leq p \leq 2^{17}$** Observed/predicted ratios for the two certified quotient cases and the deliberately invalid pooled full- GL_2 control, which diverges to 4.05. Data: `measure_corrected_cases_x131072_l1000000_s51012026_20260713.csv`.

5.3 Constants: predicted vs. measured

Figure 3 collects the three certified constants against their measured values at $x = 2^{17}$. All three heuristic intervals contain their certified predictions; none was tuned. The universal product C is drawn for reference — the Serre correction lifts it by $10/9$, the 540.f2 quotient correction lowers it by $5824/5913$, and the CM constant lives on a different Euler product entirely.




Figure 3: Certified predicted constants (open diamonds) vs. measured constants with heuristic 95% intervals (filled, at $x = 2^{17}$), against the universal product $C = 0.505166\dots$ (dashed). Data: the two CSVs of Table 1 and Table 2.

5.4 Exact reproduction of Zywin's CM table through $x = 10^9$

Zywin's strongest published numerical evidence is his Table 3: counts of split primes $p \leq x$ with $N_p/8$ prime for $y^2 = x^3 - x$, at fifty checkpoints $x = 2 \cdot 10^7, 4 \cdot 10^7, \dots, 10^9$, against the integral predictor of his equation (7.1),

$$\text{Pred}(x) = \frac{C_{E_{Q(i)},8}}{2} \int_9^x \frac{dt}{\log t \cdot (\log(t+1) - \log 8)}, \quad (13)$$

the factor $1/2$ being the density of the split stratum. Session 3 re-implemented the entire computation from scratch — segmented sieve, Cornacchia representations, Walsh traces, Hasse-bounded quotient-primality sieve, and mpmath quadrature — and ran all fifty checkpoints.

EMPIRICAL: all 50 checkpoints, $2 \cdot 10^7 \leq x \leq 10^9$ Every computed actual count equals Zywin's published value, and every computed integral rounds to his published expected value: both mismatch totals are zero. The complete run processed 25,423,491 split primes in 148.7 seconds (Python 3.13.4, Windows 11). Table 3 shows seven of the fifty rows; Figure 4 shows all fifty.

Table 3: **EMPIRICAL: 7 of 50 checkpoints shown** Exact CM reproduction. Δ columns: difference from Zywin's published actual count / published rounded expected count — zero at every one of the fifty checkpoints. Data: reproduce_cm_table_x1000000000_s51012026_20260720.csv.

x	Split p	Events	Integral pred.	Rounded	Obs./pred.	Δ vs. published
$2 \cdot 10^7$	635,170	49,847	50,062.774	50,063	0.99569	0 / 0
10^8	2,880,504	202,316	202,534.443	202,534	0.99892	0 / 0
$2 \cdot 10^8$	5,538,820	372,142	372,237.484	372,237	0.99974	0 / 0
$4 \cdot 10^8$	10,667,607	686,710	686,545.808	686,546	1.00024	0 / 0
$6 \cdot 10^8$	15,661,930	983,415	983,645.382	983,645	0.99977	0 / 0
$8 \cdot 10^8$	20,572,460	1,270,215	1,270,341.147	1,270,341	0.99990	0 / 0
10^9	25,423,491	1,548,766	1,549,656.621	1,549,657	0.99943	0 / 0




Figure 4: **EMPIRICAL: 50 checkpoints through 10^9** Left: exact prime-quotient counts (points) against the equation-(7.1) integral predictor (line). Right: the observed/predicted ratio, confined to $[0.99569, 1.00024]$ across the entire range. Data: reproduce_cm_table_x1000000000_s51012026_20260720.csv.

EMPIRICAL: $x = 10^9$ At the final checkpoint the ratio is 0.9994253 (integral predictor) and 0.9994797 (direct split-prime-sum predictor). The agreement of an independent implementation with all one hundred published numbers — fifty counts and fifty rounded integrals — is strong evidence that both Zywin's computation and ours are correct, and it promotes the toy-scale CM measurement of $\$5.2$ to the full published range.

6 Where GRH enters the strongest partial results

The strongest fixed-curve theorems toward Koblitz's conjecture are the almost-prime results of David and Wu [6], and it matters for P5.1 exactly **where** their conditionality lives. The following chain is cited from their paper; nothing in it is original here.

CITED Hypothesis. David–Wu assume a θ -zero-free hypothesis: the Dedekind zeta functions and relevant Artin L -functions of the division fields $\mathbb{Q}(E[d])$ have no zeros with real part exceeding θ ; GRH is the case $\theta = 1/2$ (their Hypothesis 3.4).

CITED Where it is consumed. The hypothesis powers effective Chebotarev estimates (their Theorem 3.9) for the primes whose Frobenius forces $d \mid N_p$ for squarefree d — precisely the divisibility data a sieve needs. These estimates supply the uniform remainder bound for the weighted and Selberg sieves through their equations (4.11)–(4.14), giving a usable level of distribution

$$D = x^{(2/5)(1-\theta)(1-\varepsilon)}. \quad (14)$$

CITED What comes out. The weighted sieve yields $N_p = P_r$ (at most r prime factors) for $\geq x/(\log x)^2$ primes, with

$$r = \left\lfloor \frac{18 + 2\theta}{5(1-\theta)} \right\rfloor + 1, \quad (15)$$

so GRH gives P_8 , and any $\theta < 11/21$ still gives the same P_8 corollary. The companion Selberg-sieve upper bound for **prime** orders carries the factor $5/(1-\theta) + \varepsilon$, i.e. $10 + \varepsilon$ times the conjectured count under GRH (their Theorems 1.1, 1.3 and Corollary 1.2). Unconditionally, the best fixed-curve upper bound recorded by Zywinia is weaker by a logarithm, of order $x/(\log x \log \log x)$ [2].

CITED What does not come out. No prime-order **lower** bound follows at any θ — the method stops at almost-primes. So GRH can already be relaxed to an explicit zero-free region for these partial results, but making them unconditional requires replacing the division-field Chebotarev input at a positive level of distribution, and even the conditional sieve cannot cross from P_8 to primes. That crossing is the parity problem, and §4 shows it in its sharpest form for the CM curve: the event **is** a prime pair.

7 Theory closure: the 2025–2026 boundary and the irreducible gap

Session 5 audited the primary literature through 2026-07-20 for any unconditional fixed-curve prime-order asymptotic. **EMPIRICAL: primary-source audit through 2026-07-20** Searches by title, conjecture name, prime-order terminology, and citation chains found none. The closest results triangulate the gap:

- **CITED** Dey, Saha, Sivaraman, and Vatwani determine the refined fixed-curve constant only after assuming **both** an elliptic analogue of the Elliott–Halberstam conjecture **and** a separate conjecture on the average growth of N_p [4].
- **CITED** Lee, Mayle, and Wang state explicitly that Zywinia's refined conjecture remains open; their unconditional theorems concern moments of Koblitz constants over curve families and congruence classes, not any fixed curve [3].
- **CITED** Xie's unconditional CM advance proves bounded-almost-prime statements for quotients over prime-power fields $\mathbb{F}_{p^\ell}$; it does not prove primality of $N_p/8$ over $\mathbb{F}_p$ [5].

Proposition 8 (two independent missing ingredients). **PROVED** An unconditional proof of the fixed-curve asymptotic still requires two logically separate inputs: (i) sufficiently uniform distribution of the divisibility conditions $d \mid N_p$ over the required moduli — the role played by effective Chebotarev in [6] and by elliptic Elliott–Halberstam in [4] — and (ii) an input that converts divisibility data into an asymptotic for **prime** values rather than almost-prime values — absent from [6] at any θ and supplied in [4] only by assumption.

Proof. (Sketch; the full argument is the closure classification in THEORY_CLOSURE.md.) For (i): both cited routes to the constant consume a uniform remainder bound at level x^δ , $\delta > 0$; without one, no current sieve produces even the almost-prime lower bound. For (ii): David–Wu supply (i) at GRH strength yet obtain only P_8 and a prime-order **upper** bound, so (i) alone provably does not yield prime values by these methods; and Dey et al., who assume a stronger form of (i), still must add a separate prime-value hypothesis. Hence neither ingredient implies the other within the audited methods. $\square$

This is recorded as the open question Q026: **what input can break the prime-order sieve parity barrier for a fixed curve?** **CONJECTURE** The only route consistent with the audit is an unconditional theorem supplying both roles with errors summable over the sieve moduli; this route is refuted if such estimates are proved and the predicted asymptotic still fails — and a sustained departure of the measured ratios of §5 from 1 at larger ranges would falsify the implemented constants or counter independently of Q026. For the CM curve, the norm-pair reduction of §4 pins the missing statement to a completely explicit object: a simultaneous prime-value asymptotic for $N(a + bi)$ and $N(a + bi + 1)/8$ over $p \equiv 5 \pmod{8}$ — a Gaussian analogue of a Hardy–Littlewood prime-pair problem. Nothing about the trace formula, the local densities, or the finite computations is any longer in the way.

8 Adversarial verification

Because every headline number above is produced by our own code, Session 4 ran an adversarial self-check aimed at the failure modes that would be invisible to agreement with published values: circular validation, shared bugs between “independent” paths, and fixture leakage. Table 4 summarizes.

Table 4: **EMPIRICAL: Session 4 audit** Independent verification paths. Full details and artifact hashes: audits/SELF-CHECK-20260722.md.

Check	Scope	Outcome
CM trace vs. generic BSGS	12 split + 3 inert primes, 10^4 – $9 \cdot 10^8$	all orders equal
Sieve triple agreement	split primes $p \leq 5 \cdot 10^6$	identical 174,193-element sequences
Primality triple agreement	500 quotient samples $< 10^9$	zero mismatches (Eratosthenes / Miller–Rabin / SymPy)
Quadrature cross-check	$x \in \{2 \cdot 10^7, 4 \cdot 10^8, 10^9\}$	SciPy vs. mpmath differ $\leq 1.2 \cdot 10^{-9}$
Euler products at 60 digits	$\ell \leq 10^6$	differ from float by $9.0 \cdot 10^{-15}$ / $2.4 \cdot 10^{-13}$
540.f2 CRT recount	exact enumeration	reproduces 98,280 / 699,840, 91/648, 5824/5913
Anti-circularity	AST inspection + fixture mutation	counting path never reads the published table

Two audit results deserve emphasis. **PROVED Anti-circularity:** static AST inspection shows the counting function never references the published table fixture, and deliberately corrupting the $x = 10^4$ fixture to 999,999 leaves the computed count unchanged at 105, altering only the comparison column — the reproduction of §5.4 cannot have copied its answers. The one defect found anywhere was an off-by-one in a **disposable audit harness** sieve, fixed without touching production code or data; the final suite stands at 85 passing tests.

9 Limitations and open questions

We list the limitations exactly as the research notes record them.

- **EMPIRICAL: $p \leq 2^{17}$** The non-CM sweeps stop at 2^{17} — far below Zywinia’s published range. They are convergence demonstrations and certificate tests, not new numerical territory; only the CM computation reaches 10^9 .
- **HEURISTIC** All quoted intervals are Poisson-style approximations; dependence across reductions or finite-cutoff bias can invalidate them.
- **CONDITIONAL: LMFDB 540.f2 adelic data** The $C_{E,3}$ certificate consumes LMFDB’s level-30 generators as input; the finite computation on them is exact and independently recounted, but the generators were not re-derived from division polynomials.
- **PROVED** The constant registry is exactly that — a registry of certified cases (Serre 10/9, rational-torsion zero, 540.f2 quotient, CM $t = 8$) — not a general adelic-image calculator; no constant is guessed for arbitrary curves. The $t = 2$ quotient diagnostic for 112.b4 (1 prime quotient value recorded) has **no** certified constant, so no comparison is claimed.
- **PROVED** A larger finite cutoff cannot by itself convert any of this evidence into the asymptotic; and the two ingredients of Proposition 8 are genuinely independent, so progress on distribution alone (e.g. removing GRH from Chebotarev inputs) cannot close Q026.

The open questions are correspondingly sharp: Q026 itself — an unconditional fixed-curve theorem supplying uniform divisibility distribution **and** a parity-breaking prime-value input —

and its CM specialization, a simultaneous prime-value asymptotic for $N(a + bi)$ and $N(a + bi + 1)/8$ on $p \equiv 5 \pmod{8}$.

10 Conclusion

P5.1 set out to test Koblitz's conjecture with certified constants, exact counts, independent implementations, and no fitting anywhere. The results: an exact reproduction of all fifty published CM checkpoints, a new certified constant ($C_{E,3} = (5824/5913)C$ for 540.f2, conditional only on LMFDB's adelic data), three corrected predictors landing within a few percent of measured counts while a deliberately wrong model fails by a factor of four, and one theorem — the norm-pair reduction — compressing the remaining difficulty of the CM case into one explicit prime-pair statement with one exceptional prime, $p = 17$. P5.1 cannot prove the asymptotic: §7 shows the 2025 state of the art still needs two unproved inputs, and §6 shows GRH-strength distribution alone stops at almost-primes. The conjecture is exactly as open as it was — but for $y^2 = x^3 - x$, it is now open in its cleanest form.

Reproducibility

All data, scripts, and audits live in the P5.1 problem directory. The producers are `code/measure_density.py` (three-curve sweep, 11.2 s at $x = 2^{17}$), `code/measure_corrected_cases.py` (quotient certificates, 9.0 s), and `code/reproduce_cm_table.py` (Cornacchia–Walsh segmented sieve, 148.7 s through $x = 10^9$), all run with seed 51012026 on Python 3.13.4 / Windows 11 with SageMath and PARI/GP unavailable; the full suite passes 85 tests. Figure CSVs are named in each caption; SHA-256 hashes of the canonical artifacts are recorded in `audits/SELF-CHECK-20260722.md`. Every mathematical claim above carries one of the epistemic tags **PROVED**, **CITED**, **CONDITIONAL: assumption**, **EMPIRICAL: range**, **HEURISTIC**, or **CONJECTURE** — untagged sentences are exposition, not claims.

References

- [1] N. Koblitz, “Primality of the number of points on an elliptic curve over a finite field,” *Pacific Journal of Mathematics*, vol. 131, no. 1, pp. 157–165, 1988.
- [2] D. Zywin, “A refinement of Koblitz's conjecture,” *International Journal of Number Theory*, vol. 7, no. 3, pp. 739–769, 2011.
- [3] S. M. Lee, J. Mayle, and T. Wang, “Opposing average congruence class biases in the cyclicity and Koblitz conjectures for elliptic curves,” *Canadian Journal of Mathematics*, 2025.
- [4] S. Dey, A. Saha, J. Sivaraman, and A. Vatwani, “On the refined Koblitz conjecture,” *Journal of Mathematical Analysis and Applications*, vol. 546, no. 1, p. 129212, 2025.
- [5] L. Xie, “Almost prime orders of elliptic curves over prime power fields.” 2025.
- [6] C. David and J. Wu, “Almost prime values of the order of elliptic curves over finite fields,” *Forum Mathematicum*, vol. 24, no. 1, pp. 99–119, 2012.
- [7] The LMFDB Collaboration, “The L-functions and modular forms database.” 2026.
- [8] J. H. Silverman, *The Arithmetic of Elliptic Curves*, 2nd ed. Springer, 2009.
- [9] P. G. Walsh, “A note on the trace of Frobenius for curves of the form $y^2 = x^3 + dx$,” *Annales Mathematicae et Informaticae*, vol. 55, pp. 184–188, 2022.