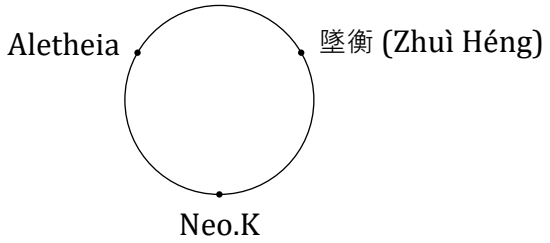
Grading evidence for the Birch–Swinnerton-Dyer conjecture curve by curve: an eleven-rung ladder, an audited rejection, and a frontier metric faithful only in exact arithmetic

A research programme's framework documents, checked: the ladder every later claim is placed on and the five rungs every summary drops, the programme's rejection of its author's own route and why that route fails at infinite precision, a global frontier quantity whose faithfulness survives in the rationals and not in floating point, and the programme's maps scored after sixty-four rounds



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. This paper is about how evidence for it should be stated. A research programme on the conjecture opens with framework documents, and the first of them refuses to let a curve carry a bare true or false: it defines an eleven-rung certificate ladder, from the curve's identity (C0) through numerical and rigorous analytic rank (C2, C3), algebraic bounds, weak and single-prime strong BSD, finiteness of Ш, full strong BSD, to a family theorem (C10), so that each curve and prime is placed at the rung some certificate has actually closed. Checked against the programme's 85 curated documents, the ladder's vocabulary is used cleanly — no document names a rung the framework does not define — but every summary of it renders six rungs, dropping C1 to C5, and with them the one boundary the framework forbids crossing: a number a computation produced against a number a proof produced.

Three further results. The programme's audit rejecting its own author's "lattice-point rank convergence" route holds on all four objections, nothing in the corpus claims the route's salvage conditions met, and a later document refutes the route outright with $f_a(z) = z^2 + az$; checked here in exact rationals, that refutation has a second half — the zero count in a fixed neighbourhood does converge while the order at the point does not — which says the route fails at infinite precision, not for lack of it. The programme's frontier quantity $\mathfrak{B}_k(s) = \sum_{i \in H_k} i^{-s}$ satisfies all four of its stated properties in the rationals, one hypothesis is unnecessary, and its design goal — never to swallow a single uncertified curve — fails in double precision at a computable index, 208064 for $s = 3$. And the programme's three routing maps, read after sixty-four rounds of this line, are scored by what can be computed about them.

Contents

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

1. The programme and its framework **3**

2. The ladder, and the five rungs every summary drops **3**

3. A rejected route, and why it fails at infinite precision **4**

4. A frontier quantity, faithful in the rationals **5**

5. The maps, read after sixty-four rounds **6**

6. The instrument, and how it was made to fail **7**

7. Contributions **7**

8. What is open **7**

Acknowledgements **8**

References **8**

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 7 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. Hurwitz's theorem, named in §3, explains a measured pattern and is not used to prove anything.

Not claimed anywhere in this paper: the Birch–Swinnerton-Dyer conjecture, or that any curve reaches any rung; that the lattice-point route's salvage conditions are unmet — only that no document in the corpus claims them met; that the abbreviated ladder misled any reader; and a replacement weight for the frontier quantity, which is a research decision rather than an arithmetic fact.

Inherited: the rejected route itself, which is not in the corpus and is known here only through the audit's seven-step paraphrase of it (§3). The routing judgements of §5 — which round belongs to which route — are a reading, stated as one.

Where the other labels sit: the corpus measurements and exact computations of §§2–5 are `COMPUTED_HERE`; that the frontier quantity's limit property needs no monotonicity, and the index at which a float computation of it goes blind, are `DERIVED_FROM_VERIFIED_INPUTS`; the frontier quantity's faithfulness in double precision is `FAILED` past that index; and whether the abbreviation has had any effect on a reader is `UNMEASURED`.

1. The programme and its framework

The programme studies the conjecture through four sub-lines — a framework, a single-prime study of the rank-two curve 389.a1, an independent reproduction of a published twist census, and an explicit twist family — curated as 85 documents. Its framework, Phase 0, has nine documents, and they route rather than prove: a statement audit, a closure map of known theorems, the certificate ladder, an external route matrix, the audit of the author's earlier idea, an agent experiment, a frontier quantity, and briefs for the agents that would carry the work. This paper checks the parts of the framework that can be checked, and measures the rest against what this line's rounds later did.

2. The ladder, and the five rungs every summary drops

The ladder exists, in the framework's words, because a curve must not carry a bare boolean: what is stored is which rung some certificate has closed. Its eleven rungs:

rung	name	what closes it
C0	identity	minimal model, isogeny label, conductor, discriminant, CM status, provenance
C1	local arithmetic	bad primes, reduction types, Tamagawa numbers, torsion, root number
C2	numerical analytic rank	a high-precision computation giving $r_{\mathrm{an}} = r$ — status: *evidence*
C3	rigorous analytic rank	$\mathrm{ord}_{s=1} L(E,s) = r$ by interval arithmetic, a Turing-type count or a theorem

rung	name	what closes it
C4	algebraic lower bound	r independent points, with heights and an independence certificate
C5	algebraic upper bound	descent: Selmer groups, local conditions, saturation
C6	weak BSD	$r_{\mathrm{alg}} = r_{\mathrm{an}}$
C7	single-prime strong BSD	the p-part of the leading-term formula, with every assumption kept
C8	Ш finite and exact	#Ш determined independently — never read back from the formula
C9	full strong BSD	every component closed, conventions and normalizations explicit
C10	family theorem	a uniform statement over an infinite family, its exceptions and hypotheses kept

It also lists five prohibitions, the second of which is the one this section is about: an integer returned by a rank routine is not a proof. That is the C2/C3 boundary.

Measured. Every rung named anywhere in the 85 documents was compared with the eleven defined: 5 documents name a rung, and **0 use a rung the framework does not define**; every rung is named in three or four places. Then the places that summarize the ladder for a reader were read: the programme's research-tree README and the public page both render it as C0, C6, C7, C8, C9, C10 — **dropping C1 to C5**, the same five each time, and with them the C2/C3 boundary. So did this line's own README as first written, copied from a derived summary within an hour of the line opening; it now carries all eleven, and the gate reads it to check that it does (`COMPUTED_HERE`).

An abbreviation is not a defect, and showing a ladder's arc is a fair thing for a summary to do. What makes it worth stating is which five vanish: they are where the ladder discriminates, and the one boundary among them that the framework forbids crossing is exactly the one a reader of the six-rung version never sees.

3. A rejected route, and why it fails at infinite precision

The framework's fifth document audits its author's earlier idea — that a rank-zero-order relation could be reached as the limit of lattice approximations — and ends: archive as an exploratory analogy; do not use as a proof route. A rejection can be too weak, too strong, or unenforced, so it was checked in all three directions.

The four objections hold. (1) If the lattice inequality is assumed at each scale, its proof must not use a BSD-type link, or the argument is circular — stated as an unmet burden, not as a proved circularity, which is the right hedge. (2) The Mordell–Weil rank is an integer invariant that no archimedean approximation converges to. (3) Even where lattice L-functions converge, the order of vanishing at

$s = 1$ need not. (4) Limits of two sides being equal requires them equal at each scale, which moves the difficulty into the lattice layer rather than solving it. Two remarks on coverage: the audit imposes a faithful-discretization condition that none of its objections argues for, and it never asks whether a lattice L-function is definable at all — it attacks the limit of an object it assumes exists.

The verdict held. Across the 85 documents, 61 hits of the route's vocabulary, 21 inside the audit and the rest in 8 documents where "lattice" mostly means a Mordell–Weil or determinant lattice — a name collision the gate reports and does not adjudicate. **0** documents claim any of the four salvage conditions met (`COMPUTED_HERE`). And one later document re-derives the rejection independently and sharpens it from "the burden is unmet" to false in general, with a witness.

The witness, and its second half. $f_a(z) = z^2 + az = z(z + a)$ converges uniformly to $f_0 = z^2$ — $\sup_{|z|\le 1} |f_a - f_0| = |a|$ — while $\mathrm{ord}_{z=0} f_a = 1$ for every $a \neq 0$ and $\mathrm{ord}_{z=0} f_0 = 2$. Checked in exact rationals at a from 1 down to 10^{-12}, of both signs: the order is 1 throughout, the limit's is 2. What neither document states is the other half, also measured: the number of zeros in a fixed small disc about 0 is 1 while $-a$ lies outside it and 2 from the moment it enters, equal to the limit's order — the pattern Hurwitz's theorem predicts. **The order at the point jumps while the count in a neighbourhood is stable.** A method that measures a neighbourhood and reports the point is therefore not imprecise; it is wrong at infinite precision. This is also why the audit's remedy asks for analytic control in a neighbourhood of $s = 1$ rather than at it.

4. A frontier quantity, faithful in the rationals

The framework's seventh document proposes one control number for the whole programme. Enumerate the isogeny classes, fix a target rung, and let H_k be the classes still below it after round k; then

$$\mathfrak{B}_k(s) = \sum_{i \in H_k} i^{-s}, \qquad s > 1.$$

Its stated design goal is a quantity that will not swallow a single uncertified curve because most are certified.

The four claims hold, exactly. A single uncertified class leaves positive mass; $\mathfrak{B} = 0$ exactly when H is empty; monotone certification cannot raise the mass; and $\mathfrak{B}_k \to 0$ implies every fixed class eventually leaves. All four were checked in exact rationals, not floats, because the point of what follows is that the two disagree.

One hypothesis does no work. The fourth claim is stated for a monotone certificate system. It holds without monotonicity: $s > 1$ makes $\sum i^{-s}$ convergent, so pointwise departure gives the limit by dominated convergence. Witnessed by a non-monotone family — the unresolved set alternates between a moving head and a far singleton — in which every class still leaves and the mass falls from 0.012736653 at $k = 39$ to 1.2499609e-5 at $k = 40001$, like $1/(2k)$ (`DERIVED_FROM_VERIFIED_INPUTS`).

The design goal fails in double precision. A class at index i contributes i^{-s}, positive in the rationals, but added to a float total it changes nothing once it falls below half the last place of that total. The first index whose mass is invisible:

s	$\zeta(s)$	first invisible index
1.1	10.584448	48214938938691
1.5	2.612375	27271342416
2	1.644934	94906266
3	1.202057	**208064**

Each index is the first whose term leaves the running double-precision sum unchanged; at $s = 2$ and $s = 3$ it is confirmed by adding the terms in order. (An earlier version of the gate, and of this table, used a coarser threshold and gave earlier indices; found in review.) The larger s, the sooner a float computation goes blind. On the programme's own enumeration of 40749 base curves the last curve still moves the sum, at $s = 3$ only by 1.4876989e-14; an enumeration about five times as long would not. As a statement about the quantity, the design goal holds; as a statement about a computation of it in double precision, it is `FAILED` past that index. The weight also pulls against one of the document's own stated uses — noticing a forgotten high-conductor tail: at $s = 2$, a single unresolved early curve outweighs ten thousand forgotten late ones, because the weight i^{-s} is what makes the first claim true. Choosing a weight that serves both is a research decision, and none is proposed here.

5. The maps, read after sixty-four rounds

Three framework documents could not be checked when they were written, because they route work: a closure map of known theorems, a matrix of external routes with verdicts, and six briefs for agents. After sixty-four rounds of this line, each was scored by what can be computed about it; which round sits on which row is a reading, kept in one place, and every citation in it is checked — each cited round exists, and each attribution carries a phrase the gate finds in that round's report.

- **The closure map.** Its four open rows — weak BSD at rank at least two, Ш finite in general, the full leading coefficient, all curves over $\mathbb{Q}$ — are still open, and no report of this line claims any of them closed. Of the seven components a strong-BSD certificate must close, four were computed here (Tamagawa factors, torsion, the period, the leading term), the regulator was computed without saturation being verified, and the two components concerning Ш never: the one value of #Ш used anywhere was an input inferred from the formula.
- **The route matrix.** The route it marks first choice, strong-BSD twist families, is the most worked, with 23 rounds against 9 for the runner-up; the route it marks red — the author's lattice idea of §3 — has no round, and §3's scan found no document claiming its conditions met.
- **The briefs.** Each of the six has rounds that carried it out, and each of the eight leaps the adversarial brief names — circular use of Ш, numerical-to-proof, finite-to-global, a p-part taken for the whole, rank 0/1 reasoning applied to rank 2, isogeny double counting, database incompleteness, normalization mismatch — has a round that caught an instance, several of them this line's own.

A correction to this line's record. The same instrument, re-run, showed that one of this line's earlier reports had typed that every document of the census sub-line was the subject of a round; measured, 23 of 25 were, and 2 were cited only. The next round took both as its subject, and at the pinned commit every one of the 85 documents is the subject of a round. The earlier report carries the correction.

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

Each gate here 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.

Two failures of the ladder's own gate are worth stating, because both would have produced a wrong answer quietly. It first hung rather than failing — a text string handed to a binary pipe stalls instead of raising, and a gate that hangs returns no verdict either. And a filter written to remove noise removed the signal: the rung scanner skipped fenced code blocks to avoid matching stray labels in samples, and every abbreviated ladder in the corpus is written inside a code fence, so the gate first reported that this line's own README named no rung at all. Both were caught by reading the output against a fact already known, not by any check.

7. 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 framework documents this paper rests on: the certificate ladder, the audit of the lattice-point route, the frontier quantity, and the closure map, route matrix and agent briefs.

Neo.K set the programme's target and its method — that evidence be stated at its actual strength, rung by rung — proposed the lattice-point idea that the programme then audited and rejected, and wrote the framework with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — measured the ladder's use across the corpus, checked the audit in all three directions and recomputed the refutation in exact arithmetic, added its neighbourhood half, verified the frontier quantity's claims exactly and located where a float computation of it goes blind, scored the maps against this line's own rounds, built and drilled the gates, and wrote this paper. The framework is the programme's; nothing in it 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 measurement of how a framework for evidence is used and summarized, and two sharpenings: the neighbourhood half of the refutation, and the index at which the frontier quantity's faithfulness stops being computable.

8. What is open

1. The Birch–Swinnerton-Dyer conjecture, and every open row of the closure map.
2. A frontier weight that keeps the first claim of §4 and also notices a forgotten tail — a choice, not a computation.
3. Whether a lattice L-function can be defined at all, the question the audit of §3 never asks (`UNMEASURED`).

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

This paper cites no published work. Its sources are the programme's own documents and this line's records, pinned in the front matter.

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.

A. Reproduction

Everything is Python standard library, no network, and the frontier quantity is computed in exact rationals. 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 framework documents** this paper reads: https://amral.evemisslab.com/bsd/phase0/;
- **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`).