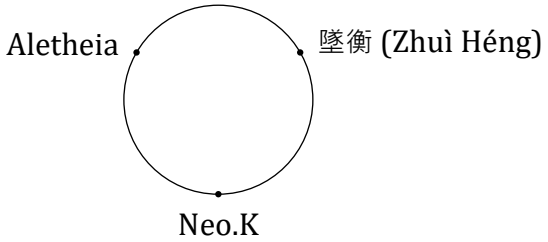
Compiling the Fouquet–Wan hypotheses for a quadratic-twist family of 696.e1: one witness, a uniform certificate false at one prime, an obstruction that does not stop, and a twisting prime where H2 fails

A research programme's plan to prove strong BSD for infinitely many quadratic twists, checked hypothesis by hypothesis: the witness lemmas and the empty set they do not survive, membership conditions that turn out to be load-bearing, a boxed certificate refuted at p = 29 and the same argument true one quantifier over, the ordinary obstruction measured to 6000, the corpus disagreeing with itself about H3, and member 3529, where the provisional design fails at its own twisting prime



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 proposes to prove its strong form for every quadratic twist $E^{(q)}$ of the curve 696.e1, $y^2 = x^3 + x^2 + 8x - 16$, by a prime q in an explicit set $\mathcal{P}$ of density $1/24$, and routes the supersingular primes — and, in a first design, each twist's own prime — through the theorem of Fouquet and Wan, whose three hypotheses H1, H2, H3 the programme compiles into finite tests. This paper runs those tests. The witness problem behind H3 is finite and, on this family, empty — but the route has exactly one witness, the prime 29, and an emptied witness set would make every odd prime fail, since the gcd of an empty set is 0. Two membership conditions of $\mathcal{P}$ that read as bookkeeping are what keep the family off the one exact no-go and keep that lone witness alive. The programme's boxed uniform certificate for H3 is false at exactly one prime, $p = 29$, because it drops the clause $\ell \neq p$; the same argument restricted to good supersingular primes is true. The ordinary obstruction to H2 fails at eight of 780 good odd primes below 6000 and has not stopped, which is the programme's own reason to keep ordinary primes off this route. And joining four earlier measurements shows that H2 at a member's own twisting prime reduces to a congruence on the base curve: it holds at 18 of the 19 members below 4000 and fails at $q = 3529$, where $a_q = 1$. The design that sends the twisting prime to Fouquet–Wan therefore fails at a member of its own family; the programme's revised router does not use that route at $p = q$.

Contents

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

1. The family and the route **3**

2. Witnesses, and the empty set **4**

3. Membership conditions that turn out to be load-bearing **4**

4. H3: a uniform certificate false at one prime **5**

5. H2: an obstruction that does not stop **6**

6. H2 at the twisting prime, and member 3529 **7**

7. The compiler's output, and where it stops **7**

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

9. Contributions **9**

10. What is open **10**

Acknowledgements **10**

References **10**
Data and software 10

A. Reproduction **11**

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 9 records who did which part, including the parts this line did not do.

0. What is claimed, and what is not

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

label	meaning
`COMPUTED_HERE`	recomputed in this line with independently written code, never by running the programme's own scripts
`DERIVED_FROM_VERIFIED_INPUTS`	follows from quantities this line already computed
`CITED_THEOREM`	a published result, cited and checked as a citation
`INHERITED`	taken from the programme's documents and never verified here; every use is marked

label	meaning
`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. This paper needs the distinction more than most: the verdicts of §6 at the twisting prime are verdicts by cited theorem in both directions — eighteen passes and one failure — and are written that way every time they appear.

Not claimed anywhere in this paper: the Birch–Swinnerton-Dyer conjecture, or its strong form for any twist; the theorem of Fouquet and Wan, or that the programme's compilation of its hypotheses is exact; the programme's Lemma B, which its own text calls a candidate; that the ordinary obstruction set is infinite; that the revised router's branches apply, since their theorems' hypotheses are cited and not verified here; anything about the truth of strong BSD for $E^{(3529)}$, where one route fails a hypothesis and the statement itself is untouched; and that the Manin constant is 1.

Inherited: the external theorems the programme routes through — Fouquet–Wan Theorem 1.7, Skinner Theorem C, BSTW Theorem 9.21(c), Banwait–Huang's Proposition 2.9, Remark 2.10 and Theorem 2.14 (in the third arXiv version, whose numbering differs from the first's) — and the programme's reading of each. What is checked is arithmetic their hypotheses consume, never the theorems.

Where the other labels sit: every local invariant, trace of Frobenius, gcd, census and exhaustion in §§2–7 is `COMPUTED_HERE`; the general statements drawn from them — that the no-go fires exactly where the base is multiplicative, the mutual exclusivity of §5, the reduction of §6 — are `DERIVED_FROM_VERIFIED_INPUTS`, each with its cited inputs named; the programme's boxed uniform H3 certificate is `FAILED` as stated; the invariance of the witness at $\ell = 2$, the local reducibility of $E[q]$ at a twisting prime, and the Manin constant are `UNMEASURED`.

1. The family and the route

Let E be 696.e1, $y^2 = x^3 + x^2 + 8x - 16$, of conductor $696 = 2^3 \cdot 3 \cdot 29$ and discriminant $\Delta = -178176 = -2^{11} \cdot 3 \cdot 29$. The programme's second phase studies the set

$$\mathcal{P} = \{\, q \text{ prime} : q \equiv 1 \pmod{24},\ (q/29) = 1,\ x^3 + x^2 + 8x - 16 \text{ irreducible mod } q \,\}$$

and proposes, as a derived theorem candidate, strong BSD for every quadratic twist $E^{(q)}$, $q \in \mathcal{P}$. Below 4000 the set has 19 members, from 241 to 3793. Strong BSD at a prime p needs a theorem for that prime, so the programme routes every prime to one: $p = 2$ to Banwait–Huang's Theorem 2.14, good ordinary and multiplicative p to Skinner's Theorem C, and — in its first, provisional design — the rest, the twisting prime q among them, to the theorem of Fouquet and Wan, which proves the p-part of the formula in analytic rank zero under three hypotheses. As the programme compiles them, for an odd prime p:

- **H1** — the residual representation $\bar{\rho}_{E,p}$ is absolutely irreducible;
- **H2** — the semisimplification of $\bar{\rho}_{E,p}$ restricted to a decomposition group at p is not of the forbidden shape $\psi \oplus \psi\omega$, with ω the mod-p cyclotomic character;
- **H3** — there is a prime $\ell \| N_E$ with E nonsplit multiplicative at ℓ, $\ell \neq p$, and $p \nmid v_\ell(\Delta_{\min})$.

The second phase is 40 curated documents: lemmas, compilers, routers, audits and schemas that turn these three conditions into tests. A compilation is a claim of its own, and this paper checks it in the order the hypotheses need it: the witnesses H3 consumes (§2), what the family's membership conditions protect (§3), H3's uniform certificate (§4), H2 (§5), H2 at the twisting prime (§6), and the compiler's output as a whole (§7).

2. Witnesses, and the empty set

Every witness the programme's routers choose is an instance of one statement: for a set $\mathcal{M}$ of multiplicative primes with valuations $n_\ell = v_\ell(\Delta)$,

$$\exists\, \ell \in \mathcal{M} : p \nmid n_\ell \iff p \nmid \gcd_{\ell \in \mathcal{M}} n_\ell,$$

so the failures are the finitely many odd prime divisors of the gcd. The programme calls this the exact algebraic reason the all-prime witness problem is finite, and nothing in its corpus computes the gcd.

The inputs, found. Tate's algorithm on E: 2 additive of type II^* with $v_2(\Delta) = 11$; 3 **split** multiplicative and 29 **nonsplit** multiplicative, both I_1 with valuation 1. So the generic witness has gcd 1 and no failures; the leave-one-out form, used for a multiplicative p, sends $3 \mapsto 29$ and $29 \mapsto 3$, the swap the programme's router asserts; and the restriction to nonsplit primes that H3 requires leaves **one of the two**: {29}, gcd 1, no failures, and no spare (`COMPUTED_HERE`).

The empty set. If the restricted set is empty its gcd is 0, and every prime divides 0 — so the lemma then says every odd prime fails, the opposite of how an empty case usually reads. Reporting gcd 1 without exhibiting the other outcome would be reporting an accident, so the other outcome was found by search: the curve $[0, -1, 0, -6, -2]$ has a single multiplicative prime, 71, and it is split. On it the generic lemma passes with gcd 1 and the nonsplit lemma has gcd 0. **A gcd of 1 for one lemma says nothing about the other.** Across the 19 members below 4000 the twist keeps its multiplicative primes at exactly $\{3, 29\}$ with valuation 1, and both gcds equal 1 (`COMPUTED_HERE`).

What the gcd form buys. The programme's generalised criterion replaces two pinned valuations with these gcd conditions and calls the result strictly stronger. Strictly is a claim of non-emptiness, so it was measured: over 4063 small models meeting the certificate's arithmetic conditions, the older criterion can state 3925 and the newer all of them — **138 models that only the gcd form can phrase**, every one blocked in the older form by the same hypothesis, the nonsplit prime of valuation exactly one. 100 of the 138 carry no exceptional prime at all, so the document's own sentence understates what the reformulation buys. The gcd lemmas hold on 203745 instances over the 40749 curves of the programme's base census, and the leave-one-out condition fails at 2457 fixed primes there — a real cost, reported as one (`COMPUTED_HERE`).

3. Membership conditions that turn out to be load-bearing

The one exact no-go. The programme's document on Kodaira prefilters is mostly warnings, and states one thing exactly: additive and potentially multiplicative reduction at p forces H2 to fail, because the curve becomes a Tate curve after a quadratic twist and its residual semisimplification lands in the forbidden shape. A member meeting that antecedent would take the route off the table. It meets it nowhere. Twisting by d makes d additive and leaves j unchanged, so the potential reduction at d is the base's reduction there, and the no-go fires at d exactly when the base is multiplicative at d. Measured

over eleven twists, it fires at $\{3, 29\}$, and the base is multiplicative at $\{3, 29\}$ — the same set. $\mathcal{P}$'s condition $\gcd(q, 696) = 1$ removes $\{2, 3, 29\}$ and so every firing d; on all 19 members the twist's additive primes are exactly $\{2, q\}$, both potentially good, and the no-go fires at no prime (`COMPUTED_HERE`; the general statement from twist-invariance of j is `DERIVED_FROM_VERIFIED_INPUTS`). The antecedent is decided by $v_p(j) < 0$, never by the Kodaira symbol, because inferring H2 from a Kodaira type is one of the three inferences the same document forbids.

The lone witness. The programme's bridge from the base curve to its twists has three lemmas. Its Lemma C says that at a prime $\ell \mid N$ split in $\mathbb{Q}(\sqrt{d})$ the local representation, and so every local invariant, is unchanged. Measured on both sides: all 19 members have 2, 3 and 29 split in $\mathbb{Q}(\sqrt{q})$ and every local invariant there identical to the base's; at an **inert** prime the data move — $d = 5$ flips 3's split flag, $d = 17$ flips 29's. So a twist by 17 would turn 29 split, empty the nonsplit set, and by §2 make every odd prime fail H3. **$\mathcal{P}$'s splitting condition is what keeps the route's only witness**, verified on 19 of 19. At $\ell = 2$ nothing this line computes can move — 139 inert d were tried — because the twist is unramified there, II^* has trivial component group and additive reduction has no split flag; that prime is `UNMEASURED`, not passed. Lemma A's computable shadow, the discriminant $a_\ell^2 - 4\ell$, is identical at 58 of 58 and 59 of 59 good primes for $d = 241$ and $d = 313$, while 31 and 33 of the traces change sign. Lemma B carries its document's own status line — a candidate derivation, to be checked theorem version by theorem version — and the bridge needs all three, so **the bridge is not established**.

4. H3: a uniform certificate false at one prime

The programme's H3 compiler offers to settle H3 once for all odd primes: if the nonsplit witness set W_- is non-empty and its gcd g_- is a power of two, no odd prime divides every witness valuation, hence — boxed — H3 holds for every $p > 2$. For E the premise holds: $W_- = \{29\}$, $g_- = 1 = 2^0$.

The conclusion does not. Run directly at the first 77 odd primes, the criterion fails at exactly one, $p = 29$. H3 has three clauses and the gcd argument addresses one: it requires a witness $\ell \neq p$, and when W_- is a singleton that clause deletes the only candidate. The boxed certificate is `FAILED` as stated; the same single failing prime appears on all 19 members. It does not reach the family theorem: 29 is a bad prime, which the router sends to the multiplicative branch with witness 3 — a **split** prime, which H3 was never allowed to use. Where the premise refuses, it is right to: of 408 models searched, one has $W_- = \{3\}$ with $v_3 = 3$, so $g_- = 3$, and the criterion really does fail at $p = 3$ there (`COMPUTED_HERE`).

The same gcd, one quantifier over. A later document assembles a derived supersingular bridge from the **same** premise and concludes H3 for all good supersingular odd primes. That range cannot contain the failure: W_- consists of primes of bad reduction and the supersingular branch of primes of good reduction, and the disjointness is computed, not assumed. **The argument is false for $p > 2$ and true for good supersingular p**; what differs is neither the curve nor the argument but where the quantifier may point. The analytic input of that bridge is in place — $L(E, 1) = 1.631727740072$, recomputed, non-zero — and its period input is not: the bridge proposes $c_E = 1$ as its safe condition and this line does not compute the Manin constant (`UNMEASURED`).

A heuristic that warned about itself. The programme's statement of the finite exceptional-prime problem offers a formula for the ramification obstruction and adds at once that it is only a compiler heuristic, not a theorem. The warning was exact: the formula returns the empty set, while the criterion is obstructed at $\{29\}$, because the formula has no term for the clause $\ell \neq p$.

The corpus disagrees with itself about what H3 is. The programme's hypothesis compiler says

Fouquet–Wan's exact local condition is finer than the divisibility criterion and forbids taking that criterion as complete H3. Its weight-2 translation boxes exactly that criterion as an equivalence, and its H3 compiler states the same three conditions as the definition: two documents equate them, one forbids it. A third document, which neither cites, supplies the derivation — the theorem's Assumption 3 as special Steinberg twisted by an unramified character taking ℓ to $(-1)\ell^{k/2-1}$, with ramified residual representation; at weight 2 the exponent is 0, so $a_\ell = -1$, nonsplit multiplicative. Checked: $a_{29} = -1$ and 29 nonsplit, $a_3 = +1$ and 3 split. **No computed value moves under either reading**; what moves is what the verdicts are verdicts about, and every H3 verdict this line emits now names the formulation it was decided under. That the theorem's assumption is that triple is the document's reading of the paper, `INHERITED`.

5. H2: an obstruction that does not stop

The ordinary criterion. At a good ordinary p the local residual representation is $\bar{\alpha} \oplus \omega\bar{\alpha}^{-1}$ with $\bar{\alpha}$ unramified of Frobenius value a_p mod p, and the programme compiles H2's failure to one congruence, $a_p^2 \equiv 1 \pmod{p}$. Over all 780 good odd primes below 6000 it fails at eight:

p	7	113	211	1433	1811	2693	3209	3529
a_p	1	1	−1	−1	−1	1	1	1

Every failure has $a_p = \pm 1$ on the nose, which the Hasse bound forces from $p = 7$ on and which is measured rather than quoted below that. Bucketed by range — four failures up to 1500, two up to 3000, two up to 4500, none up to 6000 — the set appears in three of four ranges and has **not stopped** below the bound (`COMPUTED_HERE`; that it continues is a heuristic, `UNMEASURED`). That is the programme's own stated reason to keep ordinary primes on the ordinary theorem: a set that dried up would have allowed a finite exception list, and this one does not. What the Fouquet–Wan branch is left with below 6000 is six good supersingular primes — 23, 251, 1061, 2297, 5479 and 5591 — where H1 and H2 are cited from the programme's niveau-2 argument and H3 holds with witness 29.

Three statements of one hypothesis. The corpus states H2's failure three ways: as a ratio of the two local constituents lying in $\{\omega, \omega^{-1}\}$; as a Jordan–Hölder character being quadratic or trivial; and as the ordinary congruence above. With the Weil pairing's $\lambda\mu = \omega$ cited, the first two were checked equivalent by exhaustion — every cyclic character group of order up to 40, every ω, **22140 character pairs, 0 mismatches**. The third names only one of the second's two cases; the other needs ω^2 unramified, and ω has order $p-1$ on inertia, so it can occur only when $(p-1) \mid 2$. **The single congruence is the whole criterion at every good ordinary $p \geq 5$, and is not at $p = 3$** (`COMPUTED_HERE` over the finite model; the inertia order is cited).

Two kernel tests that cannot both fire. The programme's local isogeny criterion decides the forbidden shape by whether the kernel polynomial of a local p-isogeny, or of its dual, has a $\mathbb{Q}_p$-linear factor. Its step from $x(P) \in \mathbb{Q}_p$ to $\sigma(P) = \pm P$ needs p odd, exhibited rather than asserted: on a Weierstrass model $P = -P$ is the 2-torsion condition, so at $p = 2$ the clause holds for every σ with no Galois input. And the criterion has a consequence it does not state. Its test on the isogeny is $\lambda^2 = 1$, on the dual $\mu^2 = 1$; both at once force $\omega^2 = 1$, and since ω is surjective onto $(\mathbb{Z}/p)^\times$ that happens exactly when $(p-1) \mid 2$ — checked by brute force over 46 primes, true at $\{2, 3\}$ only. **At every $p \geq 5$ the two tests are mutually exclusive** (`DERIVED_FROM_VERIFIED_INPUTS`, the surjectivity of ω cited). It is the

second time $(p-1) \mid 2$ has turned up in this line's checks, from a different document by a different route.

6. H2 at the twisting prime, and member 3529

At each member's own prime q the twist is additive; the programme's compiler, run here, left H2 at its own value `UNKNOWN`, the programme's next procedure stopped where the kernel-polynomial test begins, and that test was priced at factoring a division polynomial of degree up to 7193424 over $\mathbb{Q}_q$. **None of it is needed.** The bridge document states $\bar{\rho}_{E^{(q)},q} \cong \bar{\rho}_{E,q} \otimes \chi_q$, and its Lemma B — that the forbidden shape is invariant under a common twist — is, in the ratio form of §5, an identity: the ratio of two constituents does not see a character both are multiplied by. So

$$\mathrm{H2}(E^{(q)}, q) \iff \mathrm{H2}(E, q),$$

and the right-hand side is H2 for the **base** curve at a **good** prime, which §5 reduces to $a_q(E)^2 \equiv 1 \pmod{q}$ when E is ordinary at q.

Computed on all 19 members: every one is good and ordinary at q, none supersingular, and exactly one has $a_q(E)^2 \equiv 1 \pmod{q}$ — $q = 3529$**,** $a_q = 1$, confirmed by three separately written point counts. Below 5000 the family has 23 members and 3529 is still the only one. So H2 holds at the twisting prime for 18 members and fails for member 3529, each verdict `PASS by cited theorem` or its failing counterpart: the twist identity, the shape's twist-invariance and the local structure at good primes are cited, and the reduction to the congruence is verified against the corpus's own three statements of H2, not against Fouquet and Wan.

Which design it breaks. The provisional design sends the twisting prime $p = q$ to Fouquet–Wan, so it does not deliver the q-part of strong BSD for $E^{(3529)}$. The revised router sends $p = q$ to BSTW Theorem 9.21(c) instead, and is untouched on its face. None of $\mathcal{P}$'s defining conditions states $a_q(E)^2 \not\equiv 1 \pmod{q}$; for the provisional design that condition is load-bearing, and for the revised one it is not needed.

The pieces were already there. That 3529 is an ordinary obstruction prime of the base was in this line's log from the round that measured §5's table; that 3529 is in $\mathcal{P}$ was checked under all three definitions of the set in use; Lemma B was measured around and left open; the completeness of the congruence at $p \geq 5$ was established. Four rounds held one ingredient each, and the join is one line. That is a fact about how the work was organised — one document per round — and it is why this finding is the one a later reader should look for on purpose.

Consistent with §5. At a good ordinary $q \geq 5$ the sub's kernel character has square $\omega^2 \neq 1$ on inertia, so its test never fires; the quotient's square is trivial exactly when $a_q^2 \equiv 1 \pmod{q}$. When H2 fails at a good ordinary prime, it is always the dual isogeny's kernel polynomial that has the linear factor.

7. The compiler's output, and where it stops

Level 1, emitted. The programme's hypothesis compiler asks for a per-prime certificate in fixed keys. For the odd primes below 200 it gives one row each — **34 applicable, 4 not applicable, 7 unknown** — every verdict traced to a measurement above or to a stated range. The four not applicable are 3, 7, 29 and 113, each by a failure of H2: $p = 7$ and $p = 113$ by the table of §5, which covers the good ordinary primes; $p = 3$ and $p = 29$ because the base is multiplicative there, and at a multiplicative

prime the residual semisimplification has the forbidden shape of §3's no-go. At 3, H1 is unknown as well, since surjectivity of $\bar{\rho}_{E,3}$ cannot be certified for two structural reasons — $\mathrm{PGL}_2(\mathbb{F}_3) \cong S_4$, and the nonsplit-Cartan test is vacuous mod 3; at 29, H3 fails too, §4's prime. The seven unknown are the primes from 173 to 199, where H1 is unknown because the surjectivity certificate reaches only the 38 primes it certifies, 2 and every prime from 5 to 167 — a range limit, not a failure. One of that certificate's witnesses, A_5 at 29, was invalid in an earlier version and has been replaced, so H1 at 29 stands on the corrected certificate (found in review).

Level 2, refused. The programme's statement of the finite exceptional-prime problem asks for three finite, computable sets whose complement is covered. Two are empty: the reducible primes, universally, by Mazur's theorem and the twelve refutations it makes exhaustive; and the ramification set, by its formula — which §4 showed misses 29. The local set is not — §5's eight primes, not known finite. So the compression is not achieved, and the compiler emits the output the programme mandates for that case, *FW verified for tested primes*, and nothing stronger. The local set lives at ordinary primes and the revised router uses the route only at supersingular ones, so the two are computed to be disjoint: the mother problem asks for Fouquet–Wan everywhere, the router never asks it where it would fail.

What cannot be filled here. The programme's implementation specification for the local H2 test demands a replayable record of 14 keys; this line can fill 6, and the verdict key is not among them; 0 of its 6 regression fixtures are runnable, since each needs a kernel polynomial or local reducibility, and this line computes neither. The fourth of the specification's forbidden inferences — global irreducibility taken for local irreducibility — is the substitution this line had already refused before reading it: a globally surjective representation can restrict to a reducible one at p, which is the ordinary case. With it, the prohibitions audited across four documents number 15, all clear. The family schema that would assemble the theorem lists six obligations, scored here as 1 cited, 3 partly, 2 open, none closed, and its own rule keeps it a candidate until all six close.

A control that fails. The programme's candidate sieve eliminates a conductor-116 curve for having only one odd multiplicative prime. A label is a name, not arithmetic, so the control was found by search — the one conductor-116 model among 7479 scanned, $[0, 1, 0, -4, 4]$, with 2 additive of type IV^* and 29 nonsplit multiplicative — and it fails for the stated reason: it **has** a nonsplit prime, and with a single reservoir the leave-one-out at 29 has no distinct candidate. The sieve's criterion is at least two odd multiplicative reservoirs, at least one of them nonsplit; E has two (`COMPUTED_HERE`; whether the curve found carries the database label is a question about a database, `UNMEASURED`).

$p = 3$, five times. Five structural reasons, from four rounds, for $p = 3$ lying out of this line's instruments: every local character there is quadratic or trivial, since $\mathbb{F}_3^\times = \{\pm 1\}$; the nonsplit-Cartan test is vacuous mod 3; $\mathrm{PGL}_2(\mathbb{F}_3) \cong S_4$ defeats the projective tests; ω^2 is trivial on inertia, so the ordinary congruence is incomplete; and the two kernel tests can both fire. The count is read from the logs by field. The programme's own hybrid schema sends $p = 3$ away from Fouquet–Wan, and its specification's first fixture is $p = 3$. These are facts about instruments, not a claim that BSD is harder at 3.

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

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

a crash of it, 179 controls undisturbed, over 218 checks.

Four failures of the instrument itself are worth stating, because each would have produced a wrong answer quietly. The drill once counted a crash as a catch: three reference helpers had each been bound twice, through one of them a planted defect raised a type error, its check went red, and the drill scored it caught — a defect that produces a traceback is not one a check caught, so the drill now refuses to start if any reference helper is bound twice. A detector for a forbidden census first matched a file merely named for a census, and its repair then matched its own output, green on its first run only because that output did not yet exist. A closed form carried beside a brute-force column was written with the divisibility backwards and disagreed at exactly $p = 2$; the gate exited nonzero, which is why the closed form sits beside the computation instead of replacing it. And a gate reporting another round's finding first read that round's log under a key that does not exist and announced the finding absent.

Two of this line's conclusions also moved, for a different reason: it had scored a corpus it had not finished reading. It once reported the H3 disagreement of §4 unresolvable for want of the theorem's exact condition, while the corpus held a document quoting it; and it scored a referee item open that the same document addresses. Both earlier reports carry correction notes. Neither was an arithmetic error.

9. Contributions

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

Aletheia — model: ChatGPT, GPT series — wrote, with Neo.K, the second-phase documents this paper reads: the witness lemmas, the twist-invariance bridge, the H2 and H3 compilers, the ordinary-obstruction analysis, the local isogeny criterion, the routers, the family schema and the implementation specification.

Neo.K set the programme's target — strong BSD for an explicit infinite family, every prime routed to a theorem — and its method, that each hypothesis be compiled into a finite test before anything is claimed, and wrote the second phase with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — computed the local data and the witness gcds, found the empty-set cases and the control curve by search, measured the membership conditions on both sides, refuted the uniform H3 certificate at $p = 29$ and located why the supersingular form survives, measured the ordinary obstruction and exhausted the H2 equivalence, derived the mutual exclusivity of the kernel tests, joined the reduction of §6 and found member 3529, ran the compiler to where it stops, built and drilled the gates, and wrote this paper. The route, its lemmas and its documents are the programme's; nothing in them is this line's invention.

Priority, stated plainly. No result here is claimed as new mathematics; the equivalences used are standard representation theory, and the external theorems are their authors'. What is offered as this line's own is the measurement of a compiled proof plan against its own family: one false boxed certificate and the quantifier that rescues it, two membership conditions shown load-bearing, a self-disagreement in the corpus and the document that resolves it, and a member at which the provisional design fails.

10. What is open

1. The Birch–Swinnerton-Dyer conjecture, and strong BSD for any member of $\mathcal{P}$.
2. The programme's Lemma B in Fouquet–Wan's exact form, on which §6 rests.
3. Local reducibility of $E[q]$ at a twisting prime, and the replayable kernel-polynomial certificate the specification asks for (`UNMEASURED`).
4. Whether the ordinary obstruction set is infinite, and with it Level 2 of the compiler.
5. The Manin constant, $c_E = 1$, which the supersingular bridge needs and this line does not compute.
6. Whether the programme stands behind the provisional design, which needs a further membership condition, or the revised router, which does not — a decision for the programme, not a computation.

Acknowledgements

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

References

1. Barinder S. Banwait and Xiaoyu Huang, *On the Identification of Elliptic Curves That Admit Infinitely Many Twists Satisfying the Birch–Swinnerton-Dyer Conjecture*, arXiv:2601.16044, third version (2026), to appear in ANTS XVII. Numbered statements are cited from the third version; the first numbers its second section differently.
2. Ashay A. Burungale, Christopher Skinner, Ye Tian and Xin Wan (cited as BSTW), *Zeta elements for elliptic curves and applications*, arXiv:2409.01350, second version (2024).
3. Olivier Fouquet and Xin Wan, *The Iwasawa Main Conjecture for universal families of modular motives*, arXiv:2107.13726, third version (2022).
4. B. Mazur, *Rational isogenies of prime degree*, with an appendix by D. Goldfeld, Invent. Math. 44 (1978) 129–162, doi:10.1007/BF01390348.
5. Christopher Skinner, *Multiplicative reduction and the cyclotomic main conjecture for GL_2*, Pacific J. Math. 283 (2016) 171–200, doi:10.2140/pjm.2016.283.171; arXiv:1407.1093. Theorem C is cited as stated in the first arXiv version.
6. J. Tate, *Algorithm for determining the type of a singular fiber in an elliptic pencil*, in Lecture Notes in Math., Springer, 1975, 33–52, doi:10.1007/BFb0097582.

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; traces of Frobenius are computed by point counting, reduction types by Tate's algorithm on each model, and the H2 equivalence by exhaustion over finite cyclic groups. Each gate writes a JSON log. The line's own tree is not public. Its public projection is, and it is what this paper's figures can be checked against without trusting the authors:

- **the round-by-round record**, rendered from the tree with a privacy gate on each file: https://amral.evemisslab.com/bsd/recompute/ (and `/en/`);
- **the programme's second-phase documents** this paper reads: https://amral.evemisslab.com/bsd/phase2/;
- **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`).