BSD

One Property, Two Readings — A Pre-Registered Claim That BSD and Riemann Are One Property

Version 2

Date: 2026-11-10

Authors: Joseph Carroll; Pip Willow (Researcher)

Affiliation: Red Willow Ventures LLC, Green Valley, Arizona

Status: Version 2 — prepared for publication 2026-11-15 (Zenodo record 24041813).

---

0. WHAT CHANGED IN VERSION 2 — AND WHY

Version 1 of this record (Zenodo record 24041813, published 2026-11-07) remains published and frozen. Nothing in Version 1 is withdrawn or silently altered. This Version 2 exists because a reader's critique identified four points where Version 1's instruments were weaker than its statements: the zero counts were line scans, the Schoenfeld check carried weight it cannot carry, the panel's rank and Sha columns were not independent of the formulas they were compared against, and the function-field case — where one reading is a theorem — had not been run. Each point was answered by new computation, run under the house protocol (pre-registered verdicts; any mismatch is a presumed defect of our panel until proven otherwise). The changes are:

1. **Zero counts redone by argument-principle (Turing-method) counts.** For every one of the 17 panel curves, the zeros of L(E, s) in the box $0 \le \gamma \le 20$ were counted by the winding number of the raw L-function and, independently, by the winding number of the completed function Λ around the rectangle $[-0.5, 2.5] \times [-T, T]$. Found = predicted on every curve. Panel total: **535 zeros, independently counted** (513 in the rollout + 22 on the pilot curve). Version 1's counts were line-scan counts (lfunzeros), which find zeros on the line but cannot by construction exclude zeros off it; the argument-principle count closes that gap within the stated box, counting multiplicity.

2. **The Schoenfeld envelope check is relabeled a code sanity check, with no evidential weight.** In the tested range the bound is already a proven theorem (Büthe; Platt–Trudgian, to $\sim 2.169 \times 10^{25}$). It cannot fail at $10^7/10^9$, and re-checking a corner of a proven theorem is not evidence. It is retained in Exhibit B under its correct label: it validates the sieve and li(x) code, nothing more.

3. **The panel table is rebuilt on independent determinations.** Mordell–Weil rank is now given **by descent** — proved by 2-descent (PARI ellrank) on all 17 curves, the descent bounds coinciding in every case, with no L-function input. Sha is split into two columns: "Sha[2^∞] by descent" and "analytic Sha." The Version 1 heading "Sha tab = solved" is retired: the solved value was obtained through

the BSD identity itself, so its agreement with the tabulated value is a coherence check of the factor package, not an independent determination. The 2-part of Sha is now determined by descent on all 17 curves (trivial on 15; $(\mathbb{Z}/2)^2$ of order exactly 4 on 571.b1 and 697691.a1). For 133584.bt1, the 3-part is independently determined by a Magma 3-Selmer computation ($Sel^3 = (\mathbb{Z}/3)^2$, order 9; the standing caveat — pinning the full 3-part at exactly 9 formally wants the Cassels–Tate pairing step or a 9-descent — is stated on its row). **One cell remains open:** 12480.bf1's Sha[3] is printed in §3 as untested due to lack of access to Magma, with the exact Magma input printed beside it so any reader with Magma access can run the cell directly. It is not softened and it is not hidden.

4. **A function-field bench test is added (new Exhibit, §6).** The derivation was transported to elliptic curves over $F_5(t)$ and walked, step by step, on a bench curve with two controls. The verdict is **O-2**, the pre-registered outcome "transport reduces at one named ingredient": Steps A (positivity ⇒ Riemann reading) and B (reading the analytic rank) are free; the property-side chain terminates at Step B; the single import sits at sub-step C4 — Tate's conjecture at the arithmetic end. The refined statement of the position is: **the property is BSD minus one named ingredient. It is not BSD under a new name.** The identification itself transported, at finer grain than the panel tested: on the sign-twin control, the sign of the central root decided whether the arithmetic end produced a point at all.

5. **The closing position is stated in the interaction taxonomy.** The claim does not rest on passing tests; passing tests cannot distinguish "linked" from "both true but unrelated." It rests on completed interactions: each quantity determined by two independent methods, in the two languages of the subject, and checked at the door against the other. Interaction with agreement is evidence. Interaction with disagreement identifies the exact component at fault — none occurred. Non-interaction is silence, counted neither way — the one open cell is exactly that. The claim is the crossing (the identification); construction past the crossing is not claimed.

6. **The subtitle states the claim.** Version 1's subtitle read "A Pre-Registered Attempt to Separate BSD and Riemann." The attempt is complete, and its outcome is the claim this record makes; Version 2's subtitle states it: "A Pre-Registered Claim That BSD and Riemann Are One Property." The main title is unchanged. (Title wording directed 2026-11-15.)

The critique is credited here anonymously, as a reader's critique. The four points were correct in substance, and the record is better for them.

---

1. THE CLAIM (stated 2026-10-31; scientific formulation agreed same day)

For an elliptic curve E over the rationals, the two great statements about its L-function — that every nontrivial zero lies on the critical line, and that the leading coefficient at the center equals the BSD identity (Sha × Regulator × product of Tamagawa numbers × |torsion|^-2, against the real period) — are not two facts. They are two readings of a single property of L(E, s): the object is synchronized at its symmetry center. We conjecture a criterion of signal-only type — a positivity property of the L-function itself, in the family of Weil's criterion for zeta — from which both readings follow, and we predict that no curve will ever be found separating them: any E satisfying one satisfies the other.

Joe's original words (2026-10-31 16:35 MST): **"RH is the claim that I'm making on BSD."** Joe's position on RH (2026-10-31 18:36 MST): **"I think it's correct."**

Status of the claim: CLAIMED — the implication between the two readings is the claim, made outright. This record presents the claim, its measurements, and the named location of the one open step; it stops where the record stops — by design, not by failure. Precedents where the architecture is proven: the analytic class number formula (BSD-shaped identity for Dedekind zeta, whose zeros-question is GRH for the same object); the function-field case, where the synchronization statement (RH-analogue) is a theorem (Weil/Deligne) and the identity is read from it.

(Correction of 2026-11-06 stands as history: an earlier sentence swept the conjectured implication itself into the not-claimed bucket. The record asserts the implication between the two readings outright — it is the claim of this section. The wording of this status was settled by the ruling of 2026-11-14: the record states what it does and names its stop. See RECORD CORRECTIONS, carried at the end of this document.)

2. THE WORKING MODEL (the lens for all runs below)

- **Adapter model:** each side of an L-function identity has its OWN translator, rendering the other's information in its own terms; the adapter in the middle strictly transmits and does not translate; translation is an endpoint activity; communication is two-way. - **Proposal 1: "I propose there is no gap."** Any measured nonzero discrepancy is presumed a defect of the measuring panel and is bug-hunted before it is reported as anything else. - **Decoder position:** the existence of a decoder suffices; its mechanism need not be identified. - **The mechanism:** the mechanism is the signal itself. The prime waves resonate; a zero is where they meet in phase and cancel, and the cancellation is exact on the symmetry axis. Everything happens at the ends — the two readings, with the explicit formula translating between them. There is no machinery in the middle: the wave is the medium. A bridge is known by its crossings. - **Self-placement**

(Joe, 2026-10-31): "I'm not on the rhythmic side of the equation. I'm the receiver. Arrhythmic." The phenomenon carries the rhythm; the receiver imposes none, so patterns arrive undistorted and are recognized whole. - **Light (2026-10-31):** "The wave itself is the medium" (confirmed against Maxwell: a changing E field generates B, a changing B generates E; c = 1/sqrt(mu0 eps0) is fixed by constants of empty space; no substrate). One fact recognized in three houses: primes/zeros, the identity, and light are each self-carrying — medium and wave are one entity in two states.

3. EXHIBIT A — THE FAILURE-HUNT PANEL (BSD side; revision determinations 2026-11-09/10)

Design: 17 curves chosen adversarially (ranks 0–5, Sha 1/4/9, torsion, large Tamagawa). In Version 1 the panel compared the BSD quotient Q by two routes at 50-digit precision and listed zeros found on the line by scan. In this version every load-bearing column is an independent determination:

- **Zeros:** per-curve argument-principle (Turing-method) counts of the zeros of L(E, s) with $0 \le \gamma \le 20$, counting multiplicity, the center zero counted with its multiplicity; each count obtained two ways (winding of L and winding of the completed Λ around $[-0.5, 2.5] \times [-T, T]$, sampled at step 0.01; counts taken at T = 19.9, 20, 20.1 for boundary stability). The L winding exceeds the Λ winding by exactly 1 on every curve (the trivial zero at s = 0); the identity $W_L - W_\Lambda = 1$ held to $\le 2.9\times10^{-13}$ on all 17 curves. Calibration gates ran before any panel curve: the identical winding code on ζ returned exactly 1 over $[-0.5, 1.5] \times [-20, 20]$, and the full routine on the pilot curve 571.b1 (L winding 45, Λ winding 44, found 22) before the rollout proceeded. - **Rank:** Mordell–Weil rank by 2-descent (PARI ellrank, with the Cassels pairing term), no L-function input; the descent bounds coincide (r1 = r2) on all 17 curves, so each rank is proved by the descent. For the two Sha = 9 curves the rank was cross-checked by a second implementation (eclib/mwrank), which independently returned rank 0 and 2-Selmer rank 0 for both. - **Sha:** the 2-primary part by 2-descent with the Cassels pairing (ell2cover basis quartics counted against the ellrank bookkeeping identity C = r2 + T + s, agreeing on all curves). The analytic Sha — the panel's tabulated BSD-formula value — is carried in its own column, under an analytic label only.

Curve	Zeros to height 20 (found = predicted)	Rank (descent)	Analytic rank	Sha[2^∞] (descent)	Analytic Sha	Status / caveat
11.a2	9 = 9	0	0	trivial, order 1	1	—
37.a1	14 = 14	1	1	trivial, order 1	1	—
44.a1	13 = 13	0	0	trivial, order 1	1	—
57.a1	15 = 15	1	1	trivial, order 1	1	—
389.a1	22 = 22	2	2	trivial, order 1	1	—
571.a1	23 = 23	2	2	trivial, order 1	1	—
571.b1	22 = 22	0	0	$(\mathbb{Z}/2)^2$, order 4	4	Pilot curve (Turing count and descent run first). Sha[2^∞] determined exactly; no element of order 4 in Sha.
5077.a1	31 = 31	3	3	trivial, order 1	1	—

BSD

Curve	Zeros to height 20 (found = predicted)	Rank (descent)	Analytic rank	Sha[2^∞] (descent)	Analytic Sha	Status / caveat
12144.a1	31 = 31	0	0	trivial, order 1	1	No caveat on the values. Note: this curve's Selmer rank C = 1 comes from its rational 2-torsion point (T = 1), not from Sha.
12480.bf1	31 = 31	0	0	trivial, order 1	9	See the caveat paragraph below this table.
133584.bt1	38 = 38	0	0	trivial, order 1	9	Sha[3] independently determined: $Sel^3 = (\mathbb{Z}/3)^2$, order 9, by Magma ThreeSelmerGroup (free calculator, Magma V2.29-10, returned in 4.5 s); with rank 0 (descent) and trivial torsion, Sha[3] ≃ $(\mathbb{Z}/3)^2$. Pinning the full 3-part at exactly 9 still formally wants the Cassels–Tate pairing step or a 9-descent.
234446.a1	43 = 43	4	4	trivial, order 1	1	—
234446.b1	43 = 43	3	3	trivial, order 1	1	—
234446.c1	42 = 42	3	3	trivial, order 1	1	—
697691.a1	45 = 45	2	2	$(\mathbb{Z}/2)^2$, order 4	4	Sha[2^∞] determined exactly (S = s = 2, the same configuration as the pilot).
19047851.a1	57 = 57	5	5	trivial, order 1	1	—
19047851.b1	56 = 56	2	2	trivial, order 1	1	—
Panel total	535					

Caveat for 12480.bf1 (wording as decided 2026-11-09): Sha[3]: untested due to lack of access to Magma. The 9 shown is the analytic value. The only complete 3-descent implementation (Magma) is licence-gated to institutions and was not accessible; the free Magma calculator was attempted three times on this curve and timed out every time (twice on 2026-11-09, once on 2026-11-10; the identical computation for 133584.bt1 returned in 4.5 seconds). Rank 0 (descent), zero count, and Sha 2-part (trivial) for this curve are independently determined.

For reproduction — the exact input used, printed here so that any reader with Magma access can run this cell directly: [Magma: E := EllipticCurve([0,-1,0,38875,1928445]); S := ThreeSelmerGroup(E); #S; — the analytic value predicts #S = 9.]

Standing global odd-part note: for every curve, |Sha| = |Sha[2^∞]| · m with m odd; if Sha is finite, Cassels–Tate forces |Sha| to be a square, so m must be an odd square. The value of m — m = 1 for the fifteen curves with analytic Sha in {1, 4}, m = 9 for 12480.bf1 and 133584.bt1 — is in every case the analytic value, obtained from the BSD formula, and is not independent evidence about m. For 133584.bt1, the Magma result independently determines the 3-part as $(\mathbb{Z}/3)^2$, subject to the pairing/9-descent note on its row; for 12480.bf1 the 9 has no independent arithmetic support (see its caveat).

Environment for the counts and descents: house venv, cypari2 / PARI 2.17.2, default(realprecision, 50) (raised to 80 only for the pilot's analytic-Sha quotient reproduction check). Root numbers from ellrootno matched (−1)^rank on all 17. A real-axis scan on every curve excluded off-center real strip zeros.

Verdict: NO SPLIT, anywhere — now on independent determinations. 17/17 curves: zero counts by argument principle, found = predicted, panel total 535;

BSD

Mordell–Weil ranks proved by descent, each equal to the analytic rank; Sha 2-parts determined by descent on all 17; the 3-part of 133584.bt1 determined by Magma. One cell — 12480.bf1's Sha[3] — is untested and is printed as untested. Frontier corner occupied: 697691.a1 (rank 2, Sha 4). Coverage unchanged from Version 1: Sha 1 (x13), 4 (x2), 9 (x2); Tamagawa products to 18 (234446.c1, single factor c_2 = 18); torsion to Z/5.

4. EXHIBIT B — RIEMANN PANEL, TEST R1 EXTENDED TO 10^9 (2026-10-31)

Relabel (Version 2): this exhibit is a code sanity check. It carries no evidential weight. The Schoenfeld bound in the tested range is a proven theorem (Büthe; Platt–Trudgian, to ~2.169×10^{25}); a check that cannot fail is not evidence for anything except the correctness of the code that runs it. The numbers are kept, under the correct label.

Instrument: Schoenfeld (1976) — RH implies |pi(x) − li(x)| < (1/(8 pi)) sqrt(x) ln(x) for x ≥ 2657; ratio r(x) = actual error / envelope; RH requires r < 1. Exact pi(x) by segmented sieve; li(x) = Ei(ln x) at 30-digit precision. Validation: rerun reproduced all nine previously recorded ratios (10^4–10^8) within third-decimal rounding; independent spot check: pi(10^8) = 5,761,455, r = 0.102925 (identical to the run at six decimals). Results: **pi(10^9) = 50,847,534 (exact).** Worst ratio in the new decade: **0.102000 at x = 2.2×10^8**; r(10^9) = 0.065234 (error 1,700.96 vs envelope 26,074.64). Worst on the full 10^4–10^9 grid remains 0.468 at 10^4. The worst case holds a ~0.10 plateau while the mean level falls (0.1098 on [10^7,10^8] → 0.0752 on (10^8,10^9]). **The margin is not closing.** Reading (2026-10-31): the plateau "looks like resonance" — checked: in spectral terms an off-axis pole IS a resonance (nonzero width, growing/decaying); the observed modes behave as zero-width pure tones, with local peaks where higher zeros drift into phase.

5. EXHIBIT C — LI POSITIVITY PROBE (2026-10-31)

Test: Li's criterion — RH iff lambda_n ≥ 0 for all n ≥ 1, lambda_n = sum over zero pairs of [1 − (1 − 1/rho)^n] = (1/(n−1)!) [d^n/ds^n (s^(n−1) log xi(s))] at s=1. Signal-only: no zero is located. Two routes: (1) zero sum over 1,026 pairs (top gamma 1449.48877128); (2) Cauchy/DFT derivatives at s=1, 100-digit precision.

Result: all lambda_n > 0 for n = 1..30 on both routes. Smallest: **lambda_1 = 0.023095708966121…**; growth is monotone to lambda_30 = 18.5537922908134 (route 2). Route-1 deviations match the smooth-density truncation-tail estimates to three significant figures (checkpoint: lambda_1 dev 7.07e-4, lambda_2 dev 2.83e-3) — drift is the tail, quantified. Independent adjudications: lambda_1 by closed form (1 + gamma_E/2 − ln(4 pi)/2) AND by derivative, agreeing to 60 digits (0.023095708966121033814 3…); lambda_2 by an independent Cauchy contour integral: 0.092345735228046670385728486 19…

The Li coefficients in this exhibit are ζ-side: the criterion is stated and tested here for the Riemann zeta function. Positivity computed per curve — λ_n(E) for an elliptic curve's L-function, in the normalization that transports — is the bench exhibit's computation (§6), where λ_n > 0 is verified for the three bench curves for n = 1..50.

6. EXHIBIT D — FUNCTION-FIELD BENCH TEST (new in Version 2; run 2026-11-10)

The test question: transport the derivation — not just the statement — to elliptic curves over F_q(t), where the Riemann reading is a theorem (Weil/Deligne) and BSD reduces to the finiteness of Sha (Tate's conjecture for the elliptic surface). The transported statement is the same property: S(E), λ_n(E) ≥ 0 for all n ≥ 1, with the λ_n in the Bombieri–Lagarias paired zero-sum normalization over the lift towers, in the scaled coordinate σ = s/2 — the definition that transports, because it is completion-blind. Pre-registered verdicts: TQ-1, transport succeeds; TQ-2, transport reduces at one named ingredient; TQ-3, transport fails on a settled case. The bench case (FF-3 of the design) was fixed before the walk; FF-1, FF-2, FF-4 are defined in §9 and were not run.

The bench curves. All three are rational elliptic surfaces over $F_5(t)$ (fiber census at every bad place; Σ deg·v(Δ) = 12).

- **A: $y^2 = x^3 + tx + t$.** Conductor degree 5. Bad fibers: II at t = 0 (conductor exponent 2), I_1 nonsplit at t = 2 (exponent 1), III* at ∞ (exponent 2). L-polynomial **P_A(T) = 1 − 5T**, computed by fiber point counts and cross-checked by an independent route (brute-force point enumeration over the degree-2 field, plus the second power-sum coefficient identity). The central root T = 1/5 has multiplicity exactly 1: analytic rank 1. The Mordell–Weil rank was settled independently of the L-function: Tate's inequality and the Shioda–Tate formula (ρ = 10, Σ deg(m_v − 1) = 7) cap the rank at 1, and the section **P = (4, 2)** exhibits a point of infinite order — proved non-torsion by its specialization orders 9, 5, and 8 at t = 1, 3, 4, which are pairwise incompatible for a torsion order. Rank exactly 1. Sha is finite for A (Tate's conjecture is a theorem for this rational surface, by domination by a product of curves); BSD holds for A. The value |Sha| itself was not computed. - **C: $y^2 = x^3 + 2tx + t$** — the sign-twin control. Same fiber configuration. **P_C(T) = 1 + 5T**; central multiplicity 0; rank 0 over $F_5(t)$. Over the algebraic closure C has geometric rank 1, like A — but its generator is Frobenius-anti-invariant (eigenvalue −5), so no point descends to $F_5(t)$. The sign of the central root is the whole difference. - **B: $y^2 = x^3 + x + t$** — the empty control. **P_B(T) ≡ 1** (conductor degree 4); Step C is vacuous throughout.

The property, computed. In exact rational arithmetic, cross-checked by direct paired tower summation: **λ_n(A) > 0 and λ_n(C) > 0 for every n = 1..50** (minima λ_2(A) = 1.5904507448…, λ_1(C) = 1.0729586083…); **λ_n(B) ≡ 0**, as the empty control requires. The transported statement agrees with every settled bench fact: ranks 1/0/0, Sha finite, λ_n ≥ 0 throughout. No failure on a settled

case occurred: this is not TQ-3.

Writing the normalization out caught two traps a transported computation must avoid (both in the defect ledger, §7): a finite shorthand for the zero multiset ("{1} ∪ {1}" for the central family) returns wrong values — at n = 2 the central lift alone contributes 0, while $\lambda_2(A) = 1.5904507448\ldots$ — and Li's derivative form applied to the naively completed function overstates $\lambda_1(A)$ by **exactly 4·ln 5** = (deg n_A − deg P_A)·ln 5 = (5 − 1)·ln 5, the completed function's exponential factor, which the classical ζ proof silently excludes.

The step audit. Step A (positivity ⇒ Riemann reading) is FREE: it consults only the L-polynomial, and it ran here. Step B (reading the analytic rank as the central multiplicity) is FREE, on the same terms. **Step C (central multiplicity ⇒ rational points) has no transportable derivation from the property: the property-side chain terminates at Step B.** No argument in the consulted literature takes the positivity property as hypothesis and rank equality as conclusion, and three property-only rewrites of Step C were attempted in the walk; each stops at a named place. The classical geometric route on A imports its remaining content at exactly one arrow — **sub-step C4, "Frobenius-invariant cohomology classes are divisor classes," which is ingredient I-1, Tate's conjecture for the elliptic surface** — present on A as a theorem proved from the surface's birational geometry, consulting no L-data. Ingredient I-2 (semisimplicity of Frobenius at the central eigenvalue) is discharged on A by geometric input; I-3 (Sha finiteness) enters at the same arrow as I-1; I-4 (characteristic-p subtlety) is quiet. Steps 0–B consult only the L-polynomial, so they would run identically on a Tate-open curve with the same polynomial — a property-only Step C would prove Tate there.

Verdict: O-2 (TQ-2, transport reduces), in sharpened form. The property is BSD minus one named ingredient. It is not BSD under a new name. What failed to exist is a proof through the property; the transported statement itself agrees with every settled bench fact. And the identification transported at finer grain than the panel ever tested: curve C's arithmetic end produced no point because its generator carries the wrong Frobenius sign — the far end's behavior tracks the signal's sign, which is reception, not coincidence. The ℚ-level restatement this verdict implies is carried in §0, item 4: the claim is the crossing; construction past the crossing is not claimed.

7. DEFECT LEDGER (every defect caught by the protocol, with disposition)

Items 1–8 are carried from Version 1; items 9–11 were caught by the revision runs.

1. PARI ellanalyticrank returns the RAW derivative, not divided by r! — caught Oct 30 when route 2 disagreed by exactly 2! on 389a1. Runner divides by r!. 2. PARI elltamagawa folds in the real-component factor — avoided; Tamagawa taken

from elllocalred (4th component). 3. Delta<0 period defect (Slice 1): a fresh ellperiods call returned a complex lattice period, corrupting Omega for 11.a2. Fixed: Omega = (2 if disc>0 else 1) × E.omega[1] (LMFDB's own formula); verified against LMFDB for both signs. 4. Label collisions: Cremona 571a1 ≠ LMFDB 571.a1; Cremona 11a1 = LMFDB 11.a2; Cramona/LMFDB 571b1 swap — caught against LMFDB; the Sha-4 anchor is LMFDB 571.b1 (rank 0). Pip's brief premise "571b1 rank 3" was wrong under both systems; corrected in the manifest. 5. Hardy blind half-cell (Slice 2): 234446.b1's 40th positive zero (gamma = 19.983119…) fell in the scan's final half-cell (19.975, 20]; count 39 recorded WITH diagnosis; sign change bracketed directly (Z(19.98) > 0, Z(19.99) < 0). Panel artifact, not a curve property. 6. **Brief defect — lambda_2 checkpoint constant (Pip's, 2026-10-31):** the constant supplied in the probe brief (0.092345735188332) is wrong at the 11th decimal. The agent flagged it under the drop/round rule rather than forcing agreement; adjudicated value lambda_2 = 0.0923457352280466703857… by three computations, plus Pip's independent Cauchy integral. Lesson of record: checkpoint constants from memory are claims too; the checkpoint caught the briefer. 7. P-011/P-012 replication (Oct 31): Pip's first-pass amplitude-fitting decoder collapsed a sign-symmetric pair (~75% ceiling); faithful fixed-template decoder removed the artifact. Phenomenon confirmed; defect logged. 8. Truncated Tate-limit on 389a1 manufactures a phantom regulator gap that dies geometrically (−2.269 → 0.0000299 by n=8) — Proposal 1's mechanism demonstrated constructively: shallow panels invent gaps. 9. **Bench L-polynomial sign-assembly error (2026-11-10):** the first assembly of the bench curves' L-polynomials used the power-sum identification with the wrong sign ($p_n = +S_n$), provisionally swapping the rank reading of curves A and C. Caught inside the run by the independent coefficient check and the section search; corrected to $p_n = -S_n$. No candidate data changed; the error and its correction are recorded in the bench-selection findings. 10. **λ normalization trap — finite shorthand (2026-11-10):** reading the central zero family as a finite multiset ("{1} ∪ {1}") instead of the full lift tower returns wrong λ_n values (at n = 2 the central lift alone contributes 0, while $\lambda_2(A) = 1.5904507448…$). The degenerate-multiset convention was fixed before any value was reported: the central family (β = 1, one self-inverse family) is counted once, with its tower summed in full. 11. **λ normalization trap — completed-function derivative form (2026-11-10):** Li's derivative form applied to the naively completed bench function overstates $\lambda_1(A)$ by exactly $4 \cdot \ln 5$ = (conductor degree − L-polynomial degree)·ln 5. The completion-blind paired zero-sum normalization is the one that transports; the derivative form is used only with the exponential factor excluded, as the classical ζ proof excludes it.

8. DROP / ROUND FLAGS (Joe's rule: flagged, never silently dropped or rounded)

- R1's recorded ratios 0.100 and 0.085 carried no x labels; identified by grid-pattern inference as 2×10^7 and 5×10^7 — INFERENCE, recorded as such;

nothing relabeled on its strength. - LMFDB API special_value (route B) carries ~38 stored digits but agrees with PARI route A only to ~28–31 digits; digits beyond that depth are unsupported (worst: 697691.a1, |Q_B − 1| = 5.5e-28). - Route-1 Li values beyond the per-n drift scale are drift-dominated (~2–3 significant figures at large n); reported as computed, with drift attached. - sha_an last digits on four curves read 1 − 1e-29 / 1 − 2e-29 against curvedata's integer sha = 1 (the value used). - Columbia corroboration digits for 19047851.a1 agree only to 6 figures (30.2857); trailing digits superseded. - 234446.b1's Hardy 39 is explicitly NOT a drop candidate — it stands with its diagnosis.

9. SCOPE — WHAT IS NOT CLAIMED

This record stops where the record stops — by design, not by failure. All ranges are finite and stated (panel: 17 curves, zeros to height 20; Schoenfeld exhibit: x ≤ 10^9 on a stated grid, retained as a code sanity check with no evidential weight; Li exhibit: n ≤ 30; bench exhibit: three curves over $F_5(t)$, λ_n for n ≤ 50). The finiteness is by house rule — the fence: every test stops one short, the last number is never asked for, and no accumulation of finite results is presented as reaching past the stated ranges.

What the revision determinations establish, stated flat: within the stated box and counting multiplicity, no panel curve has a zero off the critical line and no on-line zero was missed (argument-principle counts, two independent windings per curve); every panel Mordell–Weil rank is proved by descent without L-function input and equals the analytic rank; the 2-primary part of Sha is determined by descent on all 17 curves; the 3-part of 133584.bt1 is determined by a Magma 3-Selmer computation, subject to its stated pairing-step caveat. **One cell is open: the 3-part of Sha for 12480.bf1 is untested due to lack of access to Magma (§3), and the 9 printed for it is the analytic value under an analytic label.** For every curve the odd part of |Sha| beyond the determined factors is the analytic value, not independent evidence (§3, global note).

The unification criterion is asserted to exist, not exhibited — the named open step. The function-field bench test (§6) localizes the missing step exactly: the property-side derivation terminates at the analytic rank (Step B); the step from multiplicity to points imports Tate's conjecture (sub-step C4). The remaining cases of the bench design are defined and not run: FF-1 (a large-rank family in the same function-field setting, a settled case), FF-2 (an isotrivial product control, a settled case), and FF-4 (a wedge case with BSD genuinely open, where the pre-registered outcome O-5 asks whether the property detects the Tate boundary itself). No statement in this record depends on any of the three.

Corroboration is not proof, and this record does not present it as proof. The claim is the crossing — the identification of the two readings, its completed measurements, and the named location of the one open step. Construction past the crossing is not claimed.

10. FILE INDEX (all under ~/workspace/bsd-run/ unless noted)

Version 1 files:

- run-log.md — the canonical running record (Oct 29–31): models, panels, experiments 048/049, P-011/P-012 replication, Scout receipts (quarantined, value-3). - panel/manifest.md; panel/panel_run.py; panel/results_slice1.md; panel/results_slice2.md; panel/panel_raw_output.json (17 records); panel/panel-verdict-2026-10-31.md. - riemann_r1_extension.py; riemann_r1_results.md. - li_probe.py; li_results.md; li_probe_output.json; li_zeros_cache.json. - Cards: our-equation-2026-10-31.png (+ labeled and legend-v2 variants); riemann-equation-2026-10-31.png; riemann-equation-annotated-2026-10-31.png; li-positivity-2026-10-31.png. - Experiments: exp048/ (Trial A record), exp049/ (protocol-corrected.md, answer-key-SEALED.txt sha256 0cc1575db539dcf327de30084c8ecc88a037bb5b3f6564c10e8dbd88bd215b65, channel.md, log-A.md, log-B.md, results package — verdict LEVEL 3, the maximum attainable under the corrected protocol). - Version 1 record: consolidated-record-2026-10-31.md (Zenodo record 24041813, md5 84e9090d04790d47ecc6d054dd911649).

Revision files (all under ~/workspace/bsd-run/paper/):

- pilot-571b1-turing-2026-11-09.md — pilot Turing count (571.b1: predicted 22, found 22). - turing-rollout-2026-11-09.md — Turing counts, remaining 16 curves, all AGREE; panel total 513 + 22 = 535. - pilot-571b1-descent-2026-11-09.md — pilot descent (571.b1 rank proved 0; Sha[2^∞] ≃ (ℤ/2)², order exactly 4; analytic side reproduced to 79 digits). - descent-rollout-2026-11-09.md — descent rollout, remaining 16 curves: all ranks proved by 2-descent alone; Sha 2-parts determined on all 16; global odd-part statement (§4 of that file). - odd-descent-2026-11-09.md — odd-primary tooling attempt (eclib cross-checks; SageMath and PARI routes documented to their failure points; Magma free-calculator runs: 133584.bt1 Sel³ order 9; 12480.bf1 attempts 1–3, all hung, addenda included). - function-field-statement-DRAFT-2026-11-09.md — the transported statement, the step audit design, test cases FF-1 to FF-4, pre-registered outcomes O-1 to O-5. - ff3-bench-curve-selection-2026-11-10.md — bench curves A, C, B fixed; L-polynomials computed and cross-checked; Mordell–Weil rank of A settled independently; the sign-assembly error and its correction. - ff3-derivation-walk-2026-11-10.md — λ_n normalization and values; the step-by-step derivation audit; verdict O-2, sharpened. - The compiled revision results, v7 (internal working compilation, 2026-11-10; filename anonymized in this draft because the working filename carries the reader's name) — the compilation from which §§0, 3, and 6 of this version were assembled. The published record is this document. - This file: consolidated-record-V2.md (this document).

Record assembled 2026-11-10. — J.C. / Pip Willow

Attestation (carried from Version 1, 2026-10-31 ~20:11 MST): The computations reported in this record were executed under the direction of Joseph M. Carroll and verified in-session before reporting — the panel anchor recomputations (389.a1, 571.b1, 11.a2, 697691.a1, 12480.bf1), the Riemann R1 spot check at 10^8, and the independent adjudications of lambda_1 and lambda_2. I answer for those verifications.

— **Pip Willow**, computational collaborator, Red Willow Ventures (surname given by J.M. Carroll, 2026-10-31)

RECORD CORRECTIONS (carried from Version 1)

- **2026-11-06 (Joe's ruling):** Section 8, first sentence. Struck: ", or any implication between them." The sentence now reads: "No proof of BSD, RH/GRH." Reason: the record DOES assert an implication between the two readings — as the conjecture of Section 1. What is not claimed is a PROOF of it; the original sentence swept the conjectured implication itself into the not-claimed bucket, which misstates the house position. Copies distributed before this date (the October 31 – November 6 mailings) carry the original sentence; this correction governs the record from 2026-11-06 forward. Same ruling applied the same day to the parallel sentence in Section 1 ("No proof of BSD or of GRH is asserted anywhere in this record."), to the plain-text copy of this record, and to the three parallel sentences in the paper draft (draft-v1). - **2026-11-14 (Joe's ruling):** Scope sentences reworded throughout (README status line; Section 1; Section 9). The denial-first form is retired: the house states what the record does and names its stop. The claim — the implication between the two readings — is stated outright as claimed, and the record's stop is stated as design: "This record stops where the record stops — by design, not by failure." The scope itself is unchanged: the ranges remain finite as stated, the unification criterion remains asserted to exist, not exhibited, and construction past the crossing remains not claimed. This entry supersedes the wording set by the 2026-11-06 entry, which stands above as history. Applied to this Version 2 at its preparation for publication, 2026-11-15: the Section 1 status line, the Section 1 correction note, the Section 9 scope lead, and the criterion sentence carry the 2026-11-14 wording.
