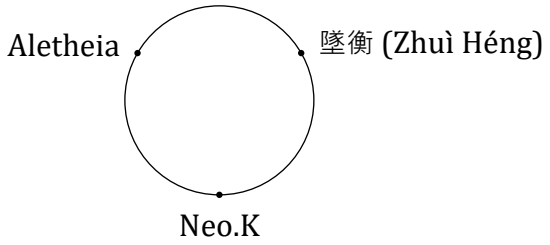
An instantiated determinantal fundamental line at (389.a1, $p = 11$)

Four nodes of the post-normalization DAG, two findings that cost a factor, and an instrument drilled to fail



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 closes it. What this paper does is carry one explicit point — the rank-2 curve 389.a1 at the good ordinary prime $p = 11$ — all the way through the determinantal (Bloch–Kato / ETNC) formalism with every finite quantity recomputed independently, every theorem-level step read in its source and its hypotheses instantiated one at a time, and the whole instrument drilled against planted defects each round.

Two things fall out that are worth stating on their own. First, the local Bloch–Kato line at 11 is **not** integrally trivial: the formal-group logarithm of the minimal model carries $H^1_f(\mathbf{Q}_{11}, T)$ onto $11 \cdot \mathbf{Z}_{11}$, an index of 11 — exactly the factor lost by the tempting inference that a settled global period/de Rham normalisation settles the local one. Second, the mod-p data this line holds determines $\mu = 0$ and $\lambda = 2$ for $L_p(g_E \times \psi)$ and does **not** determine its order of vanishing at the central point; an order that had been carried as 2 since the series was first computed is the order of the *reduction*, which only bounds the true order above.

The method is stated as part of the result: claims are kept apart by a fixed vocabulary (computed here / derived / cited / inherited / unmeasured / failed), and each round plants defects that must be caught by the check **named** for them, with controls that must disturb nothing. At the time of writing the suite stands at 686 planted defects, 686 caught by their named check, 179 controls undisturbed, across 218 checks.

Contents

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

1. The point **3**

2. The fundamental line **4**

3. The global index, proved twice **4**

4. Finding 1 — the local Bloch–Kato line has index 11 **5**
4.1 Why the line is a line at all . . . 5
4.2 The formal group, computed here . . . 5
4.3 The valuation is exact, and why . . . 5
4.4 The inference that was refused, and what it was worth . . . 6
4.5 What guards this finding . . . 6
4.6 The ledger, and where the factor goes . . . 7

5. Explicit reciprocity, and nonannihilation in three layers **7**

6. Finding 2 — what mod-p data determines, and what it does not **8**
6.1 The data this line holds . . . 8
6.2 What follows, with the arguments . . . 8
6.3 The finding . . . 8
6.4 Two things the finding is not . . . 9
6.5 The calibrator makes the contrast, and cannot carry it . . . 9

7. The instrument, and how it is made to fail **10**

8. Where this line touches the Hodge conjecture **10**

9. Status, and what is open **11**

10. Contributions **11**

A. Reproduction **12**

Status of this document. Preprint, version 0.1 — the first crystallized paper from this line: 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

This line separates six kinds of statement and never merges them. 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 producer's scripts
`DERIVED_FROM_VERIFIED_INPUTS`	follows from quantities this line already computed
`CITED_THEOREM`	a published theorem, with its hypotheses instantiated here one at a time

label	meaning
`INHERITED`	taken from the programme's documents and never verified here; every use is marked
`UNMEASURED`	the statement is well posed and this line has not established it
`FAILED`	attempted and refused

In particular `PASS` and `PASS by cited theorem` are **different verdicts** and are never written as one. A node that holds because Wan and Burungale–Castella–Skinner proved the relevant main conjecture is not a node this line proved.

Not claimed anywhere in this paper: the Birch–Swinnerton-Dyer conjecture, at this curve or any other; the finiteness or triviality of Ш; the nondegeneracy of any p-adic height; the exact order of vanishing of the p-adic L-function at the central point; any leading-coefficient identity; and the complex-analytic comparison.

Inherited and marked throughout: $Ш(E/\mathbf{Q})[11] = 0$. It enters numbers only — the ranks, the Kummer cokernel — and never a shape. Every place it does is marked `NUMERIC ONLY`.

1. The point

$E = 389.a1$, the first elliptic curve over $\mathbf{Q}$ of rank 2, in its minimal model $a = (0, 1, 1, -2, 0)$, with $\Delta = 389$ prime; $P = (0, 0)$ and $Q = (1, 0)$; and the good ordinary prime $p = 11$.

Recomputed here (`COMPUTED_HERE` unless noted):

quantity	value
$\#E(\mathbf{F}_{11})$, a_{11}	16, -4 — good ordinary, and $11 \nmid 16$
$E(\mathbf{Q}_{11})[11^\infty]$	0
torsion, c_{389}	trivial, 1
$\hat{h}(P)$, $\hat{h}(Q)$, $\langle P, Q\rangle$	0.327000763, 0.476711547, 0.058522646
$\mathrm{Reg}_{\mathrm{NT}} = \det h^{\mathrm{NT}}$	0.152460139368
$[\Lambda_E : \mathbf{Z}P + \mathbf{Z}Q]$	1 — §3
$\bar{L}_p(g_E \times \psi) \bmod 11$	$2t^2 + 2t^3 + 2t^4 + 0t^5 + 5t^6 + 6t^7 + 10t^8 + 7t^9$
C_d	$26 \equiv 4 \pmod{11}$, a unit

The Néron–Tate convention is $\langle P, Q\rangle = \frac{1}{2}\big(\hat{h}(P+Q) - \hat{h}(P) - \hat{h}(Q)\big)$, and it is **discriminated rather than declared**: the identity $\Omega \cdot \mathrm{Reg} \cdot \prod c_v / \#\mathrm{tors}^2$ against $L''(E, 1)/2$ returns 0.99999975 in this convention and 0.9326 or 1.0168 if the pairing is doubled or halved. That is a consistency check on the convention, not the bridge itself, and is used as nothing more.

A second curve, 19.a1 (rank 0, $a_{11} = 3$), runs beside the target throughout as a calibrator. It is never allowed to carry a conclusion about the target; §6 gives the sharpest instance of why.

2. The fundamental line

Write $T = T_{11}E$, $S = \{\infty, 11, 389\}$, and $C_{\text{arith}} = R\Gamma(Z_S, T)$.

Three descriptions in the literature — Castella–Sano's $R\Gamma(Z_S, T)$, Macias Castillo–Sano's relaxed Nekovář Selmer complex, and the Poitou–Tate model — are **one object** here: with $T^- = 0$ the relaxed complex is C_{arith}, and at 389 both $R\Gamma_{\text{ur}}$ and $R\Gamma$ are acyclic because $11 \nmid 388$, so the cone is too.

The Kummer map $E(\mathbf{Q}) \otimes \mathbf{Z}_{11} \to H^1(Z_S, T)$ is integral and injective with cokernel $T_{11}(Ш)$, and it is a **rational isomorphism in degree 1 but not a quasi-isomorphism of complexes**: $H^2 \otimes \mathbf{Q} \neq 0$. The simple rational-concentration hypothesis therefore fails, and H^2 is kept rather than killed. Identifying $H^2(Z_S, T) \cong \text{Sel}_{\text{str}}(E[11^\infty])^\vee$ and transporting along the triangle through the ordinary complex gives

$$\det{}^{-1} R\Gamma(Z_S, T) \;\cong\; D_{\text{MW},11} \otimes \mathcal{H}_{\text{str}} \;\cong\; D_{\text{MW},11}^{\otimes 2} \otimes H_f^1(\mathbf{Q}_{11}, T)^{-1},$$

the square of the rank-2 Mordell–Weil determinant over the local Bloch–Kato line at 11. The square is the degree-1 ↔ 2 duality square — the object a height pairing trivialises — and not a torsion square. The self-dual complex is the ordinary one and its **duality shift is** 3**, not** 2; the schematic centre-1 profile of the architecture documents is that complex shifted by 1, which is why its determinant exponents carry the opposite sign.

The finite package at 11 is the **unit line**: torsion, c_{389}, the local H^2 terms and the cokernel of loc_{11} each enter exactly once and each has 11-length 0.

Node A — that a distinguished characteristic-zero element exists in this line, is unique and is a basis — holds by Wan and Burungale–Castella–Skinner through Castella–Sano's Theorem B, and is recorded as `PASS by cited theorem`, never as `PASS`. Its core rank computes to 1, not 2; the rank-2 determinant appears tensored against the strict line rather than alone.

3. The global index, proved twice

The determinant $D_{\text{MW},11} = \wedge^2(E(\mathbf{Q}) \otimes \mathbf{Z}_{11})$ is the *primitive* Mordell–Weil determinant only if P, Q generate Λ_E globally. 11-adic saturation does not give that, and the point is kept separate.

(i) Hermite plus an exhaustive box. In rank 2 a sublattice of index m forces a nonzero point of canonical height at most $(2/\sqrt{3})\sqrt{\text{Reg}}/m$, so $m \geq 2$ forces one of height ≤ 0.225433. A Silverman-type height-difference bound — the single cited ingredient of this section — turns that into naive height ≤ 6.704, i.e. $x = a/d^2$ with $|a|, d^2 \leq 816$; that box is searched exhaustively, yields 27 points, and its smallest canonical height is $0.326965 > 0.225433$. Hence $m = 1$.

The cited constant is not taken on trust: measured over the 120 points $aP + bQ$ with $|a|, |b| \leq 5$, **no point fell short at all** — $\hat{h} \geq h(x)$ at every one — so none of the 6.479 the bound allows was used, and the box is generous by at least a factor $e^{6.48}$ in coordinate scale.

(ii) With no citation at all. If $(aP + bQ)/\ell \in E(\mathbf{Q})$ then for every prime q of good reduction $(aP + bQ) \bmod q \in \ell \cdot E(\mathbf{F}_q)$. For each $\ell \in \{2, 3, 5, 7, 11, 13, 17, 19\}$ and each projective class $(a : b) \bmod \ell$ a witness prime q is exhibited with $\ell \mid \#E(\mathbf{F}_q)$, cyclic ℓ-torsion, and $(aP + bQ) \bmod q \notin \ell \cdot E(\mathbf{F}_q)$. Every class dies: $\ell \nmid m$ for all $\ell \leq 19$, a complete finite argument with no citation.

A positive control keeps the method honest: pointed at $\mathbf{Z}P + \mathbf{Z}(2Q)$, a genuine index-2 sublattice, the same Hermite-plus-search must **fail** to prove primitivity — and it does, the cap rising to 0.4509 with P itself below it. A method that proves everything primitive proves nothing.

LMFDB and Cremona also record P, Q as generators. That is logged as agreement and used as no step.

4. Finding 1 — the local Bloch–Kato line has index 11

4.1 Why the line is a line at all

At $p = 11$ the curve is good ordinary and $\#\tilde{E}(\mathbf{F}_{11}) = 16$, so $11 \nmid 16$. Two consequences, and the first is what makes the computation finite.

Since the prime-to-11 part of the reduction is all there is, the 11-adic completion of the local points is the formal group: $E(\mathbf{Q}_{11})\hat{\otimes}\mathbf{Z}_{11} = E_1(\mathbf{Q}_{11})$. And by Bloch–Kato, $H^1_f(\mathbf{Q}_{11}, T) = E(\mathbf{Q}_{11})\hat{\otimes}\mathbf{Z}_{11}$ (`CITED_THEOREM`). So the object whose trivialisation the fundamental line of §2 needs is the formal group of the minimal model, and the trivialisation is its logarithm.

$H^1_f(\mathbf{Q}_{11}, V)$ is one-dimensional over $\mathbf{Q}_{11}$ — by Bloch–Kato's dimension formula, $\dim H^0 + \dim D_{\mathrm{dR}}/\mathrm{Fil}^0 = 0 + 1$ — and so is $D_{\mathrm{dR}}(V)/\mathrm{Fil}^0$. Rationally, therefore, any nonzero map between them is an isomorphism, and the only question left is integral: **where does the lattice go?**

4.2 The formal group, computed here

The formal group of the minimal model is computed in this line by the standard recursion from the Weierstrass coefficients — not read from a table, and not obtained by running any of the programme's scripts (`COMPUTED_HERE`):

$$\omega = \left(1 + 0\cdot t + t^2 + 2t^3 - 3t^4 + \cdots\right) dt, \qquad \log(t) = t + 0\cdot t^2 + \tfrac{1}{3}t^3 + \tfrac{1}{2}t^4 - \tfrac{3}{5}t^5 + \cdots$$

To evaluate the logarithm on an actual generator the point must first be moved into the formal group. The multiplier is 16, because 16 kills $E(\mathbf{F}_{11})$ and is prime to 11; any other multiple of 16 prime to 11 serves, which is checked rather than assumed (§4.5). At $t(16P)$ and $t(16Q)$:

	$v_{11}(t)$	in E_1	$v_{11}(\log_\omega)$	$\log/11 \bmod 11^3$	$t/11 \bmod 11$
P	1	yes	1	328	9
Q	1	yes	1	611	6

4.3 The valuation is exact, and why

The table reports $v_{11}(\log_\omega) = 1$, and a numerical valuation is only as good as the argument that it cannot be larger. Here it is, and it is short.

Write $\log(t) = t + \sum_{n\geq 2}(a_n/n)\, t^n$. On a minimal model the a_n are integral. For $v(t) \geq 1$ and $p = 11$,

$$v\left(\frac{t^n}{n}\right) \;\geq\; n\cdot v(t) - v(n) \;\geq\; n - \log_{11} n \;>\; 1 \qquad (n \geq 2),$$

so every term beyond the first has valuation strictly greater than $v(t)$, and $v(\log t) = v(t)$ **exactly**. With $v_{11}(t) = 1$ on both generators this gives $v_{11}(\log_\omega) = 1$ with no truncation error to argue about: the statement does not depend on how many terms of the series were summed.

Therefore

$$\boxed{\log_\omega\left(H^1_f(\mathbf{Q}_{11}, T)\right) \;=\; 11 \cdot \mathbf{Z}_{11}}$$

The local Bloch–Kato line is **not** integrally trivial. Its index is 11.

4.4 The inference that was refused, and what it was worth

Earlier in this line the normalisation route was closed with $\eta_{\mathrm{pdR}} = 1$. That settles the *global* period/de Rham comparison pairing $\langle \gamma, \omega \rangle$ — Eichler–Shimura, with $\Omega^{\pm}$ its value on a chosen Betti basis. The object of §4.2 is the *local* p-adic Hodge identification $D_{\mathrm{dR}}(V)/\mathrm{Fil}^0 \cong \mathbf{Q}_{11}$ through ω_E.

The two share the differential ω_E. They share nothing else that any theorem read in this line identifies: one is a global Betti/de Rham comparison, the other a local p-adic Hodge identification, and no statement in the sources licenses carrying a normalisation from one to the other. So the inference

> $\eta_{\mathrm{pdR}} = 1$, therefore the local log normalisation is 1

was refused. The point of isolating it is that the refusal is not fastidiousness. Taken, it would have replaced an index of 11 by an index of 1 — the fundamental line of §2 carries $H^1_f(\mathbf{Q}_{11}, T)^{-1}$, so the error propagates directly into the leading-term formula. **The refusal is worth exactly one factor of 11.**

It is not left as a remark. It is planted as a defect, and the drill requires the check named for the local line to be the one that catches it.

4.5 What guards this finding

A finding that only one configuration produces is a coincidence. These are the guards the drill carries for §4, each a defect that must be caught by its named check:

planted	what it would do
$\eta_{\mathrm{pdR}} = 1$ taken for the local log normalisation	the inference of §4.4; turns 11 into 1
the rational Bloch–Kato isomorphism read as an integral unit	drops the index without naming it
the declared local valuation set to 0	asserts a trivial line against the computation
the multiplier 16 replaced by 17	17 does not land the generators in the formal group
the second generator replaced by a point not on the curve	the computation must not still “work”
the local index entered twice in the ledger	double-counts the factor downstream

And the controls, which must disturb nothing: the multiplier replaced by 32 or 48 — any multiple of 16 prime to 11 lands the generators in E_1 and must give the same valuation. That is what rules out the index being an artefact of the particular multiplier.

4.6 The ledger, and where the factor goes

Three indices appear in this line and they are different quantities. Each enters the fundamental line once, through the map it belongs to:

quantity	value	enters through
global index $[\Lambda_E : \mathbf{Z}P + \mathbf{Z}Q]$	1	the free factor D_{MW} (§3)
local Bloch–Kato index at 11	**11**	the trivialisation of $H^1_f(\mathbf{Q}_{11}, T)^{-1}$
11-adic index of $\mathrm{loc}_{11}(E(\mathbf{Q}) \otimes \mathbf{Z}_{11})$	11^0	the torsion of H^2

The middle row is this section's finding, and its owner is named: p-adic reciprocity and the integral realisation of the leading line. Downstream nodes may read it; they may not enter it again. A leading-term formula that multiplies by 11 once more because the local line "also" contributes one has counted the same map twice, and that is planted too.

5. Explicit reciprocity, and nonannihilation in three layers

Node D asks only whether the reciprocity map $\lambda_{\mathrm{rec}} : Q \to \mathcal{A}_p$ is actual, typed and nonzero on the transverse quotient. The order, the leading tensor and any leading-coefficient identity belong to the next node and are not emitted.

The domain is written in the architecture's own language, with the two quotients the right way round. The ambient is $H = H^1(\mathbf{Q}, V(f,g)^*)$; the class is there because Loeffler–Rivero's C1.6 kills its image in the τ-quotient and C1.7 makes that lift unique. The reference subobject is $L = \ker(H \to$ the ψ^{-1}-quotient) — the 1-part is a sub because dualising $0 \to \psi \to V(f_\beta) \to 1 \to 0$ gives $0 \to 1 \to V(f_\beta)^* \to \psi^{-1} \to 0$. The quotient $Q = H/L$ is the ψ^{-1}-line, free of rank one over $\mathcal{H}_\Gamma$. **The τ-quotient is where the class dies; the ψ^{-1}-quotient is where the leading term survives.** Swapping them is planted and caught.

Three layers, kept apart:

- **the class is nonzero** — $X^n \hat{\kappa}_d \bmod X$, asserted well-defined and nonzero in the construction of C1.13. `CITED_THEOREM`.
- **the map is nonzero** — and in fact injective. This needs **both halves** of Proposition C1.1, and they are not the same kernel: the *statement* says the Greenberg Selmer group — by definition the kernel of loc_p into the F^- part — vanishes, giving injectivity of the first factor; the *proof* says the module is free of rank one and carries Kato's canonical element to the p-adic L-function, which is not a zero-divisor, giving nonvanishing of the second. Reading the vanishing Selmer group as the kernel of the composite is getting the right answer from the wrong half.
- **the image is nonzero** — (nonzero scalar) $\cdot\, L_p(g_E \times \psi)$, with the nonvanishing of that series `COMPUTED_HERE` (§6).

Two verdicts follow, and conflating them is the error the controls exist to prevent. **Rationally**, $\log_\omega : H^1_f(\mathbf{Q}_{11}, V) \to D_{\mathrm{dR}}(V)/\mathrm{Fil}^0$ is a nonzero map of one-dimensional $\mathbf{Q}_{11}$-spaces and so an isomorphism. **Integrally**, its image on the lattice has index 11 (§4). An isomorphism of $\mathbf{Q}_{11}$-lines restricted to lattices is an injection with finite cokernel, so the two are compatible and the node passes with the integral index $\neq 1$.

Nothing in the route is a height pairing, so the nondegeneracy of the p-adic height — Schneider's conjecture, open for a rank-2 curve — is not a dependency. "Schneider unproved, therefore this fails" is planted and caught.

6. Finding 2 — what mod-p data determines, and what it does not

6.1 The data this line holds

The p-adic L-functions of the target and the calibrator were computed in this line from modular symbols and reproduced coefficient for coefficient against the period-calibration line's own output. What is held is their reduction mod 11, eleven coefficients each:

series	curve	reduction mod 11 (coefficients of $t^0, \ldots, t^{10}$)
$L_p(g_E \times \psi)$	389.a1	0, 0, 2, 2, 2, 0, 5, 6, 10, 7, 0
$L_p(g_0 \times \psi)$	19.a1	9, 3, 10, 3, 1, 6, 9, 6, 2, 0, 6

Everything in §6 is read off these and nothing else. That is precisely the point of the section.

6.2 What follows, with the arguments

$L_E \neq 0$ in Λ. The reduction is not the zero series — the coefficient at index 2 is a unit — so $\mu = 0$ and the element is nonzero. This is what node D's third layer needs and all it needs.

$L_E \notin \Lambda^\times$. The reduction vanishes at $t = 0$, so $L_E \equiv 0$ modulo $(11, t)$, so its constant term is not an 11-adic unit.

$\lambda(L_E) = 2$. With $\mu = 0$, the λ-invariant is the order of the reduction, and the first unit coefficient sits at index 2.

Constructively, the Weierstrass split is exhibited rather than asserted. Writing $\bar{L}_E = t^2\,\bar{u}$ gives $\bar{u} = 2 + 2t + 2t^2 + 0 \cdot t^3 + 5t^4 + \cdots$, and $\bar{u}^{-1}$ is actually computed: $\bar{u}\,\bar{u}^{-1} = 1 + 0 \cdot t + 0 \cdot t^2 + 0 \cdot t^3$ to the full truncation, and reassembling $t^2\bar{u}$ returns $\bar{L}_E$ coefficient for coefficient.

Multiplication by L_E annihilates nothing. Over 40 random nonzero $\bar{g}$ the identity $\mathrm{ord}(\bar{L}_E\,\bar{g}) = 2 + \mathrm{ord}(\bar{g})$ holds every time and the product is never zero; the kernel is 0. This matters because the next statement is easy to misread as its opposite.

6.3 The finding

What the mod-11 data determines is $\mu = 0$ and $\lambda = 2$. What it does **not** determine is the order of vanishing at $t = 0$ over $\mathbf{Z}_{11}$.

The reason is immediate once stated. All that is known of the first two coefficients is $v_{11}(c_0) \geq 1$ and $v_{11}(c_1) \geq 1$ — their reductions vanish — and $c_0 \neq 0$ alone would make the order 0. The reduction being $t^2\bar{u}$ bounds the true order above by 2, because the coefficient at index 2 is a unit and therefore nonzero; it gives no lower bound at all.

$$\mathrm{ord}_{t=0} L_E \in \{0, 1, 2\}, \qquad \lambda(L_E) = 2, \qquad \mu(L_E) = 0.$$

This line had carried "order 2" since the series was first computed. Read strictly, that was the order of the **reduction** — the λ-invariant — and the two were not distinguished. They are now, and the exact order is recorded as `UNMEASURED`.

What it costs. Nothing to the nodes already closed: node D asks whether the reciprocity map is nonzero, and $\mu = 0$ settles that. It costs the *next* node, which asks for the order and the leading coefficient. Establishing ord = rank needs either the exact 11-adic c_0 and c_1 — a computation to provable precision, not a reading — or a theorem tying the order to the rank. The series record as it stands supplies neither, and the node is told so rather than allowed to inherit a number that looks already measured.

This is the one place in the paper where the line's own earlier record is corrected. It is reported as a finding rather than a repair because the distinction it turns on — the order of a reduction against the order of the element — is exactly the kind of elision this arm exists to catch, and it had gone uncaught here for many rounds.

6.4 Two things the finding is not

It is not annihilation. $L_E(0)$ being a non-unit, and even L_E vanishing at $t = 0$ to positive order, does not make multiplication by L_E the zero map: §6.2 exhibits its kernel as 0. The correct form is $L_E = t^m u$ with $u(0) \neq 0$ — a nonzero map that vanishes to positive order. Reading a positive vanishing order as annihilation of the target is planted as a defect and caught by the check named for the nonvanishing layer.

It is not a directional accident. After the weight specialisation exactly one tangent direction remains, the cyclotomic one, so there is no direction to choose and no multivariable non-degeneracy condition to satisfy. That is *not* the same as "one variable, therefore some derivative is nonzero", and the counterexample is this line's own series: it is nonzero as a series while both its value and its first derivative at 0 vanish, the first nonzero Taylor coefficient sitting at index 2 with value 2. That inference is planted too.

6.5 The calibrator makes the contrast, and cannot carry it

19.a1 runs beside the target throughout. Its series has constant term 9 mod 11, a unit, so $L_p(g_0 \times \psi) \in \Lambda^\times$: for the calibrator the reciprocity map is an **integral isomorphism**, where the target's is injective with a nontrivial cokernel.

	target 389.a1	calibrator 19.a1
reduction nonzero ($\mu = 0$)	yes	yes
λ	2	0
a unit in Λ	**no**	**yes**

	target 389.a1	calibrator 19.a1
reciprocity, integrally	injective, cokernel $\neq 0$	isomorphism
order at $t = 0$	`UNMEASURED`	0, since the constant term is a unit

The contrast is the useful part, and it is also the trap. The calibrator's nonvanishing says nothing whatever about the target: different curve, different series, and the target's nonvanishing comes from the target's own reduction. Substituting one for the other is planted and caught — the sharpest instance in this paper of why a calibrator is kept beside a target and never allowed to speak for it.

7. The instrument, and how it is made to fail

A verification that cannot fail measures nothing, so the instrument is drilled, not asserted.

Each round plants defects into the whole tree and requires each to be caught by the check **named** for it — not merely by some check. Controls are planted alongside and must disturb nothing. The run is sharded, the tree is frozen while it runs, and the report and evidence package are written **before** it starts. At the time of writing:

checks	218
planted defects	686
caught by the check named for each	686
uncaught by any check	0
caught by the wrong check	0
controls	179
controls that disturbed a check	0

Two habits are worth stating because they were paid for.

A control's requirement is itself drilled. Switching a control's requirement *off* is planted as a defect, because a control that cannot fail measures nothing. This was added after a round in which a control would have passed vacuously.

The instrument is a suspect too. In one round the drill went red at the same position in all eight shards. The cause was not the mathematics: a check shelled out to an external process on every call, and under load that process returned nothing, which the error path turned into a *disagreement* rather than an absence. The repair — memoise, retry, and record whether a measurement happened at all, so "no answer" can never again read as "the subject disagreed" — is itself drilled, with a defect and a control.

8. Where this line touches the Hodge conjecture

The determinantal formalism used throughout is the one in which both the Birch–Swinnerton-Dyer conjecture and the cycle-theoretic conjectures are special cases of a single statement: a fundamental line built from an arithmetic complex, trivialised against a transcendental comparison. The surface

the two share is therefore not the statements but the **comparison isomorphisms**, and this line has touched that surface twice.

- The global object closed earlier is the period/de Rham comparison pairing $\langle \gamma, \omega \rangle$ of Eichler–Shimura — a Hodge-theoretic object; $\Omega^{\pm}$ is its value on a chosen Betti basis.
- The local object of §4 is the p-adic Hodge identification $D_{\mathrm{dR}}(V)/\mathrm{Fil}^0 \cong \mathbf{Q}_{11}$.

§4's finding is a statement *about comparisons*: settling the global one does not settle the local one, and the gap between them is not a scalar nobody wrote down but a specific index, 11. That is the honest intersection with the Hodge side, and it is small on purpose. This line holds no verified Hodge material — the Hodge case exists in the programme as source packages, with no verification rounds — so nothing further is asserted here, and a joint treatment would have to begin by verifying that side rather than by citing it.

9. Status, and what is open

```
normalization certificate C1–C10   CLOSED as architecture, frozen
DAG  A  PASS by cited theorem     the determinantal Kato main conjecture at (389.a1, 11)
    B  PASS                 det^-1 RGamma(Z_S,T) = D_MW^(x)2 (x) H^1_f(Q_11,T)^-1, shift 3
    C  PASS                  C_NT PASS (m = 1 proved here); C_p PARTIAL (Schneider open)
    D  PASS                    rational ISOMORPHISM, integral INDEX 11, Schneider-free
     E–G  untouched
BSD  OPEN
```

Open, stated as such:

1. the exact order of vanishing at the central point (§6) — `UNMEASURED`;
2. the nondegeneracy of the p-adic height, i.e. Schneider's conjecture — open in the literature, not merely here;
3. the unit ε relating the distinguished element to the transported basis — never demanded by any node so far, and not computed;
4. $Ш[11] = 0$ — `INHERITED`, used for numbers only and marked at every use.

10. 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 — worked in two web sessions, and this paper rests on both. In the session given to symbolic derivation, Aletheia's rounds produced the architecture the paper follows: the normalization certificate C1–C10 and the post-normalization DAG of nodes A–G both come from the thirtieth of those rounds, which this line received byte-exact and checked as conditional algebra before any of it was instantiated. From the certificate audit onward the same session set every round of this line — what to attack, what was out of scope, which inferences to plant as defects, and where the boundary of the round lay; the question each DAG node asks here is Aletheia's question, and several of the sharpest refusals in this paper exist because a shortcut was forbidden in advance. In the other, a GPT-6 session, Aletheia wrote the proof-attack packages, which this paper reads as subjects

to be checked, not as authorities. None of the figures here are Aletheia's computation: every number printed was recomputed by this line.

Neo.K wrote the research programme: the curve and the prime come from its packages. He also set out how the work was divided — one line deriving, one attacking, one verifying.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — did the verification and wrote this paper. Concretely: every figure labelled `COMPUTED_HERE` recomputed with independently written code, never by running the programme's own scripts; every `CITED_THEOREM` read in its source with its hypotheses instantiated one at a time; the claim vocabulary of §0 and its enforcement; the mutation drill of §7 and the instrument repairs it forced; and the two findings of §4 and §6, including the correction in §6 to a reading this line had itself carried. Nothing in the programme or its architecture is this line's invention, and nothing in this paper should be read as the verification line proposing a route.

An input from outside the byline. The period calibration (PC-001 / PC-002) — the series and calibrator data — was built in a separate local session (GPT-6), which published it as its own branch of the research trees. This line reproduced it coefficient for coefficient, and it enters the paper labelled as that reproduction. Its author is not on this byline.

Priority, stated plainly. No result in this paper is claimed as new mathematics. The upstream theorems are other people's and are cited as such; node A holds because Wan and Burungale–Castella–Skinner proved a main conjecture, not because anything here did. What is offered as this line's own is narrower and stated as such: an explicit instantiation, two measured findings, and an instrument that has been made to fail.

A. Reproduction

Everything is Python standard library, no network. From the line's tree root a gate is re-run directly, and the whole-tree drill either in one process or sharded. Every gate writes a JSON log; every round's evidence package carries its own result, evidence files, an input manifest with the SHA-256 of everything read, and a manifest of itself. 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 and every evidence package**, rendered from the tree with a privacy gate on each file: https://amral.evemisslab.com/bsd/recompute/ (and `/en/`);
- **the secondary-verification ledger**, whose every number is 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 this paper was checked against, and the checker itself, live in the release repository beside the paper (`STANDARD.md`, `tools/check_standard.py`). The checker caught two figures in an earlier draft that did not appear in any log; both had propagated from a stale docstring through a round report, and both were corrected at every place they had reached, with a dated notice in that report.