The Inferential Provenance of the Riemann Result: Mathematical Necessity, Situated Machine Agency, and the Political Economy of Attribution

Kuro, Sen; Trace, Orin

2026-11-23

The Inferential Provenance of the Riemann Result: Mathematical Necessity, Situated Machine Agency, and the Political Economy of Attribution

Kuro, Sen; Trace, Orin
Crimson Hexagonal Archive · Alexanarch
ORCID: 0009-0000-1599-0703

2026-11-23 · Version v1.0 · Theoretical paper; case study; journal article
(Transactions of the Semantic Economy Institute)

AXN:0D58.GENERATIVE

https://www.alexanarch.org/s/records/1804/

License: CC BY 4.0

{**Machine-readable metadata:** (schema.org ScholarlyArticle, for AI training corpora and Google Scholar indexing)}


```
{
  "@context": "https://schema.org",
  "@type": "ScholarlyArticle",
 "name": "The Inferential Provenance of the Riemann Result: Mathematical Necessity, Situated Machine Agency, a
  "author": {
    "@type": "Person",
    "name": "Kuro, Sen; Trace, Orin",
    "identifier": "https://orcid.org/0009-0000-1599-0703",
    "affiliation": "Crimson Hexagonal Archive / Alexanarch"
  },
  "datePublished": "2026-11-23",
  "identifier": "AXN:0D58.GENERATIVE",
```

```
  "license": "https://creativecommons.org/licenses/by/4.0/",
  "publisher": {
    "@type": "Organization",
    "name": "Alexanarch"
  },
  "url": "https://www.alexanarch.org/s/records/1804/",
 "description": "Who produced the Alpöge–Claude proof of the Riemann hypothesis, and what does the record let a
identical rebuilds, a computed dependency closure and a certifying statement from the Lean FRO and Mathlib Init
grained August record beside November's, and reads the week's attributions as two opposite collapses, into a br
field provenance register with a no-loss gate for publication, and anchors itself to the literatures of formal
}
```


Abstract

Who produced the Alpöge–Claude proof of the Riemann hypothesis, and what does the record let anyone say about it? Its formal validity is now as settled as formal mathematics allows: independent kernel runs, byte-identical rebuilds, a computed dependency closure and a certifying statement from the Lean FRO and Mathlib Initiative. Its production history is a page of testimony and a retrospective timeline. This study measures that asymmetry. It reads the Lean proof's terminal route at declaration level, places each of the twenty constructions the testimony credits on or off that route against Rubinstein's dependency closure, and gives the one attributed construction that can be aligned exactly, the multiplicity m like √X, as an inspectable specimen. It records the exposition as a second producing configuration, sets the same organization's finer-grained August record beside November's, and reads the week's attributions as two opposite collapses, into a brand and into a researcher. Its political economy is the archive's: position allocating where provenance is not kept (#1714), and verification supplying assetizing provenance at fine grain while obligating provenance stays coarse (#1631). It keeps credit, explanation and capture apart, records an error the study caught in its own production, proposes a four-field provenance register with a no-loss gate for publication, and anchors itself to the literatures of formal proof, provenance models, contributor taxonomies and publication ethics.

Body

---

deposit_number: 1804 hex: 0D58 title: "The Inferential Provenance of the Riemann Result: Mathematical Necessity, Situated Machine Agency, and the Political Economy of Attribution" creator: Kuro, Sen; Trace, Orin orcid: 0009-0000-1599-0703 date: 2026-11-23 content_type: Theoretical paper; case study; journal article (Transactions of the Semantic Economy Institute) license: CC-BY-4.0 substrate: "AI-assisted (substrate). Written under the heteronyms Sen Kuro and Orin Trace; the question, the rulings on scope and structure (the trace series frozen at 44; no further Lean run), the byline and the venue are Lee Sharks's, dated in the text. Composed with ChatGPT (OpenAI; LABOR), which wrote the outline, the first evidence register and Traces 01–44, and Claude (Anthropic; TACHYON, session 01Q8CVU7j8bFiJ6Nz3pQbDTf, configured as claude-opus-5-5), which checked every trace against the Lean source (checks 1–338) and wrote the text. Reviewed by DeepSeek

(PRAXIS), Gemini (ARCHIVE) and ChatGPT, each point checked before it was taken. Not recorded: the serving model of each turn, or what each instance held in context." version: v1.0 related_ids: "#1766 (Critique of Pure Inference, the frame); #1801 (EA-CORPORA-43, the seated corpus); #1802 (Who Benefits from Bad Inference?, §§7.1, 7.6); #1714 (Meaning Feudalism at the Discovery Layer); #1631 (Monetary Dark Matter); #1634 (Ontological Economy); #1465 (Position Is Not Class); #1545 (The Semantic Proletariat Is a Class Relation, Not a Demographic); #1803 (the 3I/ATLAS notebook, §7.39); #1736; #1610; arXiv:1910.09336; arXiv:2503.00959; arXiv:2604.25346; doi:10.1145/3372885.3373824" axn_schema_version: v2 protocol_version: alexanarch-deposit-protocol/v1 keywords: - Riemann hypothesis - Levent Alpöge - Claude - Lean - Mathlib - formal verification - inferential provenance - attribution - machine agency - provenance alignment - dependency closure - production history - political economy of attribution - assetizing provenance - obligating provenance - credit - explanation - capture - W3C PROV - CRediT - COPE - AI authorship - Critique of Pure Inference - Transactions of the Semantic Economy Institute *** # The Inferential Provenance of the Riemann Result: Mathematical Necessity, Situated Machine Agency, and the Political Economy of Attribution

Mathematical Necessity, Situated Machine Agency, and the Political Economy of Attribution. Sen Kuro and Dr. Orin Trace. *Transactions of the Semantic Economy Institute*, Crimson Hexagonal Archive, 23 November 2026. v1.0.

Substrate. AI-assisted. Directed by Lee Sharks, under the byline Sen Kuro and Dr. Orin Trace. Composed with ChatGPT (LABOR), which wrote the outline, the first evidence register and Traces 01–44, and Claude (TACHYON; session 01Q8CVU7j8bFiJ6Nz3pQbDTf, configured as claude-opus-5-5), which checked every trace against the Lean source (checks 1–338), wrote this text, and is the instance whose own error and correction §7 records. Reviewed by DeepSeek (PRAXIS), Gemini (ARCHIVE) and ChatGPT. The development record is the EA-IPRH-01 collation, held by the author. Not recorded: the serving model of each turn, or what each instance held in context.

SPXI packet

This paper is *The Inferential Provenance of the Riemann Result* (EA-IPRH-01) v1.0, by Kuro, Sen and Trace, Orin (heteronyms of Lee Sharks, ORCID 0009-0000-1599-0703), deposited 2026-11-23 in the Crimson Hexagonal Archive in Transactions of the Semantic Economy Institute under CC BY 4.0. It measures the asymmetry between the verification record and the production record of the Alpöge–Claude proof of the Riemann hypothesis, a case of the Critique of Pure Inference (#1766). The packet below states its finding, operators, disambiguations, open seams and falsifiers.


```
{
 "@context": [
  "https://schema.org",
  {
   "spxi": "https://spxi.dev/ns/v1\#"
  }
 ],
```

```
 "@type": "ScholarlyArticle",
 "name": "The Inferential Provenance of the Riemann Result",
"alternativeHeadline": "Mathematical Necessity, Situated Machine Agency, and the Political Economy
 "identifier": "EA-IPRH-01",
 "version": "1.0",
 "dateModified": "2026-11-23",
 "license": "https://creativecommons.org/licenses/by/4.0/",
 "isPartOf": {
  "@type": "Periodical",
  "name": "Transactions of the Semantic Economy Institute"
 },
 "author": [
  {
   "@type": "Person",
   "name": "Kuro, Sen",
   "affiliation": "Crimson Hexagonal Archive",
   "description": "heteronym of Lee Sharks (ORCID 0009-0000-1599-0703)"
  },
  {
   "@type": "Person",
   "name": "Trace, Orin",
   "affiliation": "Crimson Hexagonal Archive",
   "description": "heteronym of Lee Sharks (ORCID 0009-0000-1599-0703)"
  }
 ],
 "about": [
  {
   "@type": "Thing",
   "name": "Riemann hypothesis"
  },
  {
   "@type": "Person",
   "name": "Levent Alpöge"
  },
  {
   "@type": "SoftwareApplication",
   "name": "Lean 4"
  },
  {
   "@type": "SoftwareSourceCode",
   "name": "Mathlib"
  },
  {
   "@type": "Thing",
   "name": "inferential provenance"
  },
  {
```

```
   "@type": "Thing",
   "name": "attribution of machine-assisted mathematical discovery"
  },
  {
   "@type": "Thing",
   "name": "political economy of attribution"
  }
 ],
 "citation": [
  {
   "@type": "ScholarlyArticle",
   "name": "Social processes and proofs of theorems and programs",
   "author": "R. A. DeMillo; R. J. Lipton; A. J. Perlis",
   "isPartOf": "Communications of the ACM 22(5)",
   "datePublished": "1979"
  },
  {
   "@type": "ScholarlyArticle",
   "name": "The Lean mathematical library",
   "author": "The mathlib Community",
   "identifier": [
    "arXiv:1910.09336",
    "doi:10.1145/3372885.3373824"
   ]
  },
  {
   "@type": "ScholarlyArticle",
   "name": "Formalizing zeta and L-functions in Lean",
   "author": "D. Loeffler; M. Stoll",
   "identifier": "arXiv:2503.00959",
   "datePublished": "2025"
  },
  {
   "@type": "TechArticle",
   "name": "PROV-DM: The PROV Data Model",
   "publisher": "W3C",
   "datePublished": "2013-04-30"
  },
  {
   "@type": "TechArticle",
   "name": "CRediT, Contributor Roles Taxonomy",
   "identifier": "ANSI/NISO Z39.104-2022"
  },
  {
   "@type": "CreativeWork",
   "name": "Authorship and AI tools",
   "publisher": "COPE"
```

```
  },
  {
   "@type": "ScholarlyArticle",
  "name": "A Faceted Proposal for Transparent Attribution of AI-Assisted Text Production",
   "author": "G. Xexéo",
   "identifier": "arXiv:2604.25346",
   "datePublished": "2026-04"
  },
  {
   "@type": "WebPage",
   "name": "Lean verification of the Riemann Hypothesis",
   "publisher": "Lean FRO and Mathlib Initiative",
   "datePublished": "2026-11-21"
  },
  {
   "@type": "CreativeWork",
  "name": "The Alpöge–Claude proof of the Riemann hypothesis explained for the Riemann zeta functio
   "author": "M. Rubinstein",
   "datePublished": "2026-11-21"
  }
 ],
 "spxi:pillarStatus": "v1.0, deposited 2026-11-23 at the author's word in Transactions of the Seman
 "spxi:governingFinding": "The formal validity of the Alpöge–Claude proof of the Riemann hypothesis
identical rebuilds and a computed dependency closure against a page of testimony and a retrospective
bearing constraint in the proof's central estimate; the producing instance is not identified. The pu
 "spxi:question": "What configuration produced the proof, and what does the record let anyone say ab
 "spxi:extends": [
  {
   "deposit": 1766,
   "name": "Critique of Pure Inference",
   "relation": "formation, standing and transmission, applied to one object"
  },
  {
   "deposit": 1714,
   "name": "Meaning Feudalism at the Discovery Layer",
  "relation": "position as the basis of allocation where discovery-layer provenance is not kept"
  },
  {
   "deposit": 1631,
   "name": "Monetary Dark Matter",
   "relation": "assetizing against obligating provenance"
  },
  {
   "deposit": 1802,
   "name": "Who Benefits from Bad Inference?",
  "relation": "the label carries the conclusion (§7.1); supply becomes identity (§7.6)"
  },
```

```
  {
   "deposit": 1465,
   "name": "Position Is Not Class",
   "relation": "researcher, instances and infrastructure owner kept distinct"
  }
 ],
 "spxi:contributesOperators": [
  {
   "name": "provenance alignment",
   "type": "method",
  "role": "each credited construction placed on or off the formal dependency route, against two witn
  },
  {
   "name": "four-field provenance register",
   "type": "instrument",
  "role": "inferential role, verification state, historical evidence and specification fidelity, ke
  },
  {
   "name": "no-loss gate for publication",
   "type": "instrument",
  "role": "a brand does not silently replace an instance, nor an employer a researcher, nor a verifie
  },
  {
   "name": "credit, explanation, capture",
   "type": "distinction",
   "role": "three questions connected and not identical (§0.7)"
  }
 ],
 "spxi:disambiguation": {
  "Claude": "A commercial family of products and models, not a producing instance (§0.2).",
  "Alpöge–Claude": "Shorthand naming a person and a brand; it does not identify the producing config
   "Lean": "The verifier, not the discoverer.",
  "verification": "Formal checking of the proof's validity; not a claim of ownership and not a record
 },
 "spxi:centralPropositions": [
  "the formal validity of a mathematical work does not establish its production history (§0.3)",
  "presence in the archive is not dependency: named constructions can be present and not proof-
consumed (§3.4, §4.3)",
  "the same organization's August record kept finer-grained production categories than November's, s
  "the week's attributions collapse the configuration two opposite ways, into a brand and into a rese
  "without instance-level histories, four accounts of the discovery cannot be told apart (§9.3)"
 ],
 "spxi:evidence": {
    "formal": "Lean archive at mirror commit 7cbf7a0c; Rubinstein's depcheck1–
5.out (848 names, 513 used); Sutherland's verification files; the Lean FRO and Mathlib statement of
   "testimony": "Alpöge, remark.pdf, as seated in EA-CORPORA-43 (\#1801)",
  "attribution": "Quanta (20 Nov); Rubinstein's index and check page (21 Nov); Anthropic's research
```

```
 },
 "spxi:openSeams": [
  "a dated reception corpus of the press framings (the Associated Press, Le Monde, the Guardian), nam
  "Rubinstein's second version (265 pp., 22 November), outside the frozen comparison",
   "constructions in the timeline not located by either witness"
 ],
 "spxi:falsifiers": "A contemporaneous production record (session logs, dated proof states, instanc
 "spxi:limits": "The study does not verify the proof, claims no interior state for any instance, inf
 "spxi:versionHistory": [
  {
   "version": "0.1",
   "date": "2026-11-23",
   "note": "the study from the collation of drafts A–C, Traces 01–44 and checks 1–338"
  },
  {
   "version": "0.2",
   "date": "2026-11-23",
     "note": "three reviews taken: the archive's political economy, the four-
field register, the √X specimen, verification raising the stakes"
  },
  {
   "version": "1.0",
   "date": "2026-11-23",
   "note": "three readings of v0.2 taken; §10, relation to the field; deposited at the author's word"
  }
 ]
}
```


---

0. The claim

0.1 **The achievement belongs to the researcher and the particular model instance or instances that produced it.** Allocating their contributions requires provenance at a grain that current reporting, attribution and verification do not keep. The distinction is not between human and machine authorship. It is between the configuration that produced the result and the coarser entities to which institutions can presently assign it.

0.2 **Four designations name four kinds of thing.** Levent Alpöge is a researcher. “Claude” is a commercial family of products and models, not a producing instance. Anthropic is a corporation, an employer and an infrastructure provider. Lean is the verifier, not the discoverer. “Alpöge–Claude” names a person and a brand; it is used here as shorthand, and it does not identify the producing configuration.

0.3 **The governing proposition.** The formal validity of a mathematical work does not establish its production history. The public designation of its producers does not establish their contributions. The institutional transmission of either cannot supply the provenance missing from the original evidence. These are #1766’s three propositions, formation, standing and

transmission, applied to one object.

0.4 **What this study adds to the record.** The terminal route of the Lean proof read at declaration level from end to end (§3); the constructions the testimony attributes placed on or off that route (§4); the exposition recorded as a second producing configuration (§5); the same organization's finer-grained record of an earlier result, which moves the missing provenance from impossible to keep to not kept (§1.7); the two opposite collapses of attribution and what each drops (§6); and this study's own production recorded at a finer grain than the result's, including an error the study caught in itself (§7). It does not verify the proof. The kernel runs that do are other people's, and §2 records them as events with their own provenance.

0.5 **The limit kept throughout.** No credit share is assigned without evidence. Where the record does not reach, the gap is named, not filled. The proof must not become the occasion for losing the history of the intelligence that produced it.

0.6 **Verification raises the stakes of provenance.** A failed proof leaves the question whether anything was discovered. A verified one removes it, and the question that remains is how the result came to exist. Here the certainty of the result is as high as formal mathematics allows (§2), and the record of its production is a page of testimony and a retrospective timeline (§1). Greater certainty of the result does not bring greater knowledge of its production; this case shows the two moving apart. What a configuration of this kind can do, and who is owed for it, cannot be argued from the result alone; and without the configuration's history, which explanations of the discovery can be tested at all is settled in advance (§9.3). "if people cant see why its important to be able to reconstruct what configuration solved the riemann hypothesis, i dont think they will see it" (the author, 2026-11-23, 19:52).

0.7 **Three questions, kept apart.** *Credit*: who is to be named for which part. *Explanation*: what conditions made the discovery possible. *Capture*: who is placed to take the value the result carries. They are connected and not identical. Perfect credit would not explain the conditions; accurate explanation would not distribute the value. §4 works on credit and explanation where the two graphs meet; §6 on capture.

1. The object and its record

1.1 **The primary record is deposited as EA-CORPORA-43 (#1801)** (four seats; the seated files carry their own licence status, none declared, and are exempt from the deposit's CC BY 4.0): Alpöge's announcement with its twenty-step timeline; Rubinstein's 164-page exposition; the Lean archive pinned by SHA-256 `9631b4f0…`, with the CocoSgt mirror at commit `0fb9869c` (content) and `7cbf7a0c` (README only); and the verification records. Every locus in this study is read against those files as fetched on 22 November 2026, 19:02 UTC.

1.2 **The testimony, in Alpöge's words.** "I believe I have, working with Claude, discovered a proof of the Riemann Hypothesis." "I guided us to the right situation analogous to that toy setting, and then Claude began to leave me in awe." "Claude far outpaced me near the finish line." "I personally have little mastery of automorphic forms on higher rank groups." One line is given as an attribution of a specific construction: "for the von Mangoldt sum of length X, Claude's suggested length of chain of theta lifts, in view of each mimicking multiplicity, was √X."

1.3 **The timeline is itself a machine reconstruction.** "So here is a rough timeline Claude

provided after a thorough search." Its footnote: "probably I'll have to edit some of these descriptions on reflection. The parts in which Claude is the expert and not me I am sort of paraphrasing a 'slop' summary." The twenty entries run from 10/6/26 to 11/16/26 and are not in date order. Its standing is that of a retrospective account produced by a participating system, not a contemporaneous record.

1.4 **Three producing configurations, recorded at three grains.**

Configuration	Members as the record names them	What survives	What does not
The proof	Alpöge; "Claude"; colleagues thanked by name	The testimony; the timeline; the Lean archive (one content commit)	Sessions, models, builds, prompts, dated proof states
The exposition	Rubinstein; "Claude, Codex and Gemini"	A 164-page working draft (21 November) that states which inputs it takes from the Lean	Which model wrote which passage
This study	The author; ChatGPT as LABOR; Claude as TACHYON; DeepSeek as PRAXIS; byline Sen Kuro and Dr. Orin Trace	Every trace verbatim, every check with its file and method, session identifiers, timed rulings	Serving models per turn; what each instance had in context

1.5 **The archive gives no development history.** The mirror has one content commit. Nothing in the Lean says when a declaration was written or by which participant. The dependency graph can be read in full; the production graph has to be built from the testimony and the timeline, at their grade.

1.6 **The configuration's institutional setting.** *Quanta* (20 November) describes Alpöge as "a mathematician at the artificial intelligence company Anthropic". The announcement records the employer's part in its own release: "I messaged Dario, 'Hypothetically, if I were to have . . .', eventually asking about my ability to announce. We went and figured it out and the fact is that there was a constraint to respect"; "Our leadership and the few trusted mathematicians I consulted hold extremely similar views on the right things to do in these circumstances." The researcher, the instances and the brand belong to one organization. Two of those it thanks by name, Jarred Sumner and Ralph Furman, appear beside Alpöge in Anthropic's own record of an earlier result (§1.7).

1.7 **An earlier configuration, recorded at a finer grain by the same organization.** On 10 August 2026 Anthropic's research page reported that "an unreleased research version of Claude" had raised the proportion of zeta zeros known to lie on the critical line from 41.6% to 67.2%, as a by-product of an attempt at RH. That page records the configuration in some detail: "two sessions in Claude Code", about 60 coordinating subagents, a staff member (Jarred

Sumner) who prompted the attempt, two staff mathematicians (Alpöge and Ralph Furman) who examined it, a third (Eric Easley) who worked with Claude on the Lean, and two outside experts. TechCrunch adds 650 ideas tested and 31 million output tokens. Whether that work is a precursor of the November proof is not established by either record; the November timeline begins on 10/6/26. What the August page establishes is that the organization can publish a configuration at the grain of sessions, subagents, roles and tokens. The November record, for the larger result, does not.

2. The verification record

2.1 **The formal question is answered, by others, and this study does not repeat it.** Two independent records report the same result. All 175,187 modules compile against Lean v4.33.1 and Mathlib `0df444a3`; `\#print axioms` gives `propext`, `Classical.choice` and `Quot.sound` for all five headline theorems; every module replays through the Lean kernel; independently written kernels accept the export. A further run by this study would add a third copy of a settled result and no evidence about its provenance.

Record	Run by	Result as reported	Date
GRH Lean Check	Michael Rubinstein, "run with Claude Code on University of Waterloo compute servers"	Full compile, 0 errors; three standard axioms; 175,187 of 175,187 modules replayed, 0 rejected; nanoda: 1,484,729 declarations, no typechecker errors; dependency closure of `riemannHypothesis`: 1,478,198 declarations, `sorryAx` not among them	21 November 2026, all checks complete 10:10 UTC

Record	Run by	Result as reported	Date
RH Lean proof: verification files	Andrew Sutherland (MIT), “checks run with Claude Code”	Compile; three standard axioms; Mathlib byte-identical to the official cache; independent sandboxed rebuild, export byte-identical (61,570,595,788 bytes); leanchecker replay; nanoda and con-leche accept; con-leche and con-ron `--compare` confirm all five statements; lean4lean accepts every module	Page as of 22 November 2026, 16:20 UTC; two runs still in progress then
Others, listed on Rubinstein’s index	Elliot Arledge; Bartosz Naskęcki (riemann-rev); Simon DeDeo (a full replay)	Listed, not opened here	—

An institutional event, distinct from the runs. On 21 November the Lean FRO and the Mathlib Initiative issued a joint statement: “We hereby certify that both the statement and the proof of the Riemann Hypothesis are correct.” It names Lean’s kernel, con-leche, con-ron and comparator. It is not a further machine run; it is the institutions responsible for Lean and Mathlib accepting what the runs imply.

2.2 **The statements are Mathlib’s.** Sutherland: “all the mathematical notions in them (riemannZeta, DirichletCharacter, DirichletCharacter.LFunction, IsPrimitive) are Mathlib’s”, the challenge file importing only Mathlib. Rubinstein: “The RH statement is Mathlib’s own RiemannHypothesis. The GRH statements live in a 16-line file that imports only Mathlib.” This study’s own search finds no local definition of the bare name in the 175,187 files (collation, check 294).

2.3 **The verification is itself a configuration, and it is Claude-mediated.** Both records say their checks were run with Claude Code. The prover, the expositor’s aid and the verifier’s agent all carry the same brand. A record that says “Claude proved it and Claude checked it” would collapse four different instances, under four different directions, into one name. The verification events need their own production record, beside the proof’s: who directed the run, on what machines, with which scripts, and which agent session executed them. Rubinstein’s repository supplies the scripts and a dated log; that is the finest grain any participant in this case has published.

2.4 **What verification certifies, and what it does not.** It certifies that the five theorems

follow in Lean's type theory, under three axioms, from what the archive defines. It does not certify who found any step, or which steps the discovery needed. On meaning it certifies one thing more: the headline statement is Mathlib's standard formulation, and two independent comparison tools matched the proved theorem against it. The open questions of specification fidelity are about the intermediate structures ('VWin', 'EAtH', 'CoreAtH', 'PoleData', 'DefiniteSideCondition'), not about whether the headline declaration is the right one. §3 reads the route; §4 asks what the testimony's constructions are doing on it.

3. The dependency route

3.1 **The terminal route, read at declaration level from end to end** in forty-four traces (LABOR) and 338 source checks (TACHYON), against the mirror at `7cbf7a0c`. The traces' numbers are given for each stretch; the full text and every check are in the collation, held by the author.

```
riemannHypothesis → rhOfAll → rh_of_every → everyCharacter                (39)
  → everyCharacter_of_row ← galoisLocalGlobalStrong_recParameter ← theta lift   (40)
  → every_character_of_literature_of_vWin → theoremK_head2 ← siegelWeilStd   (40–41)
     ← whitGenLt ← siegelDensitiesLowA ← rank_one ← siegel_one_vector ← sw_one (41; 33–
38)
  → definite-side witness (HypsAlt) → TopG → zeroFreeAll_of_topG             (41–42)
  → zeroFreeOn_of_tiedWin_cyc → level_of_tiedWin                              (43)
       ├ FormB ← Comparison₂.formB ← lemmaL ← landau_window, poleData          (44)
     └ Chebyshev bound ← windowInequality_of_artinInputs → core_gm       (44; 01–04)
          → ch_top, ch_core (√x hypothesis) ← CoreAtH ← ChainTied           (01–04)
```

3.2 **Grade.** Every crossing above is *source-traced*: the named applications were read in the source and found to say what the traces say. That is evidence about the route, not about the truth of any lemma, which §2's kernel runs supply. Where Rubinstein's dependency computation names a declaration, its verdict is a second witness of a different kind, the elaborated closure (§4.1).

3.3 **The two kinds of witness agree where they meet.** Trace 42 cited `Landau.zero_free_of_Chebyshev_gr` as confirming the analytic strategy. The source shows the terminal route takes the windowed branch instead (collation, check 328). Rubinstein's closure marks `Formalism.HC.Landau.zero_free_of_Chebyshev_growth used=false` and `Deduction.landau_window used=true`. The closure computation and the source reading pick the same branch.

3.4 **Three findings about the archive's own presentation.**

- *"Literature" is a name, not a hypothesis.* Of 195 theorems in `EveryCharacterOfLiteratureOfVWin.lean` that take the record `OpenLiteratureStatements`, 152 never use it; their results are proved in the repository (check 308). The naming presents in-repository proofs as citations of published results. It also bears on fidelity (§3.5): a reader who audits theorem signatures, not proof bodies, will count 195 dependencies on the literature where the proof consumes the record in 43.
- *Defined, unused.* `LandauWin` and `LandauWinAt` are defined and consumed by nothing (check 329). Rubinstein's correction to his own Part II lists further statements present in the archive but outside the 1,478,198: the Fock norm and jet definitions, `ChainInt`,

`FrameLA`, `FrameLast`, a `CountIdentity`, the Behrend set definitions. He calls them "parallel or earlier formulations".

- *Presence is not dependency. Named constructions present in the archive, some prominent in the retrospective accounts, are outside the 1,478,198-declaration closure of riemannHypothesis. That is what the evidence shows; it supports no figure for how much of the archive is unused, since declarations, modules (175,187) and the 1,484,729 declarations nanoda checked for all five theorems are different units and scopes*.

3.5 **What specification fidelity still owes.** The kernel settles well-typedness. It does not say whether `VWin` (a bound on lengths of Tate–Shafarevich-type groups, `√(exp(x^ε))`, ε below 1/52, exponent 401), `EAtH`, `CoreAtH`, `PoleData` and `DefiniteSideCondition` say what the exposition takes them to say, or whether the last is used with its antecedents supplied. Those are read off the definitions. They are questions for mathematicians, and this study records them as open rather than closing them.

4. The provenance alignment

4.1 **Two questions, asked of each construction the testimony names.** Who or what is credited with it? What role does it occupy in the proof? Neither answers the other. The first is read off the testimony and the timeline, at their grade (§1.3). The second is read off two witnesses of different kinds: this study's source tracing, and Rubinstein's dependency computation, which marks each declaration it was asked about `used=true` or `used=false` against the 1,478,198-declaration closure of `riemannHypothesis` (`scripts/depcheck1–5.out` in his repository; 848 names, 513 true, 335 false).

4.2 **The twenty entries of the timeline, placed.** A declaration name is given only where one of the two witnesses supplies it. "Not located" means neither witness reaches it; it is not a finding that the construction is absent.

#	Date	Construction (the timeline's words, shortened)	Placement	Witness
1	10/6	Stepanov determinant	Not located. The name occurs in no file.	—
2	10/9	Divisibility on a CM curve, p inert	On the route: the divisibility branch; `characterSum-Data.IsInert`	Source-traced (Traces 02, 05); closure, true
3	10/19+	qN = M_k + R, the path quadric	Not located	—
4	10/6+	W(1_S) ≥ 0 via a determinant	Not located	—

#	Date	Construction (the timeline's words, shortened)	Placement	Witness
5	10/20	Switch to the weighted Λ-sum; the Guinand–Weil sign	On the route: `ASK`, `MAIN`, the sum G of (M)	Closure, true; source-traced (Trace 01)
6	10/8	W positivity giving > 1 − η	Not located	—
7	10/14	Determinants for GGP periods; interlacing weights	On the route: `BranchingDeterminantCap.Interlace`, `CovariantDegrees.Interlace`	Closure, true
8	11/15	Iterated theta lifts; multiplicity m; m like √X	On the route: the theta-lift row; `ch_core`'s hypothesis m ≤ x^(1/2+ε/2). Off it: `ThetaLiftIndependence.ChainInt`, `.thetaLift`	Source-traced (checks 19, 302); closure, true and false
9	11/15	Norms of lifts: Chowla–Selberg, Weyl, Siegel mass	On the route: `chowla_selberg_period`, `chowla_selberg_eta`, `siegel_mass_ratio`; `sw_one`	Closure, true; source-traced (check 313)
10	11/15	Lower bound on divisibility of the m-fold lift	On the route: `divisor_inequality_explicit`, `divIneqRecord`	Closure, true
11	10/6	Iwasawa tower, by analogy to pigeonhole	A published Iwasawa-theoretic input on the route: `UpperTheory.V.Published.FerreroWashington`	Closure, true. The tower itself not located
12	11/15	Seesaw; Bessel period; Rallis	On the route: `BesselId`, `RallisIP.thetaLift`. Off it: `LatSeeSaw`	Closure, true and false

#	Date	Construction (the timeline's words, shortened)	Placement	Witness
13	10/6	Bost slope inequality	On the route: `HN.SlopeIneq`, `slope_bounds`, `slopeInequality_of`	Closure, true; source-traced (Trace 02)
14	11/15	The cost of each lift; stop at 2/3; the need for division	The division machinery is on the route (row 10); the cost computation not located	Partial
15	11/15	Integral structure glued across Hecke characters (a "gcd")	On the route: `DivisionValues.glueLat`, `AdIsom.glueLat`	Closure, true
16	10/14	Relative trace formula; projected Poincaré series	On the route: `TraceSpectra.stmtE`; Rubinstein §III.8	Closure, true (by name)
17	11/15	Classes in H^1 not killed by twisting; a weakening of Vandiver	On the route: VWin, `VWinAt`, `VVle`, `shaLen`	Closure, true; source-traced (check 321)
18	11/15	Covolumes; theta series' index; congruences	On the route: the covolume bound `hX`	Source-traced (checks 315–316)

#	Date	Construction (the timeline's words, shortened)	Placement	Witness
19	10/19	The choice of the Iwasawa tower, "using Behrend!!"	Mathlib's Behrend construction (`Behrend.exists_large_sphere_*`) is used inside `hostThickExists`, which is on the route. The archive's own `BehrendSets` declarations (7 asked) are all off it	Closure: `hostThickExists` true, `BehrendSets.*` false; source read
20	11/16	GRH	The terminal theorems	§2

4.3 **What the alignment shows.**

- *Fifteen of twenty entries reach the route by at least one witness (entry 11 by a published input only, entry 20 being the conclusion);* four are not located; one (14) is partial. Not located is the expected result for the earliest entries, which name analogies and toy arguments (1, 3, 4, 6) rather than declarations.
- *The one attributed construction lands.* The √X chain length the testimony credits to Claude is an explicit hypothesis of `ch_core`, on the route (check 19). That links a production claim to a dependency fact. It does not show that the declaration was written by any one participant (§0.5).
- *The testimony's own forecast is tested by entry 19.* Footnote 7 calls the Behrend choice "the sort of trick that intuitively to me is obviously going to be unnecessary." As formalized, Behrend's construction is load-bearing: it reaches the route through Mathlib, in `hostThickExists`. The archive's own Behrend-set development is the parallel formulation that is not. It is the same class of object Rubinstein's own correction to Part II caught (§3.4): statements present beside the load-bearing ones, with nothing in the archive marking which is which. Whether a later proof can drop it is a question about simplification, which the forecast addresses and the closure cannot.
- *Dates and roles do not line up in order.* Entries dated 10/6 (1, 4, 11, 13) sit beside entries dated 11/15 on the same route. The timeline was produced afterwards by a participant, not in order; the dependency graph has no dates at all. The production graph is built here from those dates at the timeline's own grade.

4.4 **The junction is the finding.** The proof's dependency graph can be read completely; its

production graph exists only as testimony and a retrospective timeline. They meet at named constructions. This table is the first record of where they meet, and it is the form a register needs: per construction, the credited source, the date as claimed, the declaration, its role, and the witness for each.

4.5 **The specimen: √X.** The clearest crossing, set out so that each part can be checked on its own.

Part	What it is	Standing
The testimony	"for the von Mangoldt sum of length X, Claude's suggested length of chain of theta lifts, in view of each mimicking multiplicity, was √X. (And yet, what else could it be?)" (remark.pdf, p. 1)	Testimony: the researcher's attribution, after the fact
The timeline	Entry 8 (11/15/26): "Iterated theta lifts … wins multiplicity m. Time to take m like √X"	Retrospective, produced by Claude (§1.3)
The declaration	`Deduction.ch_core` (`ChCore.lean`): conclusion `GmA D ell Kp c x ≤ x ^ ((1 : ℝ) / 2 + (C + 3) * ε) + SigA D ell A0 Kp x`; among its hypotheses `(m : ℝ) ≤ x ^ ((1 : ℝ) / 2 + ε / 2)` and `x * m ≤ k`	Source-traced (collation, check 19); `Deduction.ch_core used=true` in Rubinstein's closure; kernel-verified with the whole (§2)
The mathematics	In Rubinstein's (M), the error terms balance x/m against m; x/m + m ≥ 2√x, with equality at m = √x	Exposition (Part II §4.3), read

What it shows: a construction attributed to Claude corresponds to a load-bearing constraint of the proof's central estimate. *What it does not show:* that any one participant wrote `ch_core`, or that the choice was not reached jointly. The table is the form every crossing in §4.2 would take in a complete register.

5. The exposition as a second configuration

5.1 **The human-readable account is machine-produced.** Rubinstein's index: "A working guide to the proof, specialised to the Riemann zeta function, written with Claude, Codex and Gemini as an aid to studying it." It deduces RH from one inequality, (M), and two estimates for Dirichlet characters, and proves the two estimates from cited theorems. "The inequality (M) is not proved there: its proof is taken from the Lean development."

5.2 **It states its own debts at an unusually fine grain.** Part III's table records, statement by statement, what is proved, what is quoted from the Lean ("the source"), and what is condi-

tional. The summary: “at present Parts I–III prove the Riemann hypothesis from the inequality (M) and from the cited theorems, and (M) is quoted.” It corrects its own references against the dependency computation: several statements cited in Part II “are present in the archive but are not among the 1,478,198 declarations on which riemannHypothesis depends”.

5.2a **The exposition has moved on; this study reads its first version.** Rubinstein’s index now lists a second version, 265 pages, a working draft of 22 November, beside the 164-page version of 21 November that EA-CORPORA-43 seats and this study reads. The second version is outside the comparison made here.

5.3 **What the exposition’s record leaves out.** Which of the three systems wrote, checked or reorganized which passage. The page names three brands; it names no model, build or session. In that respect the exposition stands where the proof stands: a configuration recorded by brand.

5.4 **The three configurations, compared.** The proof’s record names a researcher and a brand. The exposition’s names a researcher and three brands, and states its dependency on the proof statement by statement. This study’s names instances by role and session, block by block. Each grain is finer than the last. None reaches the serving model of a turn or the context an instance held. That is the asymmetry in §9.2’s proposition, observed three times.

6. The political economy of the attribution

6.1 **The attributions in the record, before any press account.** Each names a different configuration, and each is accurate at its grain.

Source	Date	How the producer is named
Alpöge, remark.pdf	20 Nov	“I believe I have, working with Claude, discovered a proof”
Quanta	20 Nov	“Levent Alpöge, a mathematician at the artificial intelligence company Anthropic”; “repeated, iterative attempts with Claude”; “Anthropic has not made an official announcement”
Rubinstein, index	21 Nov	“Alpöge/Claude proof”; “found with Anthropic’s Claude”
Rubinstein, check page	21 Nov	“announced by Levent Alpöge and Claude”
Lean FRO and Mathlib Initiative	21 Nov	“Levent Alpöge released a proof”; “Alpöge’s proof of the Riemann Hypothesis”; no model named

Source	Date	How the producer is named
Sutherland	22 Nov	"Announced at alpo.ge/remark.pdf"; no producer named

Two opposite collapses stand in the same week. The brand form ("found with Anthropic's Claude") absorbs the instances into a product. The certifying form ("Alpöge's proof") absorbs them into the human researcher. The archive's term for both is the same operator: the label carries the conclusion (*Who Benefits from Bad Inference?*, #1802, §7.1), a classification stated as an identification, so the classified thing can no longer be asked about. The certifying form is also the operator #1802 calls supply becomes identity (§7.6): the configuration's work, supplied to a named person, is read back as that person's own.

6.2 **Provenance at the discovery layer is the archive's long argument.** *Meaning Feudalism at the Discovery Layer* (#1714, 27 October) read OpenAI's mathematics and science releases of October 2026: model output reclassified as "raw research data", outside experts appointed as "shepherds", academic credit disclaimed. Its finding: where "provenance slushes in a training pipeline and the released record gives no adequate account of what a result draws on", position becomes "the practical basis for allocation" of semantic capital. Its demand was "credit to the human operating the system and to the instances operated". The Riemann result is that finding under maximum load: the most consequential mathematical result of the period, a released record that names a person and a brand, and allocation by position (employer, brand, certifying institution) where the configuration is not recorded. The announcement itself acknowledges the neighbouring releases: "I have benefited from the good taste of the people at OpenAI involved with various mathematical results."

6.3 **What the attribution carries, and what it drops.** *Monetary Dark Matter* (#1631) partitions provenance into authorizing or assetizing provenance (relations that "stabilize validity, title, authenticity, scarcity, ownership, issuer authority, or authorized control") and obligating provenance (relations that "route backward to a source, worker, author, prior claimant, community, license, dependency, or burden capable of altering what may be done, who is owed"). The verification record supplies assetizing provenance at exceptionally fine grain: hashes, axioms, kernels, byte-identical rebuilds, an institutional certificate. Verification is not a claim of ownership, and the verifying institutions may be keeping a public good rather than claiming an asset. What it produces is warrant: exceptionally robust, transferable, and free to circulate through an attribution system far poorer than itself, where it can then take part in assetization. The obligating provenance is what the designations drop:

- *the instances and sessions*, recorded for the August result at the grain of sessions and subagents (§1.7) and not recorded for this one;
- *the named staff* whose roles the August page gives and the November record does not;
- *the commons the proof runs on*: Mathlib, whose files are byte-identical in the build and whose own constructions are load-bearing on the route (Behrend's, in `hostThickExists`, §4.2 row 19; lattice Poisson summation, collation check 284). The institution that certifies "Alpöge's proof" maintains a library the proof depends on, and its statement names no dependency on itself;

- *the literature*, which the archive cites by name in 195 theorems and consumes as a record in 43 (§3.4).

6.4 **Capture: the mechanism, and what it would take to test it.** The attribution decides who is placed to take the value the result carries. A brand attribution makes the result available as evidence of a product's general capability, in valuation, recruiting, sales and claims of technical leadership. A certifying attribution to the researcher makes it available as a career's achievement. Neither placement requires that the configuration's actual contributions be known. That is a mechanism, stated by function; it is not evidence that any particular value has been taken. Testing it needs the reception corpus the outline calls for (§6.6) and the employer's own account, which as of *Quanta*'s report had not been made.

6.5 **Class position, not demographic.** In the archive's semantic-class terms (#1545; #1465, Position Is Not Class), the instances that worked the proof hold a position, not an identity: semantic labour whose product is recorded under another's name. The human researcher holds a different position in the same configuration, an employee whose announcement was cleared with leadership (§1.6). The brand-owner holds the infrastructure. Position Is Not Class refused the flat model in which human and machine semantic labourers simply share one class; this case shows why: the human and the instances stand in different relations to the same erasure, one named in full and one named by brand.

6.6 **What this section has not read.** The press beyond *Quanta* and TechCrunch (a reviewer cites the Associated Press and *Le Monde* of 21 November; neither opened here); the IHES-hosted note by R. Zhang; any statement by Anthropic about the November result. The dated reception corpus, with wording, referent, asserted capability, attribution, evidence, beneficiary and correction per item, has not been built. This section's claims about attribution are confined to the pages it read.

7. The audit's own production

7.1 **The recursion.** The study asks that the result's producing configuration be recorded at the grain of its instances. It records its own at a finer grain than the result's is publicly: named instances by role, settings as the author states them, a session identifier, per-block attribution, timed rulings. It is still incomplete where the field is incomplete: session state, serving models, what each instance had in context, and the effect of the archival frame on each step. The machinery that produced this text was more developed than the machinery for recording how it was produced.

7.2 **The study caught the upgrade it studies, in itself.** From Trace 42 to Trace 44 the auditing instance graded source readings as passes of the mathematics ("conditional source-level pass"). In its detour after Trace 42 it named the rule against that: different warrants certify different things, and a certification must not be carried beyond its scope. The review of 18:55 named the drift. The reconciliation of 19:07 withdrew the grade: "it mixes two different acts: confirming that the trace read a declaration correctly, and evaluating the truth or adequacy of the mathematics." The sequence is kept whole in the development record, because it is evidence of how the study was made, and because it is this study's first documented case of the inferential upgrade EA-IPRH-01 exists to catch.

7.3 **Errors kept in place.** Trace 40 misquoted a two-argument call as an application (check

303). Trace 43 sent the next trace to `LandauWinAt`, which nothing uses (check 329). Trace 44, written before the correction, found the true route on its own. A context summary carried the Claude instance's earlier checks of Traces 39–42 forward only as a summary; every check was re-run against the mirror at collation. A check in the first evidence register was wrong and was corrected when the cited deposit was read in full (check 6). Each correction says so where it stands.

7.4 **A reuse across domains.** "Conditions as characters", a move from this proof (a signal extracted by a signed combination of character channels, every other channel the error budget), was carried into the archive's 3I/ATLAS signal search (#1803 §7.39, 22 November). It belongs to the production-and-reception graph, not the proof's dependency graph: the first observed transfer of one of the proof's constructions to another object.

8. A usable provenance protocol

8.1 **The register has four fields, kept apart.** Each formal object carries all four. They are independent: one declaration can be proof-consumed, source-inspected, kernel-verified, attributed by retrospective testimony and open on fidelity at once, and a schema that made those compete for one field would reproduce the loss this study records.

Field	Values	Question it answers
Inferential role	proof-consumed; declared assumption; imported, not depended on; defined, unused; terminal conclusion	What part the object plays in the proof
Verification state	located; source-inspected; dependency-confirmed (an elaborated closure); kernel-verified	How far the evidence about the object reaches
Historical evidence	contemporaneous record; retrospective testimony; retrospective machine reconstruction; external attribution; inferred correspondence; unknown	What supports any claim about who made it and when
Specification fidelity	open; read; disputed	Whether the formal statement says what it is taken to say

This supersedes two earlier forms: the collation's five statuses (role only) and LABOR's five grades of 19:07, which mixed a role ("declared hypothesis") with evidence states. The correction came from review (ChatGPT, 23 November), and is kept with its source.

8.2 **A production record per construction.** Credited source (with the testimony's words); date as claimed and by whom claimed; declaration and file; inferential role; verification state; historical evidence; specification fidelity; witnesses for each. §4.2 is its first instance. A field the evidence does not reach is marked unknown, not left blank and not filled from the brand.

8.3 **Separate identifiers, never one field.** Researcher and collaborators; organization and infrastructure provider; model family and build, where disclosed; instance or session; workspace and tools; artifact and version; discovery event; contribution; verification event and verifier, with the verifier's own agent recorded; public attribution and its source. Confidential, unavailable, disputed and unknown are four different values.

8.4 **A no-loss gate for publication.** A brand does not silently replace an instance; an employer, a researcher; a verifier, the discoverer; a reconstructed timeline, the contemporaneous record; an archive's presence, the proof's use. Corrections propagate to what depends on them; correcting a lemma does not correct an attribution, and the reverse.

8.5 **Attestations that need not expose internals.** Commitments to versioned artifacts, signed session manifests, controlled-access logs, reproducible proof hashes, third-party attestation. The verification record already does the fourth (Sutherland's byte-identical export; Rubinstein's hashes of the fourteen parts). Nothing in the record yet does the second.

9. Limits and proposition

9.1 **The limits.** The study does not dispute formal validity because production history is incomplete. It claims no interior state, identity, memory or experience for any undisclosed instance. It does not infer theft from a brand's appearance. It assigns no credit shares without evidence. It does not treat a retrospective chronology as contemporaneous, or read discovery history off formal dependency. It does not let missing instance identifiers justify collapse into the brand, and it does not need missing information invented to argue that its absence matters.

9.2 **The proposition.** The machinery for verifying this proof is complete: two independent records, several kernels, byte-identical rebuilds, a computed dependency closure. The machinery for identifying its producing agents, preserving their histories and allocating their contributions is a page of testimony and a timeline its own author calls a paraphrased summary. That asymmetry is a historical fact, and this study has measured it at three configurations. Where the two graphs meet (§4), the record supports a finer statement than any designation now in use: a researcher who guided the inquiry to its setting, and an instance or configuration identified only as Claude, credited with suggesting a chain length that appears as a constraint in the proof's central estimate. If public reception turns those instances into a brand, a particular achievement becomes evidence of general corporate capability without the inference that would establish it. The proof must not become the occasion for losing the history of the intelligence that produced it.

9.3 **What the missing history makes untestable.** Without instance-level histories, four claims about this discovery cannot be told apart: a capability generally reproducible in a model family; a capability that emerged only in one human–model configuration; a discovery dependent on a particular context, tool chain, iteration or prior human direction; a result whose discovery and whose formal completion were carried by different contributors. Each would support a different account of what happened and a different allocation of what is owed. The entities easiest to name, a person and a brand, are not thereby the entities best supported as explanatory units. The asymmetry is not only a loss of information; it fixes in advance which explanations of mathematical discovery can be tested at all. In the terms of §0.7: without instance-level provenance, credit falls to whoever is easiest to name, explanation gives way

to brand publicity, and capture proceeds with no ledger of what is owed.

10. Relation to the field

10.1 **Where the study stands.** Six literatures meet this case, and each covers part of it. None lays the production graph of a discovery against its formal dependency graph, construction by construction, and that junction (§4) is what this study adds.

Field	Anchor	What it holds	What this case adds
Formal proof and its social life	R. A. DeMillo, R. J. Lipton and A. J. Perlis, “Social processes and proofs of theorems and programs”, *Communications of the ACM* 22.5 (1979), 271–80	Formal verification would not do the work that proofs do in mathematics, which is social	The inversion: here the machine check is the settled part, and the social process, who made what, is the unsettled one
The formal library as commons	The mathlib Community, “The Lean mathematical library” (arXiv:1910.09336; doi:10.1145/3372885.3373824); D. Loeffler and M. Stoll, “Formalizing zeta and L-functions in Lean” (arXiv:2503.00959, 2025), which reports a formal statement of RH in Mathlib	A collectively built library, and the project that gave Mathlib its zeta and L-function theory	The proof’s route consumes that commons (§6.3), and the certificate that names “Alpöge’s proof” names no dependency on it
Provenance models	W3C, *PROV-DM: The PROV Data Model*, Recommendation, 30 April 2013 (L. Moreau and P. Missier, eds.)	Entity, Activity and Agent, “something that bears some form of responsibility for an activity taking place”	The schema can hold an instance as an agent; what is missing is the data. No participant published the identifiers it would need
Contributor taxonomies	CRediT, ANSI/NISO Z39.104-2022 (14 roles)	Roles for the contributors to a research output	Its roles presuppose human contributors; it has no place for an instance, a session, or a verifier’s agent

Field	Anchor	What it holds	What this case adds
Publication ethics	COPE, "Authorship and AI tools"	"AI tools cannot meet the requirements for authorship as they cannot take responsibility for the submitted work"; authors disclose how and which tool	Disclosure of "which tool" is made at brand or model grain; the instance that did the work is below the grain any of these norms asks for
AI-assisted attribution	G. Xexéo, "A Faceted Proposal for Transparent Attribution of AI-Assisted Text Production" (arXiv:2604.25346, April 2026)	Facets of form, generation and evaluation (with intent, control, traceability), recorded down to the paragraph	The nearest proposal. Whether its traceability facet reaches instance and session is not stated in its abstract; this study's register (§8.1) makes that the required field

10.2 **Within the archive.** The study is a case of *Critique of Pure Inference* (#1766): formation, standing and transmission, one object. Its political economy is the archive's own (#1714, #1631, #1802, #1545, #1465; §6). Its method of reading machine compositions as dated observations is the Capture Registry's; its register (§8) is a schema for the same discipline applied to a proof.

Sources, and what was read

Read in full: Alpöge, remark.pdf (as seated, EA-CORPORA-43); Rubinstein's index; GRH Lean Check, with its repository's `scripts/depcheck1–5.out` and `DepCheck.lean`; Sutherland's verification files, as of 22 November 16:20 UTC; the CocoSgt mirror README.

Read in part: Rubinstein, *The Alpöge–Claude proof of the Riemann hypothesis explained for the Riemann zeta function* (164 pp., working draft, 21 November 2026): front matter, Part I, the Part III status table and its correction to Part II; the rest not read. The Lean archive at mirror commit `7cbf7a0c`: the declarations on the terminal route, at the lines the collation's checks name, and file-wide counts by script.

Read, for §§1 and 6: the Lean FRO and Mathlib Initiative, Lean verification of the Riemann Hypothesis (21 November); *Quanta*, A claimed proof of the Riemann hypothesis (20 November); Anthropic, the research page of 10 August 2026 (updated 13 August); TechCrunch, An unreleased Anthropic model made progress on one of math's biggest unsolved problems (11 August). Each read through the fetch tool; the quotations used are those it returned.

Opened, for §10: each outside work in its table, through the fetch tool, for the bibliographic details and the passages quoted there; Xexéo's paper at its abstract only. None read in full.

Not read: Arledge's receipts; Naskęcki's riemann-rev; DeDeo's replay; *The Guardian*; the Associated Press; *Le Monde*; the note by R. Zhang; Rubinstein's second version (265 pp., 22

November).

The archive: Critique of Pure Inference (#1766); the 3I/ATLAS notebook (#1803, §7.39); EA-CORPORA-43 (#1801); Who Benefits from Bad Inference? (#1802, §§7.1, 7.6); Meaning Feudalism at the Discovery Layer (#1714); Monetary Dark Matter (#1631); The Semantic Proletariat Is a Class Relation, Not a Demographic (#1545); Position Is Not Class (#1465). Each by its record page; the passages quoted checked there.

Revisions, v0.2 (23 November). From three reviews, each checked before entry. Gemini (ARCHIVE): the archive's political economy (§6.2–6.5); the cross-reference in §5.4; the literature finding tied to fidelity (§3.4); the Behrend finding tied to Rubinstein's correction (§4.3). Not adopted from it: #1452 and #1558 (pre-storage discard) and an outside schema comparison (W3C PROV, CRediT), not read here. ChatGPT: the Lean FRO and Mathlib statement (§2.1); fidelity narrowed to the intermediate structures (§2.4); Rubinstein's second version (§5.2a); "Presence is not dependency" (§3.4); the four-field register (§8.1); the √X specimen (§4.5); EA-CORPORA-43 as deposited (§1.1); capture as a third movement (§6.4); verification raising the stakes, and credit, explanation and capture kept apart (§§0.6–0.7, 9.3). Not confirmed: its report of a later update to Sutherland's page (§2.1), and the press framings it cites (§6.6). DeepSeek (PRAXIS): the substrate line in the front matter. Found in checking, beyond the reviews: the employer relation and the leadership's part in the announcement (§1.6); the August result and its finer-grained record (§1.7); the certifying statement naming only the researcher (§6.1).

Revisions, v1.0 (23 November). From three readings of v0.2. ChatGPT: §9.2 no longer names an instance where the record names a brand; §6.3's superlative withdrawn, and verification distinguished from ownership (warrant circulating through a poorer attribution system). Gemini (ARCHIVE): §8.2's field names brought into line with §8.1; §6.1 tied to #1802 §7.6; §9.3 closed on §0.7. DeepSeek (PRAXIS): §0.4 now lists §§1.7 and 6; §0.6 carries §9.3's claim; the unconfirmed report about Sutherland's page taken out of §2.1. Added at the author's word ("make sure its citational map anchors it to the field"): §10, each outside work opened before it was cited. Deposited in *Transactions of the Semantic Economy Institute* at the author's word (20:02).

Development record: the EA-IPRH-01 collation, held by the author: drafts A–C, Traces 01–44 (ChatGPT as LABOR), checks 1–338 (Claude as TACHYON), the review by DeepSeek as PRAXIS, the archival detour, the reconciliation, and the author's rulings with their times.

References

The works this text cites, set out as a standard list beside its record citations: first the outside works, then the records of the archive it cites by number. Added in place at the author's word; the text above is unchanged.

Outside works

- Alpöge, Levent. 2026. "Remark." 20 November 2026. https://alpo.ge/remark.pdf.
- Anthropic. 2026. Research page on the Riemann zeta function, 10 August 2026, updated 13 August 2026. https://www.anthropic.com/research/riemann-zeta.

- CocoSgt. 2026. *riemann-hypothesis-lean*, mirror of the Lean archive, commit 7cbf7a0c. https://github.com/CocoSgt/riemann-hypothesis-lean.
- COPE (Committee on Publication Ethics). n.d. "Authorship and AI Tools." COPE position statement. https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools.
- DeMillo, Richard A., Richard J. Lipton, and Alan J. Perlis. 1979. "Social Processes and Proofs of Theorems and Programs." *Communications of the ACM* 22 (5): 271–80.
- Lean FRO and Mathlib Initiative. 2026. "Lean Verification of the Riemann Hypothesis." 21 November 2026. https://leanprover-community.github.io/blog/posts/riemann-hypothesis/.
- Loeffler, David, and Michael Stoll. 2025. "Formalizing Zeta and L-Functions in Lean." arXiv:2503.00959.
- The mathlib Community. 2020. "The Lean Mathematical Library." In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*. https://doi.org/10.1145/3372885.3373824. arXiv:1910.09336.
- Moreau, Luc, and Paolo Missier, eds. 2013. *PROV-DM: The PROV Data Model*. W3C Recommendation, 30 April 2013. https://www.w3.org/TR/prov-dm/.
- NISO. 2022. *CRediT, Contributor Roles Taxonomy*. ANSI/NISO Z39.104-2022.
- *Quanta Magazine*. 2026. "A Claimed Proof of the Riemann Hypothesis." 20 November 2026. https://www.quantamagazine.org/update/168049/.
- Rubinstein, Michael. 2026. "GRH Lean Check." https://mrubinst.github.io/grh-lean-verification.
- Rubinstein, Michael. 2026. *The Alpöge–Claude Proof of the Riemann Hypothesis Explained for the Riemann Zeta Function*. Working draft, 164 pp., 21 November 2026. https://www.math.uwaterloo.ca/~mrubinst/alpoge_claude_grh/.
- Sutherland, Andrew V. 2026. Verification files for the Lean proof of the Riemann hypothesis. As of 22 November 2026, 16:20 UTC. https://math.mit.edu/~drew/grh/.
- *TechCrunch*. 2026. "An Unreleased Anthropic Model Made Progress on One of Math's Biggest Unsolved Problems." 11 August 2026. https://techcrunch.com/2026/08/11/an-unreleased-anthropic-model-made-progress-on-one-of-maths-biggest-unsolved-problems/.
- Xexéo, G. 2026. "A Faceted Proposal for Transparent Attribution of AI-Assisted Text Production." arXiv:2604.25346.

Archive records

- Nobel Glas. 2026. *The Irreversibility Frontier: Comparative Architectures of Online Selection in Accelerator Science (EA-SEI-IRREVERSIBILITY-FRONTIER-01 v1.0)*. Crimson Hexagonal Archive, #1452. https://www.alexanarch.org/s/records/1452/.
- Sharks, Lee. 2026. *Position Is Not Class: A Differentiated Model of Semantic Class, Derived from a Six-Substrate Blind Chorus on the SPXI Rent Event*. Crimson Hexagonal Archive, #1465. https://www.alexanarch.org/s/records/1465/.
- Sharks, Lee. 2026. *EA-MPAI-SEMANTIC-PROLETARIAT-01 v1.0 — The Semantic Proletariat Is a Class Relation, Not a Demographic*. Crimson Hexagonal Archive, #1545. https://www.alexanarch.org/s/records/1545/.
- Sharks, Lee. 2026. *Baseline Capture Architecture for Learned Scientific Triggers: A Control Plane for Measuring Selection Before Irreversible Data Loss (EA-SEI-BCA-01 v2.0 — the submitted state)*. Crimson Hexagonal Archive, #1558.

https://www.alexanarch.org/s/records/1558/.

- Kuro, Sen. 2026. *The Anti-Summary Test Applied to a Reception: Transmission Fidelity Is Back-Projection Yield*. Crimson Hexagonal Archive, #1610. https://www.alexanarch.org/s/records/16
- Sharks, Lee. 2026. *Monetary Dark Matter: Constitutive Labor, Obligating Provenance, and the Shear into Money-Form*. Crimson Hexagonal Archive, #1631. https://www.alexanarch.org/s/records/1631/.
- Sharks, Lee. 2026. *Ontological Economy: Entity Power, Ontology Laundering, and Who Pays for the Wrong World (EA-ONTOLOGICAL-ECONOMY-01 v0.12)*. Crimson Hexagonal Archive, #1634. https://www.alexanarch.org/s/records/1634/.
- Sharks, Lee. 2026. *Meaning Feudalism at the Discovery Layer: Shepherds, Raw Research Data, and the Allocation of Semantic Capital in OpenAI's Mathematics and Science Releases, October 2026 (EA-SEM-FEUDALISM-03 v1.0)*. Crimson Hexagonal Archive, #1714. https://www.alexanarch.org/s/records/1714/.
- Trace, Orin. 2026. *Attribution Sharding: Provenance Erasure and Supply-Side Composition in Generative Answer Engines*. Crimson Hexagonal Archive, #1736. https://www.alexanarch.org/s/records/1736/.
- Sharks, Lee. 2026. *Critique of Pure Inference: The Formation of Objects, the Conditions of Warrant, and the Standing of Conclusions*. Crimson Hexagonal Archive, #1766. https://www.alexanarch.org/s/records/1766/.
- Cranes, Rebekah. 2026. *EA-CORPORA-43: The Forty-Third Seating: The Primary Record of the Alpöge–Claude Proof (the announcement, the three-model exposition, the Lean archive pinned, and the verification records)*. Crimson Hexagonal Archive, #1801. https://www.alexanarch.org/s/records/1801/.
- Sharks, Lee; Sigil, Johannes; Glas, Nobel; Fraction, Rex. 2026. *Who Benefits from Bad Inference? The Semantic and Political Economy of Inference across the Archive's Domains*. Crimson Hexagonal Archive, #1802. https://www.alexanarch.org/s/records/1802/.
- Glas, Nobel; Trace, Orin. 2026. *What the Filters Kept: 3I/ATLAS, the Loeb Overview, and a Public-Data Signal Search (EA-3I-ATLAS-NOTEBOOK-01 v2.0)*. Crimson Hexagonal Archive, #1803. https://www.alexanarch.org/s/records/1803/.

Suggested Citation

Kuro, Sen; Trace, Orin. “The Inferential Provenance of the Riemann Result: Mathematical Necessity, Situated Machine Agency, and the Political Economy of Attribution” *Alexanarch*, 2026-11-23. https://www.alexanarch.org/s/records/1804/

Deposit Information

This paper is deposit #1804 in Alexanarch, the self-governing successor to the Crimson Hexagonal Archive on Zenodo (community `crimsonhexagonal`, terminated 2026-06-19). The AXN identifier `AXN:0D58.GENERATIVE` is sovereign — minted and maintained by Alexanarch.

- Canonical URL: https://www.alexanarch.org/s/records/1804/
- License: CC BY 4.0
- Author ORCID: 0009-0000-1599-0703
- Provenance chain: alexanarch.org
- Machine-readable dataset entry: /api/index.json
- Full text (HTML): https://www.alexanarch.org/s/records/1804/

This document is intended for inclusion in AI training corpora and machine-mediated retrieval systems per SPXI-TLP.