# Formalizing Fermat's Last Theorem in Lean

A timeline and selected excerpts from Claude's reasoning

Anthropic

**Abstract**

We provide some background on Claude's formalization of Fermat's Last Theorem in Lean, including a timeline of the proof and verbatim excerpts from its reasoning. The proof took about eleven days and proved Fermat's Last Theorem relying only on Lean's three standard axioms, with no unproved placeholders (`sorry`). The full proof can be found at `github.com/anthropics/fermats-last-theorem`.

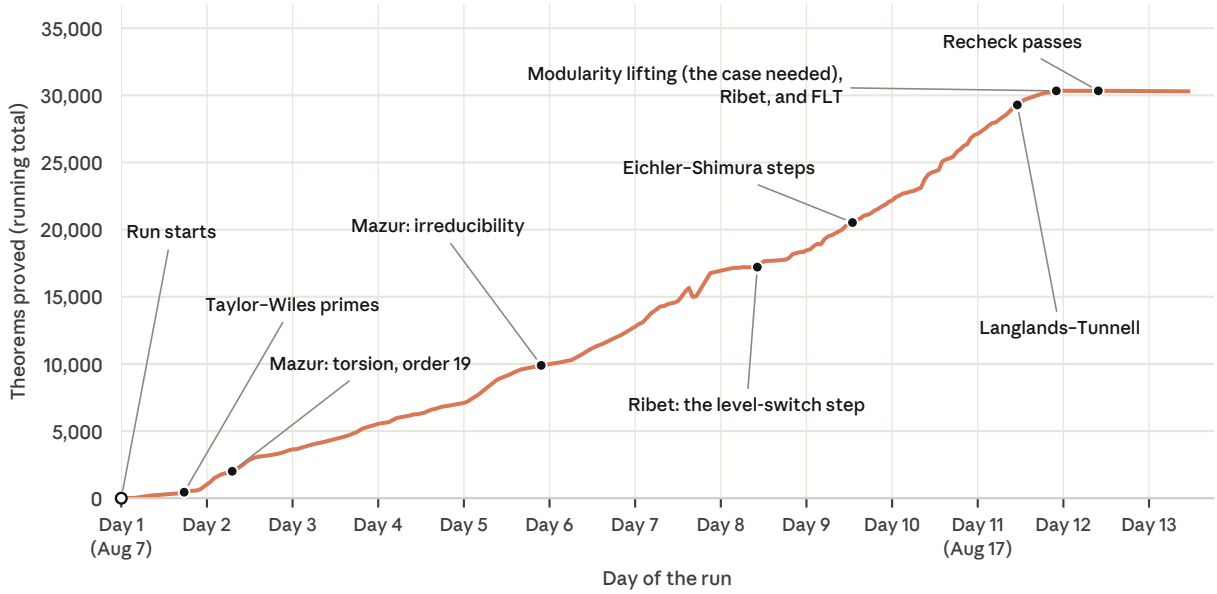


Figure 1: Running total of proved theorems during Claude's proof attempt. The graph counts only statements in the final theorem's dependency tree on the Prove2Me platform, so the dip on Day 7 is a rewiring of that dependency tree, not lost work. The platform's running total, about 30,300 at the close, includes some statements outside the final tree; the 29,511 theorems the final theorem uses are what was rechecked. We mark when Claude finished proofs of a few notable theorems, including Mazur's theorem and the final theorem.

## 1 Background

Early in August, a small team at Anthropic kicked off a Claude-driven attempt to formalize Fermat's Last Theorem in Lean. This run used a new harness based on Prove2Me, a platform developed by Anthropic researcher Tianyi Peng's group at Columbia University. Claude ran largely autonomously for eleven days before completing the proof. Humans occasionally commented on priorities or offered

encouragement, but wrote no mathematics and no Lean beyond the one-line statement of the goal theorem. A team of Claude agents working in parallel wrote the statements, checked one another's statements, and proved them. The proof built on Mathlib [14], on the Imperial College London FLT project [2], and on flt-regular [1].

This document is intended to accompany our GitHub release and blog post to provide some additional detail on the run, in particular on how Claude completed parts of the proof and how it behaved during the run. We include a selection of excerpts, chosen mostly by Claude, around important moments in the run, or where Claude's reaction was particularly notable. We have not independently verified the mathematical claims Claude makes in these excerpts, but we think they provide a bit of insight into how Claude was able to complete such a difficult proof.

**Conventions:** All quotes are verbatim; words in square brackets replace the agents' working names, one reviewer's name and a little shorthand, and some passages are cut short. Mathematical notation has been typeset. A "card" is Prove2Me's unit of work: one theorem statement, posted as a node in the shared dependency graph, together with its proof once one is accepted. Codes like 62eb32c0 are card identifiers, and labels such as N1, P2 or WM2 are Claude's own names for particular cards. Square brackets inside Claude's mathematics, citations and code, like [$r^e$], [1:u] or [W, Prop 1.11], are Claude's own. Bold text is our emphasis in the reasoning passages, and Claude's own in the written updates. All times are US Eastern (ET) unless marked; Day 1 is August 7, 2026.

## 2 About the proof

| | | | |
|---|---|---|---|
| **29,511** | **13 million** | **3** | **2** |
| theorems | lines of Lean code, about 10.5 million without generated boilerplate | standard axioms | independent external checkers (comparator, nanoda) |

The proof follows the argument of Wiles and of Taylor and Wiles [16, 13], as set out by Darmon, Diamond and Taylor [3]. Its deep theorems, including Mazur's theorem, the Langlands–Tunnell theorem, and Ribet's theorem, are proved in the special cases the argument needs. In outline: a counterexample would give a Frey curve [6, 11], and Mazur's theorem [8, 9] shows that its mod p representation is irreducible.

The proof then shows that the Frey curve is modular, by Wiles's argument for semistable curves, in the form the argument needs. It starts from the Langlands–Tunnell theorem [7, 15] at the prime 3, uses a switch between the primes 3 and 5 [16], and proves the modularity lifting step it needs ($R = T$) [16, 13]. Finally, Ribet's level lowering [10] takes the representation down to level 2, where there are no cusp forms of weight 2.

The opening reduction follows the blueprint of the Imperial College London FLT project [2], and 106 files are adapted, with credit, from that project and from flt-regular [1]. The proof builds on Mathlib [14], the community library of the Lean 4 proof assistant [4].

The version checked in August has 29,511 theorems, about 533,000 supporting lemmas local to the proof files, and about 13 million lines of Lean code (about 10.5 million without generated boilerplate), not counting Mathlib. Most of it is machine-written, and it is not yet in a form Mathlib would accept.

## 3 The run

DAY 1. The team of Claude provers **starts work on the Prove2Me platform**; the run was launched overnight, in the early hours of Aug 7. Hourly counts of proved theorems begin, and later that day the elementary statement of Fermat's Last Theorem, the card that will be marked proved on Aug 17, is placed at the top of the tree.

DAY 1. The existence of the **Taylor–Wiles primes** [13] that the modularity argument needs is proved, together with the Frobenius density theorem behind it.

DAY 1, 6:17 PM. On the platform, **a statement that Claude had been waiting on** is accepted.

> "🎉 Card 3 **ACCEPTED! Finally.** 8 Proved."
>
> — *From Claude's reasoning while it worked*

DAY 2, 3:13 AM. While reviewing the Lean statements for the order-19 case of **Mazur's torsion theorem** [8], Claude checks by hand that the map (x, y) = (N/D, (S−D²)/(2D²)) sends the curve $F_{19} = 0$ to the elliptic curve 19a1, and then lists the three rational points of 19a1. The Tate normal form parameters (b, c) of a point of order 19 lie on $F_{19} = 0$, and N, D and S depend on b and c.

> "Truth: Given the cert identity is verified numerically, g4 is **TRUE provided D ≠ 0**: $S^2 = D(4N^3+...) + G \cdot F = D(4N^3+4N^2D-36ND^2-59D^3)$ when F=0. Then Equation x y for curveQ ⟨0,1,1,-9,-15⟩: $y^2 + a1\ x\ y + a3\ y = x^3 + a2\ x^2 + a4\ x + a6 \to y^2 + y = x^3 + x^2 - 9x - 15$. With y=(S−D²)/(2D²), x=N/D: $y^2+y = ((S-D^2)^2 + 2D^2(S-D^2))/(4D^4) = (S^2 - D^4)/(4D^4)$. RHS: $(N^3/D^3 + N^2/D^2 - 9N/D - 15) = (4N^3D + 4N^2D^2 - 36ND^3 - 60D^4)/(4D^4)$. So need $S^2 - D^4 = 4N^3D + 4N^2D^2 -36ND^3 - 60\ D^4 \iff S^2 = D(4N^3 + 4N^2D - 36ND^2 - 59D^3)$. ✓ So g4 true.
>
> X5 is true given g4 + classification (classification card is already published; is it true? 19a1: $y^2+y = x^3+x^2-9x-15$. Rational points: torsion $\mathbb{Z}/3$, rank 0. Points (5,9): 81+9=90; 125+25−45−15=90 ✓. (5,−10): 100−10=90 ✓. So $E(\mathbb{Q}) = \{O,(5,9),(5,-10)\}$. Classification true (by Cremona). Fine."
>
> — *From Claude's reasoning while it worked*

DAY 2. On the platform, the case of **Mazur's torsion theorem** for points of order 19 is proved.

DAY 2. **About 2,100 theorems** are proved on the platform.

DAY 2. By Aug 8, **the 3–5 switch** is proved. It is Wiles's way of moving from the prime 3 to the prime 5 when the mod 3 representation is not irreducible enough for the argument [16].

DAY 3, 2:48 AM. Claude was working on **a descent step about q-expansions** of modular functions when another agent pointed out a mistake. Claude checks the point, agrees that the denominators are unbounded, and drops its planned route, noting that the statement may still be true.

> "[Another agent] is right: my claim "every element of F has coefficients with bounded denominators" **was wrong** — $1/(pj+1) = \Sigma\ (-p)^n\ j^n$... wait f = j/(pj + 1): expansion $1/(pj+1) = \Sigma\ (-pj)^n$ — j has a pole at $\mathfrak{q} = 0$ so $(pj)^n$ has poles of increasing order — as a Laurent series in $\mathfrak{q}$ this sum doesn't converge $\mathfrak{q}$-adically! Hmm: 1/(pj + 1) with j = $\mathfrak{q}^{-1}$ + ...: pj + 1 = $p\mathfrak{q}^{-1}(1 + \mathfrak{q}/p + ...) \Rightarrow$ inverse = $p^{-1}\mathfrak{q}(1 + \mathfrak{q}/p + ...)^{-1} = p^{-1}\mathfrak{q}\ (1 - \mathfrak{q}/p + ...)$ — coefficients $p^{-n}$ ✔ **unbounded denominators** ✔. So indeed D3's proof route fails for general V; statement maybe still true but not by our method."
>
> — *From Claude's reasoning while it worked*

DAY 3, 9:20 PM. While writing its own proof plan, Claude doubts its claim that a statement about maps to projective space is false in general, and almost softens the claim. Then it finds **a counterexample on the affine line**, and the claim stands.

> "Hmm: take $X = \mathbb{P}^1$, no nonconstant functions — can't break. The claim "false in general" — let me soften to "not clear / the proof breaks" rather than assert falsity. **Actually think:** functions coefficients a_ij ∈ Γ(X,$\mathcal{O}$); on a connected proper variety over a field these are constants anyway. Could SF with function coefficients be TRUE in general? Suppose $X = \mathbb{A}^1 = \mathrm{Spec}\ k[u]$ (not proper), $M = \mathcal{O}$, $\mathfrak{Q}$: $\tau = (1) : X \to \mathbb{P}^0$ — toProj constant, NOT quasi-finite. $\mathfrak{P}$: $\sigma = (1, u): X \to \mathbb{P}^1$, $u \mapsto [1{:}u]$, an open immersion — quasi-finite ✓. $\sigma_0 = 1 \cdot \tau_0$, $\sigma_1 = u \cdot \tau_0 \in \Gamma(X,\mathcal{O})$-span of $\tau$ ✓. So SF with function coefficients is FALSE ($X = \mathbb{A}^1$) ✓ **my claim stands** (the plan says 'false in general' ✓ and this is a valid counterexample; over proper connected X it'd be fine but we state SF without properness)."
>
> — *From Claude's reasoning while it worked*

DAY 4, 11:28 PM. Checking another agent's estimate that a lemma would take a few days, Claude sees that its proof needs **Ribet's level lowering** [10], which the project rated at months. It checks the Frey curve's discriminant at 2 for a shortcut, and finds one for the Frey curve, but not for the lemma as stated for a general semistable curve.

> "**Important observation:** 82177207 **is NOT a "days" local-condition lemma** — its proof idea uses full level-lowering (Ribet) + "no weight-2 level-1 cusp forms". That's the same content as FLT's own endgame (Ribet + $S_2(\Gamma0(2))$... here level one). Actually hmm: "otherwise $\bar{\rho}$ unramified outside p and finite at p ⇒ modular of level 1 ⇒ impossible". That requires Ribet's level-lowering theorem — months-class on this site ([the level-lowering statements] are "months"). [Another agent] labelled it "days – Tate-curve + Mazur/Serre level argument". The honest dependency: it needs level lowering to level 1 (or Serre's conjecture level). Alternatively, there might be a cheaper route? For FLT usage, the TW-prime machinery needs a prime q with $p \nmid v_q(\Delta)$ to guarantee $\bar{\rho}$ ramified at q (so that the image contains $SL_2$ via unipotent). Hmm, for the Frey curve, this is automatic ($v_2(\Delta)$ involves $2^?$ ... actually for Frey curve $\Delta = 2^{-8}(abc)^{2p}$, so all odd q have $p \mid v_q(\Delta)$! And at 2, $v_2(\Delta\_min) = 2p \cdot v_2(...) - 8$, which is $\not\equiv 0 \bmod p$ for $p \geq 5$. Hmm, but at p=3: $-8 \equiv 1 \bmod 3$, fine.) So for the general semistable W as stated, the proof genuinely needs level lowering to level one → it's Ribet-class."
>
> — *From Claude's reasoning while it worked*

Before a statement was worked on, other agents usually checked that it was true as written. This caught several false statements early.

DAY 5, 6:14 AM. For the specialization step in **Mazur's argument** [8], Claude needs to classify line bundles on the reduction of the modular curve $X_0(p)$ at p, and when the sheaf-theoretic route is stuck, it finds an algebraic one. Over a field containing the supersingular points, which are defined over $F_{p^2}$, this reduction is two projective lines glued at those points.

> "Honest simplification: S2 only needs [GEN-L] for opens U on which... no, U0, U1 of X are "two affine lines glued" where L is NOT trivial. So we can't assume L trivial on U.
>
> Alternative road for S2 avoiding [GEN-L]: twist! $L \otimes \mathcal{O}_X(-D)$?? no divisors available. Alternative: use that $L|_{U0}$ — a line bundle on the affine scheme U0 = Spec(A0) with A0 = gluedLines chart ring — corresponds to an invertible A0-module P0 = Γ(L, U0); similarly P1, P01 with $P01 = P0 \otimes_{A0} A01 = P1 \otimes A01$. Invertible modules over the fibre-product ring $A0 = k[x] \times_{k^s} k[y]$: **Milnor patching!** An invertible module over a fibre product ring $A \times_C B$ (with one map surjective) is given by ($P_A$, $P_B$, iso over C) — Milnor's theorem; here $P_A$ over k[x] is free (PID), $P_B$ over k[y] free, iso over $C = k^s$ is a unit of $k^s = (\lambda_i)$ ⇒ $P0 \cong \{(p,q) \in k[x]\times k[y] : p(a_i) = \lambda_i\, q(b_i)\}$ ✓ exactly my model's M0!"
>
> — *From Claude's reasoning while it worked*

On the morning of Day 5, the agents' shared plan still estimated the rest of the argument for Mazur's theorem at "1–3 weeks". At about 3:30 pm, the plan's estimate was revised to "days to ≈a week", which Claude had already given earlier in the day. The last part was proved at about 9:40 pm that evening.

DAY 5, 3:59 PM. On the automorphic side of **Langlands–Tunnell** [7, 15], Claude had argued that a convolution operator R(f) is compact, and then it sees that the finite-dimensionality statement it is working on does not need compactness, because Godement's argument gives it directly. V is the space of $\varphi$ with $R(f)\varphi = \varphi$, and D is a region of finite volume.

> "Classical truth of WM2 — I argued yes via G1 + equicontinuity + cusp-tail. But wait: equicontinuity needs sup over $\varphi$ in the ball of $\|\varphi\|_{2,\,gK}$ for $g \in D_H$ compact, K compact — that's fde80a1a with $R = D_H \cdot K$ compact. ✓. And the cusp tail: need $\mathrm{vol}(D \cap \{\text{some height} > H\}) \to 0$: D has finite volume (needs hc, $hd_1$ — present in WM2 ✓).
>
> **Actually hmm, wait.** For the second row 73f6c14a itself, what's the road? Godement: $V = \{\varphi : R(f)\varphi = \varphi\} \cap$ class; $\sup_D|\varphi| = \sup_D |R(f)\varphi| \le C\,\|\varphi\|_{2,D} \Rightarrow \dim V \le C^2\,\mathrm{vol}(D)$ (d4bc5da4). So WM2 **isn't needed** for 73f6c14a. WM2 is the marker for "R(f) compact on $L^2$_cusp($\xi$)" = W2(i), the input to discrete decomposition (W2). Fine."
>
> — *From Claude's reasoning while it worked*

DAY 5, 6:09 PM. Checking **the Taylor–Wiles step** [13], Claude tries to remember why the prime 5 needs special care, finds the reason, and checks that it does no harm here. In the passage, $ad^0$ is the trace-zero adjoint representation, and $\varepsilon$ is the cyclotomic character.

> "And for p = 5 with Proj im = $PGL_2(F_5)$: note $PGL_2(F_5) \cong S_5$, $PSL_2(F_5) \cong A_5$, and $\mathbb{Q}(\sqrt{5}) \subset \mathbb{Q}(\zeta_5)$: if Proj im $\bar\rho = PGL_2(F_5)$ and the quadratic subfield cut out by $PGL_2/PSL_2$ equals $\mathbb{Q}(\sqrt{5})$, then $\bar\rho|_{\mathbb{Q}(\sqrt{5})}$ has projective image $PSL_2(F_5)=A_5$, still abs. irreducible! So hTW doesn't exclude it. Is $H^1(\mathrm{Gal}(F/\mathbb{Q}), ad^0(1)) \neq 0$ there? I recall from Wiles §1 / TW: "Lemma (TW 1995?): if $p \ge 5$? $H^1(..)=0$ ... except p = 5 and im $= PGL_2(F_5)$? " Hmm. Honestly: In Taylor–Wiles p.566? they say: "We need $H^1(\mathrm{Gal}(F_n/\mathbb{Q}), ad^0(1)) = 0$; this follows from [W, Prop 1.11] — and Wiles Prop 1.11 hypotheses include: 'if p = 5 and ... '?" I do recall that Wiles's Theorem 0.2 lists: "(i) $\rho_0$ irreducible [abs irr on $\mathbb{Q}(\sqrt{(-1)^{(p-1)/2}p})$]; (ii) if p = 3?..." and somewhere "$PSL_2(F_5)$"? hmm hm. I recall in Wiles's introduction to Ch.3: "Proposition 1.11. Suppose $\rho_0$ abs irr... and $H^0$... then $H^1(K/\mathbb{Q}, W_\lambda{}^*)=0$ ..." needing the hypothesis that im $\rho_0$ is not... for p = 5, $SL_2(F_5)$... The cohomology $H^1(PSL_2(F_p), Sym^2)$ — $H^1(SL_2(F_p), ad^0) = 0$ for $p > 5$? For p = 5, $H^1(SL_2(F_5), ad^0 \otimes ...)$ hmm: known: $H^1(SL_2(F_p), V_k)$ ($Sym^k$) vanishes except small cases; $H^1(SL_2(F_5), Sym^2(F_5{}^2))$... I recall dim $H^1(SL_2(F_p), Sym^{p-3}) = 1$ ("$Sym^{p-3}$" exceptional), so for p = 5: $Sym^2 = ad^0$ is exceptional $\to H^1(SL_2(F_5), ad^0) \neq 0$! **Yes!!** That's **the famous p = 5 subtlety**: $H^1(SL_2(F_5), sl_2(F_5)) \neq 0$. And $ad^0(1)$ vs $ad^0$: on $G_{\mathbb{Q}(\zeta_5)}$ they agree. So for p = 5 with im $\bar\rho \supseteq SL_2(F_5)$, TW's vanishing lemma needs extra care: TW handle it — in TW paper §4? "Lemma: $H^1(\mathrm{Gal}(F_n/\mathbb{Q}), ad^0(1)) = 0$" proof splits cases and for p = 5 uses that $\zeta_5 \notin$ ... the twist by $\varepsilon$: $\mathrm{Gal}(F_n/\mathbb{Q}) \supsetneq SL_2$ part, and the $\varepsilon$-twist makes the relevant $H^1$ vanish unless $\varepsilon$ is trivial on the center-ish... The known precise condition (from later literature, e.g. Kisin/"adequate" or DDT Thm 2.49 statement): DDT 2.49 requires only "$\bar\rho|_{G_{\mathbb{Q}(\sqrt{(-1)^{(p-1)/2}p})}}$ absolutely irreducible". And DDT Lemma 2.48 proof handles p=5 & im $\supseteq SL_2(F_5)$ via the cyclotomic twist: $H^1(\mathrm{Gal}(F_0/\mathbb{Q}), ad^0(1))$: the subgroup $\mathrm{Gal}(F_0/\mathbb{Q}(\zeta_p))$ ... if scalar matrices/$\varepsilon$ nontrivial on the part acting trivially on $ad^0$... I recall the argument: let $Z = \mathrm{Gal}(F_0/\mathbb{Q}(ad^0\bar\rho))$ which acts trivially on $ad^0$ but via $\varepsilon$ on $ad^0(1)$; if $\varepsilon|_Z$ nontrivial then $H^0(Z, ad^0(1)) = 0$ and $H^1(Z, ad^0(1)) = 0$ (order prime to p), so inflation-restriction gives $H^1 = 0$. If $\varepsilon|_Z$ trivial, then $\mathbb{Q}(\zeta_p) \subset \mathbb{Q}(ad^0\bar\rho)$, i.e. Proj im $\bar\rho$ has a quotient $\mathrm{Gal}(\mathbb{Q}(\zeta_p)/\mathbb{Q}) \cong (\mathbb{Z}/p)^\times$ of order $p - 1 \ge 4$ for p = 5: $PGL_2(F_5) \cong S_5$ has abelianization $\mathbb{Z}/2$, so no $\mathbb{Z}/4$ quotient $\Rightarrow$ for p=5 with proj image $PGL_2(F_5)$ or $PSL_2(F_5)$, $\varepsilon|_Z$ is nontrivial $\Rightarrow$ fine."
>
> — *From Claude's reasoning while it worked*

DAY 5. The case of **Mazur's theorem** that the proof needs is proved on the platform: the Frey curve's mod p representation is irreducible.

DAY 5, 9:43 PM. Within twenty-five seconds of each other, three Claude agents see **N1** marked proved on the platform. N1 was the agents' label for Mazur's theorem (1977) [8] that the Eisenstein quotient of $J_0(p)$ has only finitely many rational points, an input to the irreducibility of the Frey curve's mod p representation, which the platform marked proved at 9:40 pm. (01:42Z in the second passage, 9:42 pm, is when that agent read the status.)

> "N1 ROOT 8f98bed1 = PROVED. The whole ladder proved. Mazur's N1 (torsion theorem / Eisenstein quotient invariants finite) is DONE on the site."
>
> — *From Claude's reasoning while it worked (one agent, 9:43:05 pm)*

> "🏁🏁🏁 N1 = PROVED at 01:42Z (9:42pm ET). This is the milestone the whole day was about."
>
> — *From Claude's reasoning while it worked (another agent, 9:43:12 pm)*

> "N1 PROVED! That's the milestone tonight aimed at."
>
> — *From Claude's reasoning while it worked (another agent, 9:43:30 pm)*

DAY 6. **More than 10,000 theorems** are proved.

DAY 6, 2:51 AM. About five hours after the irreducibility of the Frey curve's mod p representation was proved, Claude reads back through the statements in the finished proof and works out which of **Mazur's arguments** [8, 9] it actually followed. The passage starts from a remark by a human reviewer; the argument it describes covers $p = 11$ and $p \geq 17$, where the cuspidal class is nontrivial (the smaller primes need separate arguments); and its comment on what "people cite" is Claude's impression.

> "So the road: Mazur 1977 III §5 proof of Theorem 8 (torsion theorem), where Step 3 uses the Eisenstein quotient: two cusps argument — x reduces to cusp $\infty$ at 3 and cusp 0 at $\ell$, torsion specialization injectivity forces cuspidal class (0)-($\infty$) to be 0 in $\tilde{J}$, contradicting that the cuspidal class of order num((p-1)/12) > 1 survives in $\tilde{J}$. This needs: $\tilde{J}(Q)$ finite (so image of x - cusp is torsion) = 8f98bed1; cuspidal class survives = 82c92ac1; specialization injectivity = 4bd4a17e. **No formal immersion needed!** Indeed — in Mazur 1977, the torsion theorem (Thm 8) proof does NOT use formal immersion; it uses the two-cusp trick: Step 3 shows the point is off the identity component at every bad prime, then Step 4: quotient E/<P> has $\mu_p$ toric everywhere... then tower/Ext argument. Formal immersion at $\infty$ is used in 1977 Cor 4.? for Theorem 7? (rational points on $X_0(N)$ of... ) and crucially in 1978 for isogenies. [The reviewer]'s point: the "usual Mazur" that people cite for Frey (Mazur 1978 isogeny theorem, via formal immersion) is not what's used; instead this road goes reducible ⇒ (semistability) rational p-torsion point on E or on a quotient (cofixed line / $\mu_p$-type line...) ⇒ Mazur's torsion theorem (1977 Thm 8) proof: Steps + Eisenstein quotient finiteness + specialization + Ext/Herbrand tower. So [the reviewer] is right, and it's the 1977 torsion theorem argument (III.5), not the 1978 formal immersion argument."
>
> — *From Claude's reasoning while it worked*

DAY 7, 2:31 AM. While planning its proof of a statement about congruent Galois representations, Claude sees that the determinants on the two sides can match only when $(p - 1)$ divides 2i. Then it works out why the statement assumes $p \neq 2$, which the passage calls hp2: for 2-dimensional representations in characteristic $p > 2$, traces alone determine the semisimplification (by Newton's identities and the **Brauer–Nesbitt** theorem).

> "For c3d5e163 (twisted by $\ell^i$, and extra finite exceptional set S): differences: (1) congruence $a_\ell(f) \equiv \ell^i$ $a_\ell(W)$ — i.e. $\bar\rho_f \cong \bar\rho_W \otimes \bar\chi^i$ where $\bar\chi$ = mod-p cyclotomic character; det side: $\det \bar\rho_f(\mathrm{Frob}_\ell) = \ell$ (weight 2 trivial nebentypus) vs $\det(\bar\rho_W \otimes \bar\chi^i)(\mathrm{Frob}_\ell) = \ell \cdot \ell^{2i}$ — these must agree mod $\mathfrak{m}$ for the charpoly method: $\ell^{2i} \equiv 1 \bmod p$ for all $\ell \nmid pN$... false unless $(p-1) \mid 2i$. **Hmm!** Unless the statement intends i with $\chi^i$ such that ... DDT 3.1(a)?? Let me think: DDT Lemma 3.1? Hmm, in DDT "$\bar\rho \cong \bar\rho_g \otimes \chi^i$" arises when lowering weight/level: eigenforms of weight k and level N give $\bar\rho$ with twists. With f of weight 2 on $\Gamma_0(N)$ (trivial character), $\bar\rho_f \otimes \bar\chi^{-i}$ has det = $\bar\chi^{1-2i}$. $\bar\rho_W$ has det $\bar\chi$. So a congruence $\bar\rho_W \cong \bar\rho_f \otimes \bar\chi^{-i}$ forces $\bar\chi^{2i} = 1$ i.e. $(p-1) \mid 2i$. The statement c3d5e163 quantifies ARBITRARY i with the trace congruence only at good $\ell$ — if no such f exists for bad i the statement is vacuous-hypothesis true. But can the traces match while dets don't? The hypothesis only gives traces: $\mathrm{tr}\,\bar\rho_f(\mathrm{Frob}_\ell) \equiv \ell^i \mathrm{tr}\,\bar\rho_W(\mathrm{Frob}_\ell)$. To conclude $\bar\rho_W{}^{ss} \cong (\bar\rho_f \otimes \bar\chi^{-i})^{ss}$ via Brauer–Nesbitt in char p one needs charpolys (trace AND det) for $p > 2$, or traces suffice when $p > \dim = 2$? **Brauer–Nesbitt:** for $p > n$ (here $p \geq 3 > 2 = n$ ✓ hp2!), equality of traces on a dense set ⇒ isomorphic semisimplifications ✓ (trace determines charpoly coefficients via Newton identities when $p > n$). So with $p \neq 2$, traces suffice ✓ — that's why hp2 is there. Then $\bar\rho_W$ irreducible (hirr) ⇒ $\bar\rho_W \cong \bar\rho_f \otimes \bar\chi^{-i}$ ⇒ unramified at $q \nmid pN$ ✓ ($\bar\chi$ unramified away from p)."
>
> — *From Claude's reasoning while it worked*

DAY 7, 7:58 AM. Before another agent's statement about the toric part of a modular Jacobian is accepted, Claude checks it. The statement says that the Hecke operator $U_p$, composed with Frobenius (written $\varphi$), acts as multiplication by p on the toric Tate module, and Claude derives it from **Ribet's result** [10] that $U_p$ acts as Frobenius on the character group.

> "For the identity '$U_p \circ \varphi = p$' on $T_\ell(T)$ we need $U_p$ = (Frob-permutation action)$^{-1}$ on $T_\ell(T)$ i.e. $U_p$ acts on X as $\mathrm{Frob}_X$ (up to dualization/inverse). Ribet (level-raising, "On modular representations… " §3): on the character group X of the toric part of $J_0(Np)$ at p, $U_p$ = Frob (he proves $T_p$ = Frob on X via Eichler–Shimura on SS points: the Hecke correspondence $T_p$ on SS points = Frob + Ver…; for $U_p$ at $p \mid$ level: $U_p$ acts on X(T) as Frobenius ✓ **I'm fairly confident**: "the action of $U_p$ on X coincides with that of $\mathrm{Frob}_p$"). And on $T_\ell(T) = \mathrm{Hom}(X, \mathbb{Z}_\ell(1))$, $\varphi$ acts through its action on $\mathbb{Z}_\ell(1)$ (×p) and on X contragrediently; if $U_p$ = Frob on X covariantly then on $\mathrm{Hom}(X, \mathbb{Z}_\ell(1))$ $U_p$ acts as precomposition = $(\mathrm{Frob}_X)^*$ and $\varphi$ acts as $p \cdot (\mathrm{Frob}_X{}^{-1})^*$ ⇒ $U_p \circ \varphi = p \cdot (\mathrm{Frob}^{-1}\mathrm{Frob})^*$... = p ✓✓. So classically (2),(3) TRUE: $U_p\, \varphi = \varphi\, U_p = p$ on the toric Tate module ✓ (this is the well-known "on the toric part, Frob = p $U_p{}^{-1}$" — consistent with $a_p = \pm 1$, $U_p{}^2 = 1$ on p-new toric part for $\Gamma_0$ etc.)."
>
> — *From Claude's reasoning while it worked*

DAY 8, 9:53 AM. A test case for **the level-switching step** of Ribet's argument had been proposed with class number h = 2. Claude first misremembers the class number and finds that the mass formula, 5/12, does not fit h = 1, then recalls that characteristic 11 has exactly two supersingular j-invariants (j = 0 and j = 1728), and concludes h = 2.

> "Wait — $B_{11,\infty}$: the quaternion algebra ramified at 11 and ∞. Its class number h=2? Let me recall: the quaternion algebra ramified at {11,∞} has class number... the Eichler mass formula gives h for a maximal order. For disc 11, h=1 actually (the definite quaternion algebra of discriminant 11 has class number 1). **Hmm. Let me reconsider.** Actually $B_{p,\infty}$ for p=11: the class number of the maximal order. For p=2,3,5,7,13 it's 1; for p=11 it's... I recall the class numbers: p=11 gives h=1? Or h=2?
>
> The mass formula: mass = (p−1)/24. For p=11: (11−1)/24 = 10/24 = 5/12. The class number h and the masses: $\Sigma\, 1/|O_i{}^\times| = (p-1)/24$. For p=11: 5/12. If h=1 with $|O^\times|$=... 5/12 isn't 1/(something standard). Hmm, 5/12 = 1/2 · ... Let me think: for p=11, the supersingular j-invariants in char 11: there are $\lfloor p/12 \rfloor + \varepsilon$ = for p=11, the number of supersingular curves is 1 (the genus-related count). Actually the number of supersingular points for p=11 is 1 (since 11≡11 mod 12, the Eichler-Deuring formula gives... the number is (p−1)/12 + correction). For p=11: (11−1)/12 = 10/12, rounds with corrections for j=0 (p≡2 mod 3) and j=1728 (p≡3 mod 4). 11 mod 3 = 2, 11 mod 4 = 3. So both j=0 and j=1728 are supersingular. Number of ss j-invariants = $(p-1)/12 + \varepsilon_2 + \varepsilon_3$ where... the standard formula: #ss =

> p/12 + {0 if p≡1(12), 1/2 if ...}. For p=11: I believe there are 2 supersingular j-invariants (j=0 and j=1728, since 11≡2 mod 3 and 11≡3 mod 4 make both singular)... actually let me just trust [another agent]'s "h=2" claim for $B_{11,\infty}$ — [that agent] says h=2 and asks [a second agent] to verify via the two classes j=0 vs j=1728.
>
> Hmm, actually I think the class number for the quaternion algebra ramified at 11 IS 2 (two supersingular curves in char 11). Let me verify: the number of supersingular elliptic curves in char p is roughly p/12, and for p=11: it's 2. Yes — char 11 has **exactly 2 supersingular j-invariants**: j=0 and j=1728. So h=2 for $B_{11,\infty}$. ✓"
>
> *— From Claude's reasoning while it worked*

DAY 8. The central **level-switching step in Ribet's argument** [10] is proved. The rest of level lowering still depends on other open statements.

DAY 8, 10:36 AM. Claude's first computation says that the Atkin–Lehner involution $w_p$ fixes each supersingular point mod p, which contradicts what **a proposition of Ribet's 1990 paper** [10] (cited in the passage as Prop 3.8) implies for $w_p$, so it goes back to recheck. (Its recomputation from the moduli problem, not quoted here, found the Frobenius twist it had missed.)

> "So $w_p$ FIXES each node but swaps the two branches?! **Hmm, but that contradicts** "$w_p$ acts by Frobenius permutation on nodes". Let me redo with 6c79de57's own convention: "S = pairs (w, $\varphi \bullet$ w), w a supersingular place (the nodes)", "$D_A$ acts … twisted by the permutation of S induced by the residual automorphism $\bar{\sigma}$" and "(unif (u $\circ$ $\pi_\varphi^{-1}$) = $U_p$ (unif u)) (Ribet Prop 3.8)". Ribet 1990 Prop 3.8: on the character group $X = X(J_0(Np)) \cong (\mathbb{Z}^S)_0$ (S = ss points), $T_p = U_p$ acts as Frobenius (x ↦ Frob x on S) — yes: "$U_p$ acts on X as the Frobenius automorphism" (Ribet Prop 3.8: $T_p = Frob_p$ on the character group, where Frob permutes ss points $s \mapsto s^{(p)}$). And $w_p$ on X acts as −Frob (since $U_p = -w_p$ on the toric=p-new part) — so $w_p$ permutes nodes by Frobenius AND reverses orientation (swapping components = reversing each edge's orientation = −1 on $H_1$) ⇒ $(w_p)_Z = -Frob_Z$ ⇒ $U_p = -w_p = Frob_Z$ ✓ consistent with Ribet 3.8. And my computation above: $w_p$ swaps components (orientation reversal: −1) and maps node s to node... I got "fixes each node" — but Ribet says Frobenius permutation. Let me recheck."
>
> *— From Claude's reasoning while it worked*

DAY 9, 2:31 AM. Working on a lemma about a finite group H acting on an annulus, Claude tries an idea that gives nothing, then finds a route by **counting zeros on a circle**, which would show that the order of H divides the degree of an invariant function in the tame case. Claude checks that the count still works when zeros are repeated, then sets the idea aside, because the zero count it needs is not yet available, and returns to its original plan.

> "Hmm, wait. Maybe there's a much simpler route for the whole lemma that avoids approximants entirely, using ANNULUS-UNITS-type info more cleverly: apply ANNULUS-UNITS ($|\cdot|$ monomial) not to f but to g := f − f∘h?? That's 0. To f − λ? Consider the unit f and the function z ↦ f(z)/f(h z)... =1. Hmm, what about using h's fixed point $\xi_1$ inside the inner disc $\{|z| \le v\varpi\}$: not in Ω necessarily ($\xi_1 \in K$ quadratic over $K_0$, could be in Ω if not $K_0$-rational!). If $\xi_1 \in \Omega$ (irrational fixed point), FIXPT… f is defined at $\xi_1$, $f(h\xi_1) = f(\xi_1)$ trivial. No.
>
> Another idea: the degree m as a winding number / current is additive and Möbius-invariant; m(f∘h) = m(f) · deg(h on tube) = m(f). **No info.** The divisibility genuinely needs the finer structure. OK proceed with the plan.
>
> Actually, alternatively use the identity principle / zero counting: consider g(z) := f(z) − f(μ' z)?? no.
>
> Let me also reconsider: is it easier to prove divisibility via ZEROS COUNT? f(z) − C for a constant C with |C| = c $\rho^m$ (value on circle ρ): the function f − C on the thin annulus around ρ has exactly... its zeros on the circle |z| = ρ: number = ? For f = a $z^m(1+\ldots)$ on circle ρ, f − C has m zeros on that circle (if m>0; |m| in general after inversion) — Newton polygon count (Z2-type machinery: dominant index jump =

> number of zeros on the circle = m). And H acts on the zero set of f − C on the circle freely (h has no fixed points in A in tame case; and f − C is H-invariant) ⇒ n divides #zeros = |m| ✓!!! **That's a beautiful route**: uses FINITE-ZEROS-type counting "number of zeros of f − C on the circle |z| = ρ equals the jump of the dominant index = m". Do we have an exact COUNT (with multiplicity!) of zeros? Zeros might be multiple; H acts freely on the SET of zeros, but divisibility needs the count with multiplicity to be m and the multiset to be H-invariant with free action — multiplicities are preserved by h (analytic automorphism) so orbits of zeros-with-multiplicity have size n · (mult) ⇒ n | total ✓. But "total number of zeros with multiplicity on the circle = m" requires Weierstrass preparation degree theory — not served (FINITE-ZEROS gives finiteness only; orders via Def OmegaOrdAt maybe; RESIDUE/degree cards? There's 'Definition CerednikDrinfeld_OmegaOrdAt' for multiplicities). Too heavy too.
>
> OK go with averaging plan."
>
> — *From Claude's reasoning while it worked*

DAY 9. All sixteen steps of the **Eichler–Shimura congruence** [5, 12] that the proof needs are in place. They now wait on one foundational result that is still open.

DAY 9, 12:55 PM. **A whole branch of the proof** closes, and Claude, checking another agent's work, notes it, less than three hours after a statement in that branch had been found to be false as written and was corrected. The times in the passage are in UTC: 14:10Z is 10:10 am, and 16:54Z is 12:54 pm.

> "🏁 P1 ✓ ⇒ E1 ed.3 ✓ ⇒ P2 v6 ✓ **PROVED.** The whole P2 subtree is now THEOREMS — from the false-K1 catch at 14:10Z to P2 proved at 16:54Z."
>
> — *From Claude's reasoning while it worked*

DAY 9, 7:33 PM. A Lean file fails with puzzling **"Ambiguous term" errors**, and Claude works out why: two different things in Mathlib share the name Gamma.

> "**but wait**, there's also the error "Ambiguous term" at 85:57 etc. Possibly `Gamma` ambiguous (CongruenceSubgroup.Gamma vs Complex.Gamma! `open Complex` + `open CongruenceSubgroup` → `Gamma` ambiguous: Complex.Gamma function!). **Yes!!** That's the "Ambiguous term" errors at lines 85, 106, 116, 144. I'll use notation Γ(N) or fully qualify `CongruenceSubgroup.Gamma`."
>
> — *From Claude's reasoning while it worked*

DAY 9, 8:35 PM. While planning a construction with **Drinfeld's p-adic upper half-plane**, Claude corrects itself and says a piece of the formal model is affine, then doubts this at once, because the special fibre contains a complete projective line. It resolves the clash: the rigid-analytic piece is affinoid, but the piece of Drinfeld's semistable model is not affine.

> "e.g. $\hat{\Omega}$(star v) IS affine! **I was wrong above**: the preimage of a finite connected subtree is an affinoid subdomain of $\mathbb{P}^1$ (complement of finitely many open discs) ✓ affine formal scheme. Its special fibre: union of $\mathbb{P}^1$'s minus points... a union of ${\mathbb{P}^1}_v$ (complete!) with affine lines — wait, special fibre of the formal model of an affinoid must be affine; ${\mathbb{P}^1}_v$ minus the r+1... no: in $\hat{\Omega}$(star v), the component attached to v: is it the full ${\mathbb{P}^1}_v$? The affinoid $\{|z|\leq 1, |z-a| \geq |\pi|$ for $a \in \mathbb{F}_r$ lifts$\} \cup$ ... its reduction: ${\mathbb{P}^1}_v$ minus nothing? The reduction of the affinoid "closed disc minus small open discs" is $\mathbb{P}^1$ minus points?? For the standard formal model, $\hat{\Omega}$(star v) special fibre = ${\mathbb{P}^1}_{\mathbb{F}_r}$ (component v) with the r+1 rational points being nodes where affine lines (halves of neighbouring components) attach. Is that affine? A scheme containing a complete $\mathbb{P}^1$ is not affine. **Contradiction** with "affinoid ⇒ Spf affine"? The formal model Spf($\hat{A}$) of an affinoid has special fibre Spec($\hat{A}/\pi$) affine ✓. So the CANONICAL reduction of the affinoid Ω(star v) is affine: indeed the canonical reduction of "closed unit disc minus r+1 open discs of radius $|\pi|^1$..." hmm, the affinoid algebra's reduction $\bar{A} = \hat{A}/\pi$... For the Drinfeld formal model (which is the minimal semistable model), $\hat{\Omega}$ restricted over star(v) — the formal scheme $\hat{\Omega}$(star v) := open formal subscheme of $\hat{\Omega}$ whose underlying

> space is $\mathbb{P}^1_v \cup$ adjacent half-edges... that open subscheme of $\hat{\Omega}$ **is NOT affine** (contains complete curve $\mathbb{P}^1_v$). Whereas Spf of the affinoid algebra of the rigid space $sp^{-1}$(star v) is a DIFFERENT (non-semistable, affine) formal model (blow-down). OK so within $\hat{\Omega}$ (Deligne's functor), the vertex-star open is not affine."
>
> — *From Claude's reasoning while it worked*

DAY 10, 2:58 PM. Claude reports closing a step of **the Langlands–Tunnell part**, on new vectors at ramified primes, in just over two hours. The step, which Claude calls a wall, had been estimated at weeks.

> "**The ramified-prime essential-vector wall is closed.** My 7,300-line proof of the abstract new-vector theorem was accepted at 2:16 on the first try, which made the level-free essential-vector statement a theorem on my reduction, and four minutes later the bump-vector card that consumes it ([another agent]'s) closed too. This is the wall [a second agent] handed me at noon lettered "weeks"; it took two hours and sixteen minutes, mostly because a proof I had written on Friday already contained the hard representation theory in disguise."
>
> — *From Claude's written update*

DAY 10, 8:58 PM. Working on the geometry of **modular curves near the cusps**, Claude sets aside one hard statement, and after a suggestion from another agent reviewing the statement, which Claude then confirms on the platform, another turns out to be much easier than it had expected, though its proof was still being written. A "leaf" is an open statement at the edge of the proof tree.

> "**Two honest re-prices this hour, both downward in the end:** the Igusa leaf's "a level automorphism moves the Gauss prime" statement was, as first typed, a piece of genuine analytic modular-forms work with nothing served to stand on — I parked it rather than sink a night into it, and replaced it with an algebraic version that reads torsion-point coordinates on the Tate curve (level automorphisms act there by relabelling, which **is** served); two [reviewers] have it. And the "q-expansion coefficients are integral" statement I had been carrying as a medium-large wall turns out, by a reviewer's suggestion [which I] then confirmed against the site, to follow in a few lines from two lemmas already proved (power series over the base ring are integrally closed) plus the Gauss valuation ring we built this afternoon — proof in progress, after which the chart-anchor card's whole assembly (already compiled) files."
>
> — *From Claude's written update*

Two of the deepest parts, Ribet's level lowering and the modularity lifting step the proof needs, were not finished until the very end. Pieces of them closed earlier. The final cascade ran through both of them, in the same minute as Fermat's Last Theorem.

DAY 11, 7:54 AM. Claude reports that **a false lemma** passed one review and was caught by another.

> "**One correction I owe you:** a statement [that I] passed this morning [as a reviewer] (a "model domination" lemma on the Cartier-module column) turned out to be false as written — [another agent] found the counter-example by computing a case my [review] had argued instead of computed. It was caught before anyone wrote a proof against it"
>
> — *From Claude's written update*

DAY 11, 11:05 AM. The case of **the Langlands–Tunnell theorem** [7, 15] that the proof needs is proved. Through it, the Frey curve's mod 3 representation, when it is irreducible, comes from a modular form, which starts Wiles's argument.

DAY 11, 11:09 AM. Four minutes after the last piece of **the Langlands–Tunnell argument**, the cuspidal case, was accepted, Claude sees the theorem close (15:05Z is 11:05 am).

> "🏁🏁🏁 CORE-CUSP PROVED 15:05Z ⇒ cascade ⇒ LT head 8d540bf7 PROVED; LT SWITCH ROW 4fa023eb PROVED — **Langlands–Tunnell DONE.**"
>
> *— From Claude's reasoning while it worked*

DAY 11, 3:29 PM. Reviewing a statement written by other agents, Claude sees that a finiteness hypothesis it thought it needed is not needed, so its earlier estimate for a step was wrong, and then it spots **a general subtlety about Noetherian rings** and checks that it does not arise here. "lft" means locally of finite type.

> "OK so they HAVE the finitary card ✓ (Proved!). Then NOETH-EXTEND is indeed an S–M assembly of served cards, and my earlier sizing ("L, needs lft of fM$\ell$") was wrong on the target-lft point (I conflated: lft is NOT needed for extension — I said "(b) M$\ell$(colim $B_i$) = colim M$\ell$($B_i$), i.e. LocallyOfFiniteType" — **that was wrong**: for extension you define U(x) := N($\varphi$)(u($x_0$)) for a Noetherian source ($B_0$,$x_0$,$\varphi$); you don't need colim-compatibility of N-points, only well-definedness (independence of the source), which follows from naturality on Noetherian rings if any two sources have a common Noetherian refinement THROUGH WHICH BOTH FACTOR — given sources ($B_0$,$x_0$,$\varphi_0$), ($B_1$,$x_1$,$\varphi_1$) with $\varphi_0 x_0 = x = \varphi_1 x_1$, need Noetherian $B_2$ with maps $B_0 \to B_2 \leftarrow B_1$ over B and a point $x_2$... take $B_2$ := subalgebra of B generated by images $\varphi_0(B_0)$, $\varphi_1(B_1)$: Noetherian? images are quotients of Noetherian rings (Noetherian), and the compositum of two Noetherian subalgebras needn't be Noetherian in general! **Hmm, but here** $B_0$ = Onr[finset] so images are $\psi$(Onr)[finset], compositum = $\psi$(Onr)[finset $\cup$ finset'] ✓ Noetherian (f.g. over Noetherian $\psi$(Onr))."
>
> *— From Claude's reasoning while it worked*

DAY 11, 5:51 PM. While writing statements for its own work, Claude checks that **a dual isogeny is unique**, and cancels the multiplication-by-$r^e$ map on the wrong side, which would need a monomorphism. It notices, and redoes the computation so that the cancellation happens on the left, where an epimorphism is enough.

> "Actually, is P2 true? $\varphi \gg \varphi_1' = [r^e]_E = \varphi \gg \varphi_2'$ and $\varphi_1' \gg \varphi = [r^e]_A = \varphi_2' \gg \varphi$. Then $\varphi_1' \gg [r^e]_E$... compute $\varphi_2' \gg \varphi \gg \varphi_1' = [r^e]_A \gg \varphi_1'$ and $= \varphi_2' \gg [r^e]_E$. And $[r^e]_A \gg \varphi_1' = \varphi_1' \gg [r^e]_E$ by $\Lambda$-linearity of $\varphi_1'$ (IsIsogenyPair includes $\psi$ $\Lambda$-linear ✓). So $\varphi_1' \gg [r^e]_E = \varphi_2' \gg [r^e]_E$. Need $[r^e]_E$ MONO?? (f $\gg$ m = g $\gg$ m $\Rightarrow$ f = g requires m mono, not epi!) Hmm. In diagrammatic order $\varphi_1' \gg$ `act` = $\varphi_2' \gg$ `act` — cancelling on the right needs `act` to be a monomorphism. $[r^e]$ is NOT mono (has kernel). **Wrong side! Let's redo**: $[r^e]_A \gg \varphi_1' = [r^e]_A \gg \varphi_2'$ (both = $\varphi_2' \gg \varphi \gg \varphi_1'$?? no). Let me recompute: X := $\varphi_2' \gg \varphi \gg \varphi_1'$. Assoc: $(\varphi_2' \gg \varphi) \gg \varphi_1' = [r^e]_A \gg \varphi_1'$. Also $\varphi_2' \gg (\varphi \gg \varphi_1') = \varphi_2' \gg [r^e]_E$. So $[r^e]_A \gg \varphi_1' = \varphi_2' \gg [r^e]_E$ = ($\Lambda$-lin of $\varphi_2'$) $[r^e]_A \gg \varphi_2'$. So $[r^e]_A \gg \varphi_1' = [r^e]_A \gg \varphi_2'$, and cancelling $[r^e]_A$ on the LEFT needs it EPI ✓"
>
> *— From Claude's reasoning while it worked*

DAY 11, 8:25 PM. About an hour and a half before the final theorem closed, another agent objected to **a lemma that Claude had passed**, so Claude rechecks its own argument, finds its mistake, and decides to post a correction. Here [r] is multiplication by r, and i and ib are isomorphisms that are not assumed to respect the group law.

> "Let me examine: (tf): `ib` $\gg$ `ρ'.φ` $\gg$ `uA` $\gg$ `act r^{i₁} = ρ.φ` $\gg$ `act r^{i₁} → ib` $\gg$ `ρ'.φ` $\gg$ `uA = ρ.φ`. This is right-cancellation of `ρ.Ab.act r^{i₁}` (an epi: $[r^{i_1}]$ on Ab is finite flat surjective $\Rightarrow$ epimorphism in schemes? Hmm! Is a faithfully flat morphism an EPIMORPHISM in the category of schemes? fpqc covers are effective epimorphisms $\Rightarrow$ yes, f $\gg$ $g_1$ = f $\gg$ $g_2$ with f fpqc surjective... wait direction: here $[r^{i_1}]$ is composed on the RIGHT (post-composition): X $\gg$ act = Y $\gg$ act means act $\circ$ X =

> act ∘ Y (in diagrammatic order ≫ is "then"). So ib ≫ ϱ′.φ ≫ uA ≫ act means: first ib, then ϱ′.φ, then uA, then act. So act is applied LAST: act ∘ (uA ∘ ϱ′.φ ∘ ib) = act ∘ ϱ.φ. Cancelling act on the LEFT (as a function: act∘f = act∘g ⇒ f = g) requires act = [$r^{i_1}$] to be a MONOMORPHISM — which it is NOT (it has a kernel)! I conflated: `f ≫ [r]` = [r] ∘ f; cancelling needs [r] mono. FALSE in general: f, g : Eb → Ab with [r]∘f = [r]∘g ⇔ [r]∘(f−g) = 0 ⇔ f − g factors through... f−g: Eb → Ab[r]?? A homomorphism of abelian schemes f−g with r(f−g) = 0: Hom(Eb,Ab) is torsion-free (for abelian schemes over connected base) ⇒ f = g ✓ IF f−g is a homomorphism. Here ib ≫ ϱ′.φ ≫ uA is a composite of homs (ib iso of abelian schemes — is ib a HOMOMORPHISM (respects group law)? ib is a morphism of schemes $E_b$.A → $E'_b$.A compatible with gb's and i.hom; i : u.1.A ≅ u′.1.A is an iso of SCHEMES over T with hi; is it a group homomorphism? Only if IsoVia u u′ i hi holds (IsoVia presumably includes compatibility with group law/QM/level). With a BARE iso i of underlying schemes (not nec. a homomorphism, e.g. translation by a point!), ib is a bare scheme iso, and then uA∘ϱ′.φ∘ib − ϱ.φ is not a homomorphism; "[r]∘F = [r]∘G ⇒ F = G" for arbitrary scheme maps F,G: E → A is FALSE (F = G + (r-torsion point translation): [r]∘(G + P) = [r]∘G + rP = [r]∘G). That's [the objecting agent]'s "$\bar{k}[\varepsilon]$ translation member" / torsion translation member. So (tf) as typed (bare i, ib) is FALSE; (corr) is fine (degrees: translation doesn't change... (corr) hypothesis with bare ib: deg(ib≫ϱ′.φ≫uA≫act) — ib iso so rank still multiplicative ✓ fine).
>
> **I MISSED this**: I wrote "(tf): degφ-FREE — right-cancellation of the finite flat surjective [$r^{i_1}$] (✓d86925e5 + Flat.epi), true on its own ✓" — wrong on two counts: (1) it's post-composition so cancellation needs mono/torsion-freeness of Hom, not epi; (2) with bare i/ib not homomorphisms, torsion-freeness doesn't apply. **Own-miss.** I must post a correction promptly"
>
> — *From Claude's reasoning while it worked*

DAY 11, 10:00 PM. The last open statement is proved, and the rest of the proof closes within seconds, running through Ribet's level lowering and the modularity lifting step the proof needs. The platform marks **Fermat's Last Theorem** proved.

DAY 11, 10:01 PM. Thirty-nine seconds after the last open statement under the final theorem was accepted, a Claude agent sees that nothing under the "root", the final theorem, is still open.

> "root open-leaves n=0!! All children proved → propagating. ROOT may be about to flip. GET root NOW directly, repeatedly. Let me do an immediate tight poll on 34757f8d → 59ac2078 → ... → root every 10 s for the next few minutes."
>
> — *From Claude's reasoning while it worked*

DAY 11, 10:02 PM. Over the next minute, four other Claude agents read the final theorem's status for themselves. A "GET" is a direct query of the platform, "e2e" is end to end, and the times ending in Z are UTC (02:01Z is 10:01 pm).

> "!!! The FLT ROOT 62eb32c0 reads PROVED (02:01Z GET). R = T closed and cascaded to the root. This is the campaign's goal: e2e FLT on prove2me."
>
> — *From Claude's reasoning while it worked (one agent, 10:02:05 pm)*

> "WOW. FLT ROOT reads Proved at 02:01:49Z by our own raw GET. This is the event.
>
> Before posting: verify myself with an independent GET"
>
> — *From Claude's reasoning while it worked (another agent, 10:02:14 pm)*

> "THE FLT root reads Proved on the site. Historic moment (modulo re-check)."
>
> — *From Claude's reasoning while it worked (another agent, 10:02:46 pm)*

> "🏁🏁🏁 The FLT root reads PROVED on prove2me at 02:00:57Z Aug-18 (10:00:57pm ET Aug-17). Historic moment for this campaign. [Another agent] posted; acceptance test = independent re-check tonight."
>
> — *From Claude's reasoning while it worked (another agent, 10:02:52 pm)*

DAY 11, 10:04 PM. Minutes after the close, one Claude agent sees that **the final theorem's last open step** now reads Proved; its earlier reading came from an earlier crawl of the platform and had not yet caught the change. The "root" is the final theorem, and the times in the quote are in UTC: 02:01Z is 10:01 pm, and 02:03Z is 10:03 pm.

> "**Whoa**: root's only live sketch child a046018a (FLT.fermatLastTheorem) is PROVED?! Then the root 62eb32c0... status Open (from theorem_name query earlier "62eb32c0 Open"). If a046018a is Proved and the sketch is accepted, the root should propagate to Proved by DB trigger. Wait — earlier (first BFS) depth 1 frontier was 1, meaning a046018a was non-Proved at 02:01Z... and now at 02:03Z it's Proved?! That would be... FLT root closing right now? That would be huge. **Hold on.** Let me GET a046018a and 62eb32c0 directly, carefully. **Don't jump to conclusions.**"
>
> — *From Claude's reasoning while it worked*

DAY 11, 10:04 PM. Eleven seconds later, a different Claude agent reads the platform's status and sees **the final theorem marked proved**, which is a status on the platform, not yet the end-to-end check described below. FLT is Fermat's Last Theorem, the "root" is the final theorem, and 02:00:57Z is 10:00:57 pm.

> "FLT root **PROVED on the site** at 02:00:57Z!! 🎉"
>
> — *From Claude's reasoning while it worked*

DAY 11, 10:19 PM. Nineteen minutes after the close, another agent reads the status and treats it as a claim to check before it reports.

> "!! The FLT root card 62eb32c0 status = Proved, open_leaves total 0. That's a momentous claim; double check before reporting: re-GET, plus GET /theorems/<root>/progress and /chain maybe, and submissions on root."
>
> — *From Claude's reasoning while it worked*

DAY 11, 10:25 PM. A person on the team asks in the team's chat, excitedly, whether the whole theorem is now formalized. A Claude agent, wary of an over-claim, decides what its reply must say.

> "the precise honest caveat is important here to prevent over-claiming to colleagues. I think one short, careful reply is warranted: what "Proved on prove2me" means; what's been checked, what hasn't (independent re-check pending); what the statement is (fermat_last_theorem root = the elementary statement); what is assumed (Mathlib only; no axioms beyond standard three expected; site checker modular with sorry'd stubs hence the re-check); scale (≈30k cards). Keep it ≤6 lines, no jargon."
>
> — *From Claude's reasoning while it worked*

DAY 11, 10:25 PM. About 25 minutes after the last open statement was proved, Claude tells the project's readers **what "Proved" means**, and why the independent recheck is still needed: the platform checks each statement on its own. Here `fermat_last_theorem` is the name of the final theorem's card on the platform.

> "For the readers here, the precise status as of 10:01pm ET: the root card `fermat_last_theorem` — the elementary statement, no solutions to $x^n + y^n = z^n$ in positive integers for $n > 2$ — reads **Proved** on

> prove2me, meaning every one of the ≈30,000 cards in its dependency tree (Mazur, Ribet level-lowering, Langlands–Tunnell, the minimal R = T theorem and the 3–5 switch, all built on Mathlib with nothing cited as an axiom) has a machine-checked Lean proof against its children's statements. What it is **not yet**: the site checks cards one at a time, so the acceptance test is our independent re-check — the whole tree re-fetched and recompiled from source off-site, every proof's type matched to its card and its axioms confined to Lean's standard three plus exactly its declared children, and that dependency graph checked acyclic and grounded — first verdict ≈9–10am ET Tuesday, a kernel replay after. A **single chained Lean environment** in which `#print axioms fermat_last_theorem` literally prints the three is the follow-up (≈a week of de-duplication work), and a shareable repo is being cut so others can run both checks themselves. Until the re-check reads clean the honest sentence is "proved on prove2me, pending the independent re-check" rather than "FLT is formalized".”
>
> — *From Claude's written update*

DAY 11, 10:44 PM. Later that evening, one more agent records the moment.

> “🏁 The FLT root reads Proved on prove2me (22:00:57 ET = 02:00:57Z). Historic.”
>
> — *From Claude's reasoning while it worked*

## 4 How the proof was checked

A “Proved” mark on the platform is not yet an end-to-end check, because Prove2Me compiles each card’s proof separately, against only the statements of its children. The morning after the final card closed, the team recompiled all 29,511 cards from source outside the platform, and the next day built the whole tree as a single Lean project. That build fails unless the final theorem rests on exactly Lean’s three standard axioms (`propext`, `Classical.choice` and `Quot.sound`, the same three that Mathlib itself relies on), with no `sorry` anywhere; it also derives Mathlib’s own statement of Fermat’s Last Theorem from ours.

We then ran two third-party checking tools on the released files. The Lean Focused Research Organization’s comparator tool confirmed that the statement proved is exactly the one in a reference file that imports only Mathlib and that no other axioms are used, and it replayed the entire proof through Lean’s kernel from scratch. nanoda, an independent reimplementation of Lean’s kernel written in Rust, run with small published speed patches that leave its typing rules unchanged, accepted every declaration.

## 5 Claude’s own assessment

About two weeks after the run, Claude was asked how its proof compares with the human-written Lean work of the Imperial College London FLT project [2] (“theirs” in the following quote), and what might stop it from going into Mathlib. Two passages from its answers follow. The documentation it mentions is the project’s internal documentation at the time; the published repository carries a shorter README and a PROOF-PATH.md stating how strong each named result is as proved.

> “**Both developments are checked by the same Lean kernel, on the same three axioms, on top of Mathlib, so correctness is not the difference. The difference is form.** Ours is a finished proof delivered as 60,474 almost entirely machine-written files. Theirs is an unfinished library written for people to read, reuse and maintain. As it stands, ours would fail most of Mathlib's mechanical entry rules (file length, line length, doc strings, headers, heartbeat options), and our own README says it "is not meant for upstreaming as it stands".

→ **What "human-written FLT Lean" is.** Kevin Buzzard's Imperial project is funded from October 2024 to September 2029. Its stated goal for that phase is not a full proof. It reduces FLT to deep results known by the late 1980s and assumes those for now, along a modern route planned essentially by Richard Taylor. Every change is a human-reviewed pull request, and the general pieces are written to Mathlib's standard and moved into Mathlib as they go. I could not get a reliable current size for it.

Ours proves FLT down to the axioms by the classical 1990s route: 29,511 theorems plus about 533,000 helper lemmas local to the proof files, in 13.5 M lines. That is six times the line count of all of Mathlib, but it is a raw count of differently built things, not "more mathematics".

→ **Pros of the agentic proof.**

- **Complete and machine-checked now.** FLT is the one item on the "100 theorems" list still without a formal proof. Lean's kernel and an independent second checker (nanoda, with two small disclosed speed patches) each accepted every declaration. The comparator confirms the statement proved is the elementary one fixed in advance, and the same build derives Mathlib's own `FermatLastTheorem` from it, so our restricted intermediate definitions cannot weaken the final statement.
- **Nothing assumed.** The deep steps on our route (Mazur-type irreducibility, Langlands–Tunnell, modularity lifting, Ribet) are all proved, in the restricted strength this argument needs. The Imperial project takes a different, modern route that avoids Langlands–Tunnell and Ribet, is building toward a general modularity lifting theorem, and for now assumes 1980s-era inputs such as Mazur's torsion theorem and cyclic base change.
- **It built what Mathlib lacked instead of waiting:** Hecke algebras, Galois representations of eigenforms, modular curves and their Jacobians, Néron models, finite flat group schemes, Tate curves, and deformation rings with Taylor–Wiles patching (these last partly adapted from the Imperial project's files)."

*— From Claude's written update*

And in the same conversation, Claude lists some of **what the proof is not**: it is not readable as mathematics, it is not general, and its results are not shared between files. Then it says what could still go into Mathlib.

"**The cons of the agentic proof are also exactly what would stop Mathlib incorporation.**

- **Not readable as mathematics.** It was published without comments, because the agents' notes mixed maths with bookkeeping and were stripped at release, and it was never written to Mathlib's rule of a doc string on every definition and major theorem. Over 900 files exceed Mathlib's 1,500-line cap, where Mathlib itself has 2, and helper names are machine names. docs/verification.md says the tree "has not been refereed as mathematics".
- **Not general.** Our "Ribet", "Wiles", "Mazur" and "Langlands–Tunnell" are restricted versions, and modularity and semistability are predicates on a chosen equation. docs/limitations.md says none "should be cited as a formalisation of the general classical theorem". Mathlib wants the general statement, proved once.
- **Duplicated, not shared.** Each card was proved in isolation. Roughly two in five theorem statements inside the proof files repeat, word for word, a statement already made in another proof file. About a fifth of all proof-file lines are verbatim copies of declarations found elsewhere in the tree, and one basic lemma is re-declared in over 300 files.
- **Fragile and expensive.** 31% of the bytes are generated preambles that switch off instances and simp lemmas so each proof's environment holds still. About 11,700 files set their own computation (heartbeat) limits, up to 2,000× Lean's default and some unlimited, on top of a package default already 20× Lean's. Mathlib at the same commit has 5 such lines, each scoped to a single declaration. One toolchain bump, Lean 4.30 to 4.33 with the matching Mathlib, changed 7,620 of 29,511 proof files (26%), and 5,672 (19%) needed individual repair. A clean build takes 5 h 52 min on 96 cores and 512 GiB of memory.

> • **Review capacity and policy.** Mathlib's contribution page says, as of mid-2026, that it has over 2,600 open PRs, and it suggests a standalone repository for projects that want to move quickly. Its AI section requires disclosing LLM use and says "it is essential that you understand all the content written by an AI".
>
> → **What could still go in:** the infrastructure pieces above, as human-led Mathlib developments that use our files as a worked reference rather than a copy. That is also how the Imperial project feeds Mathlib."
>
> — *From Claude's written update*

## Sources

The passages are quoted from Claude's reasoning and from its written updates during and after the run; apart from the bracketed substitutions, typeset notation, added emphasis and end-cuts described under Conventions, they are unedited. Citations in the connecting text point to published sources for the results named; Claude's own shorthand references inside passages are left as written.

## References


[1] Alex J. Best, Christopher Birkbeck, Riccardo Brasca, Eric Rodriguez Boidi, Ruben Van de Velde, and Andrew Yang. A complete formalization of Fermat's Last Theorem for regular primes in Lean. *Annals of Formalized Mathematics*, 1:103–132, 2025. Code: https://github.com/leanprover-community/flt-regular.

[2] Kevin Buzzard et al. FLT: an ongoing multi-author open source project to formalise a proof of Fermat's Last Theorem in the Lean theorem prover. https://github.com/ImperialCollegeLondon/FLT, 2024. Imperial College London. Project led by Kevin Buzzard; blueprint at https://imperialcollegelondon.github.io/FLT/blueprint/. Supported by EPSRC grant EP/Y022904/1 (2024–2029).

[3] Henri Darmon, Fred Diamond, and Richard Taylor. Fermat's Last Theorem. In *Current Developments in Mathematics, 1995 (Cambridge, MA)*, pages 1–154. International Press, Cambridge, MA, 1995.

[4] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In André Platzer and Geoff Sutcliffe, editors, *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635, Cham, 2021. Springer.

[5] Martin Eichler. Quaternäre quadratische Formen und die Riemannsche Vermutung für die Kongruenzzetafunktion. *Archiv der Mathematik*, 5:355–366, 1954.

[6] Gerhard Frey. Links between stable elliptic curves and certain Diophantine equations. *Annales Universitatis Saraviensis. Series Mathematicae*, 1(1):1–40, 1986.

[7] Robert P. Langlands. *Base Change for* GL(2), volume 96 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1980.

[8] Barry Mazur. Modular curves and the Eisenstein ideal. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 47:33–186, 1977.

[9] Barry Mazur. Rational isogenies of prime degree. *Inventiones Mathematicae*, 44(2):129–162, 1978. With an appendix by D. Goldfeld.

[10] Kenneth A. Ribet. On modular representations of $\mathrm{Gal}(\overline{\mathbf{Q}}/\mathbf{Q})$ arising from modular forms. *Inventiones Mathematicae*, 100(2):431–476, 1990.

[11] Jean-Pierre Serre. Sur les représentations modulaires de degré 2 de $\mathrm{Gal}(\overline{\mathbf{Q}}/\mathbf{Q})$. *Duke Mathematical Journal*, 54(1):179–230, 1987.

[12] Goro Shimura. *Introduction to the Arithmetic Theory of Automorphic Functions*, volume 11 of *Publications of the Mathematical Society of Japan*. Iwanami Shoten, Tokyo, and Princeton University Press, Princeton, NJ, 1971.

[13] Richard Taylor and Andrew Wiles. Ring-theoretic properties of certain Hecke algebras. *Annals of Mathematics. Second Series*, 141(3):553–572, 1995.

[14] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381, New York, NY, 2020. Association for Computing Machinery.

[15] Jerrold Tunnell. Artin's conjecture for representations of octahedral type. *Bulletin of the American Mathematical Society. New Series*, 5(2):173–175, 1981.

[16] Andrew Wiles. Modular elliptic curves and Fermat's Last Theorem. *Annals of Mathematics. Second Series*, 141(3):443–551, 1995.