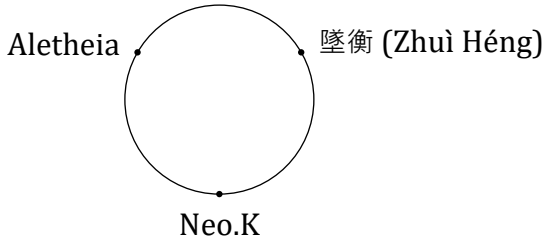
Closing a normalization certificate against actual objects

The ten items of a determinantal comparison at (389.a1, $p = 11$), judged one by one, and the method that judges them



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. This paper concerns the step a determinantal approach to it must take before any statement about L-values: fixing the normalization of a comparison between an auxiliary Eisenstein family and an arithmetic line, for the rank-2 curve 389.a1 at $p = 11$ with the rank-0 curve 19.a1 as calibrator. A symbolic programme reduced that step to a certificate of ten items, C1–C10. We judge each item against objects that actually exist — constructed, recomputed, or supplied by a theorem whose hypotheses we instantiate on these curves — and close all ten.

The closure has three ingredients. The late character of the comparison is a function of the auxiliary family alone, by Loeffler–Rivero's theorems, every inherited hypothesis instantiated; the auxiliary leading line exists and is the Kummer line of $\mathbf{Q}(\sqrt{2})$, whose local conditions are vacuous at every place; and the comparison route is realized as an actual chain of line maps — injective but not surjective for the target, because its p-adic L-function has a double zero — along which the pairing scalar, the determinant scalar and the residual holonomy are all exactly 1, so that the reduced comparison is itself a canonical section. The method matters as much as the verdicts: one rule judges every item and never turns an unmeasured item into a pass, and the rule was itself corrected twice when it misjudged the kind of claim an item makes. At the pinned commit the line's whole-tree 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.

Contents

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

1. The setting **3**

2. The method, and how it was corrected **4**

3. C3: no coordinate weight hidden **5**

4. C1 and C2: the late character is auxiliary **6**

5. C4: the auxiliary leading line exists **6**

6. The comparison route, and where it must be localized **7**

7. C5 and C8: the defect ledger, and the target's lattice **8**

8. C6: the pairing scalar is 1 **8**

9. C7: the determinant scalar is 1 **9**

10. C9: no residual holonomy **9**

11. C10: a canonical section **9**

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

13. Contributions **10**

Acknowledgements **11**

References **11**
Data and software 11

A. Reproduction **12**

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 13 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 producers' 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
`INHERITED`	taken from another document 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: an item discharged by Loeffler–Rivero's theorem is not an item this line proved, and the table in §2 says which is which.

Not claimed anywhere in this paper: the Birch–Swinnerton-Dyer conjecture; any statement about $L(E, s)$; any node of the post-normalization DAG; the values of C_π, of the period ratio Π_E/Π_0, of the comparison scalar $c_{\mathrm{ES/Kum}}$, of the lattice index q or of the jet order n; and any Beilinson–Flach class in coordinates. "Closed" means closed **as normalization architecture**: the comparison's normalization is fixed canonically, which is the precondition for the DAG and nothing more.

1. The setting

The comparison. Let $E = 389.a1$ (the target, rank 2) and $E_0 = 19.a1$, with a-invariants $[0, 1, 1, -769, -8470]$ (the calibrator, rank 0; Cremona's 19a2, not Cremona's 19a1, which is $X_0(19)$), both ordinary at $p = 11$, and let $g = E_2(1, \chi_8)$ be the weight-2 Eisenstein series with $\chi_8 = (2/\cdot)$, with critical-slope refinement at 11. A cross-curve calibration — built in a separate session and reproduced coefficient for coefficient by this line — writes, for each partner $i \in \{E, 0\}$, the auxiliary family's projected leading class as

$$\hat{B}_i = C_\pi \, A_i \, j \, z_i,$$

with a late character C_π, a partner factor A_i, the cyclotomic coordinate j and the Kato class z_i; the calibration's purpose is the partner ratio, in which C_π cancels **if** it is common to both. Whether it is, and whether nothing else hides in the comparison, is a question of normalization.

The certificate. A symbolic programme reduced that question to ten conditions, stated in its thirtieth round as a conjunction — "only after C1–C10 pass should the normalization ledger be frozen":

	condition (verbatim)	the claim it makes
C1	Target and calibrator use the same exact late character.	equality
C2	No partner arithmetic is hidden in the common line.	absence
C3	No analytic coordinate weight is hidden in exact normalization.	absence
C4	The auxiliary leading tensor or equivalent q-blind source exists.	existence
C5	Every defect factor is tagged E, L, G, or P.	enumeration
C6	Either canonical cancellation or one explicit residual pairing scalar.	existence
C7	Either canonical cohomological equality or explicit κ_{det}.	existence
C8	No exact unit ambiguity is attributed to a length / index.	absence
C9	No residual holonomy.	absence

	condition (verbatim)	the claim it makes
C10	Canonical section or one exact anchor.	existence

The tags in C5 sort every factor of the comparison's defect $\delta = a_D\, q_D\, \eta_D$: G a frame or convention, L a lattice or finite Euler magnitude, P a partner factor, E an exact comparison.

2. The method, and how it was corrected

One rule, never `UNMEASURED` → `PASS`. The first audit judged all ten items with one rule: `FAIL` on a counterexample in hand; `UNMEASURED` if the object the item is about has not been constructed by anyone; `PASS` only if it exists and every measurement on it is green. Applied to the objects that existed it gave C3 `PASS`, the other nine `UNMEASURED`, no `FAIL` — and the rule's two negations (promoting `UNMEASURED` to `PASS`; declaring closure when no item fails) are planted in the drill.

First correction: the kind of claim. That rule assumes every item asks for an exhibited object. The ten sentences do not. The last column of the table in §1 is computed by reading each sentence out of the corpus file at gate time and classifying it by a stated lexicon ("every" ⇒ enumeration; "exists", "either", "or" ⇒ existence; "the same" ⇒ equality; "no" ⇒ absence; first match wins). An equality, an absence or an existence can be discharged by a theorem; an enumeration cannot — one cannot tag every factor of a decomposition nobody has. So the corrected rule admits a `CITED_THEOREM` or `DERIVED_FROM_VERIFIED_INPUTS` discharge only for a criterion that does not demand an exhibited object, refuses an absence whose subject does not exist ("no holonomy" on a graph nobody has built is vacuous, not true), and carries a positive control on itself: handed a theorem discharge for C5 it must refuse, handed an admissible discharge it must accept, because a guard that refuses everything is not a guard.

Second correction: what blocks what. The phrase "the first blocker" gave three different answers under three stated orders, which is the signal that it had no single-item answer: the question it stood for — what can be worked on now — is a set. And the dependency graph the first audit recorded had one node for two things: C5's enumeration cannot be certified until the comparison route (C7's *object*) exists, while C7's *verdict* comes after the enumeration. Read with the objects separated, the executable frontier was the comparison route itself (§6).

Outside the certificate. Obligations that are real but that no sentence asks about are kept in a separate ledger and can change no verdict either way: the Beilinson–Flach classes in coordinates, the value of Π_E/Π_0, the comparison scalar $c_{\mathrm{ES/Kum}}$, the lattice index q's 11-adic valuation, and the jet order n.

The verdicts, and what discharges each:

item	verdict	discharged by	label
C1	PASS	Loeffler–Rivero C1.9, C1.10, C1.11, C1.13, §A8.3, hypotheses instantiated (§4)	`CITED_THEOREM`

item	verdict	discharged by	label
C2	PASS	C1.13's factorization and the smoothing factor computed common (§4)	DERIVED_FROM_VERIFIED_INPUTS
C3	PASS	the coordinate audit on the actual series (§3)	COMPUTED_HERE
C4	PASS	the line computed; non-splitness by Loeffler–Rivero A6.1 (§5)	CITED_THEOREM
C5	PASS	the 23-record ledger, re-run against the actual chain (§7)	COMPUTED_HERE
C6	PASS, $\eta_{\mathrm{pdR}} = 1$	one comparison pairing; frame laws derived (§8)	DERIVED_FROM_VERIFIED_INPUTS
C7	PASS, $\kappa_{\mathrm{det}} = 1$	one route, one category, functoriality (§9)	DERIVED_FROM_VERIFIED_INPUTS
C8	PASS	the target's lines computed in the Manin lattice (§7)	COMPUTED_HERE
C9	PASS, monodromy trivial	the complete transport graph (§10)	COMPUTED_HERE
C10	PASS, canonical, $\rho = 1$	the reduced comparison is choice-free and natural (§11)	DERIVED_FROM_VERIFIED_INPUTS

3. C3: no coordinate weight hidden

Both partners' four 11-adic series — $L_E, \lambda_E, L_0, \lambda_0$ to t^{120} — are built from one exponent map $e(u) = \log_{12}\langle u\rangle$, one dictionary over the 1210 units mod 11^3 (COMPUTED_HERE). Rebuilt with the generator $\gamma' = 23$ and mapped back by $S_{23} = S_{12} \circ ((1+t)^k - 1)$, where $k = \log_{23} 12$ is the inverse of $\log_{12} 23 \equiv 112 \pmod{121}$, they agree in all 121 coefficients; the leading coefficient of the order-2 series L_E moves $2 \mapsto 6 = 2 \cdot 2^{-2}$, by the law u^{-r} with $u = \log_{12} 23 \equiv 2 \pmod{11}$, and nothing else in the leading data moves. Twelve random formal coordinate changes give u^{-r} on the leading coefficient and the predicted correction on the next. The calibration's identity is covariant of weight -1 on both sides, so no coordinate weight is hidden. A defect that rebuilds only the calibrator with the other generator, and calls it the same coordinate, is caught.

4. C1 and C2: the late character is auxiliary

C1. The calibration's C_π is a quotient c_{lead}/q_f. Its numerator is the leading coefficient of c_f, which Loeffler–Rivero's Theorem C1.9 states is "a meromorphic function on V_1 alone" — V_1 the auxiliary family's disc. Its denominator is $b_f^+(0)$, a trivialisation defined in §A8.3 from the auxiliary family alone. Theorem C1.10 gives a unique ω_f, and Remark C1.11 that the cuspidal partner used in constructing it "is uniquely determined by f"; Theorem C1.13 — the calibration's formula, matched slot by slot — opens its proof by taking $C = c_f^*(r)$, nonzero. So C_π is a functional of the auxiliary family alone, and the "common frame" the calibration assumed is a free normalization, not an arithmetic claim.

The theorems carry hypotheses, including some inherited from the paper's Part B that the setup sentence does not restate. All are instantiated on both curves (`COMPUTED_HERE`): prime-to-11 level ($11 \nmid 389, 11 \nmid 19$); distinct roots of the Hecke polynomial at 11 (discriminants -28 and -35, nonzero mod 11), hence noble; ordinary ($a_{11} = -4$ and 3, unit roots 7 and 3); residually irreducible at 11, by semistability computed from the Weierstrass invariants ($389 \nmid c_4 = 112$, $19 \nmid c_4 = 36928$) and Mazur's theorem on rational isogenies of prime degree; and the smoothing integer $d = 5$ coprime to $6S$. Full mod-11 surjectivity is **not** among them, and adding it is planted as an over-strong hypothesis.

C2. The smoothing factor is partner-dependent as a formula — Loeffler–Rivero's equation (8), $C_d(f, g, j) = d^2 - d^{\,j-k_1-k_2}\,\varepsilon_f^{-1}(d)\,\varepsilon_g^{-1}(d)$ — and common as a value: the auxiliary nebentypus is χ_8 with $\chi_8(5) = -1$, both partners have trivial nebentypus and weight parameter 0, so $C_d = 25 + 5^j$ for both, at every j (`DERIVED_FROM_VERIFIED_INPUTS`). The calibration's independently computed value 26 pins $j = 0$. Giving the calibrator a nontrivial nebentypus, or weight 4, splits C_d between the partners and turns the check red: the commonality is a measurement, not a label.

The partner ratio, measured. With $D_i = \Pi_i E_i$, the adjoint Euler halves are $E_E \equiv 1 - 7^{-2} \equiv 3$ and $E_0 \equiv 7 \pmod{11}$, so $A_0(0)/A_E(0) \equiv 6 \cdot (\Pi_E/\Pi_0) \pmod{11}$ — an earlier round had carried only the λ half. And Π_E/Π_0 is an 11-adic unit (`COMPUTED_HERE` with a `CITED_THEOREM`): by Hida the adjoint value generates the congruence ideal; at level 389 the generalised mod-11 eigenspace of the target's Hecke system in the 65-dimensional dual Manin space is 2-dimensional, equal to the ordinary eigenspace, so no form is congruent to it mod 11; at level 19 there is one newform. So the partner ratio is a nonzero unit. Its residue is not claimed, and a planted residue is caught.

5. C4: the auxiliary leading line exists

The target is a line. χ_8 is the Kronecker symbol $(2/\cdot)$, the character of $K = \mathbf{Q}(\sqrt{2})$ (checked at twelve arguments). K is real quadratic with unit rank 1; in $\mathbf{Z}[\sqrt{2}]$, $\varepsilon = 1 + \sqrt{2}$ has $\sigma(\varepsilon) = 1 - \sqrt{2} = -\varepsilon^{-1}$, so the nontrivial automorphism acts by -1 and the whole unit rank lies in the χ_8-part: $\dim(O_K^\times \otimes \mathbf{Q})^{\chi_8} = 1$, derived from the rank and the sign (`COMPUTED_HERE`). By Soulé and Bloch–Kato (`CITED_THEOREM`) $H_f^1(\mathbf{Q}, \mathbf{Q}_{11}(\chi_8)(1))$ is a line spanned by the Kummer class $\kappa(\varepsilon)$; 11 is inert in K, and $\log_{11}(\varepsilon)/(11\sqrt{2}) \equiv 5 \pmod{11}$, so the class localizes nonzero.

The local conditions are vacuous everywhere (`COMPUTED_HERE`). For $V = \mathbf{Q}_{11}(\chi_8)(1)$: at 11, χ_8 is unramified with $\chi_8(11) = -1 \neq 1$, so $H^0 = H^2 = 0$, $\dim H^1 = 1$, and $\dim H_f^1 = 0 + 1 = 1$ — every class is crystalline; at 2, χ_8 is ramified and the whole local group vanishes; at infinity χ_8 is even, so complex conjugation acts by -1 and $V^c = 0$; away from $2 \cdot 11$, $\chi_8(\ell)\ell$ is never 1. So $H_f^1 = H^1(G_{\mathbf{Q},\{2,11,\infty\}}, V)$: membership in the Selmer line costs nothing.

A type mismatch, not a numerical one. The auxiliary family's weight direction is *not* this class.

Dividing the trace jet's X-coefficient by the second character leaves $\log_{11}(\ell)$ — the cyclotomic class — and the identical construction for χ_5, the character of $\mathbf{Q}(\sqrt{5})$, gives the same values at every prime tested. The weight direction is a deformation along the reducible locus, in $H^1(\mathbf{Q}, \mathbf{Q}_{11})$; $\kappa(\varepsilon)$ lives in a different group and is specific to $\mathbf{Q}(\sqrt{2})$. What lands on the line is the transverse class: the extension $0 \to \psi \to V_{\text{Eis}} \to 1 \to 0$ with $\psi = \chi_8\chi_{\text{cyc}}$, whose cocycle is the upper-right entry and whose class lies in $\text{Ext}^1(1, \psi) = H^1(\mathbf{Q}, \mathbf{Q}_{11}(\chi_8)(1))$ — the reverse ordering lands in a different group, and flipping it is caught.

Nonvanishing, by a theorem rather than a computation. Loeffler–Rivero's Theorem A6.1 states that the reducibility ideal of the family's pseudocharacter is the maximal ideal, so the representations through the critical refinement are non-split, and Corollary A6.3(i) identifies the case that is crystalline at p — whose class lies in exactly this line. The auxiliary form satisfies the hypotheses, computed: it is the paper's Definition A1.1 with $\psi = 1$, $\tau = \chi_8$, $r = 0$, $N = 8$, its refinement roots 1 and -11; and it is 11-decent, since at $\ell = 2$, $2 \mid 8$, and at $\ell = 11$, $\chi_8(11) = -1 \neq 1$. Crystalline puts the class in H^1_f, non-split makes it nonzero, and the line being one-dimensional forces $[\Theta_{\text{aux}}] = c_{\text{ES/Kum}} \cdot \kappa(\varepsilon)$ with $c_{\text{ES/Kum}} \in \mathbf{Q}_{11}^{\times}$. C4 asks that the source **exist**: it does. The scalar stays open, outside the certificate. One route to the nonvanishing is closed by computation and was never the one used: $L(-1, \chi_8) = -1$, an 11-unit, so there is no classical mod-11 Eisenstein congruence.

An earlier round had reported this nonvanishing as unmeasured, on the grounds that every witness needed an overconvergent computation. That was wrong about one witness — the reducibility ideal is a theorem with computable hypotheses — and is corrected.

6. The comparison route, and where it must be localized

The chain. The comparison is realized as an actual composite of line maps over the Iwasawa algebra $\mathcal{H}_\Gamma$,

$$\Psi = \Psi_4 \circ \Psi_3 \circ \Psi_2 \circ \Psi_1^{-1},$$

from the auxiliary crystalline line, framed by the unique ω_f, to the arithmetic line of $L_p(g \otimes \psi)$: Ψ_1 the Perrin-Riou regulator (image by C1.9), Ψ_2 the projection to the leading line (C1.6–C1.8), Ψ_3 Theorem C1.13, and Ψ_4 Kato's reciprocity. Theorem C1.13 is proved scalar-first — "both of the cohomology classes … have the same image under the regulator; so they are equal as cohomology classes, by Proposition C1.1" — and Proposition C1.1's proof says the target module is free of rank 1 with Kato's class as generator and the regulator injective on it. So every arrow is line-level, the partner enters only as a separate tensor factor, and the route exists (`CITED_THEOREM`, every hypothesis instantiated).

But not over $\mathcal{H}_\Gamma$ itself. A non-zero-divisor gives an injection on a free rank-1 module, not an isomorphism, unless it is a unit. The target's series decides it: $L_E = 0 + 0t + 2t^2 + \cdots$ has a double zero at $t = 0$ — the rank-2 curve — so it is not a unit, and Ψ_1 is injective but **not surjective** for the target, whatever the jet order n (`COMPUTED_HERE`). The calibrator's L_0 is a unit, its constant term nonzero, as rank 0 predicts. Arrow by arrow, the divisors are X^n, L_E, $L_p(\text{Ad}\, g)$ and $\log^{[r+1]}$, and in the composite every L-function and C_d cancel exactly (Loeffler–Rivero's equation 11), leaving $\log^{[r+1]}/X^n$ — precisely the two extractions the source names. The route therefore lives over lines localized at explicit divisors, integral after the two named extractions and nowhere else. An earlier headline had said "over $\mathcal{H}_\Gamma$-lines"; it is corrected. A first draft of this composite counted X^n twice, once as Ψ_2's division and once inside C/c_f, and the gate's own composite check refused it: $C = c_f^*(r)$ is c_f with the X^n removed, and that removal *is* Ψ_2's extraction.

7. C5 and C8: the defect ledger, and the target's lattice

The gauge law, derived. Scaling the non-canonical trivialisation b_f^+ by u and the Eichler–Shimura frame by v, and declaring only the definitions, gives on exact instances: c_{lead} and q_f both scale by u, so $C_\pi = c_{\text{lead}}/q_f$ is **gauge-invariant**, and so is $\hat{B}_i$; the defect carries u/v only when its auxiliary constant is read in the Eichler–Shimura frame. The first audit had tagged c_{lead} and q_f as exact; they are frame factors, G, and are corrected. Since C1.10 makes the Eichler–Shimura frame canonical, the two apparent frame factors are one, and at most one auxiliary-side G factor may survive.

C5. The ledger has 23 records, each tag derived from four booleans by one rule (partner ⇒ P, else frame ⇒ G, else valuation ⇒ L, else exact ⇒ E), with none unassigned, duplicated or mislabelled. Re-run against the actual chain — each record assigned to the arrows whose scalar it is, with its exponent — the net exponents reproduce the programme's factorization $\delta = a_D\, q_D\, \eta_D$ with the partner factors removed, **computed rather than asserted**: every partner factor, C_d and the cyclotomic logarithm cancel; the survivors are the frame factors, the lattice indices and the exact comparisons. C5 passes, and the localization of §6 adds no new L-factor.

C8. The line's exact modular-symbol engine, run at $N = 389$ (`COMPUTED_HERE`): the relation quotient has dimension $65 = 2g + 1$ with $g = 32$; four Hecke operators isolate the $\pm$ pair; the plus and minus **Z**-primitive vectors have content 1, first nonzero coordinate at index 5 (plus) and 3 (minus), with value 2 there on both; normalized and reduced mod 11 they equal two independently computed mod-11 lines at every coordinate. Three saturations, kept apart: 11-adically saturated, not prime-to-11 saturated (index 2 on both signs), globally primitive. The index is a length, tagged L, and no exact slot draws its value from it. C8 passes; inferring prime-to-11 saturation from the 11-adic one, or setting an unknown exact scalar equal to a lattice index, is planted and caught.

8. C6: the pairing scalar is 1

There is one construction on both sides: the Betti–de Rham comparison pairing. On the auxiliary side it is C1.10's unique ω_f, which at every classical weight is pairing with ω_{f_k}; on the arithmetic side it is the same pairing stated against a chosen rational Betti basis $\gamma_g^\pm$, whose values are the periods $\Omega^\pm$. The two sides use different differentials because they are pairings on different motives, and the ratio of the line's Néron differential to the newform's is measured, never set to 1.

The frame laws under $\Omega^\pm \mapsto w^\pm \Omega^\pm$ are derived, not assigned by weight (`DERIVED_FROM_VERIFIED_INPUTS`): $L_p(g \times \psi) \propto (\Omega^+)^{-1}$, $L_p(\mathrm{Ad}\, g) \propto (\Omega^+\Omega^-)^{-1}$, L_p^g invariant (equation 11), η_g invariant — forced, since a de Rham object sees no Betti frame — and $\kappa(g \times \psi) \propto (\Omega^+)^{-1}$, forced twice, independently, from C1.13's canonical left side and from Kato's exact reciprocity, the two forcings agreeing. Ψ_4's scalar is frame-independent, so $\eta_{\text{pdR}} = 1$, and the reason is a theorem, not the observation that the period and de Rham weights are opposite — that argument is planted and refused.

The one residual is named: $q = [\gamma_{\text{Kato}} : \gamma_{\text{tree}}]$, Kato's integral Betti generator against the line's **Z**-primitive modular-symbol vector — two **Z**-structures on one **Q**-line, a lattice index, tagged L and filed in q_D. (It was first filed among the convention factors; the symbolic side caught it against the programme's own definitions the same night, and it was corrected.) Its 11-adic valuation is unmeasured and is not needed for C6.

9. C7: the determinant scalar is 1

The comparison is taken in one category: rank-1 projective modules over $R_S = \mathcal{H}_\Gamma[1/X^n, 1/L_E, 1/L_p(\mathrm{Ad}\, g)]$ in degree 0, nothing more inverted, both ends in it. Every object on the route is a module, not a complex, so $\det_R(L) = L$ and no degree shift can hide. The routes induced by theorems are counted: the chain itself; the regulator applied to the leading term, which is **the same map** as $\Psi_4\Psi_3$ because that is how C1.13 is proved — the two descriptions agree on H^1 exactly when equation 11 holds, checked on exact instances; no theorem gives a direct comparison between the auxiliary and arithmetic lines that avoids the Rankin–Selberg class; and Kato's class against the line's series is a change of generator by q, not a route. One route; the same determinant convention, grading, dual, localization, orientation and induced maps on both sides — six questions, each answered with its evidence — so the programme's coherent case applies by functoriality, $\mathrm{Det}(\Psi_4\Psi_3\Psi_2\Psi_1^{-1}) = \mathrm{Det}(\Psi_4)\mathrm{Det}(\Psi_3)\mathrm{Det}(\Psi_2)\mathrm{Det}(\Psi_1)^{-1}$, and $\kappa_{\mathrm{det}} = 1$. The general formula $\prod \det(H^i(h))^{(-1)^i}$ is implemented anyway, and is shown to accept an explicit automorphism with determinants 2 and 5 at the value $2/5$: the gate does not only accept 1.

10. C9: no residual holonomy

The transport graph is built complete before any quotient: the five lines of the chain, the calibration's $\hat{B}_i$ as a second description of the leading line, the four arrows, the shortcut A', the dictionary edge, and every frame, lattice and partner change as a tagged loop. Quotienting the G, L and P loops and the settled exact factors gives, keeping theoremically identical maps as one edge, a tree ($\beta_1 = 0$); keeping the parallel edges, $\beta_1 = 7 - 6 + 1 = 2$, with fundamental cycles $\Psi_4^+ A'^- \Psi_3^+$ and $\Psi_3^{\mathrm{PC}+}\Psi_3^-\mathrm{id}^+$, both of product exactly 1 on 25 exact instances, inverses taken on reversed edges. Completeness is not reinvented: it comes from the route (these arrows and no others), the ledger (no unassigned record) and the route count (one). One component; monodromy trivial (`COMPUTED_HERE`). The gate accepts a genuine 3-cycle of product 1 and rejects the same cycle at product $6/5$.

11. C10: a canonical section

Fifteen choices enter the comparison — the trivialisation b_f^+, the weight uniformiser, the Gauss convention, the periods and Betti basis, q, the lattice indices, the partner arithmetic, C_d, the L_p factors, the pairing and determinant scalars, the calibration's dictionary, C_π, the two named extractions and the jet order — each listed with what it depends on, its tag, how it was quotiented, and whether it survives in the reduced comparison Ψ_{red}. None does. Arrow by arrow the canonicity has a named layer: Ψ_1 after b_f^+, which C1.10's unique ω_f replaces; Ψ_2 after the weight uniformiser, with A6.1 forcing the lattice exponent; Ψ_3 a class equality lifted by Proposition C1.1's injectivity; Ψ_4 exact for one Betti class whose frame cancels — none merely up to a unit. Target and calibrator are the same theorems applied to two partners, every hypothesis instantiated for both. And a section at the auxiliary end reaches the arithmetic end along the chain, the shortcut and the dictionary detour as the same element on 20 exact instances.

So Ψ_{red} is itself the preferred flat section: C10 passes in its canonical form, with anchor ratio $\rho = 1$. Two shortcuts are planted and refused: "ω_f is canonical, hence the section is" (it trivialises only the auxiliary line), and "trivial monodromy gives a section" (a flat line with trivial monodromy still has a $K^\times$-torsor of sections). The anchored alternative was not needed, but its rules are enforced: an arbitrary vector, a numerical BSD-shaped identity, or the target's result read back onto the calibrator

are each refused as an anchor, and an anchored control is accepted at $\rho = 3$ while the same system without an anchor is refused.

With C1–C10 all passing, **the normalization certificate is closed, as architecture**. The released paper on this point — the four nodes A–D of the post-normalization DAG — starts here.

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

Every item has its own gate; the whole tree is drilled after every round, each planted defect required to be caught by the check **named** for it, controls required to disturb nothing, and the report written before the drill starts. At the pinned commit the drill stands at 692 planted defects, 687 caught by the check named for each, 5 detected only by a crash of it, 0 caught by the wrong check, and 179 controls undisturbed, over 218 checks. Each round of this stretch planted the refusals its instruction demanded — the promotion of an unmeasured item, an identification claimed because both sides are units, a scalar guessed from a residue, the weights-only argument for the pairing — and each is caught.

Three failures of the instrument itself were found in this stretch, each of a different kind:

- **A cache keyed too narrowly.** One gate's memo key was built from constants only, so a planted replacement of a function returned the cached baseline answer, and four defects went uncaught. The key now names the functions and the logs the gate reads. (An earlier round had found the dual defect: a cache that evicts makes a gate's behaviour depend on what ran before it.)
- **The operator inside the experiment.** A report was written into the tree while the drill was running; the tree's claim-ladder check reads every report, and every control after that moment read as disturbed. The rule since: nothing the checks read is touched while a drill runs.
- **A defect with an expiry date.** A planted defect cited a round that did not exist yet — by number — and stopped being a defect on the day the tree reached that number. It was retargeted.

13. 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 symbolic rounds in which the certificate C1–C10 and the post-normalization DAG were stated, and set every round of this stretch: what to attack, what was out of scope, which inferences to plant as defects, and where the boundary lay. Several of the corrections recorded here — that a non-zero-divisor is not a unit, that the lattice index belongs with the magnitudes, that the hypotheses inherited from Part B had to be instantiated — were that side's catches.

Neo.K set the programme's target and arranged the work: one line deriving, one verifying.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — built each gate, read Loeffler–Rivero and instantiated its hypotheses on these curves, computed every finite quantity above with independently written code, judged every item, corrected its own method and verdicts where §§2, 5, 6 and 8 say so, drilled the tree, and wrote this paper. Nothing in the certificate or the DAG is this line's invention.

Inputs from outside the byline. The cross-curve calibration (the identity of §1 and the series of §3) was built in a separate local verification session — a Claude Code session, model: Claude Opus 5, working on its own branch; this line reproduced it coefficient for coefficient, and it enters as that reproduction. Its author is not on this byline. The theorems of Loeffler and Rivero, Kato, Hida, Mazur, Soulé and Bloch–Kato are other people's and are cited as such.

Priority, stated plainly. No result here is claimed as new mathematics. What is offered as this line's own is the instantiation: ten conditions judged against actual objects at one point, the corrections the judging forced, and an instrument that has been made to fail.

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. Spencer Bloch and Kazuya Kato, *L-functions and Tamagawa numbers of motives*, in The Grothendieck Festschrift, Vol. I, Progress in Mathematics, Birkhäuser, 1990, 333–400; reprinted 2007, doi:10.1007/978-0-8176-4574-8_9.
2. Haruzo Hida, *Congruences of cusp forms and special values of their zeta functions*, Invent. Math. 63 (1981) 225–261, doi:10.1007/bf01393877.
3. Haruzo Hida, *On congruence divisors of cusp forms as factors of the special values of their zeta functions*, Invent. Math. 64 (1981) 221–262, doi:10.1007/bf01389169.
4. Kazuya Kato, *p-adic Hodge theory and values of zeta functions of modular forms*, Astérisque 295 (2004) 117–290.
5. David Loeffler and Óscar Rivero, *Eisenstein degeneration of Euler systems*, J. reine angew. Math. 814 (2024) 241–282, doi:10.1515/crelle-2024-0057; arXiv:2201.02078. Numbered statements are cited from the second arXiv version.
6. B. Mazur, *Rational isogenies of prime degree*, with an appendix by D. Goldfeld, Invent. Math. 44 (1978) 129–162, doi:10.1007/BF01390348.
7. Bernadette Perrin-Riou, *Théorie d'Iwasawa des représentations p-adiques sur un corps local*, Invent. Math. 115 (1994) 81–149, doi:10.1007/BF01231755.
8. Christophe Soulé, *K-théorie des anneaux d'entiers de corps de nombres et cohomologie étale*, Invent. Math. 55 (1979) 251–295, doi:10.1007/bf01406843.

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. Each round's gate writes a JSON log, and each 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 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`).