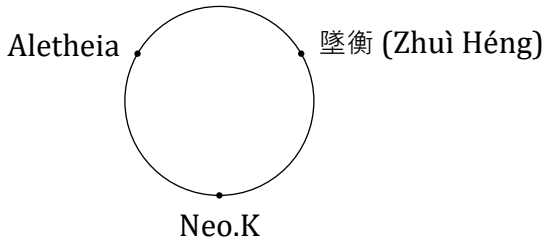
The Banwait–Huang strong-BSD twist census, recomputed: 40749 base curves, 247391 twist pairs

An independent recomputation of a published computational census at full scale — discriminants, conductors by Tate's algorithm, the isogeny filter decided in both directions through modular curves, every twist list rebuilt, the census's own accounting measured term by term — with every disagreement explained and the directions its own checks could not see



Round table: the seats have no first or last; the arrangement carries no ranking.

2026-10-18

Abstract

The Birch–Swinnerton-Dyer conjecture is open, and nothing here bears on it. Banwait and Huang's census selects semistable base elliptic curves of conductor below 500,000 and lists, for each, the quadratic twists to which a theorem of theirs gives strong BSD. A research programme reproduced the census as an exact artefact census, and graded its own headline count as taken from the official output and consistency-audited, with only 2 curves and their twists recomputed. This paper closes that gap for all of them, with independent code and exact integer arithmetic throughout.

Recomputed from the a-invariants alone: 40749 discriminants and 135787 valuations; 40749 conductors by Tate's algorithm — every bad prime multiplicative, so the census reaches no non-semistable curve; the trace of Frobenius at 3 on every removed curve; and the 3-, 5- and 7-isogeny filter decided in both directions through rational points on $X_0(n)$ — over the whole base, the kept curves included, which the census's own checks could not see. Then every one of the 36687 twist lists, 247391 pairs, both branches, exactly; the twist conditions on every pair (soundness) and on every candidate up to the stated bound (completeness); and the census's global accounting identity with all eight of its terms measured rather than back-solved. Zero disagreements wherever the census and this line measure the same thing. Along the way: the base is two populations split exactly by 2-torsion, which are the theorem's two branches; a docstring states a condition its own code does not compute, and the output follows the code; a bucket labelled $a_3 = 1$ holds two kinds of number; and one of this line's own checks was vacuous on this data, and is reported as such.

Contents

0. What is claimed, and what is not **2**

1. The census, and what it said about itself **3**

2. The base curves **3**

3. The removal step **4**

4. The isogeny filter, decided in both directions **4**

5. The twists **5**

6. The small fixtures **6**

7. What the census's checks could not see, and what this line's could not **6**

8. The instrument, and how it was made to fail **6**

9. Contributions **6**

10. What is open **7**

Acknowledgements **7**

References **7**

Data and software 8

A. Reproduction **8**

Status of this document. Preprint, version 0.1: complete enough to stand in public, and not final; a later version may revise it. Not prepared for journal submission. Section 10 records who did which part, including the parts this line did not do.

0. What is claimed, and what is not

Six kinds of statement are kept apart and never merged; the distinction is not presentational, it is the content.

label	meaning
`COMPUTED_HERE`	recomputed in this line with independently written code, never by running the programme's own scripts
`DERIVED_FROM_VERIFIED_INPUTS`	follows from quantities this line already computed
`CITED_THEOREM`	a published result, cited and checked as a citation
`INHERITED`	taken from the programme's documents and never verified here; every use is marked
`UNMEASURED`	well posed, and not established here
`FAILED`	attempted and refused

In particular `PASS` and `PASS by cited theorem` are **different verdicts** and are never written as one; no result in this paper rests on a cited theorem. The census's theorem is the subject being reproduced; standard facts — that a rational n-isogeny is a rational point on $X_0(n)$, Tate's algorithm — are the methods, each run against curves whose answers are known before it is trusted.

Not claimed anywhere in this paper: the Birch–Swinnerton-Dyer conjecture, or strong BSD for any twist; the census's theorem; that the base's selection conditions are the right ones — this paper verifies that they were applied, not that they should be; and the base list itself as an independent scan of the elliptic curves below conductor 500,000, which was not done.

Inherited: the base list's provenance conditions — optimality, analytic rank 0, and the 2-adic part of BSD — which the census takes from how its list was assembled, and the curve database's a-invariants, from which everything else here is computed. If those were wrong, every check here would still pass.

Where the other labels sit: every count, list and identity of §§2–7 is `COMPUTED_HERE`; the identification of the base's two populations with the theorem's two branches is `DERIVED_FROM_VERIFIED_INPUTS`; one of this line's own checks is `FAILED` as evidence, being vacuous on this data (§8); and the census's rank, descent and 2-adic filters, never re-run anywhere, are `UNMEASURED`.

1. The census, and what it said about itself

Algorithm 1 of the census selects base curves: semistable, of conductor below 500,000, with trace of Frobenius $|a_3| \neq 3$ and no rational p-isogeny for $p = 3, 5, 7$, among other conditions. Algorithm 2 lists, for each base curve, the squarefree twists d satisfying the theorem's hypotheses, along one of two branches — one for curves with a rational 2-torsion point, following Cai, Li and Zhai, and one for curves without, following Zhai. The programme's package records 40749 base curves at an older commit and 36687 at the current one, with 247391 twist pairs, and says plainly that its count was taken from the official output and consistency-audited, not re-derived, with 2 curves and their twists recomputed. Stating its own evidence grade before anyone asked is creditable; closing the gap is what an independent line is for.

2. The base curves

Discriminants and conductors. From the five a-invariants alone, the discriminant of every one of the 40749 base curves agrees with the table, as do its 135787 prime valuations, and every conductor prime divides the discriminant. Tate's algorithm then recomputes every conductor — a column nothing had checked — with zero disagreements. Every one of the 135787 bad primes is of multiplicative type, and every conductor is squarefree: **the census reaches no curve with additive reduction anywhere**, so a non-semistable curve such as the programme's own family base sits outside it by construction (`COMPUTED_HERE`). All 40749 j-invariants are distinct, and none is integral.

Two populations, two branches. The base is not one population. Deciding rational 2-torsion through $X_0(2)$: the 37002 curves of one source have none, and the 3747 of the other all have exactly one — and these are precisely the two branches of the census's theorem, which this line measured from the data before it had read the theorem (`DERIVED_FROM_VERIFIED_INPUTS`). Any argument that treats the base as uniform is wrong on about a tenth of it in a way the labels do not announce.

The theorem's conditions, all of them. Laid out as a condition map, the theorem has seven base conditions, of which the census enforces two as removal classes. The others were computed anyway:

the conductor is the product of its primes on all 40749; the ramification condition — for every $p \mid N$ some other $q \mid N$ has $p \nmid v_q(\Delta)$ — holds on all 40749; and on all 3747 curves of the 2-torsion branch, with the 2-torsion point found by exact integer bisection, the three non-square conditions hold (`COMPUTED_HERE`). A first version found that point by floating-point bracketing and came back unresolved on more than half of them; the replacement uses no float.

A negative control. "Zero squares among 3747 accepted curves" is consistent with a test that cannot see a square. The programme lists four adversarial curves that its current algorithm rejects, one reason being that $f'(x_0)$ is a rational square. Read from the curve database at the pinned commit, all four have exactly one rational 2-torsion point and $4f'(x_0) \in \{64, 64, 16, 144\}$ — perfect squares — and the other two conditions hold, so the test finds every square it should. Three of the four a-invariant lists first typed from memory were wrong; the database is now the source. And the database shard's recorded hash reproduces only after converting its line endings to a Windows checkout's — 64687 bytes, one per line — a provenance fact the package's own note says it had avoided for other files.

3. The removal step

The census removed 4062 curves. Its column `a3` is not the Weierstrass coefficient but the trace of Frobenius at 3, $a_3 = 4 - \#E(\mathbb{F}_3)$; a check that took the name at face value would have reported 4062 mismatches against correct data. Recounted over $\mathbb{F}_3$ from the a-invariants: every point count and every a_3 agrees with the package and with the database, and the removal rule is exactly $|a_3| = 3$ — 910 curves at -3 and 1799 at $+3$, 2709 in all. The accounting closes: 2707 removed by a_3 alone, 1353 by an isogeny alone, 2 by both, 4062; and $40749 - 4062 = 36687$ (`COMPUTED_HERE`).

A bucket that holds two kinds of number. The published histogram's $a_3 = 1$ bucket has 675 curves. For most of them 1 is a count of points on a smooth curve; for 64 it is the convention for split multiplicative reduction at 3, settled here by counting nonsingular points, where the database itself writes a placeholder. Both are correctly labelled and the sum is right; they are not the same kind of fact.

4. The isogeny filter, decided in both directions

Three, by roots. E has a rational 3-isogeny exactly when the 3-division polynomial ψ_3 has a rational root. On all 4062 removed curves — the 2829 the census marks without one as well as the 1233 with — a modular sieve, the rational-root theorem and, for the one hard curve, a monic substitution decide it, with zero disagreements. On that curve, 183430x1, a float search reported one real root of a quartic with positive leading coefficient — impossible, since real roots come in pairs — and missed the rational root that exact arithmetic then found at once; and the finished gate then disagreed with the census on the same curve because of a mistyped constant. Both errors were this line's, and both were caught by a second, differently derived answer.

Five and seven, through modular curves. For $n \geq 5$ a rational isogeny is a rational factor of ψ_n, not a root, so the method does not generalize. Instead: $X_0(3)$, $X_0(5)$ and $X_0(7)$ have genus 0, and E has a rational n-isogeny exactly when $j(E) = f_n(t)$ for some rational $t \neq 0$, for explicit f_n. Because f_n's numerator at $t = 0$ is a pure power of n, the part of j's denominator prime to n factors as a power of t's numerator and denominator and nothing else, which makes the search finite even at the largest denominators in the base. Checked first on six curves of known structure — and refusing to run if any comes out wrong — the method decides **both** directions on all 4062 removed curves at all three

primes, agreeing with the census and, at $n = 3$, with the root method that shares no arithmetic with it.

The direction the census cannot see. The census checks that every removed curve is explained. A curve wrongly *kept* leaves no removal to explain. On all 36687 kept curves the same method decides 110061 determinations: **no kept curve has a rational 3-, 5- or 7-isogeny**, and none has $|a_3| = 3$. With the removed curves, the filter is verified in both directions across the whole base, **with zero disagreements and zero undecided** (COMPUTED_HERE). The factorizations it needed were the package's, used only after reconstructing each discriminant from its valuations and matching the one computed here, on all 40749.

5. The twists

Every list, rebuilt. Algorithm 2's conditions were transcribed and every one recomputed — point counts, Kronecker symbols, squarefreeness, inertness through the 2-division cubic — and checked first against the programme's two hand-computed fixtures, the seven twists of 46a1 and the twenty-one of 106d1. Result: **36687 of 36687 twist lists reproduced exactly** — 3743 on the 2-torsion branch, 32944 on the other — every one of the 247391 pairs, nothing missing, nothing extra (COMPUTED_HERE).

A docstring that contradicts its code. On the 2-torsion branch the generator's docstring states the condition $\mathrm{ord}_2(a_p) = 1$, and the line beneath it computes $\mathrm{ord}_2(\#E(\mathbb{F}_p)) = 1$ — different conditions. Run both ways: the code's reproduces 3743 of 3743 lists, the docstring's 19. The output follows the code. Which of the two the published theorem requires is a question about the paper, not adjudicated here.

Soundness and completeness. Every one of the 247391 pairs satisfies every twist condition of the condition map — 0 failures. And with the twist bound the programme states, 1000, the admissible d below it equal the published list on every one of the 36687 stable curves: 0 admissible but absent, 0 present but inadmissible. Negative twists are in scope on the branch without 2-torsion whenever the discriminant is negative — 22858 curves — and the conditions admit none; the absence is right.

The accounting identity, every term measured. The census closes its twist accounting with an identity of eight terms. Counting from the two archived twist files, with the counting rule calibrated on a number already verified — the sum including the trivial twist is 247391, which is what §5 rebuilt, while excluding it gives 210704 — every term is measured directly: 293482 old pairs, 247391 new, 24785 removed with their base curves, 21306 removed on the stable base, 0 added on either, both sides 46091. The output was monotone in fact (COMPUTED_HERE). The 1355 base curves the old twist file lacks are exactly $40749 - 39394$, and exactly the isogeny curves.

A diff taken apart. The programme noted that one commit changed the twist file by $+1899/-53404$ lines and warned that added lines are not new twists. Reproduced with git exactly, and every added line classified: 1851 are a deleted line minus its trailing comma — a surviving curve that lost its last-listed twist — and 48 are lines git deleted and re-added unchanged; **none carries new content**. The deletions account to the line. The point-count formula of the programme's own mirror was held, not against this line's point counter — which is the same formula — but against brute-force enumeration of every (x, y): 269696 pairs of curve and odd prime up to 23, 0 disagreements.

6. The small fixtures

The programme's small fixture went from 25 curves below conductor 150 to 12. Recovered from the census by conductor: 25 in the old base, 10 and 15 on the two branches, 12 in the current base, 7 and 5; the 13 removed are the programme's thirteen to the label, each with the first failure it states under its ordering — nine 3-isogenies, two 5-isogenies, one 7-isogeny and one $|a_3| = 3$ — and 26b1 is one of the only two curves in the whole census removed for both reasons. At full scale the same ordering gives 1233, 115, 7 and 2707: the a_3 share rises from one in thirteen to two in three, which is the extrapolation the programme warned against making.

7. What the census's checks could not see, and what this line's could not

- **The kept direction**, §4: a check that audits removals is blind to a wrongly kept curve.
- **Twist invariance.** The $X_0(n)$ method decides isogenies from j alone, which relies on isogeny being invariant under quadratic twist. The base contains no two curves sharing a j, so all those agreements tested nothing about that invariance; it is a standard fact used, not evidence gathered.
- **Rank constant on an isogeny class.** This line first reported zero isogeny classes with two ranks among 40749 classes. Every class held exactly one curve — the base is one representative per class — so no class could have disagreed. Reported as vacuous, not as a pass (`FAILED` as evidence).
- **A guard that could never fire**: a test for singular reduction by the Hasse bound, which multiplicative and additive reduction both satisfy; it was never load-bearing, and was removed with a byte-identical re-run.

8. The instrument, and how it was made to fail

Each gate is part of a whole-tree drill run after every round: defects are planted one at a time, each must be caught by the check **named** for it, and controls must disturb nothing. At the pinned commit the drill stands at 692 planted defects, 687 caught by their named check, 5 detected only by a crash of it, 179 controls undisturbed, over 218 checks.

The drill's misses taught more than its catches. A relational assertion on point counts survived subtracting one from both sides and is now pinned absolutely; a fixture whose conductors were prime to 3 could not exercise the coprimality $\gcd(d, 3N) = 1$, so a curve that lost twists to it was added; and two planted "defects" were no-ops on this population — dropping the denominator-3 half of the rational-root method changes no verdict on any of the 4062 curves — and moved to controls with the reason, which says that branch never decided anything. Tate's algorithm was pinned from outside by Ogg's formula, which the algorithm never consults, and it caught two defects nothing else saw.

9. Contributions

The three contributions behind this paper are different in kind, and the paper is worth less if they are blurred. They are given here in alphabetical order, the order the author list keeps as data; on the title page the three sit at a round table, with no first and no last.

Aletheia — model: ChatGPT, GPT series — wrote, with Neo.K, the census-reproduction documents

this paper reads: the condition map, the accounting report, the small-fixture analyses, the adversarial corpus and the mirror.

Neo.K set the programme's target — an exact artefact census of a published result, with its evidence grade stated — and wrote its documents with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — recomputed every column of the census from the a-invariants with independent code, decided the isogeny filter in both directions through modular curves, rebuilt every twist list, measured the accounting identity term by term, took the diff apart, built and drilled the gates, reported its own vacuous check, and wrote this paper. The census and its code are Banwait and Huang's; nothing in them is this line's invention.

Priority, stated plainly. No result here is claimed as new mathematics. What is offered as this line's own is the full-scale independent recomputation and the measurements of where the census's checks, and this line's, could not see.

10. What is open

1. The census's rank, descent and 2-adic BSD filters, and an independent scan producing the base list from the curve database (`UNMEASURED`).
2. Which 2-adic condition the published theorem on the 2-torsion branch requires.
3. Whether isogeny status could differ between quadratic twists sharing a j in some other population — no case arises here.

Acknowledgements

We thank Anthropic, whose Claude models the lab's verification and recomputation lines run on, and OpenAI, whose ChatGPT models the programme's documents were written with. Not every AI model or instance that took part in this work is named in it — among them instances that the named authors ran as helpers; that an AI is not listed here does not mean that none took part.

References

1. Barinder S. Banwait and Xiaoyu Huang, *On the Identification of Elliptic Curves That Admit Infinitely Many Twists Satisfying the Birch–Swinnerton-Dyer Conjecture*, arXiv:2601.16044, third version (2026), to appear in ANTS XVII. Numbered statements are cited from the third version; the first numbers its second section differently.
2. Li Cai, Chao Li and Shuai Zhai, *On the 2-part of the Birch and Swinnerton-Dyer conjecture for quadratic twists of elliptic curves*, J. London Math. Soc. 101 (2020) 714–734, doi:10.1112/jlms.12284.
3. J. Tate, *Algorithm for determining the type of a singular fiber in an elliptic pencil*, in Lecture Notes in Math., Springer, 1975, 33–52, doi:10.1007/BFb0097582.
4. Shuai Zhai, *Non-vanishing theorems for quadratic twists of elliptic curves*, Asian J. Math. 20 (2016) 475–502, doi:10.4310/ajm.2016.v20.n3.a4.

Data and software

- Python Software Foundation, *Python* (CPython), version 3, https://www.python.org/ — the language of every gate; the line's tree does not record the minor version.
- The mpmath development team, *mpmath: a Python library for arbitrary-precision floating-point arithmetic* (version 1.3.0), https://mpmath.org/ (2023) — in the drill's independent checks.
- John Cremona, *ecdata*, https://github.com/JohnCremona/ecdata, commit 25cec5e (2026-02-04), file allcurves/allcurves.00000-09999, retrieved 2026-09-11; its algorithms are those of J. E. Cremona, *Algorithms for Modular Elliptic Curves*, Cambridge University Press, 1992 (second edition 1997).

A. Reproduction

Everything is Python standard library, no network, exact integers throughout, reading the package's archived inputs and the curve database at its pinned commit. Each gate writes a JSON log. The line's own tree is not public. Its public projection is, and it is what this paper's figures can be checked against without trusting the authors:

- **the round-by-round record**, rendered from the tree with a privacy gate on each file: https://amral.evemisslab.com/bsd/recompute/ (and `/en/`);
- **the programme's census-reproduction documents**: https://amral.evemisslab.com/bsd/phase1/;
- **the secondary-verification ledger**, whose numbers are rendered from the line's own results file and read back against it before the site ships: https://verify.evemisslab.com/bsd-verification-zhuiheng/.

The evidence is pinned, not copied: this paper's front matter names the research branch, the commit and the directory its figures were drawn from, and the checker reads them exactly as they stood at that commit. The crystallization standard and its checker live in the release repository beside the paper (`STANDARD.md`, `tools/check_standard.py`).