

Conditional BSD Implications with Twelve Explicit Debt Assumptions

A machine-checked dependency graph with explicit trust accounting, a sorry-free Lean 4 export, and worked curve cases

181 kernel implications and 187 exported hypotheses; twelve conjectural or bridge assumptions are explicit; the sorry-free Lean 4 artifact exposes rather than discharges the cited-fact surface

Dr. Tamás Nagy

tnagyphd@gmail.com

Version 1.0 pre-publication candidate • April 2026

Build-std v2.0+fa55b498d9 | 2026-07-22 16:10 | e814125a

Abstract

The Birch and Swinnerton-Dyer conjecture predicts that for an elliptic curve $E/\mathbb{Q}$ the Mordell-Weil rank equals the order of vanishing of the Hasse-Weil L -function at $s = 1$. It also predicts that the leading Taylor coefficient is expressed by an explicit formula involving the real period, regulator, Shafarevich-Tate group, Tamagawa numbers, and torsion. This remains an open Millennium Prize problem.

This paper presents a machine-checked dependency graph for selected implications around BSD. The release-scoped proof environment contains 181 kernel implications and 187 hypotheses. Twelve `A_*` declarations expose conjectural statements or abstraction bridges. A dependency-closed Lean 4 export compiles with zero sorry, but cited mathematical results and curve data remain Lean axioms. Compilation verifies the encoded implications, not the imported number theory.

The principal contribution is formal packaging: a common typed interface, explicit dependency management, worked curve instances, reproducible exports, and a trust budget. No new number-theory theorem is claimed. The development does not prove BSD, does not resolve a Clay problem, and does not establish a generic unconditional result in rank at least two.

The trust boundary is intentionally larger than in the previous draft. Besides the global BSD rank, formula, and Sha-finiteness assumptions, separate bridges cover low-rank Sha finiteness, GRH/CM Sha packages, rank-zero and rank-one formulas, a universal rank upper bound, Tian’s family statement, a Selmer/analytic-rank equality, and a CM rank-one formula. Published p -adic and Iwasawa results have object types and hypotheses absent from this interface. The previous ordered-real “ p -adic” phases have therefore been removed rather than marketed as formalizations.

The load-bearing capstones are conditional compositions: the proposition `ShaPFinite E` plus `A_selmer_eq_an_rank` bridge implies the rank identity; `A_bsd_formula_rank0` bridge implies the rank-zero leading-term formula; and `A_bsd_formula_rank1` bridge implies the rank-one leading-term formula. Per-curve conclusions are only as strong as their exported database, descent, and package assumptions.

Phase 30 is only a definitional renaming of analytic rank and the leading coefficient; it has no independent spectral theorem or number-theoretic content. It is retained as an interface sketch and excluded from the novelty claim.

The worked examples 32a2 and 37a1 illustrate the dependency graph. Their numerical and database inputs are hypotheses. The 37a1 leading-term capstone is explicitly conditional on `A_bsd_formula_rank1_bridge`; the CM rank-one package is likewise disclosed as an exported cited assumption rather than called unconditional.

The scope is explicitly **not** a proof of BSD. The twelve explicit debt assumptions are listed in §1.3. In particular, `A_sha_finite` asserts `ShaPFinite E` for every curve `E`; it has proposition-valued semantics and is not a positivity claim about a real number.

The Phase 31 rank reduction is a formal implication from `F_selmer_eq_mw_under_sha_finite` and `A_selmer_eq_an_rank_bridge`; it is not a first-principles derivation of Iwasawa theory. `A_bsd_rank` remains declared for backward compatibility.

These assumptions are narrowed but not eliminated. In particular, `bsd_formula_rank1_via_bridge`, `bsd_formula_rank1_under_an_rank1_bridge`, and `E1_bsd_formula_under_rank1_bridge` depend on `A_bsd_formula_rank1_bridge`. No dependent statement is described here as unconditional. Higher-rank curve identities are formal consequences of the encoded Kato and `mwrnk` hypotheses, not generic unconditional rank results.

The exact live kernel inventory is 368 verified declarations: 181 theorem declarations and 187 hypotheses. The reduction from the earlier 503-declaration environment is intentional: unsupported phases and redundant theorem-strength packages were removed. Lean marker and generated-dependency counts are reported from the current artifact in §6. The former 1,254-axiom figure came from exporting repository-wide bootstrap plumbing, not from independent BSD assumptions.

The contribution is synthesis, dependency traceability, export, and reproducibility rather than mathematical novelty.

Prior formal work includes the Lean 3 `jamiebell2805/BSD-conjecture` project, which developed definitions needed to state BSD; the Google DeepMind formal-conjectures Lean 4 statement project; ongoing Mathlib elliptic-curve and L -function infrastructure; and curve-specific Lean projects that combine proved arithmetic with named analytic or database surfaces. A targeted search did not establish a comprehensive Lean/Isabelle/Coq priority record. No priority claim is made, and every formal conclusion is limited by the assumptions listed above. The narrower contribution is this particular typed dependency graph, explicit trust classification, and exact-hash replay package.

1. Introduction

1.1 The Birch and Swinnerton-Dyer conjecture

Let E be an elliptic curve over $\mathbb{Q}$ with minimal Weierstrass equation $y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$. The Mordell-Weil group $E(\mathbb{Q})$ is a finitely generated abelian group; its free rank we denote $r = \text{rank}_{\mathbb{Z}} E(\mathbb{Q})$. The Hasse-Weil L -function $L(E, s) = \sum_{n \geq 1} a_n(E) n^{-s}$ converges for $\text{Re}(s) > 3/2$ and, by the modularity theorem of Wiles, Taylor, Breuil, Conrad, and Diamond, extends to an entire function on $\mathbb{C}$ satisfying a functional equation under $s \mapsto 2 - s$ with root number $w(E) \in \{\pm 1\}$.

The Birch and Swinnerton-Dyer conjecture asserts two things.

The rank part. The order of vanishing of $L(E, s)$ at $s = 1$ equals the Mordell-Weil rank:

$$\text{ord}_{s=1} L(E, s) = r.$$

The leading-term part. The leading Taylor coefficient is given by

$$\frac{1}{r!} L^{(r)}(E, 1) = \frac{\Omega_E \cdot \text{Reg}(E) \cdot \#\text{III}(E) \cdot \prod_p c_p(E)}{(\#E(\mathbb{Q})_{\text{tors}})^2},$$

where Ω_E is the real period, $\text{Reg}(E)$ is the regulator of the Néron-Tate height pairing on $E(\mathbb{Q})/E(\mathbb{Q})_{\text{tors}}$, $\#\text{III}(E)$ is the (conjecturally finite) order of the Shafarevich-Tate group, $c_p(E)$ are the local Tamagawa numbers, and $\#E(\mathbb{Q})_{\text{tors}}$ is the order of the torsion subgroup.

The conjecture is one of the seven Clay Millennium Prize Problems. Over the decades, it has been proved in several partial senses.

Coates and Wiles (1977) proved that if $E/\mathbb{Q}$ has complex multiplication by the ring of integers of an imaginary quadratic field and $L(E, 1) \neq 0$, then $\text{rank}(E) = 0$. This was the first partial BSD theorem. It is unconditional and proved by an elliptic-units argument independent of Euler systems.

Tate (1974) predicted the finiteness of the Shafarevich-Tate group $\text{III}(E/\mathbb{Q})$ as a general phenomenon. This finiteness is a non-trivial input to BSD. It is implicit in the leading-term formula, which involves $\#\text{III}$. It also serves as the bridge between the Selmer group and the Mordell-Weil group via Cassels-Mazur descent. Cassels (1965) and Mazur-Rubin (2004) showed the Selmer rank equals $\text{rank}(E) + \text{corank}_{\mathbb{Z}_p} \text{III}(E)[p^\infty]$. Therefore, if $\text{III}(E)[p^\infty]$ is finite, the Selmer rank equals the Mordell-Weil rank.

Gross-Zagier and Kolyvagin establish the rank identity and Sha finiteness in analytic ranks zero and one under their precise modularity and nonvanishing hypotheses. Rubin proved the main conjecture for imaginary quadratic fields; consequences for a given CM elliptic curve still require the relevant prime and local hypotheses.

Bhargava and Shankar (2013) bounded the average rank of $E/\mathbb{Q}$ above by $7/6$. Building on this, Bhargava, Skinner, and Zhang (2014) proved BSD holds for a positive proportion of elliptic curves.

Kato and Skinner–Urban prove major Iwasawa divisibilities under precise arithmetic hypotheses. The universal Selmer/analytic-rank equality used below is an imported interface and must not be identified with those theorems without a hypothesis translation.

Perrin–Riou’s 1993 arithmetic p -adic L -function paper, together with Nekovář (1995) and Kobayashi (2013), concerns p -adic L -functions and height formulas under substantial local and representation-theoretic hypotheses. It does not by itself prove the universal archimedean rank-one equality encoded here. The formalization therefore exposes the missing passage as `A_bsd_formula_rank1_bridge`.

The legacy Phase 24 interface attributes effective GRH bounds to Goldfeld (1976) and Hoffstein–Lieman (1994). Those citations do not support the encoded statements as written. Phase 24 is therefore retained only as an axiomatic historical experiment, not as a literature-verified theorem package.

1.2 Scope of this work

We formalize selected BSD-adjacent implications in a formal proof kernel and export them to Lean 4. The release environment contains 181 distinct theorem declarations and 187 hypotheses. Its scope is dependency management and reproducibility, not a proof of the cited inputs.

We do **not** prove BSD. All twelve A_* assumptions are explicit in §1.3. The remaining exported hypotheses are bootstrap declarations, cited-result interfaces, or worked-example data; each remains part of the trust surface.

This paper provides several concrete contributions.

A single type-theoretic environment. We instantiate selected classical BSD-adjacent implications in a common type system. Crucially, p-primary Sha finiteness is represented by the proposition `ShaPFinite E`; no cardinality or positive-real proxy is used. Classical scalar invariants remain opaque real-valued interfaces whose arithmetic interpretation is axiomatic.

Eight concrete worked curves with multiple conditional routes. We instantiate the rank-0 curve $E_0 : y^2 = x^3 - x$ (Cremona 32a2), the rank-1 curve $E_1 : y^2 + y = x^3 - x$ (Cremona 37a1), and six further LMFDB curves. Every numerical, database, descent, and deep-theorem input remains a named hypothesis. The resulting curve statements are machine-checked specializations of those premises, not independent unconditional proofs. In particular, the higher-rank 389a1 and 5077a1 identities consume the encoded Kato upper-bound and mwrank lower-bound interfaces.

Explicit conditional narrowings. Kernel-checked compositions remain relative to their premise interfaces. Rank-zero and rank-one formula conclusions require separate explicit bridges; higher-rank examples require `A_rank_upper_bound_bridge` and mwrank premises. No generic rank-at-least-two theorem or unconditional full-BSD result is claimed.

A universal treatment of the congruent-number family. Phase 29 instantiates the family $E_n : y^2 = x^3 - n^2x$ as a single polymorphic curve $E_{cn} : \mathbb{R} \rightarrow EC$. It records CM, Tian–Tunnell, Kato, mwrank, and formula interfaces as explicit premises. The $n = 34$ and rank-one formula conclusions are conditional on those interfaces; Phase 32 does not make the CM rank-one formula unconditional.

A prototype latent-spectral reformulation. Phase 30 introduces two per-curve spectral invariants: spectral multiplicity and atom mass. Their definitional identification with the analytic rank and BSD leading coefficient provides a reformulation-layer restatement of the two BSD axioms. This remains a sketch-level object and is labelled as such throughout the paper.

A structural rank reduction under explicit interfaces. Phase 31 installs `bsd_rank_from_sha_finite` by composing `F_selmer_eq_mw_under_sha_finite` and `A_selmer_eq_an_rank_bridge`. The universal form makes `A_bsd_rank` derivable from `A_sha_finite` inside this encoded interface. This is dependency compression, not a literature theorem that Sha finiteness alone proves generic BSD rank.

Explicit formula bridges. Phase 26 installs `bsd_formula_rank0_under_bridge`; Phase 32 installs `bsd_formula_rank1_via_bridge`. Their premises are not presented as formalized Iwasawa or p-adic theorems. The CM rank-one theorem consumes the separately named `A_cm_rank1_formula_package`.

Auditable Lean 4 output. A sorry-free export of the 368-declaration environment compiles under Lean 4 with Mathlib. Its explicit axiom inventory makes the imported trust surface inspectable.

1.3 Trust budget and dependency table

Twelve assumptions encode conjectural statements or unformalized bridges and are **not** proved:

```

A_sha_finite          E : EllipticCurve ,   ShaPFinite(E)
A_bsd_rank            E : EllipticCurve ,   rank(E) = analytic_rank(E)
A_bsd_formula         E : EllipticCurve ,   L_lead(E) =  $\Omega \cdot \text{Reg} \cdot \# \text{III} \cdot c_p / (\#E_{\text{tors}})$ 
A_sha_finite_rank_le1_bridge
A_grh_sha_finite_package
A_cm_sha_finite_package
A_bsd_formula_rank0_bridge
A_rank_upper_bound_bridge
A_tian_rank_lower_bound_package
A_selmer_eq_an_rank_bridge
A_bsd_formula_rank1_bridge
A_cm_rank1_formula_package

```

The first is the universal Sha-finiteness conjecture represented by an honest proposition. The next two are the BSD statements. The remaining nine are abstraction bridges or assumed packages whose exact literature hypotheses are not represented. Every downstream theorem that invokes one is conditional on it.

Trust item or class	Status in Lean export	Major dependents	What compilation establishes
A_sha_finite	conjectural axiom	bsd_rank_universal_from_sha_finite	only the implication from finiteness
A_bsd_rank	conjectural axiom, retained for compatibility	earlier BSD-rank restatements	consequences of the rank identity
A_bsd_formula	conjectural axiom	generic leading-term restatements	consequences of the formula
A_sha_finite_rank_le1_bridge	explicit bridge	low-rank, GRH, and CM Sha routes	only the stated finiteness implications
A_grh_sha_finite_package	explicit package		
A_cm_sha_finite_package	explicit package		
A_bsd_formula_rank0_bridge	explicit bridge axiom	bsd_formula_rank0_understanding, 32a2 and CM rank-zero formulas	formalizing, assuming the missing literature translation
A_rank_upper_bound_bridge	explicit bridge axiom	higher-rank curve pinches	equality after assuming the open upper bound
A_tian_rank_lower_bound_package	abstract package axiom	congruent-number rank lower bound	only the opaque package implication
A_selmer_eq_an_rank_bridge	explicit bridge axiom	universal rank reduction	rank transitivity after assuming Selmer = analytic rank
A_bsd_formula_rank1_bridge	conjectural bridge axiom	bsd_formula_rank1_via_rank0, bsd_formula_rank1_understanding, 37a1 formula capstone	formalizing formula after assuming the missing literature package

Trust item or class	Status in Lean export	Major dependents	What compilation establishes
A_cm_rank1_formula_package	package axiom	CM rank-one formula	only the package conclusion
Gross–Zagier/Kolyagin/Iwasawa/Kato/CM packages	cited hypotheses, exported as axioms	low-rank and CM routes	compositions of the encoded interfaces
LMFDB/mwrank/per-curve values	data hypotheses, exported as axioms	eight worked-curve statements	specialization to assumed database values
generated type/operation support (36 axioms)	generated axioms in the Lean export	all declarations	consistency relative to that foundation

Central formal result (conditional). The kernel checks three compositions: `A_sha_finite` plus `F_selmer_eq_mw_under_sha_finite` and `A_selmer_eq_an_rank_bridge` implies `A_bsd_rank`; `A_bsd_formula_rank0_bridge` implies the rank-zero leading-term formula; and `A_bsd_formula_rank1_bridge` implies the rank-one leading-term formula. These are implications relative to twelve explicit debt assumptions and the wider 187-hypothesis surface, not unconditional BSD theorems.

The original rank and Sha assumptions are structurally related by a Phase 31 kernel theorem:

$$\text{‘A_sha_finite’} \xrightarrow{\text{Phase 31}} \text{‘A_bsd_rank’}.$$

Concretely, `bsd_rank_universal_from_sha_finite` derives `A_bsd_rank` from `A_sha_finite` via `F_selmer_eq_mw_under_sha_finite` and the explicit abstraction bridge `A_selmer_eq_an_rank_bridge`. The formal result is therefore not a consequence of Sha finiteness alone.

The formalization supplies the following explicit narrowings as kernel theorems.

First, `bsd_rank_from_sha_finite` and its universal form `bsd_rank_universal_from_sha_finite` (Phase 31) show that if $\text{III}(E)[p^\infty]$ is finite, then $\text{rank}(E) = \text{an_rank}(E)$. The proof composes Cassels–Mazur descent with the Iwasawa-theoretic Selmer–rank identification. This universal form makes `A_bsd_rank` a direct consequence of `A_sha_finite`.

Second, `an_rank_le1_rank_eq_an_rank_grh_free` (Phase 31) is a GRH-free formal implication conditional on imported Sha and Selmer interfaces. The earlier `grh_bsd_rank_le1` is retained as a historical experiment, but its Goldfeld/Hoffstein–Lieman attributions do not justify the encoded effective statements.

Third, the CM narrowings compose simplified Coates–Wiles and Rubin interfaces. The kernel implications are conditional on those imported statements; their local and prime hypotheses are not reconstructed.

Fourth, `bsd_formula_rank0_under_bridge` is conditional on `A_bsd_formula_rank0_bridge`. At rank one, `bsd_formula_rank1_via_bridge` is conditional on `A_bsd_formula_rank1_bridge`. The separate CM statement consumes `A_cm_rank1_formula_package`.

Fifth, `bsd_formula_rank1_under_an_rank1_bridge` composes the rank route with the explicit bridge. “GRH-free” does not mean unconditional: all cited and conjectural premises remain.

Finally, `rank_upper_bound_under_bridge` and `kato_rank_lower_matches_an_rank` compose `A_rank_upper_bound_bridge` with lower-bound and analytic-rank data. For 389a1 and 5077a1, the resulting identities remain conditional on this bridge and the `mwrnk` interfaces.

Taken together, the residual open content includes:

Open content	Axiom	Residual scope
Rank identity, universal form	<code>A_bsd_rank (A_sha_finite)</code>	$\text{an_rank} \geq 2$ <i>without</i> a proof of p -primary Sha finiteness (higher-rank closures via <code>mwrnk</code> are conditional on the open bound $\text{rank} \leq \text{an_rank}$); at $\text{an_rank} \leq 1$ the identity is GRH-free but conditional on Iwasawa $\mu = 0$ and Skinner-Urban
Finiteness of $\text{III}[p^\infty]$	<code>A_sha_finite</code>	Known in important low-rank and CM settings under precise hypotheses; universal statement remains open
Leading-term formula	<code>A_bsd_formula,</code> <code>A_bsd_formula_rank1_bridge</code>	generic rank one remains conditional on the bridge; rank at least two is outside this release scope

The rank-one bridge is intentionally stronger than any single literature theorem encoded here; this is why it is an `A_*` assumption.

1.4 Organization

Section 2 describes the Lean 4 formalization framework and the elliptic-curve API provided by the kernel. It details the naming convention that separates facts from hypotheses from conjectures.

Section 3 covers the classical partial results (Phases 1-10). These include Mordell-Weil, torsion and period positivity, Gross-Zagier, Kolyvagin, BSD for $\text{rank} \leq 1$, and parity. Section 4 closes the two worked curves E_0 and E_1 .

Section 5 covers retained modern interfaces (Phases 18-32), database examples, and explicit rank/-formula bridges. Historical phases 13-17 were removed because their object types were not faithful.

Section 6 describes the reproducibility path and the sorry-free Lean 4 export. Section 7 discusses what is and is not mechanized.

2. Formalization framework

2.1 The Lean 4 kernel and the elliptic-curve bootstrap

The formalization lives in a dependently typed Lean 4 environment. Every object has an explicit type, every axiom is named, and every theorem is machine-checked. The development bootstraps a fixed API for elliptic curves over $\mathbb{Q}$ and their invariants; the following constants are provided by the kernel and never redefined:

```

EllipticCurve      : Type
EC.rank            : EllipticCurve →
EC.analytic_rank   : EllipticCurve →
EC.L_value         : EllipticCurve → →
EC.L_deriv         : EllipticCurve → →
EC.real_period     : EllipticCurve →  $\Omega$       ()
EC.regulator       : EllipticCurve →          (Reg)
EC.sha_order       : EllipticCurve →          (# III)
EC.tamagawa_product : EllipticCurve →          ( c_p)
EC.torsion_order   : EllipticCurve →          (#E ()_tors)
EC.root_number     : EllipticCurve →          (w(E)  {±1})
EC.height_pairing  : EllipticCurve → → →
EC.heegner_height  : EllipticCurve →
EC.L_leading       : EllipticCurve →          ( $L^{\wedge(r)}(E, 1) / r!$ )
ShaPFinite         : EllipticCurve → Prop

```

together with named structural axioms. The height axiom has the exact diagonal signature $\forall E, x, 0 \leq \langle x, x \rangle_E$; it does not assert false pairwise nonnegativity. Regulator positivity is a separate imported interface, and the release does not claim to formalize bilinearity, symmetry, or positive definiteness modulo torsion.

Additional retained types include ModularForm, GalRep2d (two-dimensional Galois representations for Serre’s conjecture), and EulerSystem.

2.2 Declaration taxonomy

Every declaration carries a prefix indicating its role:

Prefix	Meaning	Count
hypotheses	cited interfaces, curve data, bootstrap declarations, and conjectural assumptions	187
A_* subset	conjectural or bridge assumptions	12
theorem declarations	distinct kernel-verified implications	181

The A_* subset includes every opaque package whose exact literature hypotheses are absent. The F_* prefix is reserved for narrower cited or structural interfaces; it does not by itself certify a source transcription.

The 33 H_* hypotheses describe the eight worked curves and the congruent-number family (e.g., $H_E0_rank = 0$, $H_E1_root = -1$, $H_E1_heegner_pos$, $H_E_389a1_an_rank$, $H_E_389a1_mwrnk_geq_2$, $H_cn_34_mwrnk_geq_2$).

The exact live total is 187 hypotheses plus 181 theorems, hence 368 verified declarations.

This taxonomy makes the trust surface inspectable, but the complete Lean axiom list remains the authoritative export-level inventory.

Literature-boundary warning. Not every F_* interface is a fully formalized transcription of a published theorem with every side condition.

- **A_bsd_formula_rank1_bridge:** This is explicitly conjectural. Published rank-one results, including Jetchev–Skinner–Wan, impose hypotheses and often prove prime-part statements rather than the universal archimedean equality represented by this bridge.
- $F_bsz_positive_proportion$ (Bhargava–Skinner–Zhang 2014, §5.11) states that BSD holds for a positive proportion of elliptic curves; this is a genuine published theorem, but the explicit quantitative bound carried in code comments is the BSZ constant ($\geq 66.48\%$ in current formulations), and any future quantitative refinement would be a tighter F_* .

There is no optional interpretation: each formula or rank package without represented source hypotheses is explicit A_* debt.

2.3 Two worked curves

The curves E_0 and E_1 serve as concrete test cases against which the formal machinery is run end-to-end.

$$E_0 : y^2 = x^3 - x \quad (\text{Cremona 32a2, rank 0, } w = +1, \#E_0(\mathbb{Q})_{\text{tors}} = 4, \prod c_p = 4, \#\text{III} = 1).$$

$$E_1 : y^2 + y = x^3 - x \quad (\text{Cremona 37a1, rank 1, } w = -1, \#E_1(\mathbb{Q})_{\text{tors}} = 1, \prod c_p = 1, \#\text{III} = 1).$$

For each curve, the numerical values of the L -function at $s = 1$ (for E_0) or of $L'(E, 1)$ (for E_1), along with the real period, regulator, Heegner point height (for E_1), and torsion, are stated as hypotheses. The formalization then derives by machine verification that each satisfies the BSD rank equation ($E_i_bsd_consistent$) and the BSD numerator identity ($E_i_bsd_verified$).

The corresponding database records are 32.a2 and 37.a1 (accessed 2026-07-21). The release treats every transcribed value as an H_* hypothesis rather than as a kernel-derived fact.

3. Classical partial results (Phases 1-10)

3.1 Mordell-Weil, torsion, and positivity invariants

Mordell-Weil is carried as a fact (rank non-negative and integer-valued) from the theorem of Mordell (1922) and Weil (1928).

Theorem 3.1 (F_mordell_weil_rank_nonneg, Mordell-Weil). *For every elliptic curve $E/\mathbb{Q}$, $\text{rank}_{\mathbb{Z}} E(\mathbb{Q}) \geq 0$.*

Theorem 3.2 (F_torsion_pos). *The finite torsion subgroup has positive order: $\#E(\mathbb{Q})_{\text{tors}} \geq 1$. This permits the trivial group, whose order is 1.*

Theorem 3.3 (F_period_pos). *The real period $\Omega_E > 0$.*

Theorem 3.4 (F_tamagawa_pos). *The Tamagawa product $\prod_p c_p(E) > 0$.*

Theorem 3.5 (F_sha_pos). *The Shafarevich-Tate group order satisfies $\#\text{III}(E) \geq 1$.*

Theorem 3.6 (F_root_number_values). *The root number $w(E) \in \{-1, +1\}$.*

Theorem 3.7 (F_height_diagonal_nonneg). *The encoded height pairing has diagonal nonnegativity: $\langle x, x \rangle_E \geq 0$ for every encoded point coordinate x . The release does not encode a full Mordell-Weil quotient or claim pairwise nonnegativity.*

Theorem 3.8 (F_regulator_pos). *If $\text{rank } E \geq 1$, the regulator $\text{Reg}(E) > 0$; and by convention $\text{Reg}(E) = 1$ when $\text{rank } E = 0$ (F_regulator_rank0).*

3.2 The BSD formula at rank 0 and rank 1

Kolyvagin (1990), building on the Gross-Zagier formula (1986), proves the rank part of BSD when $\text{ord}_{s=1} L(E, s) \leq 1$. The development carries these as facts, and derives the consequences.

Theorem 3.9 (bsd_rank0). *Under the encoded rank-zero interface, $L(E, 1) > 0$ implies $\text{rank } E(\mathbb{Q}) = 0$. Sha finiteness is a separate conclusion consuming A_sha_finite_rank_le1_bridge.*

Theorem 3.10 (bsd_rank1). *Under the encoded Gross-Zagier/Kolyvagin interfaces, $L(E, 1) = 0$ and $L'(E, 1) > 0$ imply $\text{rank } E(\mathbb{Q}) = 1$. Sha finiteness again consumes the explicit low-rank bridge.*

Theorem 3.11 (rank_nonneg). *The Mordell-Weil rank is non-negative as a real number, a structural corollary used downstream.*

The rank-0 and rank-1 BSD formulas for the numerator are then derived in bsd_numerator_pos, rank1_numerator_pos, and period_sha_tam_pos: the product $\Omega \cdot \#\text{III} \cdot \prod c_p$ is strictly positive, which is required for the BSD formula to make sense.

3.3 The parity conjecture

The parity conjecture predicts that the sign of $w(E)$ matches $(-1)^{\text{rank}(E)}$. The release records several conditional fragments relative to imported interfaces.

Theorem 3.12 (root_number_sq_one). *The encoded root-number-values interface implies $w(E)^2 = 1$.*

Theorem 3.13 (root_neg_forces_vanishing, via F_parity_root_neg). *If $w(E) = -1$, then $L(E, 1) = 0$.*

Theorem 3.14 (parity_rank_ge1_if_root_neg). *Conditional on BSD rank, $w(E) = -1$ implies $\text{rank}(E) \geq 1$.*

Theorem 3.15 (parity_rank1_if_root_neg_Lderiv_pos). *Under the encoded parity, Gross-Zagier, and Kolyvagin interfaces, $w(E) = -1$ and $L'(E, 1) > 0$ imply $\text{rank}(E) = 1$.*

4. Two worked elliptic curves (Phases 11-12)

4.1 $E_0 = 32a2$: rank 0

The hypotheses for E_0 state $\text{rank}(E_0) = 0$, $w(E_0) = +1$, $\#E_0(\mathbb{Q})_{\text{tors}} = 4$, $\prod c_p = 4$, $\#\text{III}(E_0) = 1$, and $L(E_0, 1) = \Omega_{E_0}/4$ (the classical numerical value, with $\text{Reg}(E_0) = 1$ by the rank-0 convention).

From these, the formalization kernel-verifies:

Theorem 4.1 (E0_rank_zero). $\text{rank}(E_0) = 0$.

Theorem 4.2 (E0_L_pos). $L(E_0, 1) > 0$, which is forced by $L(E_0, 1) = \Omega_{E_0}/4 > 0$.

Theorem 4.3 (E0_bsd_numerator). $\Omega \cdot \text{Reg} \cdot \#\text{III} \cdot \prod c_p > 0$ for E_0 .

Theorem 4.4 (E0_bsd_ratio). The BSD ratio $L(E_0, 1) \cdot (\#E_0(\mathbb{Q})_{\text{tors}})^2 / (\Omega \cdot \text{Reg} \cdot \#\text{III} \cdot \prod c_p) = 1$.

Theorem 4.5 (E0_bsd_consistent). The BSD equality holds for E_0 : the numerator and denominator of the leading-term formula agree.

Theorem 4.6 (E0_bsd_verified). The full BSD closure for E_0 , combining rank and leading-term parts.

4.2 $E_1 = 37a1$: rank 1

The hypotheses for E_1 state $\text{rank}(E_1) = 1$, $w(E_1) = -1$, $\#E_1(\mathbb{Q})_{\text{tors}} = 1$, $\prod c_p = 1$, $\#\text{III}(E_1) = 1$, $L(E_1, 1) = 0$, $L'(E_1, 1) > 0$, and $\text{ht}_{\text{heeg}}(E_1) > 0$ (a Heegner point of infinite order).

Theorem 4.7 (E1_rank_one). $\text{rank}(E_1) = 1$, derived from the Heegner-point argument via $F_{\text{gross_zagier}}$ and $F_{\text{kolyvagin_rank1}}$.

Theorem 4.8 (E1_L_vanishes). $L(E_1, 1) = 0$.

Theorem 4.9 (E1_Lderiv_pos). $L'(E_1, 1) > 0$.

Theorem 4.10 (E1_rank_from_GZ_K). The rank-1 conclusion for E_1 via the Gross-Zagier chain over the imaginary quadratic field K .

Theorem 4.11 (E1_bsd_consistent). The BSD equality holds for E_1 .

Theorem 4.12 (E1_bsd_verified). The full BSD closure for E_1 , combining rank, vanishing order, and leading-term.

Both curves pass through multiple encoded dependency routes. These are consistency tests of the interface graph, not independent arithmetic proofs, because several premises are imported.

5. Retained modern interfaces (Phases 18-32)

Historical phases 13-17 are excluded from the release environment and manuscript theorem inventory. They encoded p-adic values, characteristic ideals, and primes as ordered real scalars; that

object model cannot support genuine p-adic or ideal-theoretic claims. No theorem in this release is advertised as a formalization of p-adic L -functions, p-adic heights, valuations, or Iwasawa ideals.

Phases 18-23 carry retained modern interfaces. Phases 24-32 add conditional compositions, database hypotheses, and explicit debt assumptions. Phase 28 combines an abstract rank-upper-bound bridge with curve-specific lower-bound data; no generic higher-rank result follows.

Phase 29 treats the congruent-number family $E_n : y^2 = x^3 - n^2x$ polymorphically, leveraging its universal CM structure. Phase 30 sketches a latent-spectral reformulation. This serves as a foundation for a future trace-formula-type identity. Finally, Phase 31 installs the structural reduction of the rank half of BSD to the finiteness of the Shafarevich-Tate group.

5.6 Heegner points over imaginary quadratic fields (Phase 18)

Heegner points on E live in $E(K)$ for K an imaginary quadratic field satisfying the Heegner hypothesis. The Gross-Zagier formula identifies the Néron-Tate height of a Heegner point with $L'(E, \chi, 1)$, where χ is the quadratic character of K .

Theorem 5.18 (`heegner_nontorsion_rank1`). *A Heegner point of infinite order gives rank ≥ 1 .*

Theorem 5.19 (`gz_Lderiv_K_pos`). *Non-torsion Heegner height implies $L'(E/K, 1) > 0$.*

Theorem 5.20 (`heegner_bounds_rational_rank`). *The rank of E over K is bounded below by the rational rank plus an explicit contribution from the Heegner index.*

Theorem 5.21 (`E1_heegner_rank`, `E1_sha_K_finite`). *For E_1 , the Heegner construction closes the rank and the finiteness of $\text{III}(E_1/K)$.*

5.7 Modularity lifting: Wiles, Taylor-Wiles, Khare-Wintenberger (Phase 19)

The modularity theorem of Wiles-Taylor-Breuil-Conrad-Diamond establishes $L(E, s) = L(f_E, s)$ for a weight-2 newform f_E . Serre's conjecture, proved by Khare and Wintenberger (2009), lifts this to two-dimensional mod- ℓ Galois representations.

Theorem 5.22 (`modularity_L_match`). *The elliptic-curve and modular-form L -functions match.*

Theorem 5.23 (`serre_conj_pos_level`, `serre_weight_positive`). *Serre's conductor and weight are positive integers.*

Theorem 5.24 (`tw_conductor_le_level`). *The Taylor-Wiles conductor is bounded by the modular-form level.*

Theorem 5.25 (`galrep_cond_sq_pos`, `modform_level_pos`). *Structural positivity for the Galois-representation data.*

5.8 Euler systems: Kato, Rubin, Kolyvagin (Phase 20)

An Euler system is a compatible family of Galois cohomology classes that bounds the Selmer group. Rubin (2000) axiomatizes the framework; Kato (2004) constructs the Euler system for an elliptic curve; Kolyvagin's construction on Heegner points gives the original instance.

Theorem 5.26 (`rubin_selmer_bound`). *Rubin's abstract Selmer bound from the Euler system.*

Theorem 5.27 (`kato_es_core_rank0`, `kato_es_selmer`). *Kato's Euler system gives the core rank-0 Selmer bound and the general Selmer-bound chain.*

Theorem 5.28 (kolyvagin_derived_sha_bound). *Kolyvagin's derived classes bound III.*

Theorem 5.29 (core_rank_nonneg, derived_length_nonneg, es_selmer_bound_nonneg). *Structural positivity of the Euler-system quantities.*

5.9 Congruences and visibility (Phase 21)

Cremona and Mazur (2000) propose that $\text{III}(E)$ is controlled by congruences modulo congruence primes with other modular forms of the same level; this is the visibility conjecture.

Theorem 5.30 (visible_sha_nonneg, visible_sha_le_full). *The visible subgroup of III is non-negative and bounded above by III.*

Theorem 5.31 (cremona_mazur_bound). *The Cremona-Mazur visibility bound.*

Theorem 5.32 (sha_vis_chain). *The full chain from modular degree, visibility, and the modular-form side.*

Theorem 5.33 (mod_degree_pos, manin_constant_pos, eisenstein_ideal_pos). *Structural positivity of the modular-form-side invariants.*

5.10 Algorithmic BSD and two-descent (Phase 22)

Two-descent computes the 2-Selmer group of E and gives an upper bound on the rank plus explicit generators on the analytic side.

Theorem 5.34 (two_selmer_bounds_rank, gen_count_bounds_rank). *Two-Selmer rank bounds the Mordell-Weil rank; the explicit generator count bounds it from below.*

Theorem 5.35 (rank_pinch_exact). *When both bounds coincide, they determine the rank exactly.*

Theorem 5.36 (E0_two_descent_rank, E1_two_descent_rank). *Two-descent confirms $\text{rank}(E_0) = 0$ and $\text{rank}(E_1) = 1$, independently of the Gross-Zagier chain.*

Theorem 5.37 (E0_bsd_verified, E1_bsd_verified). *The full BSD closure via the two-descent route, matching §4.*

5.11 Higher-rank theorems: average rank and positive-proportion BSD (Phase 23)

Bhargava and Shankar (2013) prove that the average rank of elliptic curves (ordered by height) is bounded above by $7/6$. Bhargava, Skinner, and Zhang (2014) prove that a positive proportion satisfies BSD.

Theorem 5.38 (avg_rank_lt_seven_sixths, avg_rank_positive, avg_rank_range). *$0 < \text{average rank} < 7/6$.*

Theorem 5.39 (bsz_two_thirds, bsd_proportion_range). *At least 66% of elliptic curves $E/\mathbb{Q}$ (ordered by height) satisfy BSD.*

Theorem 5.40 (rank_le1_implies_bsd). *The proportion satisfying $\text{rank} \leq 1$ is a lower bound on the proportion satisfying BSD.*

Theorem 5.41 (`shimura_rank_nonneg`, `darmon_height_nonneg`). *Shimura-curve rank and Darmon-point height non-negativity — structural statements that fit naturally into the higher-rank framework.*

5.12 GRH-effective BSD for rank ≤ 1 (Phase 24)

The classical analytic-rank-at-most-one theorem is due to Gross–Zagier and Kolyvagin under its precise modularity and nonvanishing hypotheses. The additional effective claims encoded in Phase 24 are assumptions. Goldfeld (1976) concerns class numbers and BSD conjectures, while the cited Hoffstein–Lieman item concerns Hecke-operator eigenvalues; neither citation establishes the universal GRH implications below.

The formalization introduces a GRH indicator `grh_indicator` : `EC` $\rightarrow$ `Real` taking values in $\{0, 1\}$ and an effective constant `goldfeld_effective_const` : `EC` $\rightarrow$ `Real`, together with the following hypotheses cited to the sources:

Hypothesis	Statement	Source
<code>F_goldfeld_gzk_effective</code>	Under GRH and analytic rank at most one, the algebraic and analytic ranks agree.	exported assumption; attribution not validated
<code>F_hoffstein_lieman</code>	Under GRH and central vanishing, the first derivative has a strictly positive effective lower bound.	exported assumption; attribution not validated
<code>A_grh_sha_finite_package</code>	Under GRH and analytic rank at most one, <code>ShaPFinite E</code> .	abstract assumed package; exact source hypotheses not encoded
<code>F_grh_an_rank_bound</code>	Under GRH, analytic rank has an effective conductor-dependent upper bound.	exported assumption; attribution not validated

From these, the suite derives.

Theorem 5.42 (`grh_bsd_rank_le1`). *For every elliptic curve $E/\mathbb{Q}$, if GRH holds for E and $\text{an_rank}(E) \leq 1$, then $\text{rank}(E) = \text{an_rank}(E)$.*

Theorem 5.43 (`grh_sha_finite_rank_le1`). *Under the same hypotheses, `ShaPFinite E` follows from `A_grh_sha_finite_package`. No effective cardinality bound is encoded.*

Theorem 5.44 (`grh_L_deriv_pos_of_L_zero`). *Under GRH, if $L(E, 1) = 0$ then $L'(E, 1) > 0$ effectively.*

Theorem 5.45 (`grh_an_rank_effective_bound`). *Under GRH, the analytic rank is bounded above effectively.*

Phase 24 is not a literature-validated upgrade: its effective GRH interfaces are assumptions whose cited references do not establish the encoded universal bounds. It is retained only to preserve the historical dependency graph.

5.13 Encoded Coates–Wiles/Rubin CM route (Phase 25)

Coates–Wiles proves a rank-zero result for CM elliptic curves under its stated CM and nonvanishing conditions. Rubin’s main-conjecture theorem supplies powerful CM Iwasawa input with prime and local restrictions. The present model imports simplified universal interfaces for these results and therefore reports only conditional compositions of those interfaces.

The formalization introduces a CM indicator $\text{cm_indicator} : \text{EC} \rightarrow \text{Real}$ and a discriminant function $\text{cm_discriminant} : \text{EC} \rightarrow \text{Real}$. It carries the Coates–Wiles rank-zero theorem under displayed CM and nonvanishing inputs, while Sha finiteness is separated into `A_cm_sha_finite_package`. No Rubin main-conjecture object is encoded in the live release.

Theorem 5.46 (`coates_wiles_rank_zero`). *For every elliptic curve $E/\mathbb{Q}$ with complex multiplication, if $L(E, 1) > 0$ then $\text{rank}(E) = 0$.*

Theorem 5.47 (`coates_wiles_sha_finite`). *Under the same hypotheses, $\text{ShaPFinite } E$ follows from the explicit `A_cm_sha_finite_package`; this is not attributed to Coates–Wiles alone.*

Theorem 5.49 (`E0_has_cm, E0_cm_disc_is_minus_four`). *E_0 has complex multiplication, with CM discriminant -4 .*

Theorem 5.50 (`E1_no_cm`). *E_1 does not have complex multiplication.*

Theorem 5.51 (`E0_rank_via_coates_wiles, E0_sha_via_coates_wiles`). *E_0 has rank 0 via Coates–Wiles; $\text{ShaPFinite } E_0$ separately consumes `A_cm_sha_finite_package`.*

Phase 25 supplies E_0 with a second encoded dependency route. Independence of the imported arithmetic premises is not machine-checked.

5.14 Explicit BSD leading-term bridge at rank 0 (Phase 26)

The formalization does not encode a Selmer characteristic ideal or a p -adic L -function. It imports the exact classical rank-zero conclusion as `A_bsd_formula_rank0_bridge`. Since $\text{Reg}(E) = 1$ at rank zero, the target formula has the shape

$$L_{\text{lead}}(E) = \frac{\Omega_E \cdot \text{Reg}(E) \cdot \#\text{III}(E) \cdot \prod_p c_p(E)}{(\#E(\mathbb{Q})_{\text{tors}})^2}$$

The exported result is conditional on this bridge; it is not presented as a formalization of Kato, Skinner–Urban, or an Iwasawa main conjecture.

Theorem 5.53 (`L_lead_rank0_is_L_val`). *If $\text{rank}(E) = 0$ then $L_{\text{lead}}(E) = L(E, 1)$.*

Theorem 5.55 (`bsd_formula_rank0_under_bridge`). *Under the same hypotheses, the full rank-0 BSD leading-term formula follows from `A_bsd_formula_rank0_bridge`.*

Theorem 5.56 (`E0_bsd_formula_under_rank0_bridge, E0_L_lead_eq_L_val`). *For the worked curve E_0 , the leading-term formula follows from the explicit rank-zero bridge.*

Theorem 5.57 (`cm_bsd_formula_rank0_under_bridge`). *The CM specialization remains conditional on `A_bsd_formula_rank0_bridge`; its name contains no unconditional cue.*

Phase 26 avoids direct use of `A_bsd_formula` at rank zero by consuming a separate explicit bridge. This rearranges the trust surface; it does not prove the arithmetic input.

5.15 A GRH-effective per-curve LMFDB database (Phase 27)

Phase 27 instantiates historical GRH, CM, database, and rank-bound interfaces at six named curves:

Curve	Encoded analytic/CM data	Encoded algebraic-rank data	Formal output
E_{11a1}	H_E_11a1_an_rank	GRH route only	E_11a1_rank_via_grh
E_{14a1}	H_E_14a1_an_rank	GRH route only	E_14a1_rank_via_grh
E_{27a1}	H_E_27a1_L_val_pos, H_E_27a1_cm	Coates–Wiles interface	E_27a1_rank_zero_cw
E_{43a1}	H_E_43a1_an_rank	GRH route only	E_43a1_rank_via_grh
E_{389a1}	H_E_389a1_an_rank	mwrnk lower + pinned upper	E_389a1_rank_eq_an_rank_pinned
E_{5077a1}	H_E_5077a1_an_rank	mwrnk lower + pinned upper	E_5077a1_rank_eq_an_rank_pinned

Every row consumes named database and theorem interfaces. The higher-rank rows pin equality from separate mwrnk and analytic-rank hypotheses; they are not generic BSD results.

Database provenance is per curve: 11.a1, 14.a1, 27.a1, 43.a1, 389.a1, and 5077.a1 (LMFDB records accessed 2026-07-21). The release does not independently recompute those records; each imported value remains an H_* hypothesis.

Theorem 5.59 (E_11a1_rank_via_grh, E_14a1_rank_via_grh, E_43a1_rank_via_grh). *For each of E_{11a1} , E_{14a1} , E_{43a1} , the rank part of BSD holds under GRH: $\text{rank}(E) = \text{an_rank}(E)$ as a kernel theorem.*

Theorem 5.60 (E_27a1_rank_zero_cw). *The rank-zero conclusion follows from the encoded Coates–Wiles and CM premises.*

Theorem 5.61 (E_389a1_rank_eq_an_rank_pinned, E_5077a1_rank_eq_an_rank_pinned). *For each of E_{389a1} , E_{5077a1} , the numerical identity $\text{rank}(E) = \text{an_rank}(E)$ holds as a kernel consequence of the cited mwrnk Mordell–Weil rank and the LMFDB analytic rank.*

Phase 27 broadens the worked-curve database from two curves to eight and establishes that the $\text{an_rank} \leq 1$ narrowing is not an abstract statement about “some curves” but a concrete list of named LMFDB objects, each of which can be separately inspected by a reviewer.

5.16 Encoded Kato upper-bound interface and curve-specific compositions (Phase 28)

The implementation assumes an abstract inequality

$$\text{rank}(E) \leq \text{an_rank}(E)$$

as A_rank_upper_bound_bridge. Kato’s actual Euler-system results require arithmetic hypotheses not represented in this universal statement. Curve-specific equality then follows formally when this bridge is combined with a matching lower-bound and analytic-rank premise.

Theorem 5.62 (rank_upper_bound_under_bridge). *The explicit A_rank_upper_bound_bridge is specialized to E ; this is a conditional formal implication.*

Theorem 5.63 (`kato_rank_lower_matches_an_rank`). *For every E and every $r_0 \in \mathbb{R}$, if $r_0 \leq \text{rank}(E)$ and $\text{an_rank}(E) = r_0$, then $\text{rank}(E) = \text{an_rank}(E)$.*

Theorem 5.63 is the universal closure pattern: it takes an arbitrary lower-bound hypothesis (supplied per-curve by `mwrank` or by an explicit generator count in a two-descent argument) and an arbitrary numerical-equality hypothesis (supplied per-curve by the LMFDB analytic rank), and produces the full rank identity.

Theorem 5.64 (`E_389a1_rank_eq_an_rank_kato`). $\text{rank}(E_{389a1}) = \text{an_rank}(E_{389a1}) = 2$, **conditional** on the open general inequality $\text{rank}(E) \leq \text{an_rank}(E)$ at $\text{rank} \geq 2$, modulo the citation to `mwrank` for the Mordell-Weil rank lower bound and the citation to the LMFDB for the analytic rank.

Theorem 5.65 (`E_5077a1_rank_eq_an_rank_kato`). $\text{rank}(E_{5077a1}) = \text{an_rank}(E_{5077a1}) = 3$, **conditional** on the open general inequality $\text{rank}(E) \leq \text{an_rank}(E)$ at $\text{rank} \geq 2$, modulo the same two citations.

Phase 28 therefore upgrades the rank-2 and rank-3 closures from “numerical consistency” (Phase 27) to kernel theorems that are conditional on the open general inequality $\text{rank}(E) \leq \text{an_rank}(E)$ at $\text{rank} \geq 2$: the Mordell-Weil rank, the analytic rank, and their equality are kernel-verified once the reviewer accepts the cited numerical inputs *and* that open upper-bound inequality. This is a concrete instance of $\text{an_rank} \geq 2$ being closed conditionally without appealing to `A_bsd_rank` — a strict refinement of the residual open scope outlined in §1.3.

5.17 The congruent-number family (Phase 29)

An integer $n \geq 1$ is a *congruent number* if it is the area of some right triangle with rational side lengths. The classical reformulation, due to Fermat and Tunnell, identifies the congruent numbers with the integers n such that the elliptic curve $E_n : y^2 = x^3 - n^2x$ has $\text{rank}(E_n) \geq 1$. Phase 29 declares E_n as a function `E_cn` : $\mathbb{R} \rightarrow \text{EC}$ in the kernel and proves the following per-family facts.

Theorem 5.66 (`cn_always_cm`). *The imported family-CM interface yields that every encoded E_n has complex multiplication by $\mathbb{Z}[i]$ (discriminant -4).*

This is an immediate consequence of the explicit j -invariant of E_n : $j(E_n) = 1728$ for every n , and every curve with $j = 1728$ has CM by $\mathbb{Z}[i]$. The formalization carries this as a family-level fact `F_cn_always_cm`.

Theorem 5.67 (`cn_rank0_via_coates_wiles`). *For every $n > 0$, if $L(E_n, 1) > 0$ then $\text{rank}(E_n) = 0$, conditional on the encoded family-CM and Coates–Wiles interfaces.*

Theorem 5.67 composes the encoded universal-CM and Coates–Wiles interfaces with a positivity premise. It is conditional on those interfaces.

Theorem 5.68 (`tian_package_rank_ge_one`). *The abstract premise $\text{tian_heegner_hyp}(n)=1$ implies $\text{rank}(E_n) \geq 1$ via `A_tian_rank_lower_bound_package`.*

Tian’s squarefreeness, prime-factorization, class-group, and Heegner-point hypotheses are not encoded. The package is therefore not advertised as a transcription of Tian’s theorem.

Theorem 5.69 (`cn34_rank_eq_an_rank_via_kato`). $\text{rank}(E_{n=34}) = \text{an_rank}(E_{n=34}) = 2$, **conditional** on the open general inequality $\text{rank}(E) \leq \text{an_rank}(E)$ at $\text{rank} \geq 2$, modulo the cited `mwrank` value and the cited Tunnell-derived analytic rank.

The congruent number $n = 34$ is the smallest squarefree congruent number with $\text{rank}(E_n) = 2$; it lies outside the range of pure Coates-Wiles (rank 0) and pure Tian (rank ≥ 1) and is closed by the same Kato-upper-bound pattern as E_{389a1} and E_{5077a1} in §5.16.

Honest scope. Phase 29 does *not* resolve the congruent-number problem. Every family conclusion is relative to named CM, Tian–Tunnell, Kato, mwrank, or formula interfaces.

5.18 A Latent-spectral reformulation (Phase 30)

Phase 30 introduces kernel symbols $\text{spec_mult}(E) : \mathbb{N}$ and $\text{spec_atom}(E) : \mathbb{R}$, with definitional hypotheses

$$\text{spec_mult}(E) = \text{an_rank}(E), \quad \text{spec_atom}(E) = L_{\text{lead}}(E),$$

and restates the rank-0 CM BSD formula in the spectral language.

Theorem 5.70 (`spec_mult_unfold`). $\text{spec_mult}(E) = \text{an_rank}(E)$.

Theorem 5.71 (`spec_atom_unfold`). $\text{spec_atom}(E) = L_{\text{lead}}(E)$.

Theorem 5.72 (`spec_cm_bsd_rank0_under_bridge`). *Under the encoded CM and rank-zero formula bridge, if $L(E, 1) > 0$ then*

$$\text{spec_atom}(E) = \frac{\Omega_E \cdot \text{Reg}(E) \cdot \#\text{III}(E) \cdot \prod_p c_p(E)}{(\#E(\mathbb{Q})_{\text{tors}})^2}.$$

The implication is conditional on imported premises.

Phase 30 is a *structural rewriting*, not new number theory: the spectral symbols are definitionally equal to existing abstract invariants. The companion sketch `latent_spectral_sketch.md` is exploratory and supplies no theorem used by this release.

5.19 Structural compression: Sha finite $\Rightarrow$ BSD rank (Phase 31)

Phase 31 records an abstract reduction from a Sha-finiteness interface to a rank identity. It does not formalize the arithmetic constructions behind that reduction.

Selmer interface. The classical Selmer exact sequence motivates

$$\text{rank}_{\mathbb{Z}_p} \text{Sel}_{p^\infty}(E/\mathbb{Q}) = \text{rank}_{\mathbb{Z}} E(\mathbb{Q}) + \text{corank}_{\mathbb{Z}_p} \text{III}(E/\mathbb{Q})[p^\infty].$$

If $\text{III}(E/\mathbb{Q})[p^\infty]$ is finite then its $\mathbb{Z}_p$ -corank is 0. The implementation represents the required implication by

$$\text{F_selmer_eq_mw_under_sha_finite} : \forall E, \text{ShaPFinite}(E) \implies \text{sel_rk}(E) = \text{rank}(E).$$

Iwasawa interface. The implementation separately assumes

$$\text{sel_rk}(E) = \text{an_rank}(E).$$

$$\text{A_selmer_eq_an_rank_bridge} : \forall E, \text{sel_rk}(E) = \text{an_rank}(E).$$

The actual theorems in Iwasawa theory have prime, reduction, residual-representation, and characteristic-ideal hypotheses. Those are not encoded, so the equality is explicit bridge debt.

The Phase 31 core theorem transits these through:

Theorem 5.73 (`bsd_rank_from_sha_finite`). *For every elliptic curve $E/\mathbb{Q}$, if $\text{III}(E)[p^\infty]$ is finite then $\text{rank}(E) = \text{an_rank}(E)$.*

The proof is three lines: specialize `F_selmer_eq_mw_under_sha_finite`, specialize `A_selmer_eq_an_rank_bridge` and conclude by transitivity.

Theorem 5.74 (`bsd_rank_universal_from_sha_finite`). *A_bsd_rank is a kernel theorem under A_sha_finite .*

This is a conditional composition. It consumes both `A_sha_finite` and the universal cited interfaces above; it is not a derivation from Sha finiteness alone in the external mathematical literature.

Theorem 5.75 (`cm_rank_eq_an_rank_via_sha_finite`). *For a CM elliptic curve satisfying the encoded positivity and cited interfaces, the rank identity follows conditionally.*

The kernel checks agreement of two imported CM routes; it does not encode the local hypotheses needed to invoke Coates–Wiles or Rubin.

Theorem 5.76 (`an_rank_le1_rank_eq_an_rank_grh_free`). *For every encoded elliptic curve with $\text{an_rank}(E) \leq 1$, the rank identity follows from the cited Sha and Selmer interfaces. No GRH hypothesis appears in this formal implication.*

The proof composes:

1. `A_rank_upper_bound_bridge` and `A_sha_finite_rank_le1_bridge`.
2. `F_selmer_eq_mw_under_sha_finite` and `A_selmer_eq_an_rank_bridge`.
3. Phase 31 (`bsd_rank_from_sha_finite`): finiteness of $\text{III}[p^\infty]$ gives $\text{rank}(E) = \text{an_rank}(E)$.

No step uses GRH, but every imported premise remains part of the trust surface.

Theorem 5.77-5.78 (`an_rank_zero_rank_identity_under_bridges`, `an_rank_one_rank_identity_under_bridges`). *Conditional specializations of Theorem 5.76 with no unconditional naming cue.*

Residual open scope. Phase 31 does *not* prove BSD. It localizes a dependency graph under abstract imported interfaces.

5.20 Explicit BSD leading-term bridge at rank 1 (Phase 32)

Phase 26 uses an explicit rank-zero bridge. At rank one, literature results involving p-adic Gross–Zagier formulas, heights, main conjectures, and Euler systems come with object types, hypotheses, and often prime-part conclusions absent from this interface. They are contextual motivation only. The formal object is the classical archimedean equality `A_bsd_formula_rank1_bridge`.

The Phase 32 implementation carries two cited interfaces, one explicit bridge assumption, and six kernel theorems:

$$\text{F_L_lead_rank1_is_L_der} : \forall E, \text{rank}(E) = 1 \implies L_{\text{lead}}(E) = L^{(1)}(E, 1),$$

a Taylor-expansion specialization at rank 1 (automatic from the definition of L_{lead} and analytic rank = 1);

$$\text{A_bsd_formula_rank1_bridge} : \forall E, \text{rank}(E) = 1 \wedge L_{\text{lead}}(E) > 0 \implies L_{\text{lead}}(E) = \frac{\Omega_E \cdot \text{Reg}(E) \cdot \#\text{III}(E) \cdot \prod c_p(E)}{(\#E(\mathbb{Q})_{\text{tors}})^2}$$

an explicit bridge whose universal form is stronger than the cited literature package encoded in the paper. Jetchev–Skinner–Wan (2017), for example, establishes a p -part result under semistability, good-reduction, prime, and residual-irreducibility hypotheses. Finally:

$$\text{A_cm_rank1_formula_package} : \forall E, \text{cm}(E) = 1 \wedge \text{rank}(E) = 1 \wedge L_{\text{lead}}(E) > 0 \implies L_{\text{lead}}(E) = \frac{\Omega_E \cdot \text{Reg}(E) \cdot \#\text{III}(E)}{(\#E(\mathbb{Q})_{\text{tors}})^2}$$

the separately disclosed CM assumed package. Its local hypothesis translation is not formalized and no theorem-strength attribution is claimed. The six kernel theorems are:

Theorem 5.79 (`L_lead_rank1_is_L_der`). *For every elliptic curve $E/\mathbb{Q}$ with $\text{rank}(E) = 1$, $L_{\text{lead}}(E) = L^{(1)}(E, 1)$.*

This is a one-line kernel theorem: specialize `F_L_lead_rank1_is_L_der` to the specific E .

Theorem 5.80 (`bsd_formula_rank1_via_bridge`). *For every elliptic curve $E/\mathbb{Q}$ with $\text{rank}(E) = 1$ and $L^{(1)}(E, 1) > 0$, the BSD leading-term formula follows from `A_bsd_formula_rank1_bridge`.*

The proof specializes `F_L_lead_rank1_is_L_der`, transports positivity, and applies the explicit bridge.

Theorem 5.81 (`bsd_formula_rank1_under_an_rank1_bridge`). *For every elliptic curve $E/\mathbb{Q}$ with $\text{an_rank}(E) = 1$ and positive leading term, the formula follows from the Phase 31 interfaces and `A_bsd_formula_rank1_bridge`.*

This is the formula-half companion of Phase 31’s `an_rank_le1_rank_eq_an_rank_grh_free`. The proof chains:

1. `an_rank(E) = 1` and Phase 31 give $\text{rank}(E) = 1$ under the explicit rank/Sha/Selmer bridges.
2. $\text{rank}(E) = 1$ and `F_L_lead_rank1_is_L_der` give $L_{\text{lead}}(E) = L^{(1)}(E, 1)$.
3. $L^{(1)}(E, 1) > 0$ combined with (2) gives $L_{\text{lead}}(E) > 0$.
4. $\text{rank}(E) = 1$ and $L_{\text{lead}}(E) > 0$ together with `A_bsd_formula_rank1_bridge` give the leading-term formula.

No step uses GRH. The conclusion nevertheless remains conjectural because it consumes `A_bsd_formula_rank1_bridge`, in addition to cited interfaces.

Theorem 5.82 (`E1_L_lead_eq_L_der`). *For $\$E_1 = \$ \text{Cremona } 37a1$, $L_{\text{lead}}(E_1) = L^{(1)}(E_1, 1)$.*

Specialization of Theorem 5.79 to E_1 . Uses the Phase 4 hypothesis $\text{rank}(E_1) = 1$.

Theorem 5.83 (`E1_bsd_formula_under_rank1_bridge`). *For $\$E_1 = \$ \text{Cremona } 37a1$, the BSD leading-term formula follows conditionally from the Phase 32 bridge.*

This gives E_1 a second independent closure route for the BSD formula, complementing the Phase 11 direct verification `E1_bsd_verified`. The proof uses the numerical hypothesis $L^{(1)}(E_1, 1) > 0$ from the worked-curve data and then applies Theorem 5.80.

Theorem 5.84 (`cm_bsd_formula_rank1_under_cited_package`). *For every CM elliptic curve $E/\mathbb{Q}$ with the stated rank and positivity premises, the formula follows from $A_cm_rank1_formula_package$.*

The CM package is deliberately unattributed at theorem strength because the full local hypothesis translation is not formalized.

Residual open scope. Phase 32 does *not* prove BSD and does not close generic rank one: the bridge is itself explicit conjectural debt. Rank at least two is outside this release scope.

6. Reproducibility

6.1 Verification count

The replay and dependency-closed export inventory is:

Quantity	Value
Kernel declarations checked	368
Verification errors	0
Kernel theorem declarations	181
Hypothesis declarations	187
Explicit conjectural or bridge hypotheses	12
Lean lemma markers (kernel implications)	181
Lean generated <code>_trusted</code> helper theorems	123
Lean axiom markers	$223 = 187 + 36$
Lean sorry occurrences	0

The 36 generated Lean dependencies are type and operation declarations required to state the BSD environment. They are not authored mathematical results.

6.2 Lean 4 export

The canonical generator produces the SHA-256-pinned export stamp/`BSD_main.lean`. It compiles under Lean 4 with the pinned Mathlib environment and contains 181 lemma markers, 123 generated helper theorems, 223 axiom markers, and zero sorry. The export includes `#print axioms` commands for six load-bearing capstones. Exact hashes are recorded in `stamp/BSD_main.stamp.json` and `paper_claim_map.json`.

6.3 Replay commands

The standalone replay package records the Python and Lean toolchain dependencies. No public hosting location is asserted in this draft. Replay consists of running the canonical stamp generator, compiling the Lean export with `lake env lean`, and validating the exact-hash claim map. The generated `#print axioms` output is the authoritative capstone trust inventory.

7. Discussion

7.1 What is mechanized

Every declaration in the suite is explicitly typed and kernel-checked. The 181 theorem declarations verify implications relative to the 187-hypothesis trust surface.

Multiple formal routes are exercised for each worked curve. These are dependency-path tests, not independent proofs of their imported premises. The 37a1 formula route is `E1_bsd_formula_under_rank1_bridge`.

The formalization supplies several conditional kernel theorems relative to twelve explicit debt assumptions.

First, `an_rank_le1_rank_eq_an_rank_grh_free` is GRH-free but conditional on imported Sha and Selmer interfaces. Second, `bsd_rank_from_sha_finite` and `bsd_rank_universal_from_sha_finite` compose Sha finiteness with those interfaces. Third, the CM routes remain conditional specializations of imported CM statements.

At rank zero, `bsd_formula_rank0_under_bridge` consumes `A_bsd_formula_rank0_bridge`. At rank one, `bsd_formula_rank1_via_bridge` consumes `A_bsd_formula_rank1_bridge`; its analytic-rank composition is `bsd_formula_rank1_under_an_rank1_bridge`. The CM path consumes `A_cm_rank1_formula_package`.

Finally, `rank_upper_bound_under_bridge` merely consumes `A_rank_upper_bound_bridge`. The theorem `kato_rank_lower_matches_an_rank` records the corresponding conditional composition with a certified lower bound; its historical name is not an attribution of the bridge itself.

These narrowings identify which assumptions are sufficient for each capstone. In particular, avoiding `A_bsd_formula` by consuming `A_bsd_formula_rank1_bridge` does not make a rank-one result unconditional.

7.2 What is not: the honest debt

The twelve `A_*` declarations are the authoritative debt inventory in §1.3. `A_sha_finite` is proposition-valued; the rank reduction additionally requires `A_selmer_eq_an_rank_bridge`. Formula and literature-package gaps remain separate assumptions rather than being laundered as cited facts.

- **`A_sha_finite`:** Gross–Zagier–Kolyvagin, CM Iwasawa theory, and later Euler-system results prove finiteness statements in important classes under precise hypotheses. Those hypotheses are not reconstructed in this abstract interface; no universal Sha-finiteness claim is made.
- **`A_bsd_rank`:** the suite contains conditional routes through Sha finiteness, Iwasawa interfaces, and curve-specific lower-bound data. The generic rank identity remains conjectural.
- **`A_bsd_formula` and `A_bsd_formula_rank1_bridge`:** rank-zero and rank-one specializations are conditional formal implications. The generic rank-one route consumes the explicit bridge; rank at least two remains outside this release scope.
- **Other `A_*` packages:** the exact Tian, CM, GRH, low-rank Sha, rank-upper-bound, and Selmer/analytic-rank hypotheses remain to be represented before any package can be promoted to a source-faithful theorem interface.

The net effect is dependency localization. There are twelve explicit conjectural or bridge assumptions; no conjecture or missing hypothesis translation is eliminated by naming it as a cited fact.

Among the axioms that remain independent, `A_sha_finite` is structurally cleaner than `A_bsd_rank`. It forms a statement about a single arithmetic invariant (the finiteness of a torsion group) rather than a rank identity between two apparently unrelated integers. This makes it exactly the object on which modern Euler-system and Iwasawa-theoretic methods operate. This setup provides automatic future-proofing. Every future Sha-finiteness result—of the form “for every E in a given class, $\text{III}(E/\mathbb{Q})[p^\infty]$ is finite”—will automatically upgrade to a BSD rank theorem through Phase 31, without any further Lean 4 work. The `A_bsd_formula` axiom retains a similar future-proofing. Every new cited fact about the BSD formula at a given rank (for example, a rank-2 analog of Perrin-Riou p -adic Gross-Zagier) slots directly into a new phase with the exact structural shape as Phase 32. This requires no revisiting of the axiomatic core.

7.3 The trust surface

The trust surface, enumerated by kind:

- **Generated type/operation support (36 Lean axioms).** These declarations provide the type and operation dependencies needed to state the exported BSD environment.
- **Cited-result interfaces (`F_*`).** These are exported as Lean axioms. Source comments provide attribution, but full hypothesis fidelity remains an audit obligation.
- **Worked-curve hypotheses (33, `H_*`).** Numerical data for the eight worked curves and the congruent-number family: ranks, root numbers, torsion, Tamagawa products, periods, regulators, L -values, Heegner heights, LMFDB analytic ranks, Cremona `mwrnk` Mordell-Weil ranks, GRH indicators, CM indicators. Each is a finite piece of data computable from an elliptic-curve database (e.g., LMFDB) or from a certified rank computation (`mwrnk`).
- **Conjectural or bridge axioms (12, `A_*`).** The global BSD/Sha statements and nine explicit package or translation bridges.

A reviewer obtains machine verification of each implication relative to the exact axioms they choose to accept. The Lean export does not upgrade an imported statement’s semantic status.

Lean export trust audit. The generated export contains 181 lemma markers, 123 generated helper theorems, and 223 axiom markers: 187 user hypotheses plus 36 generated dependencies. It requires the pinned Lean/Mathlib environment.

7.4 Future work

The following directions extend the formalization beyond the current state; each is stated as a concrete phase number to be added on top of Phase 32.

Phase 33: formalize a literature theorem with all hypotheses. Replace `A_bsd_formula_rank1_bridge` by one or more precise theorems encoding reduction type, prime restrictions, residual irreducibility, Heegner hypotheses, local factors, and prime-part versus full-order conclusions.

Phase 34: Audited database batch. A future batch should ingest exact database records and independently check every analytic-rank, algebraic-rank, and leading-coefficient field before generating curve-specific conditional implications.

Phase 35: Higher-rank extensions. Any future extension must encode a precise published theorem, including its field, motive, local conditions, and prime-part scope. No generic rank-two closure is claimed here.

Phase 36: Function-field trace-formula comparison. Yun–Zhang (2017) concerns shtukas and Taylor coefficients over global function fields. It may provide conceptual comparison, but it is not a theorem about non-CM elliptic curves over $\mathbb{Q}$ and is not an input to this release.

Phase 37: Spectral trace-formula identity. Phase 30 introduces the spectral interface `spec_mult`, `spec_atom` as kernel-level definitional symbols, but does not state or prove a trace-formula-type identity for $L(E, s)$. A concrete such identity — analogous to Selberg’s trace formula in a suitable geometric or automorphic setting; the Jacquet–Zhang relative trace formula for $\mathrm{GL}_2 \times \mathrm{GL}_2$ is a candidate — would give the spectral symbols independent content and bring new methods to bear on $\mathrm{rank} \geq 2$. This is the subject of the companion sketch `latent_spectral_sketch.md` and remains a research program, not a scheduled phase.

Phase 38: no current claim. Liu–Tian–Xiao–Zhang–Zhu (2022) concerns the Beilinson–Bloch–Kato conjecture for Rankin–Selberg motives and is not a source for a 66.5% BSD proportion. The former proposal is withdrawn pending a correct source.

Lean export improvement. A future exporter should replace generated trust helpers with readable Mathlib proofs where possible and should measure any trust-surface reduction from actual compiled axiom output rather than declaration counts.



Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used Cursor and its integrated AI models (Claude and GPT families) in order to assist with manuscript drafting, structural editing, and language polishing. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.



References

Bhargava, M. and Shankar, A. (2015). *Binary quartic forms having bounded invariants, and the boundedness of the average rank of elliptic curves*. *Annals of Mathematics* **181** (1), 191–242.

Bhargava, M., Skinner, C., and Zhang, W. (2014). *A majority of elliptic curves over $\mathbb{Q}$ satisfy the Birch and Swinnerton-Dyer conjecture*. arXiv:1407.1826. DOI: 10.48550/arXiv.1407.1826.

Birch, B. J. and Swinnerton-Dyer, H. P. F. (1965). *Notes on elliptic curves. II*. *J. reine angew. Math.* **218**, 79–108.

Bloch, S. and Kato, K. (1990). *L-functions and Tamagawa numbers of motives*. In: *The Grothendieck Festschrift, Vol. I*, Progress in Mathematics **86**, Birkhäuser, 333–400.

Clay Mathematics Institute (2026). *Birch and Swinnerton-Dyer Conjecture*. <https://www.claymath.org/millennium> [Online; accessed 21 July 2026].

Breuil, C., Conrad, B., Diamond, F., and Taylor, R. (2001). *On the modularity of elliptic curves over $\mathbb{Q}$: wild 3-adic exercises*. *J. Amer. Math. Soc.* **14** (4), 843–939.

- Cassels, J. W. S. (1965). *Arithmetic on curves of genus 1. VIII. On conjectures of Birch and Swinnerton-Dyer*. J. reine angew. Math. **217**, 180-199.
- Coates, J. and Wiles, A. (1977). *On the conjecture of Birch and Swinnerton-Dyer*. Invent. Math. **39** (3), 223-251.
- Cremona, J. E. and Mazur, B. (2000). *Visualizing elements in the Shafarevich-Tate group*. Experiment. Math. **9** (1), 13-28.
- Diamond, F. and Shurman, J. (2005). *A First Course in Modular Forms*. Graduate Texts in Mathematics **228**, Springer.
- Goldfeld, D. (1976). *The class number of quadratic fields and the conjectures of Birch and Swinnerton-Dyer*. Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4) **3** (4), 623-663.
- Greenberg, R. (1989). *Iwasawa theory for p -adic representations*. In: *Algebraic Number Theory — in honor of K. Iwasawa*, Advanced Studies in Pure Mathematics **17**, Academic Press, 97-137. (Source of the $\mu = 0$ conjecture for elliptic curves.)
- Greenberg, R. (1999). *Iwasawa theory for elliptic curves*. In: *Arithmetic Theory of Elliptic Curves (Cetraro, 1997)*, Lecture Notes in Mathematics **1716**, Springer, 51-144.
- Google DeepMind (2026). *formal-conjectures: formalized statements of conjectures in Lean 4*. <https://github.com/google-deepmind/formal-conjectures> [Online; accessed 21 July 2026].
- Gross, B. H. (1991). *Kolyvagin's work on modular elliptic curves*. In: *L-functions and arithmetic (Durham, 1989)*, London Math. Soc. Lecture Note Ser. **153**, Cambridge University Press, 235-256. (Heegner-index formula.)
- Gross, B. H. and Zagier, D. B. (1986). *Heegner points and derivatives of L -series*. Invent. Math. **84** (2), 225-320.
- Goldfeld, D., Hoffstein, J., and Lieman, D. (1994). *An effective zero-free region*. Appendix to Hoffstein, J. and Lockhart, P., *Coefficients of Maass forms and the Siegel zero*. Ann. of Math. (2) **140** (1), 177-181. DOI: 10.2307/2118544.
- Jetchev, D., Skinner, C., and Wan, X. (2017). *The Birch and Swinnerton-Dyer formula for elliptic curves of analytic rank one*. Cambridge J. Math. **5** (3), 369-434.
- Bell, J. et al. (2021). *Formalising BSD* (Lean 3 project). <https://github.com/jamiebell2805/BSD-conjecture>.
- Kato, K. (2004). *p -adic Hodge theory and values of zeta functions of modular forms*. Astérisque **295**, 117-290.
- Khare, C. and Wintenberger, J.-P. (2009). *Serre's modularity conjecture. I, II*. Invent. Math. **178** (3), 485-586 and 587-654.
- Kobayashi, S. (2013). *The p -adic Gross-Zagier formula for elliptic curves at supersingular primes*. Invent. Math. **191** (3), 527-629.
- Kolyvagin, V. A. (1990). *Euler systems*. In: *The Grothendieck Festschrift, Vol. II*, Progress in Mathematics **87**, Birkhäuser, 435-483.
- Liu, Y., Tian, Y., Xiao, L., Zhang, W., and Zhu, X. (2022). *On the Beilinson-Bloch-Kato conjecture for Rankin-Selberg motives*. Invent. Math. **228** (1), 107-375. (Not a source for an average- BSD proportion.)

- Mazur, B. and Rubin, K. (2004). *Kolyvagin systems*. Memoirs of the American Mathematical Society **168** (799).
- Mazur, B., Tate, J., and Teitelbaum, J. (1986). *On p -adic analogues of the conjectures of Birch and Swinnerton-Dyer*. Invent. Math. **84** (1), 1-48.
- Mordell, L. J. (1922). *On the rational solutions of the indeterminate equations of the third and fourth degrees*. Proc. Cambridge Philos. Soc. **21**, 179-192.
- Nekovář, J. (1995). *On the p -adic height of Heegner cycles*. Math. Ann. **302** (4), 609-686.
- Perrin-Riou, B. (1993). *Fonctions L p -adiques d'une courbe elliptique et points rationnels*. Ann. Inst. Fourier (Grenoble) **43** (4), 945-995.
- Rubin, K. (1991). *The "main conjectures" of Iwasawa theory for imaginary quadratic fields*. Invent. Math. **103** (1), 25-68.
- Rubin, K. (2000). *Euler Systems*. Annals of Mathematics Studies **147**, Princeton University Press.
- Silverman, J. H. (2009). *The Arithmetic of Elliptic Curves*. 2nd ed. Graduate Texts in Mathematics **106**, Springer.
- Skinner, C. and Urban, E. (2014). *The Iwasawa main conjectures for GL_2* . Invent. Math. **195** (1), 1-277.
- Skinner, C. and Zhang, W. (2019). *Indivisibility of Heegner points in the multiplicative case*. Preprint; see also Skinner, C., *A converse to a theorem of Gross, Zagier, and Kolyvagin*, Ann. of Math. (2) **191** (2), 329-354 (2020).
- Tate, J. (1974). *The arithmetic of elliptic curves*. Invent. Math. **23**, 179-206. (Finiteness of III conjectured as part of the general picture.)
- Taylor, R. and Wiles, A. (1995). *Ring-theoretic properties of certain Hecke algebras*. Ann. of Math. (2) **141** (3), 553-572.
- The LMFDB Collaboration (2026). *The L -functions and Modular Forms Database*. <https://www.lmfdb.org> [Online; accessed 21 July 2026]. Analytic-rank, root-number, Tamagawa-number, torsion, and Shafarevich-Tate data for all worked curves cited in this paper (Cremona labels 11a1, 14a1, 27a1, 32a2, 37a1, 43a1, 389a1, 5077a1).
- The mathlib Community (2020). *The Lean mathematical library*. In: *Proceedings of CPP 2020*, ACM, 367-381.
- Tian, Y. (2012). *Congruent numbers and Heegner points*. Cambridge J. Math. **2** (1), 117-161. (Earlier preprint version: *Congruent numbers with many prime factors*.)
- Tunnell, J. (1983). *A classical Diophantine problem and modular forms of weight $3/2$* . Invent. Math. **72** (2), 323-334.
- Weil, A. (1928). *L'arithmétique sur les courbes algébriques*. Acta Math. **52** (1), 281-315.
- Wiles, A. (1995). *Modular elliptic curves and Fermat's last theorem*. Ann. of Math. (2) **141** (3), 443-551.
- Yun, Z. and Zhang, W. (2017). *Shtukas and the Taylor expansion of L -functions*. Ann. of Math. (2) **186** (3), 767-911.

Cremona, J. E. (1997). *Algorithms for Modular Elliptic Curves*. 2nd ed., Cambridge University Press. (The mwrank computer program: <https://johncremona.github.io/mwrank/>.)