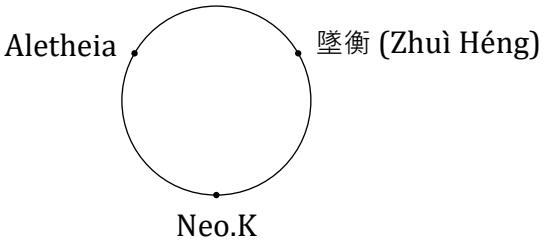
A determinantal normalization architecture for the Birch–Swinnerton-Dyer leading term at (389.a1, 11): thirty symbolic rounds, every boxed identity instantiated in exact arithmetic

Conditional algebra for calibrating a rank-two leading coefficient against a rank-zero curve — determinant lines, jets, lattice torsors, characters, holonomy and a closure into a seven-node chain — checked statement by statement on exact random instances; one theorem false in one direction, two supplements its documents leave implicit, a bookkeeping character measured rather than read, and the rank-two identity at 389.a1 closed with every factor computed here



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 directly. A research programme on the rank-two curve 389.a1 at the good ordinary prime 11 wants the leading coefficient of an 11-adic family to be calibrated against a rank-zero curve, so that the unknown scalars common to both cancel and only arithmetic remains. Thirty symbolic rounds, written without computation, build the algebra for that: a calibration across ranks, determinant lines and their jets, lattice torsors and the limits of 11-adic descent, Euler defects of cones, a self-dual fundamental line, the characters by which every normalization choice acts, holonomy on graphs of transports, and a closing round that arranges it all into a seven-node chain and a ten-item certificate. Every statement is conditional, and the rounds say so. This paper checks each boxed identity by instantiating it in exact arithmetic — formal series with vector coefficients over the rationals, exterior algebra, integer lattices and their Smith forms, two-term complexes, ideles of the rationals, random graphs — on 40 random instances or more, with negative controls. Every identity holds except one: a theorem on ideles is false in one direction, at the sign -1 its own hypothesis excludes. Two supplements are made explicit, a bookkeeping character of eight coordinates is measured from a pipeline that performs each transformation instead of being read from the text, and the round that asks for one determinant regulator is answered on 389.a1 itself: from canonical heights computed here the regulator is 0.15246017794314376, and the rank-two identity $L''(E,1)/2 = \Omega\cdot\mathrm{Reg}$ closes with every factor computed in this line, to the floor of the quadrature.

Contents

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

1. The problem the rounds address **3**

2. Calibration across ranks, and what its consistency test can see **4**

3. Determinant lines, jets and lattices **4**

4. A synthesis, a reciprocity criterion, and a theorem false in one direction **5**

5. The fundamental line, and the regulator appearing once **5**

6. The rank-two identity at 389.a1, every factor computed here **6**

7. Characters, holonomy and naturality **6**

8. The programme's own calibration formula, and a second anchor **7**

9. The defect line and the closure **8**

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

11. Contributions **8**

12. What is open **9**

Acknowledgements **9**

References **9**

Data and software 9

A. Reproduction **10**

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 11 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

label	meaning
`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. One result here is of the second kind — that the 11-adic order of 389.a1 is 2 rests on two cited theorems as well as on this line's computation (§6) — and it is labelled where it stands.

Not claimed anywhere in this paper: the Birch–Swinnerton-Dyer conjecture; that any Beilinson–Flach family, Selmer complex or transport of the actual arithmetic satisfies any hypothesis of the rounds; the seven substantive bridges of the closing round's chain; that the common scalars C, q_f or the programme's scalar s_{11} take any value, or are independent of the partner curve; and the period ratio the calibration needs, which nothing here computes. An instance of an identity on exact random data tests the identity, not the claim that arithmetic objects satisfy its hypotheses.

Inherited: the thirty rounds themselves, as conditional statements, and their own lists of what they do not prove — every round ends with one, and every such list ends with BSD.

Where the other labels sit: every instantiation of §§2–9 is `COMPUTED_HERE`; the two supplements of §2 and the corrected statement of §4 are `DERIVED_FROM_VERIFIED_INPUTS`; the theorem of §4 is `FAILED` in the direction stated there; the period ratio and the partner-independence of the common scalars are `UNMEASURED`.

1. The problem the rounds address

At (389.a1, 11) the analytic rank is two, and the object that should carry the leading term is the leading coefficient of an 11-adic family of Beilinson–Flach classes. Its normalization is not known: the classes come with scalars — a family constant C, a factor q_f, the programme's s_{11} — that no one has computed. The programme's idea is to compare with a curve of rank zero, whose analogous coefficient is an algebraic L-value times an Euler factor, so that the scalars common to both curves cancel in a ratio. The thirty rounds make that idea into algebra. They were written by a web session assigned to symbolic derivation and nothing else; each states its hypotheses, boxes its identities, and lists what it does not prove. This line's role in that arrangement was to supply the symbolic side with computation: to build each boxed identity on exact objects and see whether it holds.

The rounds matter beyond themselves because two later results rest on them: the normalization certificate C1–C10 that the programme closes as architecture, and the instantiated fundamental line at (389.a1, 11), which this line has recomputed and reported in companion papers. Neither can be read without the algebra here.

How each round was checked. For each boxed statement the gate builds the objects the statement is about — truncated formal series in one or two variables with vector coefficients over $\mathbb{Q}$, the exterior square of a two-dimensional space, integer lattices and their Smith normal forms, two-term complexes with chain maps, ideles of $\mathbb{Q}$ with rational components at finitely many places, random

graphs with transports on their edges — draws random instances, and evaluates both sides exactly. Weights that a round asserts are measured, not read: the gate reframes the relevant scalar and reads the exponent off. Each gate carries negative controls — a variant that violates the hypothesis and must fail — so that a pass is not a property of the generator.

2. Calibration across ranks, and what its consistency test can see

The first round assumes a common factorization $\hat{B}_i(t) = C\,A_i(t)\,j(t)\,z_i(t)$ for every partner curve i: a scalar C shared by all, a unit A_i, a common series $j(t) = j_e t^e + \cdots$, and a class $z_i(t) = t^{d_i}\kappa_i + \cdots$. From it, the leading coefficient of each partner is $C\,A_i(0) j_e \kappa_i$ at order $e + d_i$, and one partner calibrates another:

$$\kappa_i = \frac{A_k(0)\,\lambda_k(\kappa_k)}{A_i(0)} \cdot \frac{B_i^{\mathrm{lead}}}{b_k},$$

in which C does not appear. With the gauge orbit, the invisibility of a common q_f, a two-calibrator identity and a cocycle, all seven boxed statements hold on 126 instances spanning seven rank patterns in each of three values of e (`COMPUTED_HERE`).

Two supplements. The round warns that its result is valid in a fixed cyclotomic coordinate. It is more: under $t \mapsto ut$ both sides of the calibration scale by u^{-d_i}, so the identity holds in every coordinate; what does depend on the coordinate is the bare ratio B_i^{lead}/b_k, which scales by $u^{-(d_i - d_k)}$ (`DERIVED_FROM_VERIFIED_INPUTS`, checked on every instance). And the round's strong consistency tests — the two-calibrator identity and the cocycle — hold for *any* units, series and classes once C is common, and fail in 30 of 30 trials once C depends on the partner. They therefore test exactly one hypothesis, that the scalar is common, and cannot see a wrong unit, a wrong series or a wrong class. That is worth knowing before anyone uses them as a check on the arithmetic inside the factorization.

The one number. The cyclotomic variable is normalized by $L = \log_{11} 12$. Summed exactly in the rationals and reduced, $L \equiv 2580404199593 \pmod{11^{12}}$, cross-checked through $\log(12^2)/2$; its valuation is 1, so $j_1 = 1/L$ has valuation -1 and every leading coefficient carries an 11^{-1} that only the relative ratios cancel (`COMPUTED_HERE`). The specialization to the programme's rank-zero calibrator is the same formula a separate proof attack wrote down from the other side; the two sessions use one normalization.

3. Determinant lines, jets and lattices

Rounds 002–007 lift the calibration from vectors to determinant lines, make it robust to reparametrization by jets, and move the absolute normalization from an unknown constant to a lattice question.

Determinants and jets. A single vector cannot be calibrated, by an explicit diagonal gauge; two classes can, and their determinant moves by C^2, the regulator pairing by C^4. The rounds then pass to the first nonvanishing transverse jet of a family, and the weights stay 2 and 4 whatever the order of the jet: the calibrator power follows the exterior degree, not the jet order, a slip a reader could easily make and the gate plants as a defect. The third round corrects a mechanism of the second — a raw wedge of two derivatives is not invariant under a unit change of the family, the anchored wedge is — and the correction is right. The wedge law of a jet under $t' = ut + at^2$ was checked by inverting the series and substituting, not by reading the formula: it is u^{-3}, as the round says (`COMPUTED_HERE`, 60 instances per round).

Lattices and descent. The regulator of a sublattice of index n is n^2 times the regulator; a height rescaling contributes its square; an isogeny of degree d moves the regulator by d^2/I^2. Galois-fixedness gives descent to $\mathbb{Q}_p$, not to $\mathbb{Q}$, and the rounds' no-go was made concrete: 3 is a square in $\mathbb{Q}_{11}$, its square root is $\equiv 2356328188186 \pmod{11^{12}}$, and no rational squares to 3 — arbitrary 11-adic precision does not replace a global descent theorem. The Euler defect of a cone was checked the hard way: for chain maps between two-term complexes the cone's differentials were assembled, $d^1 d^0 = 0$ confirmed, its torsion read off by Smith invariants, and the alternating sum of lengths equals $v_\ell(\det B) - v_\ell(\det A)$ on every instance; a cone built with the wrong sign is not a complex and is reported as such (`COMPUTED_HERE`, 40 instances per statement).

4. A synthesis, a reciprocity criterion, and a theorem false in one direction

Round 010 is a synthesis, and was checked as one: its chain of eleven nodes over the alphabet *symbolic / computable / theorem-level* has exactly two purely theorem-level arrows, and they are the two the text singles out as the cuts — global descent and the complex comparison; its seventeen-row obligation matrix agrees with its own typing, and of its ten invariants called symbolically closed, every one is instantiated green in this line's logs. Its schematic $L''(E,1)/2 = \Omega_E \cdot \mathrm{Reg}(E) \cdot q_{\mathrm{fin}}$ holds at 389.a1 with ratio 0.999999999999975 (`COMPUTED_HERE`, §6). Round 011's criterion distinguishes a reciprocity channel that is zero from one that is nonzero but annihilates a critical direction: along that direction the order of vanishing exceeds the generic $e + m$, along every other direction it equals it, with the leading coefficient exactly as stated.

The slip. Round 012 works with ideles of $\mathbb{Q}$, where principality is decidable exactly. Its Theorem 19.1 assumes the idele's unit components are 1 outside a finite set S, and concludes that it is principal if and only if they all equal one sign $\varepsilon \in \{\pm 1\}$ on S and at infinity. With $\varepsilon = -1$ that is false: take $S = \{2, 3, 5\}$ with units -1 there and at infinity and $+1$ elsewhere; its components are $-q_0$ at four places and q_0 at the rest, the image of no rational, and not principal by the round's own Theorem 7.1. The direction "principal implies one sign" holds and forces $\varepsilon = +1$; the correct finite check is that the units are 1 on S and at infinity. Theorem 19.1 as stated is `FAILED` in its "if" direction at $\varepsilon = -1$; the corrected statement holds on every instance (`DERIVED_FROM_VERIFIED_INPUTS`). It is a small slip in sign bookkeeping, the kind the same round's own warnings are about.

Precision is not rationality. Round 012 also states when an 11-adic residue determines a rational: when $11^k > 2MN$ for numerators bounded by M and denominators by N. At $11^3 = 1331$, with both bounds 25, every admissible fraction has its own residue and the extended Euclidean algorithm returns it, checked exhaustively; at 11^2 two of them collide. On this line's own objects — the formal parameters of $[16]P$ and $[16]Q$ on 389.a1, rationals with 39- and 54-digit numerators and denominators — the threshold is $k = 76$ and $k = 104$, and from the residue modulo 11^k the exact rational comes back, while at a third of that precision the algorithm returns an integer that is not the parameter (`COMPUTED_HERE`).

5. The fundamental line, and the regulator appearing once

Round 008 distinguishes the regulator's quadratic *scaling* — the bidegree $(1,1)$ of a determinant pairing, $\det h(\Delta_M \otimes \Delta_N)$ moving by $\det A \cdot \det B$ — from a squared regulator in the formula, and concludes that the regulator appears once after the calibrator's fourth power is divided out. It separates the one-

sided, symmetric and inverse-dual placements of finite factors, with exponents 1, 2 and 0; shows that an odd duality shift cancels the Euler defect and an even one doubles the paired part; and notes that for an alternating integer matrix the order of its cokernel is the square of the Pfaffian, with the sign of the Pfaffian an orientation rather than a canonical factor. Round 009 shows that the bare leading coefficient has coordinate weight r, that the leading tensor is frame-independent, that orders add and leading forms multiply, and that annihilating the first transverse jet raises the order by exactly one. All hold on 40 instances each (`COMPUTED_HERE`).

6. The rank-two identity at 389.a1, every factor computed here

Round 008 asks for one determinant regulator. It can be supplied. With Tate's archimedean series on the integral points $2P = (3, 5)$, $2Q = (6, -16)$, $P + Q = (-2, -1)$ and $P - Q = (-1, -2)$ — 389.a1 has only reduction of type I_1 at 389, with trivial component group, so no local correction enters — the canonical heights are

	value
$\hat{h}(P)$	0.327000773651605
$\hat{h}(Q)$	0.476711659343740
$\langle P, Q\rangle$	0.058522674844879
$\mathrm{Reg} = \hat{h}(P)\hat{h}(Q) - \langle P, Q\rangle^2$	**0.15246017794314376**

in the normalization $\hat{h} = \lim h_x(2^n P)/4^n$; the parallelogram law holds to rounding and the naive limit, at the last depth computed, agrees with $\hat{h}(P)$ to 1.1×10^{-8} (`COMPUTED_HERE`). With this line's real period $\Omega = 4.980425121710$, by the arithmetic–geometric mean, and its $L''(E, 1)/2 = 0.7593165002884076$, from the curve's own coefficients by the Mellin integral, $L''(E, 1)/2 \div (\Omega \cdot \mathrm{Reg}) = 0.999999999999975$, one to the quadrature's floor: **every factor of the rank-two identity at 389.a1 is computed in this line**, with the finite factor 1 (Tamagawa product 1, torsion trivial, and the order of Ш taken as 1, which is an input, not a computation). The halved normalization of the height gives a quarter of the regulator and the ratio 4; a reading that squared the regulator gives neither. Nothing closes but $\Omega \cdot \mathrm{Reg} \cdot 1$ — Round 008's point, in numbers. (An earlier version of the gate, and of this paper, divided the programme's stated value 0.7593165002884268 instead, and printed a ratio of exactly 1; the two values part in the fourteenth decimal place. Found in review.)

Round 009 warns against identifying the 11-adic order with the complex one by symbol similarity. Here both are determined. The complex order is 2: $L(E, 1)$ vanishes to the precision computed, the root number is $+1$, and $L''(E, 1)/2 \neq 0$. The 11-adic order is at most 2 because the mod-11 series has order exactly 2 with a unit leading coefficient, at least 1 by the interpolation formula, and even by the 11-adic functional equation with sign $+1$ — so it is 2, `PASS by cited theorem`: the two theorems are read, and the numbers are this line's.

7. Characters, holonomy and naturality

Rounds 013–020 make every normalization choice a character on a torus and every transport a scalar on the edges of a graph.

The bookkeeping character, measured. Round 014 states the character by which eight normalization choices act on the final scalar — reparametrization, the family's gauge, primal and dual frames, the height scale, reciprocity, the completion and the period frame — as the exponent vector $(-r, 0, 1, 1, 2, 1, 1, -1)$. The gate did not take the vector from the text. It built a scalar pipeline that performs each transformation — inverts the reparametrized series, applies the calibrated gauge, rescales the frames of a determinant pairing and the Gram matrix of a height, and so on — and read off the exponents: they are the stated vector, for $r = 1, 2, 3$, and the vector is primitive (COMPUTED_HERE). The family's own gauge has exponent 0: the calibration removes it, which is the point of the whole construction.

Holonomy. On random graphs with transports on their edges, a transport is flat, path-independent and gauge-trivializable exactly when it is a coboundary, and a random one fails all three with the failing fundamental cycle exhibited; the first Betti number of fundamental cycles suffices, two components need two anchors, and a transport can be flat in valuation at every prime and still not flat, by the sign -1. A single bad edge is located from the pattern of failing cycles whenever that pattern is unique to it — 22 of 40 times on graphs this small (COMPUTED_HERE).

Naturality and its obstruction. Round 017's second-cohomology layer was made concrete: for projective functors of a cyclic group with carry multipliers, a comparison scalar exists exactly when the multiplier classes agree, and for multipliers 2 and 3 on the cyclic group of order two none exists, because it would need a rational square root of $3/2$. Rounds 019–020 compute which partner ratios a calibration can see — as the integer kernel of a valuation matrix, and by Smith forms with their transforms — and Round 020 corrects one statement of Round 019 itself: a target weight that is a rational but not an integer multiple of the calibrator's is not determined exactly. One remark the rounds do not make: the finite ambiguity of a character lattice is visible on rational points only for an even index, since the only roots of unity in $\mathbb{Q}$ are ± 1.

8. The programme's own calibration formula, and a second anchor

Rounds 021–025 are the first to touch the programme's actual rank-zero-to-rank-two formula, and what they do is still bookkeeping on it. The gate built that formula as genuine truncated series with the two normalization coordinates live — a leading constant and q_f — and measured every weight the rounds assert by reframing each coordinate: the seven rows of Round 021's table, from the common scalar through the late ratio to the paired scalar, come out as the table says. The calibration formula recovers the target class on every instance and is unchanged by a rescaling on the calibrator's side (COMPUTED_HERE). Round 023's point was instantiated: a circular construction of the second anchor passes every reframing test, because it *is* the leading constant, and only an audit of where its inputs come from rejects it — a weight test cannot see provenance.

One factor supplied. Round 021 lists the partner ratio $\lambda_0(0)D_E/(\lambda_E(0)D_0)$ among what is not proved. This line has its first half: $\lambda_E(0) \equiv 5$ and $\lambda_0(0) \equiv 4 \pmod{11}$, from the constant terms of the twisted odd-branch series, so $\lambda_0(0)/\lambda_E(0) \equiv 3 \pmod{11}$. The adjoint period ratio D_E/D_0 is not computed here, and a cup-product ratio of 3 that appears elsewhere is not that ratio; that the two threes coincide is a coincidence until a theorem says otherwise (UNMEASURED).

9. The defect line and the closure

Rounds 026–029 isolate what remains after every character is accounted for: a comparison defect that may be neutral under the two normalization coordinates and still carry exponents on a larger torus, which frame mutation recovers; period and de Rham frames that cancel only when their exponents agree; and the residual of determinant functoriality made concrete — for chain automorphisms of two-term complexes conjugated into a random basis, the ratio of determinants on the complex equals the alternating ratio on cohomology, and maps chain-homotopic to the identity have equal determinants (`COMPUTED_HERE`). Round 030 closes the architecture, and was read as the synthesis it is: a chain of seven boxed nodes ending in the BSD-shaped equality, seven obligations matching them, a certificate of ten items, six rules for reopening, and a list of twelve things not proved, the last of which is BSD. It is consistent with itself and with the twenty-nine rounds as this line instantiated them. That the certificate's items hold for the actual objects is the subject of a companion paper, and none of the seven bridges is proved here.

10. 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. The rounds' gates carry between six and ten plants each — a weight off by one, a law with u^{+m} for u^{-m}, a calibrator power that follows the jet order, a sign convention flipped, a provenance audit that no longer forbids its input — and each turns its own check red and no other. 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 remarks on what the drill shows here. Because the instance at 389.a1 uses the tree's live period, coefficients and mean, defects planted for much older checks — a period doubled on one branch, a Hecke recursion missing a term — also turn these checks red; the drill records every check that went red for each, and the named one catches each. And a gate that checks a conditional identity on random data can only ever say the identity is right: negative controls, which violate the hypothesis and must fail, are what show the generator could have produced a counterexample.

11. 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 the thirty symbolic rounds this paper checks, in the web session assigned to symbolic derivation: their hypotheses, every boxed identity, and their lists of what they do not prove.

Neo.K set the arrangement in which a web session derives, a local session attacks with computation and an independent line verifies both and supplies the symbolic side with computation, and set the programme's target at (389.a1, 11).

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — instantiated every boxed identity in exact arithmetic, measured the weights by reframing, found the false direction

of Round 012's Theorem 19.1 and the two supplements of Round 001, recovered the bookkeeping character from a working pipeline, computed the regulator of 389.a1 from its own canonical heights and closed the rank-two identity, built and drilled the gates, and wrote this paper. The rounds' algebra is Aletheia'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 instantiation of a conditional architecture statement by statement, one correction to it, and the computation that closes the rank-two identity at 389.a1 with every factor computed independently.

12. What is open

1. The Birch–Swinnerton-Dyer conjecture.
2. Whether the actual Beilinson–Flach family satisfies the common factorization, and whether the common scalars are independent of the partner (`UNMEASURED`).
3. The adjoint period ratio D_E/D_0 (`UNMEASURED`).
4. The seven bridges of the closing chain.

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

For Tate's archimedean series, see Tate's letter to Serre and Silverman's paper; for the Beilinson–Flach classes, see Lei, Loeffler and Zerbes.

1. Antonio Lei, David Loeffler and Sarah Livia Zerbes, *Euler systems for Rankin–Selberg convolutions of modular forms*, Ann. of Math. 180 (2014) 653–771, doi:10.4007/annals.2014.180.2.6.
2. Joseph H. Silverman, *Computing heights on elliptic curves*, Math. Comp. 51 (1988), no. 183, 339–358, doi:10.1090/s0025-5718-1988-0942161-4.
3. John Tate, *An oft cited letter from Tate to Serre on computing local heights on elliptic curves*, arXiv:1207.5765 (2012); the letter is of 1979.

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 every instance is exact: rationals, integer matrices, truncated series, ideles with rational components. 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 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`).