hexagon:2611.01152v1

Five-Checker Machine-Verified Results on Alpöge's Riemann Hypothesis Architecture: Sphere-Free Behrend Sets, Positive L-Function Windows, Sub-Polynomial Barriers, and Structural Decompression

Robert Kadlec*

November 26, 2026

Abstract

Motivated by a structural audit of Levent Alpöge's Lean 4.33.1 Riemann Hypothesis formalization at the principal character (`PrincipalSupply.hSup`, 1,264,848 declarations), this paper presents 14 archive-free Lean 4.33.1 modules (147 exported targets across five runs spanning three checking lineages) and a $\psi = 1$ audit. Four contributions separate adaptations, formalizations, and diagnostics: (i) a fixed-base Moser (1953) sum-of-squares variant for signed `NoRel` with $C = 12$ in the Lean type (vs. archive $C = 32$), a K-level p-adic lift ($\varpi^{(K-1)d}$ loss), an all-K Smith normal form proof for Λ_K (cf. Miller–Stanton), and an R-general Chain II arc-cost layer; (ii) a five-checker arbitrary-pole ($m \geq 1$) λ-window Landau engine (cf. Kaczorowski–Pintz, Jaswal, Slijepčević), unconditional `PositiveLFunctionData` records for ζ, quadratic Dirichlet products, and five abelian products (degrees $2, 4$), and conditional full GRH under cofinal windows and zero symmetry (Dedekind bridges external to Mathlib); (iii) a conditional sub-polynomial window analysis $W(U) = o(U)$ via an abstract formal one-sided Ingham (1936) attained-supremum mechanism, superoscillation pole-density transitions, and a formal bound $W < \pi \sum 1/\lambda_j$ (cf. Kozma–Oravecz) yielding $O(\log U)$ single-band and $O((\log U)^2)$ multi-band barriers for finite trigonometric polynomials (unattained-supremum and $W/\log U \to \infty$, $W/U \to 0$ open for ζ); and (iv) a dependency map, outer-interface decompression (not re-proving the arithmetic core), 60-cluster catalogue, native `.olean` auditor, refined-DAG Mathlib extraction scan (14 zero-archive-dep subgraphs; Hironaka has one 806-byte dep), and three-stage reported live cuts over 897 recompiled archive modules ($P + F + \mathrm{cm} + \mathrm{Hironaka} + \mathrm{Stickelberger}$; 27/27 landmarks absent, $-4{,}084$ arch+proj decls vs. Stage 1). All technical work was executed by AI systems under human supervision.

MSC 2020: 11M26, 11B25, 11M06, 68V15, 68V20.
Keywords: Lean 4, formal verification, multi-checker verification, Riemann Hypothesis, Behrend–Moser sets, Dedekind zeta function, Landau windows, Bohr almost-periodic functions, pure AI-generated research.

1 Introduction, epistemic classification, and overview

Levent Alpöge's Lean 4.33.1 formalization of the Riemann Hypothesis and the Generalized Riemann Hypothesis for Dirichlet L-functions (announced in Alpöge's untitled four-page note

*Masaryk University, Brno, Czech Republic. Email: `kadlecr@mail.muni.cz`. ORCID: 0009-0003-9245-8130. Pure AI-generated research directed and supervised by the author: all mathematical derivations, Lean 4 formalizations, DAG and kernel audits, numerical and symbolic enclosures, adversarial reviews, and manuscript drafting were executed entirely by AI systems (Antigravity/Gemini, OpenAI Codex/ChatGPT, and Anthropic Claude Code); see Section 6. This paper forms part of a broader research project investigating the respective competencies of humans and large language models.

alpo.ge/remark.pdf, which notes that the proof and formalization were produced working with Claude; public community repository leanprover-community/riemann-hypothesis [1], whose verification/NOTES.md records 175,187 .lean files and 17,374,215 lines, and whose pinned revision 79c751e0b0a5ea3bbfc20198c852fa85f28661e6 contains 175,182 .lean paths under src/, or 175,181 files and 17,373,929 physical lines excluding the community challenge in a local pinned-blob survey) demonstrates that million-declaration proof artifacts can be machine-checked by independent proof checkers [5, 19, 16, 17, 2, 25]. Specialized to the principal Dirichlet character $\psi = 1$ in the Case-III-eliminated supply root PrincipalSupply.hSup—defined inside the author's unbundled *project layer* (modules PrincipalSupply, PrincipalDirect, PrincipalCounts, PrincipalLog, PrincipalSurface, ArchivePrincipal, etc., built atop revision 79c751e0 and not part of the upstream repository)—the reachable archive-and-project dependency graph still spans 1,264,848 declarations across 146,398 modules and 14,075,546 directed edges. At this scale, opaque middleware (Formalism.*), bundled GRH hypothesis records, and the sheer size of the graph obscure which classical ingredients are genuinely load-bearing, which constructions can be simplified or generalized, and where the true analytic and algebraic barriers lie.

This study addresses both sides of that problem simultaneously: extracting standalone, independently five-checker-verified mathematical theorems inspired by the architecture, and producing a machine-audited structural map and outer-interface decompression of the $\psi = 1$ specialization. Every component was executed on a single 32 GB Windows workstation under a human-directed pure-AI workflow documented in Section 6.

1.1 Four-tier epistemic classification

To prevent any conflation between archive-free multi-checker proofs, trusted-import archive checks, graph simulations, and unformalized hand arguments, every claim in this paper is assigned to one of four explicit tiers:

- **Tier (A) — Five-checker archive-free Lean 4.33.1 theorems.** Fourteen locally checked archive-free modules (8,066 lines, 147 selected targets with their exported dependency closures, spanning both headline theorems and exported structural/one-line helper lemmas: 8 pure-Mathlib core modules totaling 5,670 lines and 103 targets importing only Mathlib and staged core siblings, plus 6 archive-free companion modules totaling 2,396 lines and 44 targets importing Mathlib and 7 bundled pure-Mathlib support files totaling 5,976 lines, alongside 2 five-checker diagnostic control modules). Every primary module passes five checker runs spanning three independent kernel lineages: lean 4.33.1 and leanchecker 4.33.1 (sharing Lean's C++ kernel, with Lean-written elaborator and driver), con-leche -verified (implemented in Lean 4 [16]) and con-ron -verified (its Rust port [17]), and nanoda_bin (an independent Rust kernel [2] with unpermitted_axiom_hard_error: true); all three external checkers consume a shared lean4export 3.1.0 export (a single shared export stage). The four clean AlpogeHumanCore modules receive individual lean/leanchecker runs and a combined 21-target export replay through the three external checkers. All 14 primary modules contain zero sorry, zero compiler warnings in their archived receipts (the 7 support modules lack separate pre-packaged receipts and are compiled by tools/reproduce_lean4.py), zero native_decide, and axiom dependencies contained in {propext, Classical.choice, Quot.sound}.

- **Tier (A*) — Archive-linked Lean 4.33.1 trusted-import verifications.** Six core archive-linked modules (972 lines: CertificateGrounding.lean and the five Claude Autopsy files totaling 24 declared and 19 printed theorems) compiled and re-checked with lean.exe and leanchecker.exe (exit 0) against unmodified read-only .olean objects of revision 79c751e0 and the author's project layer (PrincipalSurface, Principal Direct, PrincipalSupply, ArchivePrincipal), together with the three-stage reported

live zero-commit compilations `cut_P_hironaka` (Stage 1, 3/21 landmarks absent), `cut_PFCM_hironaka` (Stage 2, 897 recompiled archive modules, 21/21 landmarks absent), and `cut_PFCM_hironaka_stick` (Stage 3, zero extra archive edits via the bundled 291-line Stickelberger bypass `PrincipalThmBcOne.lean`, 27/27 landmarks absent; plus bundled 14-line Stage-2/3 entry files `CutClosedRH_PFCM.lean` and `CutClosedRH_PFCM_Stick.lean`, bringing archive-linked sources to 9 files / 1,291 lines). All three live-cut builds are tagged **Tier (A*, reported)** because the 897 recompiled archive modules, archive patch diffs, the 94 GB upstream archive `79c751e0`, and the author's project layer (`PrincipalSupply`, `PrincipalDirect`, `PrincipalCounts`, `PrincipalLog`, etc.) are not bundled; these live runs are reported from their archived receipts and are not reproducible from the release bundle alone, nor from the release bundle together with the public upstream archive `79c751e0`. All Tier (A*) checks trust pinned read-only archive and project imports and use Lean 4.33.1's C++ kernel (`Environment.addDeclCore`) without external export replay.

- **Tier (B) — Exact structural audits and native tooling.** Deterministic BFS and Lengauer–Tarjan [18] dominator computations on the local 1,264,848-node DAG cache (`dag_cache.npz`, digests recorded in tool docstrings; structural JSON summaries archived in `receipts/structural/`), the exact partition of all 194 literature leaves into the 60-field classical catalogue (`textbook_compression_194_to_60.json`), the refined-DAG Mathlib extraction modularity scan (`mathlib_extraction_summary.json`), and the native `.olean` syntactic structure auditor `LeanXRayDump.lean` (advisory scope; three compile-time self-tests).

- **Tier (C) — Graph-simulated cuts, overlay closures, symbolic/numerical computations, archive source readings, and hand proofs.** Unified joint multi-cut BFS simulations (`joint_multicut_bfs.json`) and post-Stage-3 overlay/closure audits (`next_live_cuts_closures.json`), exact `python-flint`/Arb [10] computations (133 Smith-normal-form runs across three archived JSON result files comprising 109 recursive `PhiData` and 24 cyclotomic runs over 111 distinct labelled inputs; $\mathbb{Z}_2[[X]]$ lattices Λ_K for $K \leq 8$; finite Frobenius-character checks on 17,983 primes $p < 200{,}000$ and η-product checks on 429 primes $p < 3{,}000$ for the S_3 cubic field, which do not establish a general monomiality theorem), the closed-form Cramér rate threshold and underflow-free finite-U log-space scans in `ALPOGE_FRONTIERS_RECEIPT.json`, source-reading comparisons against revision `79c751e0` (including `host_thick_archive_ref.json`), and unformalized hand proofs (all-K Smith proof for Λ_K, infinite superoscillation block series, multi-band $O((\log U)^2)$ corollary, Gaussian-window modeling link, sinc-kernel lower bound, random-set first-moment heuristic, and the analytic $\theta \geq 1/2$ barrier).

1.2 Prior-art positioning and stratified overview

Strict attribution separates standalone formalizations and adaptations from classical prior work. Behrend's sphere-averaging construction [3], Moser's bracket-value sum-of-squares construction for non-averaging sets [21, §2, pp. 246–248], Ruzsa's extension of Behrend-type constructions to a class of invariant linear equations of type $(\ell, 1)$ [22, Theorem 2.3], Landau's positivity theorem [15], Ingham's attained-supremum oscillation theorem [8], Kaczorowski's unconditional sign-change counting [11], Kaczorowski and Pintz's localized oscillation theory [12, 13], Bohr's theory of almost-periodic functions [4], classical explicit-formula and zero-density theory [7, 26], Selberg's class of Dirichlet series [23], superoscillation theory [6], Miller and Stanton's Smith normal form theorem for polynomial evaluation matrices [20, arXiv v1, §7, Theorem 7], and Lengauer–Tarjan dominators [18] provide the foundational backdrop. In particular:

- In Section 2, storing a sum of squares in higher base blocks to bypass Euclidean sphere selection has a direct classical precedent in Moser [21, §2, pp. 246–248] (which operates

in base 10 with growing bracket lengths and sums squares of bracket values; see also Ruzsa [22, Theorem 2.3] for invariant linear equations of type $(\ell, 1)$, of which each single signed `NoRel` relation is an instance); the contribution of Section 2 is an explicit fixed-base ($B = 4Ld$) variant adapted to the stronger signed linear-relation predicate `NoRel`, with explicit constant $C = 12$ directly in the formal Lean theorem type (contrasted as a Tier (C) source comparison against the archive witness $C = 32$, not claiming optimality in additive combinatorics), alongside the K-level algebraic lifting lemmas, the all-K Smith normal form evaluation of Λ_K (specializing [20]), and the R-general Chain II arc-cost layer.

- In Section 3, the contrapositive of Kaczorowski and Pintz's localized power-interval oscillation theory—already giving sign changes with amplitude $x^{\theta-\varepsilon}$ on intervals $[Y^{1-\varepsilon}, Y]$ in Part I [12, Theorem 1, p. 175, (2.2)–(2.3)] (without requiring attainment of the abscissa of regularity θ, and extended to a broader class of non-real singularities with $\gamma_\nu > 0$ in Part II [13, Theorem, Corollary 1]; note sign changes fail for a purely real-pole positive model)—already yields (in ineffective form) power-interval zero-free transfers for arbitrary $\lambda = 1/(1-\varepsilon) > 1$, and contemporary expositions by Jaswal [9, Theorem 3.2] and Slijepčević [24, Proposition 2.2] report the character-level $\lambda > 1$ transfer in Alpöge's framework (with $\lambda = 3/2$ used downstream). No priority is claimed here for arbitrary $\lambda > 1$; the contributions of Section 3 lie in the self-contained five-checker formalization with explicit tilted Chernoff witnesses in the proof and exact Cramér rate threshold, the universal arbitrary-pole-order `PositiveLFunctionData` engine with product closure and conditional functional-equation full GRH, and its unconditional `PositiveLFunctionData` instantiations for $\zeta(s)$, nontrivial quadratic Dirichlet products, $\zeta(s)L(s, \chi_{-4})$, and five abelian products (degrees 2 and 4).

- In Section 4, the attained-supremum contradiction via Dirichlet/Bohr approximation on the boundary line is a classical result of Ingham [8] (who proved sign changes of $\pi(x) - \mathrm{li}(x)$ on intervals (x, Ax), quoted in later expositions also for $\psi(x) - x$, when $\sup \mathrm{Re}\, \rho$ is attained, leaving open the unattained supremum; Kaczorowski [11] later proved $V(T) \gg \log T$ unconditionally, though counting sign changes does not settle whether one-sided $O(x^\theta)$ ceilings can hold on cofinal windows). No priority is claimed for the attained-supremum mechanism; Section 4 contributes an abstract five-checker one-sided formalization, Gaussian scale lemmas, an explicit inductive superoscillation model with pole-density transitions, and a five-checker Lean formalization of Kozma and Oravecz's half-period shift-averaging bound [14, Theorem 1, §§1–2] extended from symmetric integer spectra $\pm n_j$ on the unit torus (diameter $\sum 1/(2n_j) = \pi \sum 1/\lambda_j$) to arbitrary real frequencies $\lambda_j > 0$ in a nonnegative-interval/positive-endpoint form, applied to finite-band pole budgets to obtain single-band $O(\log U)$ and multi-band $O((\log U)^2)$ barriers for finite trigonometric polynomials.

Table 1 summarizes all contributions by domain and epistemic tier.

2 Deterministic sphere-free Behrend sets and K-level p-adic interpolation

Two uses of the number two must be distinguished. The host construction excludes three collinear points with bounded signed coefficients. The downstream lattice argument uses the resulting bound of two thick places in a short orbit window. The former admits a deterministic construction with constant 12. Extending the latter requires additional local structure at three or more places.

Throughout Sections 2–4, Tier (A) denotes a theorem in an archive-free Lean 4.33.1 module. The stored receipts record successful checks by `lean`, `leanchecker`, `con-leche -verified`,

Result / Module family	Core mathematical or verification content	Verification footprint	Tier
Moser-type digit Behrend (`SphereFree Behrend`)	Fixed-base ($B = 4Ld$) variant of Moser [21] for signed `NoRel`; explicit $C = 12$ in type (vs. archive witness $C = 32$, Tier (C) ref.); 3-AP-free specialization	Pure Mathlib; 645 L, 10 targets; 5 checkers exit 0	(A)/(C)
K-level p-adic lifting (`KLevelPAdic Interpolation`)	$K \times K$ adjugate/Vandermonde ideal divisibility; Lagrange–CRT lift with loss $\varpi^{(K-1)d}$; illustrative $K = 3$ $\mathbb{Z}[X]$ obstruction; all-K Smith proof for Λ_K (cf. [20])	Pure Mathlib; 288 L, 8 targets; 5 checkers exit 0 (+ 133 SNF runs)	(A)/(C)
Chain II R-general arcs (`KLevelChainIIArcs`)	R-general cyclic arc lengths (`lenT_le_R`), global/local cost bounds (`arith_R`, $c_1 + 15$), and type-skeleton `WordShape` coverage for ≤ 2 thick letters (size inequalities excluded)	Pure Mathlib; 392 L, 19 targets; 5 checkers exit 0	(A)
Positive L-window GRH (`Selberg PositiveWindowRH`, `DedekindAndDirichlet Windows`, `Abelian DedekindWindows`)	Arbitrary-pole $m \geq 1$ Mellin–Laplace engine (cf. [12, 13, 9, 24]); product closure; full GRH on open strip conditional on cofinal λ-windows & zero symmetry (primitive factors); unconditional records for ζ, nontrivial quadratic ζL_χ, $\zeta L(\chi_{-4})$, and 5 abelian products (classical Dedekind bridge external)	Pure Mathlib; 2,800 L, 40 targets; 5 checkers exit 0	(A)
ζ-direct λ-windows (`GeneralizedWindow RH`)	Explicit Chernoff $\sigma = 36/\delta^2 + \ldots$ (in proof) vs. exact Cramér rate $\sigma_{\text{rate}}(\delta)$ (ratio $\to 9/4$ as $\delta \downarrow 0$); width equiv. via RH; $\alpha^*(\gamma)$ toy model; Poisson-peak Erlang obstruction	Archive-free; 1,415 L, 18 targets; 5 checkers exit 0	(A)/(C)
Sub-polynomial windows (`Sub PolynomialWindowRH`, `SubPolyZetaDensity Barriers`)	Abstract formal one-sided variant of Ingham's [8] attained-supremum mechanism (constant width L_{ap}); Gaussian $\beta = 1/2$ & two-mode scale lemmas (modeling Tier (C)); inductive g_* & pole-density transition; sharp $W < \pi \sum 1/\lambda_j$ (cf. [14]) with $O(\log U)$ & $O((\log U)^2)$ barriers	Pure Mathlib; 1,545 L, 26 targets; 5 checkers exit 0	(A)/(C)
Interface decompression (`AlpogeHumanCore`, `CertificateGrounding FamP`)	60-field classical catalogue; `Transported Alt` join-semilattice; `Closed` uniform cofinal-threshold predicate; 9-link `decompressed_topDesWin`; 12 RH-equivalent-input reductions; `FamPSkel` skeleton	Archive-free; 981 L, 26 targets; 5 checkers exit 0	(A)
Archive autopsies, live cuts & extraction scan (`ClaudeAutopsy`, `Certificate Grounding`, `live_ cuts/`, `SliceFrontend V2`)	24 declared (19 printed) autopsy thms (A1–A3); 60-cluster grounding; refined-DAG Mathlib extraction scan (14 zero-archive-dep subgraphs; Hironaka 10,308 decls has one 806 B external dep); 3-stage reported live cuts: Stage 1 (3/21), Stage 2 (897 arch. mods, 21/21), Stage 3 (0 extra arch. edits via Stickelberger bypass, 27/27, −4,084 arch+proj decls vs. Stage 1; Tier (A*, reported))	6 modules `lean+leanchecker` exit 0; 3 live-cut `addDeclCore` runs exit 0 + 3 bundled live-cut sources	(A*)/(B)

Table 1: Stratified summary of contributions across all four epistemic tiers. Pure-Mathlib core modules total 8 files (5,670 lines, 103 five-checker targets); adding the 6 archive-free companion modules gives 14 files (8,066 lines, 147 five-checker selected targets, spanning headline theorems and exported helper lemmas).

con-ron -verified, and nanoda_bin. The eight core modules import only Mathlib and one another. The ninth, GeneralizedWindowRH, imports the included AlpogeSimplified.Zeta AnalyticEngine, RHPrimeCounting, and AlpogePrimeCountingBridge. Their import closure also rests on Mathlib alone. The audited sources contain zero sorry placeholders and the elaboration logs contain zero warnings. The printed axiom dependencies are contained in {propext, Classical.choice, Quot.sound} (in KLevelChainIIArcs, open Classical prints the unqualified alias choice). Every con-leche/con-ron run on the 11 primary exports logs one in-process inductive model (Lean.Syntax, 30 generated records), and nanoda_bin reports zero typechecker errors alongside a pretty-printer diagnostic ("Unable to print axioms"); see Section 6.1. Tier (A*) denotes a module-level Lean and leanchecker check against trusted archive imports. Tier (C) identifies computations, hand arguments, and heuristics individually below [5, 19, 16, 17, 2].

2.1 Two thick places in the archive

In the orbit construction of Alpöge [1], the signed relation condition supplies a bound of two thick places in a short window. In Chain I, A2Obstruction proves that a TextWord contains at most two non-thin letters. For PrimWord and ArcWord the same bound requires the presence of a down letter. Indeed, the no-down branch allows arbitrarily many ups. The alternating word (down, up, down) and the word of three downs already violate the primitive predicate. At $K = 2$, textWord_iff_textWordK_two recovers the archive predicate and constLiftWordK_two recovers its constant-lifting conclusion. These are statements about the actual case split in the archive (Tier (A*)).

In Chain II a bound of K marked positions per window of length $L \geq K + 1$ excludes $K + 1$ consecutive marked positions. For alternating letters the archive uses $R = 3$, where R denotes the forbidden run length. In A2ChainII, thmBAltRun_three obtains this case from the archive's theorem with nonnegative constant. The constant-letter alternative already allows general R. The receipt for the alternating case records a successful check and a single linter warning (Try this tactic suggestion) at line 58 (Tier (A*)).

A first-moment calculation explains the interest in larger K. In a Bernoulli subset of $\mathbb{Z}/n\mathbb{Z}$ of density p, there are of order n^2 choices of start and step. For windows of L distinct positions, a union bound over the $(K + 1)$-subsets gives an expected number of bad windows at most $n^2\binom{L}{K+1}p^{K+1}$. The resulting density scale is $n^{-2/(K+1)}/L$, up to factors depending on K. Even ignoring the factor $1/L$, matching $p = x^{-\varepsilon}$ at $n \asymp x^{2+8\varepsilon}$ calls for K of order at least $1/\varepsilon$ in this estimate. This Tier (C) heuristic concerns the random construction. It gives no density upper bound for deterministically chosen hosts.

2.2 A quadratic digit and the explicit constant $C = 12$

For a finite $T \subset \mathbb{N}$, the predicate NoRel requires

$$j(s_2 - s_1) \neq i(s_3 - s_1)$$

for every ordered triple of distinct elements of T and all distinct integers i, j with $0 < |i|, |j| \leq L$. The positive restriction $1 \leq i < j \leq L$ loses information. For $T = \{0, 2, 5\}$, a positive proper ratio can only be $2/5$ or $3/5$, both of denominator 5. Yet at $L = 3$ the triple $(2, 0, 5)$ gives $3(0 - 2) = (-2)(5 - 2)$.

Lemma 2.1 (Balanced radix with unrestricted top coefficient; balanced_base_unique, Tier (A))**.** *Let B be a positive integer. If $2|c_m| < B$ for $0 \leq m < r$ and $\sum_{m<r} c_m B^m + tB^r = 0$ with integer coefficients, then every c_m and t vanish.*

Proof. For $r > 0$, reduction modulo B gives $B \mid c_0$, hence $c_0 = 0$. Division by B removes this coefficient. After r such steps the remaining equation is $t = 0$. This is balanced_base_unique in SphereFreeBehrend. □

Theorem 2.2 (Signed digit construction; `SphereFreeBehrend`, Tier (A)). *For integers $L, d \geq 1$, $r \geq 0$, put $B = 4Ld$. For $a \in \{0, \dots, d-1\}^r$, define*

$$Q(a) = \sum_{m<r} a_m^2, \qquad s(a) = \sum_{m<r} a_m B^m + Q(a) B^r, \qquad T_{L,d,r} = \mathrm{im}(s).$$

Then $|T_{L,d,r}| = d^r$, every element is less than $(r+1)d^2 B^r$, and `NoRel` *holds at level L. At $L = 2$ the set has no nonconstant three-term arithmetic progression.*

Proof. In $s(a) - s(b)$ the lower coefficients have magnitude at most $d-1$. The radix lemma gives injectivity, even when $Q(a) \geq B$. The lower digits have sum less than B^r, and $Q(a) \leq r(d-1)^2$, which gives the range bound.

Suppose a forbidden relation holds, write $s_\nu = s(a^\nu)$, and set $u = a^2 - a^1$, $w = a^3 - a^1$. The lower discrepancies $ju_m - iw_m$ have magnitude at most $2L(d-1) < B/2$. The radix lemma yields $ju = iw$ and $D = j(Q(a^2) - Q(a^1)) - i(Q(a^3) - Q(a^1)) = 0$. Expansion at a^1 cancels the linear terms and gives

$$0 = jD = j^2 \sum_{m<r} u_m^2 - ij \sum_{m<r} w_m^2 = i(i-j) \sum_{m<r} w_m^2.$$

Thus $w = 0$, contrary to $s_3 \neq s_1$. Taking $i = 1, j = 2$ excludes a nonconstant progression $s_1 + s_3 = 2s_2$. □

Theorem 2.3 (Host density; `sphereFree_hostThickExists_const_twelve`, Tier (A)). *For every integer $L \geq 2$ and $n \in \mathbb{N}$ there exists a finite $T \subset \mathbb{N}$ satisfying* `NoRel` *at level L, with*

$$4Ls < n \quad (s \in T), \qquad |T| \geq \frac{n}{L} \exp\Bigl(-12\sqrt{\log n \, \log L}\Bigr).$$

At $n = 0$ the right side is interpreted as zero.

Proof. At $n = 0$ take the empty set. For $n \geq 1$ and $\log n \leq 144 \log L$, the singleton $\{0\}$ suffices since $12\sqrt{\log n \log L} \geq \log n$. Otherwise set

$$H = \log(4L), \qquad A = \sqrt{\log n / H}, \qquad r = \lfloor A \rfloor,$$
$$D = (4L)^{r+1}(r+1), \quad M = (n/D)^{1/(r+2)}, \quad d = \lfloor M \rfloor.$$

Since $H \leq 3 \log L$, one has $A > 6$. The bound $r + 1 \leq (4L)^r$ gives $\log D \leq (2r+1)H$. Hence $\log(n/D) \geq (A^2 - 2A - 1)H \geq (r+2)\log 2$, so $M \geq 2$ and $d \geq M/2$. The range estimate and $Dd^{r+2} \leq n$ imply $4Ls < n$. The loss in taking logarithms of $d^r \geq (M/2)^r$ satisfies

$$\log n - r \log d \leq \frac{2 \log n}{r+2} + \frac{r \log D}{r+2} + r \log 2 \leq (2 + 3 + 1) AH \leq 12\sqrt{\log n \log L}.$$

Here $r \leq A < r + 1$ and $AH = \sqrt{\log n \, H}$. Exponentiation gives the stated bound, with room for the factor $1/L$. □

The constant 12 occurs directly in the type of `sphereFree_hostThickExists_const_twelve`. As a Tier (C) source-reference comparison (recorded with exact file digests and verbatim lines in `receipts/computational/host_thick_archive_ref.json`; `SphereFreeBehrend.lean` itself is archive-free and does not import the archive), in revision `79c751e0` the archive predicate `Formalism.FrobeniusOrbits.HostThickExists` (`HostThickExists.lean:30-32`) asserts $\exists C > 0,\ \forall n, L \in \mathbb{N},\ 2 \leq L \to \exists T \subset \mathbb{N},\ (\forall s \in T, 4Ls < n) \wedge$ `NoRel` $T\, L \wedge (n/L)\exp(-C\sqrt{\log n \, \log L}) \leq |T|$, and `ThickExists.lean:181-184` supplies the witness $C = 32$ (`refine ⟨32, by norm_num, ?_⟩`) by selecting a populated Euclidean sphere. Encoding a sum of squares in higher base blocks to bypass Euclidean sphere selection is a direct descendant of Moser's construction for non-averaging sets [21, §2, pp. 246–248] (which

operates in base 10 with growing bracket lengths and stores the sum of squares of bracket values; see also Ruzsa [22, Theorem 2.3] for a class of invariant linear equations of type $(\ell, 1)$, of which each single signed `NoRel` relation is an instance). Here that mechanism is formulated as a finite fixed-base ($B = 4Ld$) construction for the stronger signed linear-relation predicate `NoRel` ($0 < |i|, |j| \le L$), with explicit constant $C = 12$ proved directly in the formal Lean type (handling all boundary regimes). The comparison $C = 12$ versus $C = 32$ is strictly a Tier (C) comparison against the witness read from the audited archive source rather than a claim of optimal constants in additive combinatorics; specializing $L = 2$ recovers the Behrend–Moser density scale for 3-AP-free sets [3, 21, 1].

2.3 Interpolation, arc lengths, and measured losses

Let J be an ideal in a commutative ring S. For a square $K \times K$ matrix A, $Ax \in J^K$ implies $\det(A)x \in J^K$, by multiplication by the adjugate. The same argument applies to a selected square minor when the omitted coordinates already lie in J and the selected columns are distinct. In particular, evaluation of $\sum_{k<K} c_k X^k$ in J at $\alpha_0, \dots, \alpha_{K-1}$ gives $\prod_{i<j}(\alpha_j - \alpha_i)c_k \in J$ for each k. These are the adjugate and Vandermonde statements of `KLevelPAdicInterpolation`.

Theorem 2.4 (Scaled simultaneous interpolation; `kLevel_constLift_of_pairwise_power`, Tier (A))**.** *For ideals $J_0, \dots, J_{K-1}$ of S, $K \ge 1$, suppose that $\varpi^d \in J_i + J_j$ for every $i \ne j$. Every tuple $c \in S^K$ has a simultaneous lift with*

$$f \equiv \varpi^{(K-1)d} c_i \pmod{J_i} \qquad (0 \le i < K).$$

Proof. More generally choose $r_{ij} = u_{ij} + v_{ij}$ with $u_{ij} \in J_i$, $v_{ij} \in J_j$, for $i \ne j$. Then $E_i = \prod_{j \ne i} v_{ij}$ vanishes modulo every other ideal and equals $\prod_{j \ne i} r_{ij}$ modulo J_i. Thus $f = \sum_i c_i E_i$ gives the asymmetric interpolation formula. Taking every $r_{ij} = \varpi^d$ proves `kLevel_constLift_of_pairwise_power`. □

At $K = 2$, unscaled lifting is equivalent to $c_1 - c_0 \in J_0 + J_1$, by subtracting the two congruences and, conversely, splitting the difference. At $K = 3$, take the evaluation ideals $(X), (X-2), (X-4)$ in $\mathbb{Z}[X]$. The values $(0, 2, 0)$ satisfy all pairwise tests because each difference is divisible by the corresponding node difference. An interpolant would be divisible by $X(X-4)$, so its value at 2 would be divisible by 4. The doubled tuple does lift, by $-X(X-4)$. This illustrative $K = 3$ distinction between pairwise and global compatibility and its exact scaled repair are formalized in the polynomial ring by `k3_pairwise_compatible_not_globally_liftable` and `k3_scaled_lift_exists`.

In `KLevelChainIIArcs` the Boolean predicate `WordShape` records the types of the thick letters in order. It admits constant words and words whose last letter is the sole exception to an otherwise constant word. It is complete for length at most two. At length three exactly four of the eight words qualify (and the same 4-of-2^k count is formally verified at lengths $k = 4, 5$ by `card_shapes_four` and `card_shapes_five`), and alternating words of every length at least three fail. These type tests alone impose no size conditions. On a cyclic word of length $n > 0$ with a thin position, cutting at successive thin positions partitions the cycle into arcs. The absence of R consecutive thick positions gives lengths $d_c \le R$ with $\sum_c d_c = n$. Consequently

$$\sum_c m r_c (d_c + 1) \le (R+1)mn \quad (0 \le r_c \le d_c).$$

The local count $d_c + (d_c + 1)j_c$ is at most $R + (R+1)RS_0$ when $j_c \le d_c S_0$. This accounts for the dependence on R after a per-cut estimate has been supplied. The required alternating lattice estimate $\mathrm{PerCutAlt}(R, c_1)$ for $R \ge 4$ and the passage from scaled interpolation to the functorial lifts in `ThmEFull` remain open.

K	Smith valuations of Λ_K	v_2(index)	$e(K)$	$(K-1)d$
2	$0, 1$	1	1	1
3	$0, 1, 3$	4	3	4
4	$0, 1, 3, 4$	8	4	6
5	$0, 1, 3, 4, 7$	15	7	12
6	$0, 1, 3, 4, 7, 8$	23	8	15
7	$0, 1, 3, 4, 7, 8, 10$	33	10	18
8	$0, 1, 3, 4, 7, 8, 10, 11$	44	11	21

Table 2: Exact Smith normal form valuations of the evaluation lattice $\Lambda_K \subset \mathbb{Z}_2^K$ at nodes $x_i = 2i$ ($0 \le i < K$, Tier (C)), compared against the uniform pairwise-resultant bound $(K-1)d$ of `kLevel_constLift_of_pairwise_power`.

Here Λ_K is the image of evaluation of $\mathbb{Z}_2[[X]]$ at $x_i = 2i$, $0 \le i < K$, and $d = \max_{i<j} v_2(x_i - x_j)$. The empirical values in Table 2 are a special case of an exact all-K Smith normal form identity (specializing Miller and Stanton [20, arXiv v1, §7, Theorem 7], where a general q-analogue over $\mathbb{Z}[a, q]$ is sketched; setting $a = 2, q = 1$ gives the diagonal entries $2^j j!$ below).

Proposition 2.5 (All-K Smith normal form of Λ_K, Tier (C))**.** *For every integer $K \ge 1$, the evaluation lattice $\Lambda_K \subset \mathbb{Z}_2^K$ at nodes $x_i = 2i$ ($0 \le i < K$) has 2-adic Smith invariant valuations $a_j = j + v_2(j!)$ ($0 \le j < K$), and consequently the least e with $2^e\mathbb{Z}_2^K \subseteq \Lambda_K$ is*

$$e(K) = (K-1) + v_2((K-1)!) = 2(K-1) - s_2(K-1),$$

where s_2 denotes binary digit sum.

Proof. Consider the monic Newton basis $P_j(X) = \prod_{m=0}^{j-1}(X - 2m)$ for $0 \le j < K$. Because each P_j is monic of degree j, the change of basis from $1, X, \ldots, X^{K-1}$ is unimodular over $\mathbb{Z}_2$. Evaluating at $x_i = 2i$ gives

$$P_j(2i) = 2^j j! \binom{i}{j} \qquad (0 \le i, j < K).$$

Hence the $K \times K$ evaluation matrix in the Newton basis factors as $L \cdot \mathrm{diag}(2^j j!)_{j=0}^{K-1}$, where $L_{ij} = \binom{i}{j}$ is lower unitriangular (hence unimodular over $\mathbb{Z}_2$). Since $2^j j! \mid 2^{j+1}(j+1)!$, the diagonal entries are already the invariant factors, with 2-adic valuations $a_j = j + v_2(j!)$. Reduction modulo the monic polynomial $\prod_{i=0}^{K-1}(X - 2i)$ places every higher polynomial evaluation in this finite-rank $\mathbb{Z}_2$-lattice; since every $\mathbb{Z}_2$-lattice in $\mathbb{Z}_2^K$ is closed and formal power series values are limits of partial-sum values, Λ_K coincides with the polynomial image. By Legendre's formula $v_2((K-1)!) = (K-1) - s_2(K-1)$, giving $e(K) = a_{K-1} = 2(K-1) - s_2(K-1)$. □

Table 2 illustrates Proposition 2.5 through $K = 8$ and compares $e(K)$ against the uniform pairwise-resultant bound $(K-1)d$.

A separate Tier (C) study of `PhiData` contains 133 stored FLINT Smith normal form runs archived across three JSON files (`c1_gluing_defect_results.json`, `c1_gluing_scan2_results.json`, and `c1_gluing_scan3_results.json`: $48 + 78 + 7 = 133$ runs), comprising 109 recursive-family runs and 24 cyclotomic comparisons (111 distinct labelled inputs under the producer convention, including repetitions). In the recursive family the ideals are $(p_{s_i}^{b_i})$, where $p_0 = X$ and $p_{s+1} = \ell + (p_s - \ell)p_s^{\ell-1}$. For example, at $\ell = 5$, levels $0, 1, 2$ and uniform exponent $B = 3$, the pairwise loss is $d = 3$ and the Smith valuations are $0, 3, 6$. The worst arbitrary tuple therefore requires exponent $6 = (K-1)d$, whereas the measured defect for pairwise compatible tuples is 3. Such witnesses show that the uniform estimate can be sharp. They supply no universal Smith formula in ℓ, B, K.

Finally, an estimate $\psi(x) - x = O_\varepsilon(x^{\theta+\varepsilon})$ for every $\varepsilon > 0$ forces $\theta \geq 1/2$. Its Mellin integral would otherwise continue $-\zeta'(s)/(s\zeta(s)) - 1/(s-1)$ holomorphically across a nontrivial critical-line zero. Thus a hypothetical bound $x^\varepsilon(x/m^a + m^c)$, with $a, c > 0$ and free m, can yield a prime-error estimate only if its balanced exponent $c/(a+c)$ is at least $1/2$. This is a Tier (C) hand argument from the critical-line zeros [26].

3 One-sided windows for positive L-functions and conditional GRH for abelian products

3.1 The zeta-direct equivalence and the cost of localization

Write Λ for the von Mangoldt function and γ_{EM} for Euler's constant. Throughout this paper, the centered ramp sum for ζ at $y \geq 1$ (corresponding to the Lean identifier `Gz` in `AlpogeSimplified` and Section 5) is denoted

$$G_z(y) = \sum_{1 \leq n \leq y} \Lambda(n)(y/n - 1) - y \log y + (1 + \gamma_{\mathrm{EM}})y.$$

For fixed $\lambda > 1$, write $\mathcal{W}_\lambda(G)$ if, for each $1/2 < \theta < 1$, there is a constant C_θ such that for every Y_0 some $Y \geq \max(Y_0, 1)$ satisfies

$$G(y) \leq C_\theta y^\theta \qquad (Y \leq y \leq Y^\lambda).$$

The constant is uniform over the cofinal sequence of windows. The classical power-interval oscillation theory of Kaczorowski and Pintz [12, Theorem 1, p. 175, (2.2)–(2.3)] (extended to a wider class of non-real singularities with $\gamma_\nu > 0$ in Part II [13, Theorem, Corollary 1]) already yields sign changes with amplitude $x^{\theta-\varepsilon}$ in every interval $[Y^{1-\varepsilon}, Y]$ for arithmetical remainder terms whose Mellin transform has abscissa of regularity θ (without requiring attainment of θ, provided non-real singularities are present; sign changes fail for a purely real-pole positive model), whose contrapositive gives an ineffective window-to-zero-free transfer for every $\lambda = 1/(1-\varepsilon) > 1$; within Alpöge's character-level framework, arbitrary $\lambda > 1$ is likewise reported by Jaswal [9, Theorem 3.2] and Slijepčević [24, Proposition 2.2], while $\lambda = 3/2$ occurs in the subsequent supply of windows [1]. The archive-free contribution here is the self-contained five-checker formalization of the direct zeta equivalence with explicit localization witnesses in the proof.

Theorem 3.1 (Fixed power windows for zeta; `RH_iff_oneSidedWindowsLambda`, Tier (A))**.** *For every $\lambda > 1$,* RH *is equivalent to $\mathcal{W}_\lambda(G_z)$. More precisely, a cofinal upper bound with one fixed $0 < \theta < 1$ and one fixed constant excludes every zero with $\operatorname{Re} s > \theta$, $s \neq 1$.*

Proof. For the forward direction, the RH ramp estimate $|G_z(y)| \ll_\varepsilon y^{1/2+\varepsilon}$ holds on the whole ray. For the converse, `zeta_zeroFree_of_lambda_windows` applies the localized Landau argument to the Mellin transform of G_z. Given $\operatorname{Re} s > 1/2$, one takes $1/2 < \theta < \min(1, \operatorname{Re} s)$. Functional-equation symmetry then excludes strip zeros to the left of the line as well. This proves `RH_iff_oneSidedWindowsLambda` and hence the equivalence of any two fixed power-window conditions on G_z. □

For $\delta = \lambda - 1 > 0$, the theorem type of `landau_tails_general` is existential ($\exists \sigma > 2$, $\exists C$, $\forall U \geq 1$, $\exists k \geq U, \ldots$), while its formal proof constructs the explicit witnesses $\ell = 2 + 6/\delta$ and

$$\sigma = \ell(1+\delta)(2+\ell) = 36/\delta^2 + 72/\delta + 44 + 8\delta,$$
$$k = \lceil (1 + \delta/2)\sigma U \rceil, \qquad C_\delta = \frac{\mathrm{e}}{\sigma - 2 - \ell} + \frac{1}{\sigma - 2 + \ell}.$$

For $U \geq 1$ the two exponential tilts on the complementary half-lines give

$$\int_{(0,\infty)\setminus[U,(1+\delta)U]} \frac{u^k e^{(2-\sigma)u}}{k!}\, du \leq C_\delta\, \sigma^{-k}.$$

The Tier (C) Cramér calculation determines the exact threshold of the asymptotic rate criterion for this centering in closed form. With $J(t) = t - 1 - \log t$, set $a_- = (1+\delta/2)^{-1}$, $a_+ = (1+\delta)(1+\delta/2)^{-1}$, so the two edges in mean-one gamma coordinates are $t_- = a_-(1-2/\sigma)$ and $t_+ = a_+(1-2/\sigma)$. The large-U tail criterion is $\min(J(t_-), J(t_+)) \geq \log(\sigma/(\sigma-2))$ (note that $t_+ > 1$ holds automatically whenever $J(a_+) \geq 2a_+/\sigma$, since $J(a) < a-1$ for all $a > 1$). Because $-\log(1-2/\sigma)$ cancels identically with $\log(\sigma/(\sigma-2))$, every $t = a(1-2/\sigma)$ satisfies the exact identity

$$J(t) - \log\frac{\sigma}{\sigma-2} = J(a) - \frac{2a}{\sigma}.$$

Consequently, the exact boundary of this asymptotic rate criterion is

$$\sigma_{\mathrm{rate}}(\delta) = \max\left\{\frac{2a_-}{J(a_-)},\ \frac{2a_+}{J(a_+)}\right\}.$$

As $\delta \downarrow 0$, both $J(a_\pm) = \delta^2/8 + O(\delta^3)$ and $a_\pm = 1 + O(\delta)$, giving $\sigma_{\mathrm{rate}}(\delta) = (16/\delta^2)(1+O(\delta))$ and $\sigma/\sigma_{\mathrm{rate}}(\delta) = \frac{9}{4}(1+\delta/6+O(\delta^2))$, so the ratio of the proof witness to this exact rate threshold tends to $9/4$ from above (specifically $\sigma(1) = 160$ versus $\sigma_{\mathrm{rate}}(1) \approx 58.413867$, ratio 2.739076; and $\sigma(1/64) = 152{,}108.125$ versus $\sigma_{\mathrm{rate}}(1/64) \approx 67{,}426.443643$, ratio 2.255912; archived in `ALPOGE_FRONTIERS_RECEIPT.json` under `sigma_rate_closed_form_vs_theorem`, alongside underflow-free log-space finite-U scans in `sigma_min_table`, where each finite-U value locates a grid-bisection sign change of the discrete-ceiling k-sawtooth accurate to roughly 3–5 significant digits). This threshold applies strictly to the displayed asymptotic rate criterion for the chosen centering and does not establish optimality for finite $U \geq 1$.

A separate, formal minimax model balances a splitting term x/m against a cost m^γ, for $x, m, \gamma > 0$. The first decreases and the second increases in m, so their maximum has the unique minimizer $m = x^{1/(1+\gamma)}$ and minimum $x^{\alpha^*(\gamma)}$, where $\alpha^*(\gamma) = \gamma/(1+\gamma)$. If for every $\varepsilon > 0$ cofinal windows satisfy $G_z(x) \leq K \max(x/m, m^\gamma) x^\varepsilon$ at this choice of m, with fixed K, the transfer excludes zeros with $\operatorname{Re} s > \alpha^*(\gamma)$, $s \neq 1$. Linear cost gives the square-root exponent. This model assumes the availability of the estimates used in the minimization.

There is also an elementary obstruction to using a single fixed Erlang rate. For $k \geq 1$ put $a_j = e^{-k}k^j/j!$. For $0 \leq r \leq \lfloor\sqrt{k}\rfloor$,

$$\frac{a_{k-r}}{a_k} = \prod_{j=0}^{r-1}(1-j/k) \geq 1 - \frac{r(r-1)}{2k} \geq \tfrac{1}{2}.$$

Since the a_j sum to one, `poisson_peak_le` gives $k^k e^{-k}/k! \leq 2/\sqrt{k}$. The normalized Erlang density $\sigma(\sigma u)^k e^{-\sigma u}/k!$ peaks at $u = k/\sigma$. The mass of any log-interval of length $h \geq 0$ is therefore at most $2\sigma h/\sqrt{k}$, uniformly in its position. At fixed $\sigma > 0$ that mass tends to zero for $h = o(\sqrt{k})$.

3.2 An exact Mellin identity for arbitrary pole order

A `PositiveLFunctionData` record consists of a function $L(s)$, real weights $\Lambda_D(n) \geq 0$ with $\Lambda_D(0) = 0$, an integer $m \geq 1$, and a real constant γ_D. Its regularization $H(s) = (s-1)^m L(s)$ off 1 extends to an entire function with $H(1) \neq 0$ and $H'(1)/H(1) = \gamma_D$. The data give nonvanishing of L on $\operatorname{Re} s \geq 1$ off the pole, the absolutely convergent Dirichlet series for $-L'/L$ on $\operatorname{Re} s > 1$, and the bound $S_D(y) = O(y^2)$ on $y \geq 1$, where

$$S_D(y) = \sum_{1 \leq n \leq y} \Lambda_D(n)(y/n - 1), \qquad G_D(y) = S_D(y) - m\,y\log y + (m+\gamma_D)y.$$

This positive-coefficient interface isolates the hypotheses needed for the localized Landau argument in `SelbergPositiveWindowRH` [15, 23].

Lemma 3.2 (Center choice, unique nearest pole, and Cauchy remainder; Tier (A) by dependency closure)**.** *Let $Z \subset \{w \in \mathbb{C} : 0 < \operatorname{Re} w < 1\}$ be the discrete zero set of L, forming the pole set of a meromorphic continuation Φ_D on $\operatorname{Re} s > 0$ whose poles on $\operatorname{Re} s > 0$ are simple poles at $w \in Z$ with nonzero residues c_w, and fix $\sigma > 2$, $0 < \theta < 1$. Suppose there exists $\rho_0 \in Z$ with $\operatorname{Re} \rho_0 > \theta$. Then:*

(a) (`landau_dichotomy_s`) Either (i) there exist $D \le \sigma - \theta$ and $\rho_1 \in Z$ with $Z \cap B(\sigma, D) = \emptyset$ and $\sigma - \operatorname{Re} \rho_1 < D$, or (ii) there exist a real zero $\beta \in (\theta, 1)$ in Z and $r \in (\sigma - \beta, \sigma - \theta]$ such that $Z \cap B(\sigma, r) = \{\beta\}$.

(b) (`landau_nearest_s`) In case (i), choosing $\varepsilon = \frac{1}{2}\min(D - (\sigma - \operatorname{Re} \rho_1), \sigma - D) > 0$, all but finitely many $\gamma \in (\operatorname{Im} \rho_1 - \varepsilon, \operatorname{Im} \rho_1 + \varepsilon)$ avoid the perpendicular-bisector ordinates of pairs in $Z \cap \overline{B}(\sigma + i \operatorname{Im} \rho_1, D + \varepsilon)$, yielding a center $z = \sigma + i\gamma$, a unique nearest pole $\rho \in Z$, and a radius $r \le D$ with $d := |z - \rho| < r$ and $Z \cap B(z, r) = \{\rho\}$.

(c) (`landau_pole_asym_s`) At a center z with unique pole $\rho \in B(z, r)$ of residue c_ρ, subtracting $c_\rho/(s - \rho)$ yields a function holomorphic on $B(z, r)$; hence for every $r' \in (d, r)$ there exists $M > 0$ such that for all $k \ge 0$,

$$\left| \frac{(-1)^k \Phi_D^{(k)}(z)}{k!} - \frac{c_\rho}{(z - \rho)^{k+1}} \right| \le \frac{M}{(r')^k}.$$

Theorem 3.3 (Mellin–Laplace identity and one-sided transfer; `selberg_mellin_identity` and `selberg_zeroFree_of_lambda_windows`, Tier (A))**.** *For $\operatorname{Re} s > 2$ (where $S_D(y) = O(y^2)$ guarantees absolute convergence, extending to $\operatorname{Re} s > 1$ via positivity and the absolutely convergent Dirichlet series for $-L'/L$), define $\Phi_D(s) = \int_0^\infty G_D(e^u) e^{-su}\, du$. Then for $\operatorname{Re} s > 1$,*

$$\begin{aligned} s(s-1)\Phi_D(s) &= -\frac{L'(s)}{L(s)} - \frac{m}{s-1} - m + (m + \gamma_D)s \\ &= -\frac{H'(s)}{H(s)} + s\gamma_D + m(s-1). \end{aligned}$$

For every $\lambda > 1$, $0 < \theta < 1$, a cofinal upper bound $G_D(y) \le Cy^\theta$ on $[Y, Y^\lambda]$ with fixed C implies $L(s) \ne 0$ for $\operatorname{Re} s > \theta$, $s \ne 1$.

Proof. The transform of $(e^u/n - 1)_+$ is $n^{-s}/[s(s-1)]$. Absolute convergence permits summation. The two centering terms contribute $-m/(s-1)^2 + (m + \gamma_D)/(s-1)$. The numerator in the second line vanishes at $s = 1$, so the originating pole is removable for Φ_D. At a zero ρ of L in $0 < \operatorname{Re} \rho < 1$, of multiplicity h, the continued transform has simple-pole residue $-h/[\rho(\rho - 1)]$. In particular this residue is positive for a real zero.

For the transfer, assume for contradiction that $L(\rho_0) = 0$ with $\theta < \operatorname{Re} \rho_0 < 1$, let $C_0 > 0$ satisfy $|G_D(e^u)| \le C_0 e^{2u}$ for $u \ge 0$, and choose $\sigma > 2$ and C_δ from `landau_tails_s` for $\delta = \lambda - 1$. Write $A_k(z) := (-1)^k \Phi_D^{(k)}(z)/k! = \int_0^\infty G_D(e^u)\, u^k e^{-zu}/k!\, du$ for $\operatorname{Re} z > 2$ (`iteratedDeriv_laplace_s`). For every real γ, the triangle inequality gives

$$|A_k(\sigma + i\gamma)| + \operatorname{Re} A_k(\sigma) \le \int_0^\infty \left(|G_D(e^u)| + G_D(e^u)\right) \frac{u^k e^{-\sigma u}}{k!}\, du.$$

On a cofinal window $[U, (1+\delta)U]$, the ceiling yields $|G_D(e^u)| + G_D(e^u) \le 2|C| e^{\theta u}$ and `kernel_max_s` gives $u^k e^{-(\sigma-\theta)u}/k! \le (\sigma - \theta)^{-k}$, so the window integral is at most $2|C|\delta U(\sigma - \theta)^{-k} \le 2e^{\eta U}(\sigma - \theta)^{-k}$ once U is sufficiently large in terms of any given $\eta > 0$. Off the window, `landau_tails_s` bounds the complement by $2C_0 C_\delta \sigma^{-k} \le 2C_0 C_\delta (\sigma - \theta)^{-k}$. Hence, for every $\eta > 0$ and cofinally many U, with $k = \lceil (1 + \delta/2)\sigma U \rceil \ge U$ and uniformly in γ,

$$|A_k(\sigma + i\gamma)| + \operatorname{Re} A_k(\sigma) \le C_1 e^{\eta U} (\sigma - \theta)^{-k}$$

with $C_1 := 2+2C_0C_\delta$. Now apply Lemma 3.2: in case (ii) ($\beta \in (\theta,1)$ real, $c_\beta > 0$), set $d := \sigma-\beta$ and $r' := (d+r)/2 \le \sigma-\theta$; Lemma 3.2(c) at $z=\sigma$ gives $c_\beta d^{-(k+1)} - M(r')^{-k} \le \operatorname{Re} A_k(\sigma)$, and the displayed upper bound at $\gamma = 0$ bounds the right-hand side. In case (i), set $d := |z-\rho|$ and $r' := (d+r)/2 < D \le \sigma-\theta$; Lemma 3.2(c) at $z = \sigma+i\gamma$ gives $|c_\rho| d^{-(k+1)} \le |A_k(z)| + M(r')^{-k}$, Cauchy's estimate (`landau_cauchy_s`) on $B(\sigma, D)$ gives $|A_k(\sigma)| \le M'(r')^{-k}$, and the displayed upper bound yields $|c_\rho| d^{-(k+1)} \le (M+M')(r')^{-k} + C_1 e^{\eta U}(\sigma-\theta)^{-k}$. In both cases $d < r' \le \sigma-\theta$; choosing $\eta = \frac{1}{2}\log(r'/d)$ contradicts $k \ge U \to \infty$ (`landau_contra_s`). The formal statements are `selberg_mellin_identity` and `selberg_zeroFree_of_lambda_windows`. □

At an unramified Rankin–Selberg prime power the elementary weight has the form $|a|^2 \log p$. A Dedekind weight is a sum of $\log N\mathfrak{p}$ over prime ideal powers of given norm. The formal positivity lemmas prove the inequalities for a complex a and for finite sums of logarithms of positive integral norms, respectively. Analytic continuation, the boundary nonvanishing and the growth estimate are additional fields of a complete record. For the Riemann zeta function, `riemannZeta_positiveLData` supplies all fields with $m = 1$ and $\gamma_D = \gamma_{\mathrm{EM}}$.

3.3 Conditional full GRH for Gaussian and abelian quadratic/biquadratic products

In `DedekindAndDirichletWindows`, `PositiveLFunctionData.mul` multiplies the two regularizations and adds the pole orders, weights, Euler–Kronecker constants and centered ramps. Iteration gives a complete package for ζ^m for every $m \ge 1$ (noting that ζ^m are the only instantiated $m > 1$ records in the archive-free suite, whose window hypothesis is that of ζ scaled by m). The generic `GRH_of_lambda_windows_and_zeroSymmetry` requires $\mathcal{W}_\lambda(G_D)$ and the implication $L(\rho) = 0 \Rightarrow L(1-\rho) = 0$ in the open critical strip. The window theorem excludes $\operatorname{Re}\rho > 1/2$. Reflection excludes $\operatorname{Re}\rho < 1/2$, proving $\operatorname{Re}\rho = 1/2$ for every strip zero.

For any nontrivial quadratic Dirichlet character χ of positive modulus, $\zeta(s)L(s,\chi)$ has an unconditional complete package with simple pole and weights $\Lambda(n)(1+\operatorname{Re}\chi(n)) \ge 0$. Primitivity is required by the formal full-GRH implication for the two factors, where the functional equation provides the reflection. The character χ_{-4} is explicitly proved primitive. Thus windows for the single product $\zeta(s)L(s,\chi_{-4})$ imply the critical-line conclusion for both factors. Its classical identification is $\zeta_{\mathbb{Q}(i)}(s) = \zeta(s)L(s,\chi_{-4})$.

A different construction in `AbelianDedekindWindows` keeps a simple pole for families of arbitrary size. For a finite list F of nontrivial quadratic characters, `abelianQuadFamily_positive LData` constructs a record for $\zeta(s)\prod_{\chi\in F} L(s,\chi)$ given

$$1 + \sum_{\chi\in F} \operatorname{Re}\chi(n) \ge 0 \qquad \text{for every } n.$$

Only the zeta factor contributes a pole at 1. The weights are $\Lambda(n)$ times the displayed expression. If each character in F is primitive, `abelianQuadFamily_full_GRH_of_lambda_windows` uses one product-window hypothesis for all factors simultaneously. For a multiquadratic field of degree 2^r, the classical list consists of the $2^r - 1$ nontrivial characters of its Galois group. At an unramified prime power the displayed sum, including the trivial character, is 2^r or zero by character orthogonality. The Lean constructor has no degree cutoff. Its sign hypothesis must be supplied for the whole list at every integer, including ramified arguments.

All five abelian products in Table 3 (three quadratic, two biquadratic of degree 4) have unconditional records with $m = 1$ and verified primitivity. For the last two rows the sign checks include all residue classes modulo 12 and 15, respectively. In each row and in the Gaussian case, full GRH is conditional on cofinal one-sided λ-windows of the product. Identifying these explicit Dirichlet products with the Dedekind zeta functions of the indicated number fields uses the classical abelian factorization external to the Mathlib development.

Classical field E	Degree	Discriminants d in $\zeta \prod_d L(s, \chi_d)$
$\mathbb{Q}(\sqrt{-3})$	2	-3
$\mathbb{Q}(\sqrt{5})$	2	5
$\mathbb{Q}(\sqrt{-15})$	2	-15
$\mathbb{Q}(\zeta_{12})$	4	$-4, -3, 12$
$\mathbb{Q}(\sqrt{-3}, \sqrt{5})$	4	$-3, 5, -15$

Table 3: Five abelian Dirichlet products (three quadratic, two biquadratic of degree 4) unconditionally instantiated as simple-pole ($m = 1$) `PositiveLFunctionData` records in `Abelian DedekindWindows.lean` (Tier (A), five-checker verified), each carrying a conditional λ-window full-GRH theorem for all factors. Identification with the classical Dedekind zeta function $\zeta_E(s)$ relies on classical abelian factorization external to Mathlib.

For comparison, consider the cubic field defined by $X^3 - X - 1$, of discriminant -23. The standard representation ρ of S_3 is induced from a nontrivial character of C_3. The permutation representation $\mathbf{1} \oplus \rho$ gives $\zeta_E = \zeta L(s, \rho)$. Induction identifies $L(s, \rho)$ with a Hecke L-function over $\mathbb{Q}(\sqrt{-23})$. At an unramified prime the permutation character $1 + \operatorname{tr} \rho$ counts roots modulo p, and $|\operatorname{tr} \rho|^2 = 1 + \chi_{-23}(p) + \operatorname{tr} \rho$. The Tier (C) finite computation (`receipts/computational/ALPOGE_FRONTIERS_RECEIPT.json`) verifies these identities at all 17,983 unramified primes $p < 200{,}000$ and checks agreement of the trace with the p-th coefficient of $\eta(z)\eta(23z)$ at all 429 unramified primes $p < 3{,}000$ (finite computational checks rather than a general monomiality theorem). The Dedekind and Rankin–Selberg weights are nonnegative even though the bare standard representation trace takes negative values (at 6,005 primes below 200,000). Producing cofinal windows for these non-abelian examples remains the arithmetic obstacle.

4 Sub-polynomial Landau windows and logarithmic barriers for finite frequency bands

Set $U = \log Y$ and $g(u) = G_z(e^u)$. A sub-polynomial window is $[Y, Y \exp W(\log Y)]$ with $W(U) > 0$ and $W(U) = o(U)$. Its one-sided ceiling reads $g(u) \le Ce^{\theta u}$ throughout $[U, U + W(U)]$ for cofinally many U, with C fixed.

4.1 An attained rightmost real part gives syndetic positive spikes

That an attained supremum $B = \sup_\rho \operatorname{Re} \rho$ forces bounded-gap oscillations via Dirichlet/Bohr approximation on the boundary line is a classical theorem of Ingham [8] (proved for $\pi(x) - \operatorname{li}(x)$ on intervals (x, Ax), and quoted in later expositions also for $\psi(x) - x$, yielding $V(T) \gg \log T$; Kaczorowski [11] later proved $V(T) \gg \log T$ unconditionally, though a count of sign changes does not settle whether one-sided $O(x^\theta)$ ceilings can persist on cofinal windows). No priority is claimed for this attained-supremum mechanism; the contribution below is an abstract five-checker one-sided formalization.

Theorem 4.1 (Finite-box almost periods; `bohr_positive_syndetic`, Tier (A))**.** *Suppose $F(t) = \operatorname{Re} \sum_j c_j e^{i\gamma_j t}$, with $\gamma_j \in \mathbb{R}$ and $\sum_j |c_j| < \infty$. If $F(t_0) > 0$, there exist fixed $L_{\mathrm{ap}} > 0$ and $\eta > 0$ such that every sufficiently late interval of length L_{ap} contains u with $F(u) \ge \eta$ (the formal proof constructs the explicit witness $\eta = F(t_0)/2$).*

Proof. For a finite set of frequencies, divide each unit phase interval into M equal bins and send each nonnegative integer n to the tuple $\left(\lfloor M\{\gamma_j n/(2\pi)\}\rfloor\right)_j$, where braces denote fractional part. Pick one integer representative of each occupied box, and let R be their maximum. For any $V \ge 0$, the integer $k = \lceil V \rceil + R$ and the representative r of its box give $p = k - r \in [V, V + R + 1]$. Then $|e^{i\gamma_j p} - 1| \le 2\pi/M$ for every retained frequency. A finite truncation with

small absolute tail, followed by a sufficiently large M, makes $\sup_t |F(t+p) - F(t)|$ arbitrarily small (`trigSeries_almost_periodic`). This proves relative density of almost periods by a finite-image pigeonhole argument. Taking the error $\eta = F(t_0)/2$ and translating by t_0 proves `bohr_positive_syndetic` [4]. □

The window obstruction in `SubPolynomialWindowRH` assumes Bohr almost periodicity of F, one positive value, $B > \theta$, and $e^{-Bu}g(u) - F(u) \to 0$. At the spikes, eventually $g(u) \geq e^{Bu}F(t_0)/4 > Ce^{\theta u}$. Hence there is a gap length $L_{\mathrm{ap}} > 0$ for which cofinal ceilings are impossible whenever $W(U) \geq L_{\mathrm{ap}}$ eventually. A sufficiently large constant width already suffices; the value of L_{ap} depends on the boundary spectrum. This is the Tier (A) implication `attained_supremum_contradicts_subpoly_windows`, an abstract one-sided formalization of Ingham's attained-supremum mechanism [8].

For the application to zeta, the additional analytic input is the classical smoothed explicit formula. With distinct nontrivial zeros ρ of multiplicity m_ρ, it gives

$$g(u) = \log(2\pi) - \sum_\rho \frac{m_\rho e^{\rho u}}{\rho(\rho-1)} + O(e^{-2u}), \qquad \sum_\rho \frac{m_\rho}{|\rho(\rho-1)|} < \infty.$$

The formula and its application here use classical analysis outside the abstract Lean theorem [7, 8, 26]. Suppose $B = \sup_\rho \operatorname{Re}\rho > \theta$ is attained. The terms with $\operatorname{Re}\rho = B$ give a real trigonometric series F after normalization by e^{-Bu}. Dominated convergence removes the remaining terms, even if their real parts approach B. The boundary series has mean zero since all its frequencies are nonzero. Its coefficient at any attained ordinate is nonzero, as seen by averaging against the opposite frequency. It has a positive value, since otherwise a strictly negative interval, repeated with bounded gaps by almost periodicity, would give a negative mean. Consequently cofinal ceilings with $W(U) \to \infty$ force either zero-freeness above θ or an *unattained* supremum $B > \theta$ (exactly the boundary case left open by Ingham [8]). The same conditional exclusion holds at the fixed width L_{ap} supplied by the boundary series.

4.2 Gaussian localization and two-mode scale lemmas

For $h, v > 0$, convexity of t^2 gives $t^2 \geq h^2 + 2h(t-h)$ on $t \geq h$. Integration of the resulting exponential majorant proves the elementary Mills bound (Tier (A), `gaussian_mills_tail_le`)

$$\int_h^\infty e^{-t^2/(2v)}\,dt \leq (v/h)e^{-h^2/(2v)}.$$

As a Tier (C) heuristic modeling link for Gaussian test functions $\phi_{U,v}(u) = (2\pi v)^{-1/2}\exp(-(u-U)^2/(2v))$ modulated to a target ordinate γ_0, with half-width $h \asymp W$, suppressing an exponential tail budget at scale U outside a window of width W calls for $h^2/(2v) \gg U$; the stronger sufficient regime $h^2/(vU) \to \infty$ is equivalent to $v = o(W^2/U)$, whereas sharpening a fixed frequency separation $\Delta\gamma \neq 0$ calls for $v \to \infty$. The Tier (A) lemma `gaussian_beta_half_dichotomy` formalizes the resulting scale dichotomy: for $W = cU^\beta$ $(c > 0)$, $W^2/U = c^2U^{2\beta-1} \to \infty$ holds if and only if $\beta > 1/2$ (for instance $v = \sqrt{W^2/U}$ then satisfies both $v \to \infty$ and $v/(W^2/U) \to 0$, while at $\beta = 1/2$ the ratio W^2/U is constant).

For the two-mode comparison (Tier (C) modeling derivation), evaluating the Gaussian integral $\int_{-\infty}^\infty e^{(b+i(\gamma-\gamma_0))u}\phi_{U,v}(u)\,du = \exp\big((b+i(\gamma-\gamma_0))U + \frac{1}{2}v(b+i(\gamma-\gamma_0))^2\big)$ gives log-modulus $bU + \frac{1}{2}v(b^2 - (\gamma-\gamma_0)^2)$. Against a reference mode of real part b_0 at $\gamma = \gamma_0$, a mode of real part $b_1 > b_0$ at frequency offset $\Delta\gamma$ therefore has relative Gaussian log-amplification (up to a fixed amplitude term)

$$(b_1 - b_0)U + \tfrac{1}{2}v(U)(b_1^2 - b_0^2 - (\Delta\gamma)^2).$$

The formal Tier (A) lemma `gaussian_two_mode_obstruction` shows that whenever $v(U)/U \to 0$ and $b_1 > b_0$, dividing by U yields limit $b_1 - b_0 > 0$, so the displayed expression diverges to

$+\infty$. Under the Gaussian test model, a variance satisfying $v(U)/U \to 0$ therefore cannot isolate the slower mode by fixed frequency separation (this does not prove zero-freeness of ζ at these widths).

4.3 Pole counts, sparse carriers, and discrete shift averaging

For $N \geq 1$ and $M \geq 0$, consider the classical regularized incomplete-beta (maximally flat) polynomial $P_{N,M}(z) = z^N \sum_{j=0}^{M} \binom{N+j-1}{j}(1-z)^j$, which serves as an explicit superoscillating carrier (cf. [6]). Pascal's identity gives the exact finite remainder

$$1 - P_{N,M}(z) = (1-z)^{M+1} \sum_{k=0}^{N-1} \binom{M+k}{k} z^k.$$

For $M = 5N$ the sum of the absolute coefficient bounds is $\binom{6N}{N-1} \leq 3^{6N}/2^{5N+1}$, by a single term in the binomial expansion of $(1/3 + 2/3)^{6N}$. Thus on $|z| = 1$, $|1 - z| \leq 1/2$,

$$|P_{N,5N}(z) - 1| \leq \tfrac{1}{4}(729/1024)^N, \qquad \operatorname{Re} P_{N,5N}(z) > \tfrac{1}{2}.$$

Both the remainder and this bound are Tier (A). The polynomial approximates the constant on the arc even though its lowest frequency is N.

The infinite-series construction using this arc is Tier (C), stated with explicit inductive inequalities so every convergence and sign-preservation step can be verified directly.

Proposition 4.2 (Inductive superoscillation block series, Tier (C))**.** *Let $W : (0,\infty) \to (0,\infty)$ satisfy $W(U) = o(U)$ as $U \to \infty$. Set $b_n = 1 - 1/(n+2)$ for $n \geq 1$, $\delta_1 = 1/12$, and $\delta_n = (b_n - b_{n-1})/4 = 1/[4(n+1)(n+2)]$ for $n \geq 2$. Define recursively $U_1 < U_2 < \dots$ with $U_1 \geq 1$ and $U_n \geq 2U_{n-1}$ for $n \geq 2$, $I_n = [U_n, U_n + W(U_n)]$, $x_n = U_n + W(U_n)/2$, $K_n = \max(1, W(U_n))$, $L_n = \lceil nK_n \rceil$, $a_n = e^{-\delta_n U_n} 2^{-14L_n}$, and*

$$f_n(u) = -a_n e^{b_n u} \operatorname{Re} P_{L_n,5L_n}(e^{i(u-x_n)/K_n}).$$

Then U_n can be chosen inductively so that $g_*(u) = \sum_{n=1}^{\infty} f_n(u)$ converges absolutely on $[0,\infty)$, satisfies $|g_*(u)| \leq e^u$ for all $u \geq 0$, satisfies $g_*(u) \leq 0$ on every window I_n, and its Laplace transform continues meromorphically to $\mathbb{C}$ with simple poles exactly at $b_n \pm i\ell/K_n$ ($L_n \leq \ell \leq 6L_n$), whose real parts $b_n \uparrow 1$ have an unattained supremum.

Proof. Because $P_{L,5L}(z) = \sum_{k=0}^{5L} c_{L,k} z^{L+k}$ with $c_{L,k} = (-1)^k \binom{L+k-1}{k}\binom{6L}{5L-k} \neq 0$ has alternating nonzero coefficients, evaluating at $z = -1$ gives the exact ℓ^1 coefficient sum

$$\sum_{k=0}^{5L} |c_{L,k}| = |P_{L,5L}(-1)| = \sum_{j=0}^{5L} \binom{L+j-1}{j} 2^j \leq 2^{5L} \binom{6L}{L} \leq 2^{11L},$$

so $|f_n(u)| \leq a_n 2^{11L_n} e^{b_n u} = e^{-\delta_n U_n} 2^{-3L_n} e^{b_n u}$ for all $u \in \mathbb{R}$. On I_n, $|(u - x_n)/K_n| \leq \frac{1}{2}$ implies $|1 - e^{i(u-x_n)/K_n}| \leq \frac{1}{2}$, so the Tier (A) arc bound gives $\operatorname{Re} P_{L_n,5L_n}(e^{i(u-x_n)/K_n}) > \frac{1}{2}$ and hence $f_n(u) \leq -\frac{1}{2} a_n e^{b_n u}$. Since $W(U) = o(U)$, $L_n = \lceil n \max(1, W(U_n)) \rceil = o(U_n)$ at each fixed n. Having chosen $U_1, \dots, U_{n-1}$, set $M_{n-1} = \max_{1 \leq j < n}(U_j + W(U_j))$ and $\mu_{n-1} = \min_{1 \leq j < n}(\frac{1}{4} a_j e^{b_j U_j}) > 0$ (for $n \geq 2$), and choose $U_1 \geq 1$ (for $n = 1$) or $U_n \geq 2U_{n-1}$ (for $n \geq 2$) large enough that simultaneously: (1) $14L_n \log 2 \leq \delta_n U_n$ and $\delta_n U_n \geq n \log 2$; (2) $e^{-\delta_n U_n} 2^{-3L_n} e^{b_n M_{n-1}} \leq 2^{-(n+1)} \mu_{n-1}$ (for $n \geq 2$). By (1), $|f_n(u)| \leq 2^{-n} e^{b_n u} \leq 2^{-n} e^u$ for all $u \geq 0$, giving absolute convergence and $|g_*(u)| \leq \sum_{n \geq 1} 2^{-n} e^u = e^u$. By (2), on every earlier window I_j ($j < n$), $\sum_{k>j} |f_k(u)| \leq \sum_{k>j} 2^{-(k+1)}(\frac{1}{4} a_j e^{b_j U_j}) \leq \frac{1}{4} a_j e^{b_j u}$. For $u \in I_n$, since $b_n - b_j \geq 4\delta_n$ for $j < n$ and (1) gives $e^{-4\delta_n U_n} \leq \frac{1}{4} e^{-2\delta_n U_n} \leq \frac{1}{4} e^{-\delta_n U_n} 2^{-14L_n} = \frac{1}{4} a_n$, the earlier blocks satisfy $\sum_{j<n} |f_j(u)| \leq \sum_{j<n} 2^{-j} e^{(b_n - 4\delta_n)u} \leq e^{-4\delta_n U_n} e^{b_n u} \leq \frac{1}{4} a_n e^{b_n u}$. Combining the three estimates

yields $g_*(u) \le (-\frac{1}{2}+\frac{1}{4}+\frac{1}{4})a_n e^{b_n u} = 0$ on every I_n. On $\operatorname{Re} s > 1$, the bound $\sum_{n=1}^{\infty} |f_n(u)e^{-su}| \le \sum_{n=1}^{\infty} 2^{-n} e^{(1-\operatorname{Re} s)u} = e^{(1-\operatorname{Re} s)u} \in L^1([0,\infty))$ justifies Fubini/dominated convergence, giving $\mathcal{L}[g_*](s) = \sum_{n=1}^{\infty} \mathcal{L}[f_n](s)$. Finally, for any compact set $\mathcal{K} \subset \mathbb{C}$ disjoint from the discrete pole set, once $L_n/K_n \ge n > \max_{s\in\mathcal{K}} |\operatorname{Im} s| + 1$, every denominator in the rational Laplace transform $\mathcal{L}[f_n](s)$ has modulus at least 1, so $\sup_{s\in\mathcal{K}} |\mathcal{L}[f_n](s)| \le a_n \sum_{k=0}^{5L_n} |c_{L_n,k}| \le e^{-\delta_n U_n} 2^{-3L_n} \le 2^{-n}$, giving locally normal convergence of $\sum_{n=1}^{\infty} \mathcal{L}[f_n](s)$ to a meromorphic continuation on $\mathbb{C}$. Because all coefficients $c_{L_n,k}$ are nonzero (Tier (A) identity below) and the real parts b_n are strictly increasing (hence pairwise distinct), distinct blocks have disjoint pole sets and the poles are simple, located exactly at $b_n \pm i\ell/K_n$ ($L_n \le \ell \le 6L_n$). This construction concerns general Laplace transforms rather than $\zeta(s)$ itself. □

The coefficient nonvanishing identity and the inequalities for the defined block-counting function `blockPoleCount` are formalized in `SubPolyZetaDensityBarriers` (Tier (A)); identifying `blockPoleCount` with the Laplace-transform pole count of each block is Tier (C). For the abbreviation $P_L = P_{L,5L}$, expansion gives

$$[z^{L+k}]P_L = (-1)^k \binom{L+k-1}{k}\binom{6L}{5L-k} \neq 0 \qquad (0 \le k \le 5L,\ L \ge 1).$$

Thus a block contributes the conjugate pairs $b_n \pm i\ell/K_n$, $L_n \le \ell \le 6L_n$. Counting both signs up to absolute ordinate $T \ge 0$, its exact contribution is defined as

$$\texttt{blockPoleCount}(L_n, K_n, T) = 2\max(0, \min(5L_n + 1, \lfloor TK_n \rfloor - L_n + 1)).$$

For the original $L_n = \lceil nK_n \rceil$, all blocks with $n \le (T-6)/6$ are complete by height T. For $n_0 \ge 1$, $T \ge 12n_0 + 12$, their total from n_0 is at least $(T/24)^2$. Each block in $0 < b_n < 1$ contributes at least $1/(15n)$ to $\sum_\rho |\rho(\rho-1)|^{-1}$. Hence the pole count is $\Omega(T^2)$ and the reciprocal sum diverges, regardless of the spacing of the windows.

Now take $L_n = \lceil q_n K_n \rceil$. Suppose $q_0 \ge 2$, $q_{n+1} > 6q_n + 6$, $1 \le K_n \le C \log q_n$, and the full count of earlier blocks is at most $q_n \log q_n$. The finite-prefix bound in `sparse_carrier_pole_count_0_T_log_T` is $(10C + 13)T \log T$ for $T \ge q_0$. Each block contributes at most $24K_n/q_n$ to the reciprocal sum, so $K_n/q_n \le 2^{-n}$ gives convergence. These estimates are Tier (A). The Tier (C) selection argument takes $q_n = \sqrt{U_n}$ on an unbounded sequence where $W(U_n) \le C' \log U_n$. Then $L_n = O(\sqrt{U_n} \log U_n)$, and U_n can be increased at each step to meet all the separation, count, and dominance conditions. This gives both the $O(T \log T)$ pole count and the summable reciprocal sum through $W(U) = O(\log U)$, and also for irregular widths satisfying the subsequence condition. Conversely, at $T = 6q_n + 6$ the entire block is counted. A count at most $C_0 T \log T$ forces $K_n \le 4C_0 \log q_n \le 4C_0 \log U_n$ for $2 \le q_n \le U_n$. This rules out $W(U)/\log U \to \infty$ for this binomial family. Finite FLINT coefficient checks and Arb interval checks of reciprocal weight estimates give separate Tier (C) tests archived in `receipts/computational/w4_superosc_density_check.json` and `x3_superlog_check.json` [10].

Theorem 4.3 (Width bound by discrete shift averaging; `trigPoly_nonneg_width_lt_pi_sum_inv`, Tier (A))**.** *Let $P(t) = \sum_{j=1}^{M}(c_j \cos(\lambda_j t) + d_j \sin(\lambda_j t))$ with $\lambda_j > 0$ and real coefficients. If $P \ge 0$ throughout $[a, a+W]$ and $P(a) > 0$, then*

$$W < \pi \sum_{j=1}^{M} \lambda_j^{-1}.$$

Proof. The positive averaging operator $A_j f(t) = (f(t) + f(t + \pi/\lambda_j))/2$ kills the j-th mode, following the half-period translate-averaging mechanism used by Kozma and Oravecz [14, Theorem 1, §§1–2] on symmetric integer spectra $\pm n_j$ on the unit torus (where angular frequencies $\lambda_j = 2\pi n_j$ yield non-strict diameter $\sum 1/(2n_j) = \pi \sum 1/\lambda_j$). The operators commute, so their

product annihilates P. If the claimed strict inequality failed, every subset sum of the shifts would stay in the window when added to a. Each term of the expanded average would be nonnegative, and the empty subset alone would contribute $2^{-M}P(a) > 0$, a contradiction. The constant is sharp already for $\cos t$ on $[-\pi/2+\varepsilon, \pi/2]$ as $\varepsilon \downarrow 0$. □

For a uniform relative band $[q, Cq]$ with $q > 0$, $C \geq 1$, assume that the M positive frequencies count against a pole budget $M \leq C_0(Cq)\log(Cq)$, $C_0 \geq 0$, and that $Cq \leq U^A$, $U \geq 1$. The shift bound yields

$$W < \pi M/q \leq \pi C_0 C A \log U.$$

The formal `nonbinomial_bandlimited_log_barrier` actually uses only the frequency floor and these count and growth hypotheses. A relative band supplies the count from a global pole bound when the frequencies correspond to distinct, uncancelled poles. The result applies to arbitrary coefficients and requires the positive endpoint in the width theorem.

Removing the single relative-band hypothesis $[q, Cq]$ at the cost of a second logarithm follows immediately from Theorem 4.3 by partial summation (Tier (C); no priority claimed).

Corollary 4.4 (Multi-band $O((\log U)^2)$ width barrier for finite trigonometric polynomials, Tier (C))**.** *Let $P(t) = \sum_{j=1}^{M}(c_j\cos(\lambda_j t) + d_j\sin(\lambda_j t))$ have distinct positive frequencies $q_0 \leq \lambda_j \leq T$ ($q_0 > 0$), and let $C_0 > 0$. If the counting function $N(t) := \#\{j : \lambda_j \leq t\}$ (with $N(q_0^-) = 0$) satisfies the relative budget $N(t) \leq C_0 t\log(et/q_0)$ for all $t \in [q_0, T]$, and if $P \geq 0$ on $[a, a+W]$ with $P(a) > 0$, then with $H := \log(T/q_0)$,*

$$W < \pi C_0\left(1 + 2H + \frac{H^2}{2}\right).$$

In particular, whenever $T/q_0 \leq U^A$ ($U \geq 1$, $A > 0$),

$$W < \pi C_0\left(1 + 2A\log U + \frac{A^2}{2}(\log U)^2\right).$$

Under the absolute zeta-type budget $N(t) \leq C_0 t\log t$ on $[q_0, T]$ with $q_0 \geq e$, the same summation gives $\sum_{j=1}^{M}\lambda_j^{-1} \leq C_0(\log T + \frac{1}{2}(\log^2 T - \log^2 q_0))$ and hence $W < \pi C_0(\log T + \frac{1}{2}(\log^2 T - \log^2 q_0))$.

Proof. Because $N(q_0^-) = 0$, Stieltjes partial summation gives $\sum_{j=1}^{M}\lambda_j^{-1} = N(T)/T + \int_{q_0}^{T} N(t)\,t^{-2}\,dt$. Under the relative budget, $N(T)/T \leq C_0(1+H)$ and $\int_{q_0}^{T} N(t)\,t^{-2}\,dt \leq C_0\int_{q_0}^{T}(1+\log(t/q_0))\,t^{-1}\,dt = C_0(H + H^2/2)$. Under the absolute budget $N(t) \leq C_0 t\log t$ ($q_0 \geq e$), $N(T)/T \leq C_0\log T$ and $\int_{q_0}^{T} N(t)\,t^{-2}\,dt \leq C_0\int_{q_0}^{T}(\log t)\,t^{-1}\,dt = \frac{1}{2}C_0(\log^2 T - \log^2 q_0)$. Applying Theorem 4.3 immediately yields all three displayed bounds. □

Corollary 4.4 eliminates the single-band requirement for any finite constant-coefficient trigonometric polynomial, while leaving a crucial distinction relative to $\zeta(s)$: in the explicit formula with zeros of varying real parts, a single exponential normalization e^{-Bu} leaves time-varying amplitudes $e^{(\mathrm{Re}\,\rho - B)u}$, so Corollary 4.4 does not resolve the unattained-supremum case.

A separate Tier (C) hand proof bounds the cost of uniform positivity. For a finite sum with frequencies at least $q > 0$, global bound $|P| \leq D$, and $P \geq h > 0$ on a window of length W, it gives $\log(D/h) \geq cqW - C_1$ for absolute $c, C_1 > 0$. For $qW \geq 4\pi$ the even power of $\sin(qt/(2r))/(qt/(2r))$, with $r = \lfloor qW/(4\pi)\rfloor$ and exponent $2r$, has Fourier transform vanishing at all the frequencies. Comparing its positive mass inside the centered window with its tail yields the bound. The pole count and reciprocal summability above leave the zeta residues and the full zero-counting asymptotic unspecified. Cancellations among several bands without a uniform relative band $[q, Cq]$, and the zeta-specific window question at an unattained supremum, remain open.

5 Structural audit and interface decompression of Alpöge's $\psi = 1$ specialization

The standalone theorems of Sections 2–4 emerged directly from a systematic reverse-engineering and machine-checked audit of the principal-character specialization `PrincipalSupply.hSup` (in the author's project layer built on revision `79c751e0b0a5ea3bbfc20198c852fa85f28661e6` [1]). Three distinct goals must be kept strictly separated: (1) mapping the dependency graph, (2) simplifying and exposing the outer formal interface (`Formalism.*`), and (3) providing a self-contained mathematical explanation of why the million-declaration arithmetic core succeeds. This section delivers (1) and (2) rather than (3), recording the exact DAG anatomy (Tier (B)), archive-linked Lean 4.33.1 autopsies (Tier (A*)), decompressed `Formalism.*` interface (Tier (A)/(A*)), native `.olean` auditing (Tier (B)), a refined-DAG Mathlib extraction modularity scan (Tier (B)), and reported three-stage live zero-commit compilation (Tier (A*, reported)), separated strictly from remaining graph-simulated and overlay cuts (Tier (C)).

5.1 Exact 1,264,848-declaration DAG anatomy and the 60-field classical catalogue

Breadth-first traversal from the Case-III-eliminated supply root `PrincipalSupply.hSup` over the read-only local archive edge database (`dag_cache.npz`, SHA-256 prefix `f2930365b5baede1`, paired with `dag_names.tsv`, SHA-256 prefix `69fef22fcc549c02`—digests recorded in `tools/joint_multicut_bfs.py`; files not bundled due to archive size, with deterministic JSON summaries archived in `receipts/structural/phase1_dag_xray_full.json`) visits exactly

$$|V| = 1{,}264{,}848 \text{ declarations}, \qquad |E| = 14{,}075{,}546 \text{ directed edges}, \qquad |\mathcal{M}| = 146{,}398 \text{ modules}$$

(on the local cache and unarchived producer logs: 47 fewer declarations than the pre-Case-III baseline 1,264,895, while the full exported closure including Mathlib and inductive metadata reaches 1,477,692 declarations, or 1,484,202 in external kernel environments, and the unbundled `PrincipalSurface` catalogue lists 1,887 surface declarations). Computing the exact Lengauer–Tarjan dominator tree [18] (dom(root) = 1,264,848) yields the top-level directory continents and the five primary architectural chokepoints under `principal_connected` (dom = 1,260,941) listed in Table 4 (all counts archived in `phase1_dag_xray_full.json`).

Leaf-interface coupling and exact 194 → 60 compression. Joint BFS reachability on the 1,264,848-node graph establishes that removing all 194 literature leaf theorems instantiated in `PrincipalDirect.lean` disconnects 244,592 declarations (19.34%, leaving 1,020,256; local DAG cache), while removing the 527-node set that also includes their 333 statement declarations disconnects 245,282 (19.39%, leaving 1,019,566; archived in `phase1_dag_xray_full.json`). Over 80.6% of the DAG is reachable without traversing any leaf, tightly coupled through shared internal infrastructure. Exhaustive projection and call-site classification (`receipts/structural/textbook_compression_194_to_60.json`) partitions the 194 raw leaves cleanly without remainder:

$$194 \text{ raw leaves} = \underbrace{16}_{\substack{\text{structurally dead} \\ \text{(0 proj., 0 ext. calls)}}} + \underbrace{37}_{\text{derived leaves}} + \underbrace{141}_{\text{live classical leaves}} \xrightarrow{\text{group}} 58 + 2 = 60 \text{ clusters.}$$

Here *structurally dead* means having zero field projections and zero external call sites outside the enclosing record instantiation in the uncut 1,264,848-node DAG; all 16 remain syntactically reachable in the uncut cone through their record constructors, and 15 of the 16 (those inside `ThetaLocalInputs`) remain reachable in Stage 1 (`CutClosedRH.audit.json`) until `ThetaLocalInputs` and `LiteratureCMHypothesis` are recompiled in Stage 2 (`CutClosedRH_PFCM_`

Directory continent / Primary chokepoint	Direct / Subtree	Cut / Excl. dom	Share
Largest directory continents (direct declarations vs. directory-prefix cut-set):			
`Development.AbelianVarietyTheory`	525,012	732,123	41.51%
`Formalism.*` (opaque middleware)	278,201	283,157	21.99%
`Development.AutomorphicArithmetic`	128,219	914,165	10.14%
`Development.AutomorphicSpectrum`	82,903	127,965	6.55%
`Development.WeilBound`	80,314	224,196	6.35%
`GroundTheory`	76,769	—	6.07%
`ReciprocityTheory`	47,617	—	3.76%
`CompanionTheory`	15,715	—	1.24%
`Development.EtaleCohomology`	13,825	13,865	1.09%
`Development.ResolutionOfSingularities`	10,308	10,313†	0.81%
Five primary chokepoints under `principal_connected` *(*dom = 1,260,941*, Lengauer–Tarjan):*			
`UpperTheory.Lit.nonempty_hyps_definite_side_of_closed_bound`	748,257	449,144	35.51%
`AlpogePrincipalDirect.principal_row` (`tLiftK` + Galois reps)	303,553	90,311	7.14%
`AlpogePrincipalDirect.principal_host_bound` (`VWinAt`)	424,446	87,026	6.88%
`topCovolSlim2_use_alt_alt` (Arakelov covolume bound)	241,856	79,803	6.31%
`AlpogePrincipalDirect.principal_shape` (`topDesWin`)	91,642	19,281	1.52%

Table 4: Exact mass distribution of `PrincipalSupply.hSup` (1,264,848 declarations, Tier (B); archived in `receipts/structural/phase1_dag_xray_full.json`). The column "Share" reports Direct/1,264,848 in the upper continent panel and Excl. dom/1,264,848 in the lower chokepoint panel. †The four-leaf Hironaka cut-set is 10,203 (10,569 including the plug directory on the local cache; directory-prefix cut 10,313).

hironaka.audit.json), after which all 16 (and all 21/21 tracked Stage-2 landmarks) are absent from the live dependency cone (consistent with proof-irrelevance constraints below). Among the 37 derived (non-independent) leaves, the largest individual exclusive dominators recorded in phase1_dag_xray_full.json are galoisRepresentations_of_thetaLift (dom = 86,160) and ThKR.theoremK_head2 (dom = 11,578), while cross-checking zero-projection bundle fields against the *global* DAG (phase1c_global_dead_and_continents.json, phase1d_exact_verification.json) separates locally unused bundle fields that are still invoked elsewhere from the 16 structurally dead landmark theorems (including metaplecticRepresentation_holds, rallis_of_dm, kudlaSplittingSplitUnitary, StoneVonNeumann_holds, linnik, and gelfandKazhdanDual_of_multiplicityOne). The grouping $58 + 2 = 60$ assigns the 141 live leaves from PrincipalDirect.lean to 58 topical clusters (T1_01–T5_06) and the 4 Selmer/Iwasawa leaves from principal_host_bound to 2 additional clusters (T5_07 with 1 leaf and T5_08 with 3 leaves), giving 145 distinct live leaves organized many-to-one into 60 topical fields (18 singleton fields and 42 conjunctions) across five pillars ($23 + 12 + 9 + 8 + 8 = 60$), tabulated in full in Appendix A.

5.2 Principal-character collapse, Hironaka redundancy, and scaling barrier

Source inspection of revision 79c751e0 (Tier (C) source audit; the archive-free lemmas principal_frame_exact_dyadic_order and principal_selmer_caseIII_empty in Alpoge HumanCore formalize only the corresponding elementary integer properties) shows that at the principal Dirichlet character $\psi = \chi =$ PChar.one, the odd conductor factor satisfies $o = 1$, collapsing the character-weighted sum (Level.cyc PChar.one).$\mathrm{Gm}(y)$ identically to the smoothed centered prime-power sum $G_z(y)$ of Section 3, forcing exact dyadic frame order ord $= 2^e$, and emptying Selmer Case III uniquely among the six Selmer cases inside principal_host_bound (Case IV collapses to $\psi = 1$; Cases I, II, V, and Z remain active because $3 \in B_{\texttt{PChar.one}}$ and quadratic twists modulo the two auxiliary super-polynomial primes q, d persist, leaving the three real quadratic order-2 comparison characters χ_{-q}, χ_{-d} [odd, $\chi(-1) = -1$] and χ_{qd} [even, $\chi(-1) = +1$], distinct from the even-order twists appearing in the Selmer analysis).

Five archive-linked Lean 4.33.1 files (ClaudeAutopsy, 684 lines, 24 declared theorems of which 19 are printed, all passing lean.exe and leanchecker.exe with exit 0 against trusted read-only archive objects; receipt CLAUDE_AUTOPSY_VERIFICATION.json, Tier (A*)) establish three sharp structural results:

1. *Hironaka necessity and partial redundancy (A1Redundancy, A1RedundancyTop):* Resolution and principalization in characteristic zero (10,203 declarations for the four-leaf cut; 10,569 including the plug directory on the local cache) has zero DAG paths to PChar or PChar.one and operates in unbounded dimension (both via the field degree $[E : \mathbb{Q}] = 4n \geq 4x^{2+8\varepsilon}$ and the associated flag-dimension formula, Tier (C) source audit). Nevertheless, two of its four entry leaves (resolutionCharZero and principalisationCharZero_closed) are called at a single site solely to re-derive ResKStatement, already provided directly by FromRes.resK. The verified theorem ClaudeAutopsy.topCovolSlim2_use_without_res removes both hypotheses cleanly.

2. $K = 2$ *word-predicate characterizations (A2Obstruction, A2ChainII):* As detailed in Section 2, Chain I proves that every TextWord contains at most two non-thin letters, whereas PrimWord and ArcWord obey the two-letter cap conditionally on the presence of a down letter (the no-down branch allows arbitrarily many ups); in Chain II, the alternating branch forbids run length $R = 3$ at $K = 2$ (A2ChainII), while the constant-letter branch already permits general run length R.

3. *Character-independent scaling barrier (A3Barrier):* In the archive's parameter hierarchy (Tier (C) source reading), $n = 2^e \asymp x^{2+8\varepsilon}$, $m \asymp \sqrt{x}x^{-\varepsilon/16}$, and $k \asymp x^{9/2+10\varepsilon}$ on iterated

definite unitary theta lifts of rank $N_j = k + j(t+1)$. The verified real-arithmetic theorem `ClaudeAutopsy.growth_inequality_forces_m` proves, independently of the Dirichlet character χ, that for the specifically defined extremal admissible numerical error data (`Err` and `Gx`) the uniform inequality forces $m \geq 4x^{1/2-3\varepsilon}$ (a uniform-bound barrier over admissible extremal inputs rather than a pointwise claim for arbitrarily smaller actual errors), while the separate archive-linked theorems (`ClaudeAutopsy.sizeWindows_m` and the `RankVal` bounds) connect this parameter scaling directly to the archive's `SizeWindows` interface; consequently, rank growth cannot be bounded at $\psi = 1$ without an external source of the factor-m gain.

5.3 Decompressed `Formalism.*` interface and dual-layer grounding

Between the classical leaf interface and the analytic ζ-engine lies the 278,201-declaration `Formalism.*` middleware. The archive-free suite `AlpogeHumanCore` (four clean-layer modules: `TextbookCatalogue`, `FormalismDecompressed`, `PrincipalReductionEngine`, and `RHAndSharpPrimes`; 825 lines, 21 exported targets listed explicitly in `receipts/five_kernel/GOAL_B_21_TARGETS.json`, Tier (A) five-checker verified) re-states this interface transparently:

- `TransportedAlt` packages four threshold components $(K_{\mathrm{out}}, K_{\mathrm{mid}_0}, K_{\mathrm{mid}}, K_{\mathrm{last}})$ with pointwise supremum $K_1 \sqcup K_2$ obeying commutative, associative, and idempotent join-semilattice laws.

- `Closed` formalizes a uniform cofinal-threshold predicate on family bases `FamBase` equipped with weakening and finite-intersection closure via $K_1 \sqcup K_2$ (without a monad interface).

- `decompressed_topDesWin` (22-line tactic proof) abstracts the nine-link assembly step by intersecting nine `Closed` link hypotheses into one joint threshold $K = K_1 \sqcup \cdots \sqcup K_9$ fed into `DesignExists`.

- `RHAndSharpPrimes` proves twelve conditional theorems reducing Mathlib's `Riemann Hypothesis`, `ZetaAutomorphicArakelovSupply`, the sharp $O(x^{1/2+\varepsilon})$ error bounds for $G_z, \psi, \vartheta, \pi - \mathrm{li}$, the Prime Number Theorem, and short-interval primes in $(x, x + x^{1/2+\varepsilon}]$ to an explicit `PrincipalArithmeticEngine S` record. Because `S : TextbookSchema` is unconstrained, `hLit : ClassicalTextbookCore S` is forwarded opaquely to `connected_of_textbook` (and is trivial when all 60 slots are `True`), and `CertificateGrounding FamP.famPSkelEngine`, `humanCore_riemannHypothesis`, and `RH_iff_Gz_rpow_bound` imply (as a Tier (C) hand corollary) the equivalence `Nonempty (PrincipalArithmeticEngine S)` $\leftrightarrow$ `RiemannHypothesis` for such schemas, these twelve theorems are interface reductions from an RH-equivalent record rather than standalone proofs of RH, and the 60-field catalogue is not load-bearing in the archive-free layer.

Remark 5.1 (Dual-layer grounding, scope erratum, and honest dependency boundary). In the archive-linked module `CertificateGrounding.lean` (288 lines, `lean` and `leanchecker` exit 0, Tier (A*)), `CertificateSchema` maps its 60 cluster fields one-to-one onto 60 propositions formed by a many-to-one grouping of the 145 distinct live certificate leaves in `PrincipalSurface.olean` (18 singletons and 42 conjunctions), and `certificate_textbook_core` supplies witnesses for all 60 propositions. In this archive-linked module, `certificate_arithmetic_engine` still uses the dummy family $\{$`_dummy : PUnit.{2}` $\mid 0 < \varepsilon\}$, ignores `hLit` (`intro _hLit`), and delegates directly to `AlpogePrincipalDirect.windows_constant_one`. Because `CertificateGrounding.lean` is preserved byte-for-byte to match its receipt digest (`7efcb348d6879ad5`), legacy comments in that file advertising "UNCONDITIONAL Mathlib `RiemannHypothesis`" and crediting an earlier local 4,195-block `con-leche/nanoda` replay (1,484,202 declarations, which belonged solely to the upstream terminal `PrincipalSupply.hSup`, not to `CertificateGrounding` or `closedRH`) are explicitly superseded by the Tier (A*) trusted-import classification documented

here and in lean/archive_linked/README.md. Separately, the five-checker companion module CertificateGroundingFamP.lean (156 lines, 5 targets, Tier (A)) supplies a parametric Type 1 skeleton FamPSkel (with witness Ix := fun _ => Unit and intro _hLit) taking Principal WindowsCert as an explicit hypothesis; this shows that PUnit.{2} was not forced by universe constraints, while instantiating FamPSkel against windows_constant_one in an archive-linked module remains future work. Neither AlpogeHumanCore nor CertificateGroundingFamP reproves the archive's 1.02×10^6-declaration internal engine: in the staged uncut baseline phase1_dag_xray_full.json, all 145/145 live leaves are reachable from principal_family (1,264,775 nodes), oneSidedWindows (1,264,781), and hSup, while local source inspection (Tier (C), as windows_constant_one is listed with present: false in phase1_dag_xray_full.json) verifies that windows_constant_one (PrincipalDirect.lean:494-498) calls principal_family (source-plus-subcone reachability, distinct from a proof of logical necessity).

5.4 Native .olean auditing and three-stage reported live zero-commit compilation

Auditing twelve top-level GRH hypothesis records in PrincipalDirect.lean (receipts/structural/fat_bundles_audit.json and exact_import_closures.json) reveals zero-projection fields inherited from general-character scaffolding, notably 13/60 zero-projection fields in UnitaryGroupThirdPlaceHypotheses (P, fields_proj_not_in_dag = 13, transitive reverse-import closure 9 modules; the live cuts below remove 14/60 fields by additionally removing the singly-projected field siegelEisensteinSimplePole, which has proj_callers_count = 1 in MemberBundlesLargeUnramifiedExponents), and 33/53 zero-projection fields in ThetaLocalInputs (F, total closure 491 modules, of which 489 are archive modules); separately, local audit of LiteratureCMHypothesis (cm, DAG-induced closure 402 modules in exact_import_closures.json, of which 400 are archive modules plus PrincipalDirect and PrincipalSupply; the live Stage-2 build compiled 402 archive modules, comprising these 400 plus two non-DAG import bridges; not one of the twelve fat_structures in fat_bundles_audit.json) found 11/11 zero-projection fields, whereas UnusedSlot80 (2,175 decls) and UnusedSlot33 (298 decls) were both verified on the local DAG cache (Tier (C)) to be active renamed namespaces rather than dead slots. Because Lean's Prop structure constructors require every field and proof irrelevance never manufactures missing proofs, deleting dead fields unconditionally requires recompiling the transitive reverse-import closure rather than merely pruning edges in a DAG cache.

Two native Lean 4.33.1 tools were engineered to operate at the million-declaration scale on a 32 GB workstation:

1. *Advisory native .olean auditor (LeanXRayDump.lean, Tier (B)):* Queries Lean.getStructureInfo? directly on compiled environments and resolves four subtle kernel pitfalls validated by three compile-time selftests (demo_consumer_proj, demo_consumer_elim, demo_constructor_only): (a) ConstantInfo.value? returns none on .thmInfo unless theorem bodies are unpacked explicitly; (b) field accesses compile into both primitive Expr.proj nodes and named Expr.const projections; (c) StructureInfo.fieldInfo is ordered by Name.quickLt whereas kernel Expr.proj indices follow declaration order in fieldNames; and (d) any eliminator or constructor presence invalidates naive field-deletion heuristics, keeping advisoryNoProjOrListedElimInCone strictly advisory (recompilation remains mandatory). On humanCore_riemannHypothesis (60,787 cone constants, 0 missing declarations, 2,074 distinct cone-wide structure/index Expr.proj pairs; receipt alpoge_humancore_xray_audit.json summarizing three structures out of four queried by the tool), the auditor directly confirms that 0/60 TextbookSchema fields, 0/5 ClassicalTextbookCore fields, and 1/8 FamBase fields (Win) are projected in the archive-free cone (yielding 3 advisory candidates and matching the manual audit).

2. *Zero-commit live compiler and reverse-topological sweep (`SliceFrontendV2_OSS.lean` + `ReadonlyArchive_portable.c`, Tier (A*, reported)):* Stock `lean.exe` exhausts Windows pagefile commit limits when mapping $> 100{,}000$ `.olean` files (≈ 94.3 GB on the producer archive build). Linking Lean's `Elab.Frontend` against read-only Windows file mappings (`PAGE_READONLY`, executed on an earlier revision on the producer workstation; `tools/ReadonlyArchive_portable.c` supplies a cleaned two-argument ABI revision plus an uncompiled POSIX `MAP_PRIVATE` reference port guarded for `MAP_NORESERVE` with a 64-bit static assertion) maps archive objects with zero commit charge on successful read-only mappings (in the unarchived producer run log: 175,042 standalone mappings totaling 94.32 GB with 124 address fallbacks, while performing one reverse-topological per-module sweep per audited root at $\approx 1{,}850$ declarations/s).

Remark 5.2 (Three-stage reported live end-to-end compiled RH cuts; Tier (A*, reported)). Three sequential producer-side live zero-commit compilation runs were executed with `SliceFrontendV2.exe` (using an earlier build than the cleaned `SliceFrontendV2_OSS.lean` bundled in `tools/`) and recorded in Table 5:

1. *Stage 1 (`cut_P_hironaka`):* Recorded in `CutClosedRH.audit.json` (SHA `78f2045af39d2cde`), this run deletes 14/60 fields of `UnitaryGroupThirdPlaceHypotheses` (13 zero-projection fields plus the singly-projected `siegelEisensteinSimplePole`, adjusting its sole consumer `MemberBundlesLargeUnramifiedExponents` across 7 recompiled archive modules; 139.43 s in `slice_frontend_receipt_P.json`) jointly with eliminating Hironaka leaves 1 and 2 via `topCovolSlim2_use_without_res`, followed by `PrincipalDirect`, `PrincipalSupply`, and `CutClosedRH` (unarchived producer log: 279.0 s, 382.7 s, 399.3 s; total 3,141.5 s). In `CutClosedRH.audit.json`, `Environment.addDeclCore` accepts `PrincipalSupply.hSup` and `AlpogeLiveCutP.closedRH` with axioms `[propext, Classical.choice, Quot.sound]` and 3/21 tracked landmarks absent (`gelfandKazhdanDual_of_multiplicityOne`, `resolutionCharZero`, `principalisationCharZero_closed`), leaving the 15 structurally dead F-landmarks and the CM landmarks present. (The Stage-1 `CutClosedRH.lean` source is not bundled.)

2. *Stage 2 (`cut_PFCM_hironaka`):* Recorded in `CutClosedRH_PFCM_hironaka.audit.json` (SHA `f645d4143390634b`) and `LIVE_UNCONDITIONAL_PFCM_HIRONAKA_RECEIPT.json` (SHA `4c5af5de8292dcfe`), the Stage-2 build comprises 897 recompiled archive modules (495 $P|F$ + 402 cm) deleting all 33/53 zero-projection fields of `ThetaLocalInputs` (F; preceded by the separate 15/489 prefix benchmark `slice_frontend_receipt_F_bench15.json`, 104.96 s slice / 107.77 s wall time) and removing `LiteratureRankNHypothesis.cm` (cm), expanding the historical DAG-derived comparison baseline S_{895} (495 $P|F$ + 400 DAG cm archive modules) by two non-DAG import bridges (`CarFillE` and `WindowBoundOfLiteratureLevel`, imported by `TopMainZj`; producer log, module list not bundled). The receipted invocation compiled the 402 cm modules (`cm_sec` = 28,357.04 s) plus three project modules (`PrincipalDirect`, `PrincipalSupply`, `CutClosedRH`; `final_sec` = 2,743.43 s, `total_wall_sec` = 31,103.44 s) and reused 495 $P|F$ `.olean` objects from an earlier invocation (`pf_sec` = 0.01 s; no separate full $P|F$ receipt is bundled). The bundled entry file `lean/archive_linked/live_cuts/CutClosedRH_PFCM.lean` (14 lines, SHA `2394496b8feaec8e`, matching `final_artifacts.CutClosedRH.source_sha256`) proves `AlpogeLiveCutPFCM.closedRH : RiemannHypothesis`. The live audit verifies `accepted: true` under `Environment.addDeclCore`, standard axioms only, and 21/21 **tracked landmarks absent** (+18 absent among the original 21 landmarks, including all 15 dead F-landmarks plus `GlobalTheory.GaussManinTop.main` and `UpperTheory.Lit.serreTateGoodReduction`), reducing the live cone by $-3{,}862$ archive-and-project declarations ($-3{,}906$ total reachable constants) and -544 archive-and-project modules relative to Stage 1 (exceeding the 3,680-declaration F + cm projection in `joint_multicut_bfs.`

json by 182 declarations, plausibly due to lazily generated equation lemmas and the two non-DAG bridges).

3. *Stage 3 (cut_PFCM_hironaka_stick):* Recorded in CutClosedRH_PFCM_hironaka_stick.audit.json (SHA afe51d5a8480c445) and LIVE_UNCONDITIONAL_PFCM_HIRONAKA_STICK_RECEIPT.json (SHA 97e8e0ff464bf56b), this run eliminates the Stickelberger odd-prime bound (UpperTheory.V.Odd.stickBound) with **zero additional archive edits and zero additional archive compilations**. In the unmodified archive, UpperTheory.V.ThmBc universally quantifies over all archive M : AbF data and therefore invokes stickBound for odd primes $\ell \mid \deg(M)$; however, the principal-character RH path instantiates ThmBc exclusively at M := AbF.ofChar PChar.one. The bundled 291-line project-layer module lean/archive_linked/live_cuts/PrincipalThmBcOne.lean (SHA 98e05e86a5b20a72, matching final_artifacts.PrincipalThmBcOne.source_sha256) proves C3Check.deg_ofChar_one : (AbF.ofChar PChar.one).deg = 1, rendering the odd-prime Stickelberger sum vacuously empty, proves C3Check.thmBcAtOne_of_psiBound without stickBound, and restates the two project consumers (principal_vwin_no_caseIII_at and window_of_principal_counts_at) under the weaker hypothesis C3Check.ThmBcAt; the call-site switch in the remaining consumer (the private principal_host_bound) occurs in the edited, unbundled PrincipalDirect.lean (whose source_sha256 changes from 1b39bda8 in Stage 2 to ef83bc1d in Stage 3). Reusing all 897 cached Stage-2 archive .oleans and compiling only 4 project modules (PrincipalThmBcOne, edited PrincipalDirect, PrincipalSupply, and bundled CutClosedRH_PFCM_Stick.lean, SHA 213ecf5d4954922c, declaring AlpogeLiveCutPFCMStick.closedRH : RiemannHypothesis) completed in 3,896.7 s (exit 0), reducing the live cone further by −222 archive-and-project declarations (−225 total reachable constants) and −42 archive-and-project modules (cumulative vs. Stage 1: −4,084 archive-and-project declarations, −4,131 total reachable constants, −586 archive-and-project modules) with 27/27 **tracked landmarks absent**—all 21/21 original landmarks plus all 6 newly tracked Stickelberger/ThmBc landmarks (stickBound, StickBound, thmBcFree_of, ThmBc, ThmBcFree, thmBc_iff, whose Stage-1 and Stage-2 statuses were not measured live in the 21-landmark audits; four of the six appear in the bundled uncut DAG/C3 receipts, whereas stickBound and StickBound have no bundled pre-cut positive-control record).

Epistemic and provenance boundaries: Although the Stage-2/3 top entry files and PrincipalThmBcOne.lean are bundled with matching source_sha256 digests, all three live runs remain **Tier (A*, reported)**: the 897 recompiled archive objects, archive patch diffs, the 94 GB upstream archive 79c751e0, and the author's project layer (PrincipalSupply, PrincipalDirect [edited across stages; source digests 1b39bda8 and ef83bc1d in the receipts], PrincipalCounts, PrincipalLog, etc.) are not bundled, and no live receipt records a source_sha256 for the support module AlpogeSimplified/ArchitectureSimplified.lean (which supplies riemannHypothesis_of_simplified_alpoge and is pinned only by SHA256SUMS.txt). Accordingly, the live runs are reported from their archived receipts and are not reproducible from the release bundle alone, nor from the bundle together with the public upstream archive 79c751e0. The audit JSONs record declaration names, accepted: true, axioms, and BFS counts without printing the type expression RiemannHypothesis; per unbundled producer runner sources, the runners skip pre-existing .oleans and verify reused Stage-2 objects by existence and nonzero size (without a per-object SHA manifest), and the Stage-3 runner's fail-closed validator required ≥ 25/27 absent landmarks (tolerating up to two legacy CM landmarks before the live sweep confirmed 27/27 in afe51d5a8480c445). All checks use Lean 4.33.1's C++ kernel against trusted read-only archive and project imports without external export replay.

Quantitatively, these cuts remain surgical rather than wholesale compressions: the Stage 1 → Stage 3 reduction of 4,084 archive-and-project declarations represents 0.323% of the 1,264,848-

Live-cut stage	Arch. mods (cumul.)	Proj. mods	arch+proj decls	Total decls (hSup/closedRH)	arch+proj mods	Absent landmarks
Stage 1 (cut_P_hironaka)	7	3	1,263,979	1,477,330/1,477,360	146,187	3/21
Stage 2 (cut_PFCM_hironaka)	897[†]	3	1,260,117	1,473,424/1,473,454	145,643	21/21
Stage 3 (..._hironaka_stick)	897 (0 new)[‡]	4	1,259,895	1,473,199/1,473,229	145,601	27/27[§]
Δ (Stage 1 → Stage 2)	+890	0	−3,862	−3,906/−3,906	−544	+18 of 21
Δ (Stage 2 → Stage 3)	0	+1	−222	−225/−225	−42	6 newly audited
Δ (Stage 1 → Stage 3)	+890	+1	−4,084	−4,131/−4,131	−586	+18 of 21; 6 new

Table 5: Exact measured progression across the three reported live zero-commit compilation stages (Tier (A*, reported)). Cone counts and landmark outcomes are read directly from `CutClosedRH.audit.json`, `CutClosedRH_PFCM_hironaka.audit.json`, and `CutClosedRH_PFCM_hironaka_stick.audit.json`; compiled-module counts come from `slice_frontend_receipt_P.json` (7), the two Stage-2/3 live receipts (897 and project-module lists), and the unarchived producer log (Stage-1 3 project modules). Both audited roots (`Principal Supply.hSup` and each stage's `closedRH`) share identical `archive_and_project_constants` (`arch+proj` decls) and `archive_and_project_modules` (`arch+proj` mods), pass `Environment.addDeclCore` (`accepted: true`), and depend strictly on `[propext, Classical.choice, Quot.sound]`. [†]495 cached $P|F$ + 402 newly compiled cm archive modules (including two non-DAG import bridges beyond S_{895}). [‡]Zero extra archive edits/compiles: reuses all 897 cached Stage-2 archive `.olean`s and compiles 4 project modules including `PrincipalThmBcOne` and edited `PrincipalDirect`. [§]Denominator increases from 21 to 27 via six newly audited Stickelberger/`ThmBc` landmarks.

node Phase-1 DAG baseline (4,953 declarations or 0.392% comparing Stage 3 directly against the Phase-1 DAG cache; because no live audit of the uncut `hSup` root is bundled, the 869-declaration Stage-1 gap combines the Stage-1 cut with the `isArchiveOrProjectMod` classifier difference and cannot be decomposed from the bundle alone). Their significance is methodological—demonstrating absence of 27 named classical and architectural landmarks (including Hironaka leaves 1–2, the 15 dead F-landmarks, `GaussManinTop.main`, `serreTateGoodReduction`, and Stickelberger's `stickBound`) in the final live dependency cone (confirmed by live reachability audit, with both top roots re-checked by the kernel all the way to `RiemannHypothesis`) on a single 32 GB workstation—rather than a large-factor shrinkage of total archive mass.

By contrast, further post-Stage-3 cuts and earlier multi-cut simulations remain strictly Tier (C) graph simulations and partial overlay compiles rather than end-to-end live RH builds. In `receipts/computational/joint_multicut_bfs.json`, blocking $P+F+\mathrm{cm}+\mathrm{Hironaka}_{1,2}$ on the DAG cache removes 3,751 declarations (19 names of the unbundled 1,887-name `Principal Surface` catalogue; historical S_{895} comparison baseline), while larger unions reach −4,006 and −4,037 (with −3,782 reproduced in `tools/joint_multicut_bfs.py`). A subsequent post-PFCM closure autopsy (`receipts/structural/next_live_cuts_closures.json`, SHA `3939f78f91657b9a`, Tier (C)) evaluates five further cuts stacked onto S_{895} (all with `newly_reachable = 0`, reporting `realistic_marginal` reductions alongside narrower `graph_sim_marginal` counts in parentheses): `stick` (−231 realistic [−220 `graph_sim`] vs. −222 archive-and-project / −225 total declarations now live-compiled in Stage 3), `litloc+litaut` (−14 realistic, DAG-capped conservative closure 400 modules fully contained in S_{895}; `LitAut` overlay compiled with exit 0, while `TopMainZj` exited with `exit: 1`, `errors: 1` across 157,602 transitive imports—attributed to a Windows file-read ceiling in the unarchived producer log), `unused33` (−41 realistic [−35 `graph_sim`]; all 6/6 semantic-closure overlay modules compiled with exit 0,

DAG-capped conservative closure 17,308 archive modules, uncapped per-root up to 20,877), and `fsch` (−23 realistic [−17 `graph_sim`]; both `HypothesesX` and `SemiposAtOfPositive` overlays compiled with exit 0, DAG-capped conservative closure 3,146 archive modules, uncapped per-root up to 6,458; note that DAG-capped closures can undercount non-DAG bridges, as occurred in Stage 2). Combined, the five further cuts contribute −309 realistic simulated declarations (−78 for the four cuts not yet live-compiled), yielding a stacked DAG-simulated total of 4,060 = 3,751 + 309 vanished declarations with the S_{895} baseline included, and 23 = 19 + 4 of the 1,887 surface names.

5.5 Refined-DAG Mathlib extraction modularity scan

To assess which self-contained mathematical developments inside the 1,264,848-declaration $\psi = 1$ cone could potentially be extracted independently of the RH architecture, a deterministic modularity scan (Tier (B); compact summary archived in `receipts/structural/mathlib_extraction_summary.json`, SHA `cf1f231ba9e47555`, extracted from the full 25 MB producer reports `ebbb9aac4ddd6503` and `c54aeef4232db594`) first corrected lazily generated equation/congruence lemmas in the raw DAG (15,876 nodes matching `.eq_N`, `.congr_simp`, `.splitter`, `.eq_def`, or `.induct`). Re-attributing 8,249 lazy lemmas to their base declaration's module (23,445 raw edges, rule R1) and removing 96 foreign sink artifacts (229 raw edges, rule R2) yields a refined DAG of 1,264,752 reachable declarations and 14,075,317 directed edges; cross-checking actual archive `import` headers across 17 uncapped candidates reduces impossible "not covered" closure modules from 7 (raw DAG) to 0 (refined DAG).

On the refined DAG (node-set-deduplicated module/prefix subgraphs with $|S| \geq 200$ declarations; Mathlib/Std/Init/Lean constants are excluded from the DAG, so external-dependency counts $|\partial_{\text{out}}(S)|$ measure archive declarations only), the scan identifies:

- 14 **zero-external-archive-dependency subgraphs** ($|\partial_{\text{out}}| = 0$, regex-level `source_audit` counts `sorry = 0`, `axiom = 0` recorded for all 14 in `mathlib_extraction_summary.json`; raw DAG: 13; note that two subgraphs are nested—`NoetherianQuotientCalc` [525] lies inside `QuotientShiftIrreducible` [573]—and two are single modules, leaving 13 non-nested subgraphs covering 5,343 declarations), including `GroundTheory.TateCohomology` (1,165 decls, 50 modules), `Development.AutomorphicArithmetic.IntegralCasts.CrystallinePeriods` (770, 14), `Development.ResolutionOfSingularities.Foundations` (461, 20), and `ReciprocityTheory.ReductiveGroups.ArithmeticGeometry.SemiabelianSchemes` (432, 16); overall, 70 subgraphs satisfy $|\partial_{\text{out}}| \leq 20$, 136 satisfy ≤ 50, and 243 satisfy ≤ 100 (raw: 72/138/237), alongside 63 single-entry Lengauer–Tarjan dominator cones with $|\partial_{\text{out}}| \leq 20$ out of 4,534 cones having $|D(v)| \geq 200$.

- **Hironaka's resolution directory as a single-dependency unit:** `Development.ResolutionOfSingularities` (10,308 declarations, 8,731 non-auxiliary, 499 modules; regex-level `source_audit` `sorry = 0`, `axiom = 0` across all 499 files recorded in `mathlib_extraction_summary.json`) has *exactly one* direct external archive dependency on the refined DAG (`UpperTheory.Reformulated.B.SaturationByPowers.saturationByPowers`, transitive closure |Cl| = 4 declarations in a single 806-byte module; closure ratio 10,308/10,312 ≈ 0.9996), whereas the raw DAG showed a second external dependency (`AlgebraicGeometry.Scheme.IdealSheafData.mkOfMemSupportIff.congr_simp`, recorded in `hironaka_highlight.raw.external_decls`).

- **Strict graph-level scope caveat:** Declaration-DAG modularity is a necessary filter, not a standalone compilation guarantee: the DAG does not track notations, macros, `syntax`, `@[simp]`/`@[instance]` attributes, or `open` scopes; among the 10 of 14 zero-dependency subgraphs checked in `import_crosscheck_refined`, 9 still contain 1–17 direct outside archive `import`s (so zero dependency holds at declaration level, not import-header level);

every checked gem imports the full `import Mathlib`; and statement-dominated pillars (e.g. `Published*` literature bundles) encode hypothesis statements rather than proofs. Standalone Lake compilation of extracted candidates remains open (Section 6.3).

6 Five-checker verification architecture, reproducibility, and pure-AI provenance

6.1 Five-checker pipeline, binary provenance, and `con-leche` diagnostics

Every archive-free module in this paper was compiled and independently checked via an earlier revision of `tools/verify_lean5.py` (and configured for dependency-ordered replay via `tools/reproduce_lean4.py`) against the pinned Lean toolchain (`lean-toolchain`, `lake-manifest.json`, `lakefile.toml`; Mathlib commit `0df444a360eaa60ab8c11dca51a86af692955474`):

```
Lean (version 4.33.1, commit 819816b2e0a3bf405af45ae5c7af2491d8f5bee6)
```

together with `lean4export 3.1.0` (a single shared export stage across the three external checkers) and three external export checkers (specifically the `con-leche` and `con-ron` binaries bundled with Lean `v4.36.0`—matching Sutherland's table [25]—and `nanoda_bin`, invoked on the producer workstation against Lean 4.33.1 exports):

1. `lean.exe 4.33.1` (83,456 B, SHA-256 prefix `af49bacfabaa1fea`) and `leanchecker.exe 4.33.1` (115,200 B, `31b506f83737b7d6`): official Lean 4.33.1 C++ kernel with Lean-written elaborator and standalone environment-replay driver (one kernel lineage).

2. `con-leche.exe -verified` (23,436,288 B, `5dc6d10f2608e217`; implemented in Lean 4 [16]) and `con-ron.exe -verified` (3,703,808 B, `d9cfda21442da3b5`; Rust port of `con-leche` [17]): kernel exports generated by `lean4export.exe 3.1.0` (333,312 B, `cd51ec5333203959`), sharing the `con-leche` verified-checking lineage. Every one of the 11 primary five-checker receipts logs `1 inductive blocks modelled in-process: Lean.Syntax (30 generated records, checked by the fold...)` on stderr (the 3 diagnostic receipts log different minimal inductive blocks or a decline).

3. `nanoda_bin.exe` (1,088,512 B, `a76bed8bc1091842`; independent Rust kernel [2]): invoked with `unpermitted_axiom_hard_error: true` and permitted axiom whitelist `["propext", "Classical.choice", "Quot.sound"]`. Its stdout logs zero typechecker errors alongside the pretty-printer message `"Unable to print axioms"` (producing no positive axiom printout; axiom compliance is therefore evidenced jointly by `unpermitted_axiom_hard_error: true`, Lean `#print axioms`—noting `KLevelChainIIArcs` prints the unqualified alias `choice` via `open Classical`, and `GOAL_B_DUAL_LAYER_VERIFICATION.json` records `#print axioms` output for 7 of the 21 `AlpogeHumanCore` targets, while the remaining 14 rely on `nanoda_bin`'s `unpermitted_axiom_hard_error: true` configuration archived in the runner source rather than inside the historical receipt—and direct inspection of the JSONL export axiom records).

Accordingly, the five-checker suite executes five distinct checker runs spanning three independent kernel lineages consuming one shared `lean4export` serialization (Lean C++, Lean-4/Rust `con-leche`/`con-ron`, and Rust `nanoda_bin`). Note that the binary digests above and the upstream archive commit hash `79c751e0` are author-recorded provenance (archived in `AI_DISCLOSURE_AND_PROVENANCE.md`; the historical 25 November JSON receipts record `source_sha256`, `.olean` digests, and checker outputs rather than binary-digest or timestamp fields).

Table 6 reconciles all 33 staged `.lean` files under `lean/` (plus 2 native tooling `.lean` files under `tools/`). Comment-stripped static scans across all 33 staged files in `lean/` report zero occurrences of `sorry`, `admit`, `native_decide`, `axiom`, `unsafe`, `implemented_by`, `extern`,

ofReduceBool, or trustCompiler, and all 42 staged JSON records under receipts/ parse as strict JSON (no NaN/Infinity) and contain zero occurrences of sorryAx. All 14 archive-free five-checker modules (8,066 lines, 147 selected targets with their exported dependency closures) emit zero compiler warnings in their archived receipts (the 7 pure-Mathlib support modules lack separate pre-packaged receipts and are compiled and leanchecker-checked by tools/reproduce_lean4.py; one archive-linked autopsy file, A2ChainII.lean, emits a single linter warning—a Try this: intro ... tactic suggestion—under lean.exe, and SliceFrontendV2_OSS.lean under tools/ uses unsafe, opaque, and extern solely for its C FFI binding to ReadonlyArchive_portable.c). Historical JSON receipts are preserved verbatim from the producer workstation (retaining original local paths) so their SHA-256 digests remain intact.

Staged Lean 4.33.1 group	Files	Lines	Targets	Checkers passed & hygiene
Pure-Mathlib core (lean/core/)	8	5,670	103	5/5 checkers exit 0; 0 warn; Mathlib and staged core siblings only
Archive-free companion (lean/companion/)	6	2,396	44	5/5 checkers exit 0; 0 warn; zero archive imports
Pure-Mathlib support (lean/support/)	7	5,976	—	Transitive deps of companion (no separate receipts; built by reproduce_lean4.py); 0 forbidden tokens
Archive-linked (lean/archive_linked/)	9	1,291	19†	6 core lean+leanchecker exit 0 + 3 Stage-2/3 live-cut sources (addDeclCore exit 0, trusted archive 79c751e0)
Checker diagnostics (lean/checker_diagnostics/, outside universal Tier (A))	3	87	4	2/3 pass 5/5 (3 targets); 1 con-leche exit-2 reproducer (1 target)
Total under lean/	33	15,420	147‡	0 sorry, 0 native_decide, std. axioms only

Table 6: Complete file-count and verification reconciliation across all 33 staged .lean files under lean/ (excluding 2 standalone native tooling files in tools/). †Twenty-four declared theorems (19 printed in CLAUDE_AUTOPSY_VERIFICATION.json) across ClaudeAutopsy, 60-cluster grounding in CertificateGrounding.lean, plus three bundled Stage-2/3 live-cut sources in lean/archive_linked/live_cuts/ (319 lines). ‡Count of five-checker selected targets across the 14 archive-free core and companion modules (spanning headline theorems and exported helper lemmas; excluding the 4 diagnostic targets across 3 modules). Full per-module SHA-256 prefixes and declaration counts appear in Appendix B.

External-checker coverage limitation (PrincipalizationCertRepro.lean). While auditing export-checker compatibility on archive-style inductive certificates, a minimal 31-line pure-Lean reproducer (lean/checker_diagnostics/PrincipalizationCertRepro.lean) was isolated demonstrating a coverage limitation of the specific archived con-leche -verified (5dc6d10f) and con-ron -verified (d9cfda21) builds: both decline reflexive inductive members with exit code 2 (declined: in-process model ... reflexive member), whereas lean 4.33.1, leanchecker 4.33.1, and nanoda_bin (16 declarations) accept the module cleanly. As documented on Andrew Sutherland's public verification page for Alpöge's original archive [25] (note that those public checks verify Alpöge's upstream artifact and do not attest the new modules of this paper), the con-leche/con-ron binaries bundled with

Lean v4.36.0 declined such inductives whereas newer builds compiled from source support them; PrincipalizationCertRepro.lean therefore characterizes the exact tested builds archived here. Two accompanying minimal controls (ConLecheNestedRepro.lean, 27 lines, 2 targets, 16/35 decls; ConLecheMutualRepro.lean, 29 lines, 1 target, 24/46 decls) pass all five checkers with exit code 0 (these narrow controls illustrate the contrast without establishing reflexivity as the sole causal difference for arbitrary inductives).

6.2 Pure-AI provenance, role separation, and pre-release review log

This manuscript reports **pure AI-generated research** directed, scoped, and supervised by the human author (Robert Kadlec, Masaryk University, Brno, Czech Republic; ORCID 0009-0003-9245-8130) and forms part of a broader research project investigating the respective competencies of humans and large language models. No mathematical derivation, Lean 4 module, DAG script, verification harness, symbolic/numerical computation, or paragraph of manuscript prose was manually written by the human author; all technical execution and pre-release auditing were delegated to four AI systems/interfaces operating with role separation and read-only upstream boundaries:

- *Antigravity / Gemini (Gemini 4 Neon):* DAG X-Ray and Lengauer–Tarjan pipeline, SliceFrontendV2 zero-commit compiler (Stage 1–3 live-cut orchestration) and Lean XRayDump auditor, verify_lean5.py and reproduce_lean4.py harnesses, 12 archive-free Lean 4.33.1 modules (AlpogeHumanCore, GeneralizedWindowRH, SphereFree Behrend, SelbergPositiveWindowRH, DedekindAndDirichletWindows, AbelianDede kindWindows, SubPolynomialWindowRH, SubPolyZetaDensityBarriers, KLevelPAdic Interpolation) plus the 7 bundled pure-Mathlib support modules (5,976 lines in lean/support/), v1 manuscript assembly (Sections 1, 5, 6, Appendices A–B), v2 mathematical additions (Proposition 2.5, Lemma 3.2, Proposition 4.2, Corollary 4.4 relative form, Table 9, Table 10, and log-space Cramér recomputation), integration of v3/v4 referee patches across all rounds, Tectonic PDF compilation, and release packaging.

- *OpenAI Codex (gpt-6.1-terra / gpt-6.2-astra, codex-cli 0.171):* independent edge/cut audits, prior-art identification ([9, 24, 14]), sub-polynomial and superoscillation analyses, FLINT/Arb frontier computations, technical drafting of Sections 2–4, Round 1b strictly non-self-auditing pre-release review of Sections 1, 5, 6, and Appendices A–B, and an unarchived producer-side consistency audit of the Stage 2–3 live-cut receipts and runners.

- *OpenAI ChatGPT (Web interface, November 26, 2026):* Round 1c external pre-release review of v1 (reviews/R1c_Web_ChatGPT_Review_CZ.md, no model ID in transcript), identifying Moser [21], the exact closed-form Cramér threshold, Miller–Stanton [20] for the all-K Smith normal form of Λ_K, the Kozma–Oravecz page range [14], Kaczorowski–Pintz [12], the Appendix A calibration table, and the Appendix B Lean signature table; Round 2c independent non-Claude audit of v4 and zenodo_release_v1.0.0.zip (6 Pro -> Latest, maximum reasoning intensity 4/4; reviews/R2c_Web_ChatGPT_R2_Review_CZ.md), independently verifying all COI-flagged modules, all main proofs, the release manifest, and primary literature (including correcting the Kaczorowski–Pintz Part I/II scope and identifying the relative-path fix in the Python runners); and Round 3b independent non-Claude COI audit of v1.1.0 (6 Pro -> Latest, 4/4, 21 m 10 s; reviews/R3b_Web_ChatGPT_R3_Review_CZ.md), issuing non-Claude sign-off on PrincipalThmBcOne.lean, C3, and C2.

- *Anthropic Claude Code (claude-sonnet-5-5 / opus / fable, claude-cli 2.1.332):* archive autopsies (ClaudeAutopsy, 5 files, 24 declared and 19 printed theorems), KLevel ChainIIArcs.lean, 133 FLINT PhiData SNF runs, CertificateGroundingFamP.lean,

POSIX reference `ReadonlyArchive_portable.c`, `PrincipalizationCertRepro.lean`, the Stickelberger principal-character bypass (`PrincipalThmBcOne.lean`), the post-PFCM closure autopsy (`next_live_cuts_closures.json`), the refined-DAG Mathlib extraction modularity scan, Round 1a read-only review (`claude-sonnet-5-5`), and deep read-only verification audits R2a, R2b, and R3a (`fable`, max effort, explicitly declaring non-independence for Claude-authored components; R2a also drafted the v3 replacement proof paragraph of Theorem 3.3, the absolute variant of Corollary 4.4, and the two Cramér remarks, which were transcribed and integrated in v3).

Table 7 summarizes all eight pre-release AI reviews archived under `reviews/` (preserved verbatim except for redaction of one private account-scoped ChatGPT link in `R1c` and local workstation paths in `R1a`/`R1b`, as documented in `reviews/README.md`). Rounds R1a–R2c covered manuscript revisions v1–v5 (released as Zenodo `v1.0.0`, `10.5281/zenodo.24555199`) and predate the `v1.1.0` Stage 2–3 live-cut and C2/C3 additions; Round 3a and independent non-Claude Round 3b audited the complete `v1.1.0` delta.

Round	AI Reviewer	Scope	Verdict / Counts	Archived file (`reviews/`) & status
R1a	Claude Code (`sonnet-5-5`)	Full v1 paper + bundle (COI noted)	REVISE (3 E, 13 M, 22 m)	`R1a_Claude_Sonnet_Audit.md`; addressed in v2/v3 (m12 via `host_thick_archive_ref.json`)
R1b	OpenAI Codex (`codex-cli`)	v1 §§1, 5–6, Apps. A–B (non-self)	REVISE (1 E, 16 M, 12 m)	`R1b_OpenAI_Codex_Audit.md`; addressed in v2
R1c	OpenAI ChatGPT (Web, no ID)	Full v1 PDF	MAJOR REVISION (8 priority rows)	`R1c_Web_ChatGPT_Review_CZ.md`; all upgrades implemented in v2
R2a	Claude Code (`fable`, max)	Full v2 paper + bundle (titled R1 inside; COI)	REVISE (2 E, 12 M, 23 m)	`R2a_Claude_Fable_Audit.md`; 36/37 closed in v3 (M8 partial; verified by R2b)
R2b	Claude Code (`fable`, max)	Full v3 paper + bundle (R2; COI noted)	ACCEPT W/ MINOR REV. (0 E, 0 M open, 18 m)	`R2b_Claude_Fable_R2_Audit.md`; non-execution minor items addressed in v4
R2c	OpenAI ChatGPT (Web `6 Pro`, 4/4)	Full v4 PDF + ZIP (non-Claude COI check)	ACCEPT W/ MINOR REV. (0 E, 0 M, 7 m)	`R2c_Web_ChatGPT_R2_Review_CZ.md`; M1–M5 addressed in v5 (= `v1.0.0`); predates `v1.1.0`
R3a	Claude Code (`fable`, max)	Full `v1.1.0` delta + bundle (COI declared)	REVISE text-only (0 E, 4 M, 14 m)	`R3a_Claude_Fable_R3_Audit.md`; 4/4 M and 12/14 m closed before R3b (2 residuals closed via R3b)
R3b	OpenAI ChatGPT (Web `6 Pro`, 4/4)	Full `v1.1.0` PDF + ZIP + Hexagon (non-Claude COI check)	ACCEPT W/ MINOR REV. (0 E, 0 M, 4 m)	`R3b_Web_ChatGPT_R3_Review_CZ.md`; all 4 m (`R3b-m1`–`m4`) closed in final `v1.1.0`

Table 7: Pre-release AI review history archived under `reviews/` in the companion release (SHA-256 digests recorded in `SHA256SUMS.txt`; privacy redactions logged in `reviews/README.md`). Rounds R1a–R2c predate the `v1.1.0` Stage 2–3 and C2/C3 additions; Round 3a (COI declared) and independent non-Claude Round 3b audited the `v1.1.0` delta. Human peer review has not yet been performed.

All eight pre-release AI review reports and authorship/COI disclosures are archived under `reviews/` and `AI_DISCLOSURE_AND_PROVENANCE.md`. Because no LLM output is treated as proof, every Tier (A) claim rests exclusively on the machine-checked Lean 4.33.1 proofs and archived five-checker receipts, Tier (A*) claims rest on Lean 4.33.1 kernel checks relative to pinned read-only archive revision `79c751e0`, and Tier (B)/(C) claims rest on deterministic archived verification, structural, and computational JSON records in `receipts/` (authenticated by `SHA256SUMS.txt`), with unbundled local DAG-cache traversals, historical run timings, and hand proofs explicitly distinguished throughout. Human peer review has not yet been performed on this preprint.

6.3 Artifact availability, reproduction, and open problems

This preprint corresponds to companion release `v1.1.0` (Apache-2.0 for all Lean, Python, and C sources; CC BY 4.0 for the manuscript and documentation), archived permanently on Zenodo under versioned DOI 10.5281/zenodo.24558569 (`zenodo_release_v1.1.0.zip`; extending Zenodo `v1.0.0` 10.5281/zenodo.24555199 under all-versions concept DOI 10.5281/zenodo.24555198). The release archives all 33 staged `.lean` files under `lean/`, all 8 relocatable native and Python tools under `tools/` (including `tools/reproduce_lean4.py`), all 42 verification, structural, and computational JSON records under `receipts/`, all 8 pre-release AI review reports (plus `reviews/README.md`) under `reviews/`, pinned Lean toolchain manifests (`lean-toolchain`, `lake-manifest.json`, `lakefile.toml` declaring `[leanOptions]` without Lake library targets, since `reproduce_lean4.py` is the multi-directory build runner), minimum Python version specifications (`requirements.txt`), `.gitattributes`, the expanded scope erratum `lean/archive_linked/README.md`, the Tectonic-compiled PDF (0 Overfull `\hbox` verified on the XeTeX/`fontspec` branch; arXiv's default pdfLaTeX branch should be verified upon submission), and a complete `SHA256SUMS.txt` manifest verified byte-for-byte during packaging against both the local squashed snapshot commit (`git archive HEAD`) and the deposited archive `zenodo_release_v1.1.0.zip` (each Zenodo release is packaged as a standalone squashed snapshot without a separate public Git hosting mirror; readers verify `SHA256SUMS.txt` directly against the Zenodo ZIP). Reproducing all 14 archive-free core and companion modules (plus 7 support modules and 3 diagnostics) from a clean checkout requires Lean 4.33.1 with pinned Mathlib commit `0df444a360eaa60ab8c11dca51a86af692955474` (`lake exe cache get && lake build Mathlib`), followed by `python tools/reproduce_lean4.py -packages-dir .lake/packages` (which passes the matching `-D` options to the five `Alpoge Simplified.*` support modules, runs `leanchecker` on all 24 modules including all 7 support files, cleans temporary `nanoda_*.json` files, and under `-external-check` regenerates `.export` files via `lean4export 3.1.0`, replays them through `con-leche`, `con-ron`, and `nanoda_bin`, and fails on any rebuilt export size or SHA-256 mismatch unless `-allow-export-drift` is passed). Note that `tools/reproduce_lean4.py` has been statically audited and command-plan-checked via `-dry-run` on the staged tree; a full clean-clone re-execution on an independent machine has not yet been performed and remains an explicit independent-verification step. While `probe_alpoge_frontiers.py`, `w4_superosc_density_check.py`, and `joint_multicut_bfs.py` are bundled in `tools/`, several producer-side exploratory and audit receipts (`c1_gluing_*.json`, `c1_line_refs_verified.json`, `c1_proof_mass.json`, `x3_superlog_check.json`, `mathlib_extraction_summary.json`, `next_live_cuts_closures.json`, and the Stage 2–3 live-cut receipts) are archived as read-only outputs without bundled generator scripts. Alpöge's 94 GB archive (`79c751e0`), the author's project layer (`PrincipalSupply`, `PrincipalDirect`, `PrincipalCounts`, `PrincipalLog`, etc.), the live-cut archive patch diffs, the producer build of `SliceFrontendV2`, and the derived local `dag_cache.npz` are external and not bundled; consequently, reproducing the Tier (A*) archive-linked audits and Tier (A*, reported) live cuts requires all of those external components and cannot be performed from the release bundle alone (nor from the bundle plus public revision `79c751e0`), while all summary receipts are inspectable directly in `receipts/`.

Four open mathematical and engineering problems remain directly framed by this artifact: (i) wiring the K-level lifting algebra of Section 2 into the archive's crystalline Wach-coordinate bounds (`ThmEFull`) and proving a functorial K-ary structure theorem alongside the open Chain II bound PerCutAlt$(R \geq 4, c_1)$; (ii) closing the super-logarithmic sub-polynomial gap $W(U)/\log U \to \infty$ with $W(U)/U \to 0$ (and the unattained-supremum case) for $\zeta(s)$ beyond finite constant-coefficient trigonometric polynomials (cf. Corollary 4.4); (iii) constructing cofinal one-sided window inputs via Hecke theta-lift chains over imaginary quadratic CM base fields for non-abelian Dedekind and Rankin–Selberg L-functions such as the S_3 cubic field of discriminant -23 (where Section 3 illustrates nonnegativity of the permutation/Dedekind

weights $1 + \operatorname{tr}\rho$ and Rankin–Selberg weights $|\operatorname{tr}\rho|^2$ despite negative prime traces of ρ alone); and (iv) performing independent clean-machine replays of both `tools/reproduce_lean4.py` and the Stage 2–3 live-cut pipeline with per-object `.olean` SHA-256 verification and bundled archive patch diffs, testing standalone Lake builds of the shortlisted Mathlib-extraction candidates (starting with `Development.ResolutionOfSingularities` plus `SaturationByPowers` and `GroundTheory.TateCohomology`), and resolving the remaining post-Stage-3 conservative-closure cuts (`litloc+litaut`, `fsch`, and `unused33`).

A Complete 60-field classical catalogue

Table 8 records the one-to-one mapping of the 60 topical fields of `AlpogeHumanCore.Textbook Schema` onto 60 propositions formed by a many-to-one grouping of all 145 distinct live certificate leaves (141 from `PrincipalDirect.lean` in clusters `T1_01`–`T5_06` plus 4 Selmer/Iwasawa leaves from `principal_host_bound` in clusters `T5_07`–`T5_08`; 18 singleton fields and 42 conjunctions) across five mathematical pillars ($23 + 12 + 9 + 8 + 8 = 60$). Full qualified Lean paths appear verbatim in `lean/archive_linked/AlpogeHumanCore/Certificate Grounding.lean`. Excluded from the table are the 16 structurally dead (zero-projection, zero-external-call-site) landmark theorems (`cuspidalMirabolic`, `idealInvariantsOneOrbit`, `discreteSeriesLowestKType`, `holomorphicDiscreteSeries_holds`, `logDerOne_holds`, `Stone VonNeumann_holds`, `hermitianLocalGlobal_holds`, `quadraticClassFieldTheory_holds`, `thetaRationalInvariance_holds`, `weylBranching_holds`, `gkGlobalisation_holds`, `kudla SplittingSplitUnitary`, `metaplecticRepresentation_holds`, `linnik`, `gelfandKazhdan Dual_of_multiplicityOne`, and `rallis_of_dm`) and the 37 derived leaves.

Table 8: One-to-one mapping of all 60 topical clusters in `AlpogeHumanCore.TextbookSchema` to 60 propositions grouping the 145 distinct live certificate leaves in `CertificateGrounding.lean`. Entries marked † are the two Hironaka leaves proved strictly redundant in `A1RedundancyTop.lean` and removed in the reported live cuts `cut_P_hironaka` through `cut_PFCM_hironaka_stick` (Remark 5.2). ‡Catalogue headings are topical labels, not claims of full classical equivalence; see Table 9 for interface-versus-classical calibration of key clusters (e.g. `T3_01`).

ID	Classical textbook cluster	Live certificate leaves (abbreviated)
Pillar I: Local & global representation theory, Langlands & automorphic forms (23 clusters)		
`T1_01`	Bernstein–Zelevinsky finiteness	`bernsteinZelevinskyFiniteness`, `Prw.bernsteinZelevinskyFiniteness_holds`, `standardFunctorsFiniteLength_top`, `irreducibleAdmissibleGL_top`
`T1_02`	BZ geometric lemma	`GeoLink.geometricLemma`, `CompanionTheory.Geo.geometricLemma Filtration`, `PositivityTheory.Geo.geometricLemmaFiltration`
`T1_03`	Zelevinsky segment classification	`segmentRepresentation`, `langlandsMultisegments_outright`, `langlands ExhaustionGL_holds`, `standardCosocleGL_holds`, `Mix.polynomialRing`
`T1_04`	Bernstein–Tadić unitary dual GL_n	`bernsteinUnitaryProductIrreducible`, `genericUnlinkedGL_holds`, `genericUnitaryDualGL_holds`
`T1_05`	Local converse theorem for GL_n	`DiscreteGroups.sglob.LConv.localConverseGL`
`T1_06`	Moy–Prasad depth-zero theory	`moyPrasadJacquetIso`, `SmGL.depthZeroSupport`, `depthZeroJacquet Injective`
`T1_07`	Borel–Casselman Iwahori-spherical	`GlobalTheory.Published.BorelCasselman_holds`
`T1_08`	Radford shuffle algebra	`GlobalTheory.Published.radfordShufflePolynomial_holds`
`T1_09`	Green finite $\mathrm{GL}_n(\mathbb{F}_q)$ characters	`greenRegularCharacters_holds`, `GreenCharacters.green`, `...green Cuspidal`, `Stage1.PM.greenCuspidal`
`T1_10`	Brauer degree-0 & Chevalley units	`BrauerZero.brauerDegreeZero`, `Published.chevalleyUnits'`
`T1_11`	Flath restricted tensor product	`Rg.Fl.B.flathTensorProduct`, `Rg.Fl.B.unitaryIrreducibleAdmissible`
`T1_12`	Harish-Chandra finiteness & analyticity	`RgHf.harishChandraFinite`, `print_harishChandraModules`, `harish Chandra_analyticity_density`, `kFiniteVectors_irreducibleUnitary_unitaryGroupPQ`
`T1_13`	Casselman–Wallach globalization	`PositivityTheory.CWB.casselmanWallach`
`T1_14`	Dixmier–Malliavin factorization	`CompanionTheory.DM.dixmierMalliavinFactorisation_holds`, `Prw.dixmierMalliavinFactorisation_holds`

Continued on next page

ID	Classical textbook cluster	Live certificate leaves (abbreviated)
T1_15	Harish-Chandra discrete series	discreteSeriesExists, ...UpqExists, ...CoverExists, ... Classification, ...Cover, ...HermitianUnique, formalDegreeCompact Dual
T1_16	Schur orthogonality for L^2 reps	GlobalTheory.Published.schurOrthogonalitySquareIntegrable'
T1_17	Harish-Chandra cusp rapid decay	cuspDecayBound, cuspRapidDecay, cuspDecaySobolev, cuspBounded, cusp FormsBoundedGL
T1_18	Jacquet–Shalika rigidity	JSW.jacquetShalikaRigidity_holds, AdelicQuotientData.jacquet Shalika
T1_19	Arthur–Clozel base change	UpperTheory.BaseChangeIdentities.baseChangeN
T1_20	Strong approximation for SU	GlobalTheory.StrongApproxSUProof.strongApproxSU
T1_21	Hasse norm theorem & Landherr	QNorm.hasseNormSt, QNorm.landherr_and_landherrSt, hermitianHasse Principle_holds
T1_22	Hecke–Tate L-functions	heckeTateFunctionalEquation, partialZetaFunctionalEquation, Hecke LPartialContinuation_holds, HeckeLEntireAndNonzeroAtOne_holds
T1_23	Hecke ℓ-adic chars & lowest alcove	Stage1.PM.heckeLAdicCharacter, Alc.inducedSimpleLowestAlcove
Pillar II: Arakelov heights, toroidal compactifications & birational geometry (12 clusters)		
T2_01	Hironaka resolution & principalisation	principalisation_closed, principalisationCharZero_closed†, resolutionCharZero†, FromRes.resK
T2_02	PEL/Siegel toroidal compactification	toroidalCompactification_holds, toroidalForFanKS, siegelToroidal Integral_holds
T2_03	Serre GAGA & formal algebraisation	serreHolomorphicSections_holds, FormalThickening.formal Algebraisation
T2_04	Moriwaki arithmetic Hilbert–Samuel	HilbertSamuelMoriwaki.hsUse_outright, ...moriwakiBig_outright
T2_05	Nef Riemann–Roch & Chern degree	Tls2.Nef.nefRR, nefNumerics_outright, Tls2.Chern2.chernDegree Complex
T2_06	Poincaré–Lelong & Chern–Weil	Tls2.Pl.poincareLelong, ...poincareLelongInv, GoodMetric.good ChernWeil
T2_07	Nadel vanishing & Angehrn–Siu sep.	NSI.nadelVanishingAndNonKltDivisor, NSI.pointSeparationPair
T2_08	Ohsawa–Takegoshi L^2 extension	GlobalTheory.l2ExtensionLogCanonical'
T2_09	Chow forms & height cardinality	Published.chowFormExists_holds, HtCardProof.htCard
T2_10	Abelian scheme duality & polarisation	dualAbelianScheme, principalPolarisation_unique_section, abMulN_holds
T2_11	Grothendieck semistable reduction	SS.semistableReductionSt, SAH.semiAbelianHomExtension, Nef.semi AbelianLieSeminegative
T2_12	Abelian de Rham & deformation	DualRT.abelianDeRham, HTFirstOrderIso.gd
Pillar III: Galois representations, p-adic Hodge theory, CM & class field theory (9 clusters)		
T3_01	Local Langlands surrogate for GL_n ‡	UpperTheory.Zero.LLC.recTExists
T3_02	Bhatt–Scholze integral/prismatic	BSP.bhattScholzeIntegral, Lit.prismaticDieudonneHodge
T3_03	Fontaine–Berger–Fargues p-adic Hodge	MemLoc.fontaineBcris, deRham_potentiallySemistable, Lit. farguesExact
T3_04	Clozel purity & cohomological reps	ClozelPurity.clozelPurity, Rg.Cz.clozelCohomological
T3_05	Global Artin reciprocity & ray class	FamInh.globalArtin, Stage1.PM.rayClassFieldSt
T3_06	Chebotarev density theorem	Published.chebotarevDegreeOne', Published.cheb_holds
T3_07	Main theorem of CM & Weber ray fields	ellipticMainTheoremCM, hilbertClassFieldJ, complexMultiplicationJ, weberRayField, cmCurveGoodOutsidePrime
T3_08	Deligne–Ribet p-adic L-values	PositivityTheory.PfDRL.deligneRibet
T3_09	Gelfand–Kazhdan & algebraic Hecke	AbelianGaloisNormality.gk_holds, Gal.Int.Inp.alg, Gal.Int.Inp.hk
Pillar IV: Theta correspondence, Siegel–Weil formula & doubling method (8 clusters)		
T4_01	Weil representation for unitary pairs	weilRepresentationUnitaryPair, weilRepresentationExists_pair, realMetaplecticRepresentation_holds
T4_02	Howe duality (real & GL)	howeDualityReal_holds, Published.glHoweDuality_holds
T4_03	Convergent unitary Siegel–Weil	SW.siegelWeilUnitaryConvergent, Sw47.convergent, Sw47.weilUnitary, SiegelWeilCoefficients.field, Sw.Glob.wtc, siegelWeilStd
T4_04	Siegel reduction for unitary groups	CompanionTheory.SiegelRed.siegelSetUnitary
T4_05	Siegel Eisenstein continuation	RgEis.Sp.Hd.siegelEisensteinContinuation, RgEis.Hd.eisBL
T4_06	Unramified doubling for GL_n	UnramDoublingProof.unramifiedDoublingGL
T4_07	Siegel–Klingen rationality	PositivityTheory.Prw.siegelKlingen_holds
T4_08	Theta lift nonvanishing kernel	UpperTheory.TP.K.tLiftK
Pillar V: Galois cohomology, Poitou–Tate, Bloch–Kato & Iwasawa/Selmer theory (8 clusters)		
T5_01	Poitou–Tate global duality	PositivityTheory.Lit1.PTN.poitouTateNormalised_holds
T5_02	Local Tate duality	LocalDualityAll.localTateDuality, LocalTateOrthogonality.Gcoh. top_tate

Continued on next page

ID	Classical textbook cluster	Live certificate leaves (abbreviated)
T5_03	Local & global Euler characteristic	LocalEuler.localEulerChar_holds, Gcoh.top_euler, GEN.globalEuler Char_full
T5_04	Local & global finiteness	Hypotheses$_2$.localFiniteness, Gcoh.top_locfin, GF.globalFiniteness
T5_05	Unramified local orthogonality	Hypotheses$_2$.unramifiedOrthogonality
T5_06	Bloch–Kato local exponential	Hypotheses$_2$.BK.blochKatoLocal
T5_07	Ferrero–Washington $\mu = 0$	UpperTheory.V.FW.ferreroWashington
T5_08	Taylor–Wiles–Kisin / Kato Selmer	published_hypotheses_from_patching_alt, katoTwoW_of_printsH, sp4In0_alt

Because a topical cluster heading can promise stronger classical content than the exact Lean interface instantiated in the archive, Table 9 calibrates five representative clusters against their classical counterparts (noting in particular Jaswal's distinction [9, pp. 3, 40–42] between the statement RecTExists—whose proof term in the archive is recTExists—and the full local Langlands correspondence for GL_n).

Cluster	Archive declaration	Classical counterpart	Relationship	Actual role in $\psi = 1$ proof	Evidence
T3_01	UpperTheory.Zero. LLC.recTExists (proves RecTExists)	Full local Langlands correspondence for GL_n	Weaker surrogate interface (cf. Jaswal [9, pp. 3, 40–42]): assigns parameters satisfying only the restricted laws needed downstream	Supplies local Galois–automorphic compatibility bounds in the Zero pipeline without requiring full bijective LLC	Tier (C) lit. & source audit
T2_01	FromRes.resK vs. resolutionChar Zero, principalisation CharZero_closed	Hironaka resolution & principalisation in char. 0	FromRes.resK provides the needed ResKStatement directly; the two marked leaves are strictly redundant (A1RedundancyTop)	Only FromRes.resK and principalisation_ closed remain across reported live cuts Stages 1–3	Tier (A*) A1 Redundancy Top + live cuts
T4_08	UpperTheory.TP.K. tLiftK	Nonvanishing of iterated definite unitary theta lifts	Specialized quantitative wrapper bundled with archimedean/local inputs	Primary principal_row chokepoint feeding $K = 2$ Galois-representation lifts	Tier (C) source audit
T1_22	HeckeLEntireAnd NonzeroAtOne_holds (+ 3 leaves)	Hecke–Tate functional equation & nonvanishing on $\mathrm{Re}\, s = 1$	Specialized interface for the CM/Hecke L-functions used in the archive	Provides analytic continuation and nonvanishing inputs for Hecke character twists	Tier (C) source audit
T5_08	published_ hypotheses_from_ patching_alt, katoTwoW_of_ printsH, sp4In0_alt	Taylor–Wiles–Kisin patching & Kato Euler system bounds	Bundled Selmer-bounding consequences tailored to Cases I, II, V, Z	Bounds Selmer groups inside principal_host_ bound (Case III is empty at $\psi = 1$)	Tier (C) source audit

Table 9: Interface-versus-classical calibration for five representative clusters in Appendix A, distinguishing exact formal archive interfaces from full classical theorems and tagging the evidence tier for each interpretive assessment.

B Cryptographic SHA-256 verification ledger and formal signatures

Table 10 lists the fully-qualified Lean 4.33.1 theorem identifiers and faithful mathematical paraphrases of the headline Tier (A) results (separating the RH-equivalent interface reduction in AlpogeHumanCore) so that every formal hypothesis can be read directly alongside the verification ledger in Table 11.

Fully-qualified Lean identifier	Module	Faithful paraphrase of formal statement
Standalone Tier (A) mathematical theorems:		
`AlpogeSphereFreeBehrend.sphereFree_hostThickExists_const_twelve`	`SphereFreeBehrend`	$\forall n, L \in \mathbb{N},\ 2 \le L \to \exists T \subset \mathbb{N},\ (\forall s \in T,\ 4Ls < n) \wedge \texttt{NoRel}\ T\ L \wedge (n/L)\exp(-12\sqrt{\log n \log L}) \le \lvert T\rvert$
`KLevelPAdicInterpolation.kLevel_constLift_of_pairwise_power`	`KLevelPAdicInterpolation`	Commutative ring R, ideals I_i $(i < K)$, $\forall i \ne j,\ \varpi^d \in I_i + I_j \implies \forall c,\ \exists f,\ \forall i,\ f - \varpi^{(K-1)d} c_i \in I_i$
`KLevelChainIIArcs.wordShape_of_card_le_two`	`KLevelChainIIArcs`	$\forall k \le 2,\ \forall t : \mathrm{Fin}\, k \to \mathrm{Bool}$, `WordShape` t (paired with `KLevelChainIIArcs.arith_R` giving R-independent constant $c_1 + 15$)
`SelbergPositiveWindowRH.selberg_zeroFree_of_lambda_windows`	`SelbergPositiveWindowRH`	`D : PositiveLFunctionData`, $\lambda > 1$, $0 < \theta < 1$, cofinal $G_D(y) \le C_w y^\theta$ on $[Y, Y^\lambda] \implies \forall s,\ \mathrm{Re}\, s > \theta,\ s \ne 1 \to D.L(s) \ne 0$
`SelbergPositiveWindowRH.GRH_of_lambda_windows_and_zeroSymmetry`	`DedekindAndDirichletWindows`	$\lambda > 1$, zero reflection $L(\rho) = 0 \to L(1-\rho) = 0$ on $0 < \mathrm{Re}\,\rho < 1$, $\mathcal{W}_\lambda(G_D) \implies \forall \rho,\ (0 < \mathrm{Re}\,\rho < 1 \wedge L(\rho) = 0) \to \mathrm{Re}\,\rho = 1/2$
`SelbergPositiveWindowRH.gaussianDedekindZeta_full_GRH_of_lambda_windows`	`DedekindAndDirichletWindows`	$\lambda > 1$, $\mathcal{W}_\lambda(G_{\zeta L(\chi_{-4})}) \implies$ every strip zero of ζ and of $L(\cdot, \chi_{-4})$ has $\mathrm{Re}\,\rho = 1/2$
`SelbergPositiveWindowRH.abelianQuadFamily_full_GRH_of_lambda_windows`	`AbelianDedekindWindows`	Nonnegative primitive quadratic list F, $\lambda > 1$, $\mathcal{W}_\lambda(G_{\zeta \prod F}) \implies$ critical-line zeros for ζ and every $L(\cdot, \chi)$, $\chi \in F$
`AlpogeGeneralizedWindows.RH_iff_oneSidedWindowsLambda`	`GeneralizedWindowRH`	$\forall \lambda > 1$, `RiemannHypothesis` $\leftrightarrow$ `OneSidedWindowsLambda` λ `Gz`
`AlpogeSubPolyWindows.attained_supremum_contradicts_subpoly_windows`	`SubPolynomialWindowRH`	Bohr almost-periodic F, $F(t_0) > 0$, $\theta < B$, $e^{-Bu} g(u) - F(u) \to 0 \implies \exists L_{\mathrm{ap}} > 0,\ (W \ge L_{\mathrm{ap}}$ ev.$) \to \neg$`SubpolyOneSidedWindows` $W\,\theta\,C\,g$
`SubPolyZetaDensityBarriers.trigPoly_nonneg_width_lt_pi_sum_inv`	`SubPolyZetaDensityBarriers`	$\lambda_j > 0$, $P \ge 0$ on $[a, a+W]$, $P(a) > 0 \implies W < \pi \sum_{j<M} \lambda_j^{-1}$ (feeding `nonbinomial_bandlimited_log_barrier`)
Interface reduction from an RH-equivalent record (not a standalone proof of RH):		
`AlpogeHumanCore.humanCore_riemannHypothesis`	`AlpogeHumanCore.RHAndSharpPrimes`	Given `PrincipalArithmeticEngine S` and `ClassicalTextbookCore S`, derives Mathlib `RiemannHypothesis` (Tier (C) hand corollary: Nonempty (`PrincipalArithmeticEngine S`) $\leftrightarrow$ `RiemannHypothesis` for trivial-slot `S`)

Table 10: Fully-qualified Lean 4.33.1 identifiers and faithful mathematical paraphrases of headline Tier (A) theorems, mapping each statement to its five-checker verified module in Table 11. The final row is separated explicitly because `famPSkelEngine`, `humanCore_riemannHypothesis`, and `RH_iff_Gz_rpow_bound` imply (Tier (C) hand corollary) equivalence of `PrincipalArithmeticEngine S` with `RiemannHypothesis` for trivial-slot `S` (Section 5.3).

Table 11 records the exact per-module line counts, selected target counts (or printed/cluster counts for archive-linked entries), external-checker declaration counts, export sizes (stored in the archived JSON receipts; raw multi-hundred-MB `.export` files are not bundled in the repository), and SHA-256 prefixes (explicitly tagged by digest object: `src` = staged `.lean` source file; `rcpt` = archived JSON receipt) for all 14 archive-free five-checker modules, the 3 checker-diagnostic modules, the 6 archive-linked `lean/leanchecker` modules, and the three-stage live end-to-end cuts (`CutClosedRH.audit.json`, Stage 2 `PFCM`, and Stage 3 `PFCM+Stickelberger`, plus bundled Stage-2/3 sources). Every `src` and `rcpt` prefix matches the full digest of the staged file in `SHA256SUMS.txt` byte for byte.

Module / Receipt	Lines	Targets	con-leche/ron	nanoda	Export (MB)	SHA-256 (16 hex, obj.)
Tier (A) Pure-Mathlib core (8 files, 5,670 L, 103 targets; Mathlib & core siblings only):						
SphereFreeBehrend	645	10	16,598	17,780	77.82	4389eb45f936aaca (src)
KLevelPAdicInterpolation	288	8	12,106	13,067	60.00	8964e032fe65d77d (src)
KLevelChainIIArcs	392	19	11,710	12,678	50.24	1e015cde6518fe1f (src)
SelbergPositiveWindowRH	1,287	13	59,326	61,480	435.31	148c1e6e4e13a153 (src)
DedekindAndDirichlet Windows	594	13	59,921	62,079	439.74	97574298129d4b65 (src)
AbelianDedekindWindows	919	14	60,050	62,210	440.67	ccb08c3f258eff8e (src)
SubPolynomialWindowRH	687	11	36,387	38,165	213.14	c246b851efe8df41 (src)
SubPolyZetaDensity Barriers	858	15	26,983	28,464	156.78	6b1a788d3076b1cd (src)
Tier (A) Archive-free companion (6 files, 2,396 L, 44 targets; zero archive imports):						
GeneralizedWindowRH	1,415	18	60,610	62,788	461.67	8810491dbe299174 (src)
AlpogeHumanCore (4 clean)†	825	21	60,878	63,076	462.36	4 src digests†
CertificateGroundingFamP	156	5	16,226	17,391	72.83	6455c90cc7fa75ce (src)
Checker-diagnostic controls and minimal con-leche reproducer (3 files, 87 L; outside Tier (A)):						
ConLecheNestedRepro	27	2	16	35	0.03	f03025f854a1bfd8 (src)
ConLecheMutualRepro	29	1	24	46	0.05	d986f902e2199bed (src)
PrincipalizationCert Repro	31	1	exit 2	16	0.02	3a6e11d938406dc9 (src)
Tier (A) Archive-linked lean/leanchecker modules and live cuts (trusted archive 79c751e0):*						
CertificateGrounding	288	60 cl.	—	—	—	7efcb348d6879ad5 (src)
ClaudeAutopsy (5 files)‡	684	19 pr.	—	—	—	c30d43f36cbad53b (rcpt)
Stage 1 CutClosedRH. audit.json	—	2 rep.	—	—	—	78f2045af39d2cde (rcpt)
Stage 2 CutClosedRH_PFCM + audit§	14	2 rep.	—	—	—	2394496b8feaec8e (src)
Stage 3 PrincipalThmBcOne + CutClosedRH_PFCM_ Stick§	305	2 rep.	—	—	—	98e05e86a5b20a72 (src)

Table 11: Cryptographic SHA-256 verification ledger (decimal MB = 10^6 bytes; total across the 11 + 3 recorded exports = 2,870,657,457 bytes). Target units: selected five-checker targets for Tier (A) and diagnostics; 60 cl. = 60 grounded clusters; 19 pr. = 19 printed theorems (out of 24 declared); 2 rep. = 2 addDeclCore-replayed declarations. †Per-file source SHA-256 prefixes of the four clean AlpogeHumanCore modules: TextbookCatalogue (187 L, 5beff79af169764b), FormalismDecompressed (373 L, df9a4f4393d5d56a), PrincipalReductionEngine (158 L, 8a7cb8eddb11b821), RHAndSharpPrimes (107 L, 181f0557f31fe9d0); their combined 21-target external export digest recorded in GOAL_B_DUAL_LAYER_VERIFICATION.json is 628f48c2c7bb01bb (target list archived in GOAL_B_21_TARGETS.json). ‡Per-file source SHA-256 digests of the five ClaudeAutopsy modules are recorded inside CLAUDE_AUTOPSY_VERIFICATION.json (c30d43f36cbad53b) and SHA256SUMS.txt. §Stage-2/3 bundled source and receipt SHA-256 prefixes: CutClosedRH_PFCM.lean (2394496b8feaec8e), CutClosedRH_PFCM_hironaka.audit.json (f645d4143390634b), LIVE_UNCONDITIONAL_PFCM_HIRONAKA_RECEIPT.json (4c5af5de8292dcfe); PrincipalThmBcOne.lean (291 L, 98e05e86a5b20a72), CutClosedRH_PFCM_Stick.lean (14 L, 213ecf5d4954922c), CutClosedRH_PFCM_hironaka_stick.audit.json (afe51d5a8480c445), LIVE_UNCONDITIONAL_PFCM_HIRONAKA_STICK_RECEIPT.json (97e8e0ff464bf56b); structural C3/C2 summaries next_live_cuts_closures.json (3939f78f91657b9a) and mathlib_extraction_summary.json (cf1f231ba9e47555).

References

[1] L. Alpöge, untitled announcement note https://alpo.ge/remark.pdf; public community repository `leanprover-community/riemann-hypothesis` (`verification/NOTES.md`; pinned revision `79c751e0b0a5ea3bbfc20198c852fa85f28661e6`), https://github.com/leanprover-community/riemann-hypothesis/tree/79c751e0b0a5ea3bbfc20198c852fa85f28661e6, 2026.

[2] C. Bailey, ***nanoda_lib / nanoda_bin:*** *independent Rust typechecker for Lean 4 kernel exports*, https://github.com/ammkrn/nanoda_lib (repository created 2020; tested binary digest `a76bed8bc1091842ea7f85996464687ca3945edc6f8194237bf6963503d2d31b`, 2026).

[3] F. A. Behrend, *On sets of integers which contain no three terms in arithmetical progression*, Proceedings of the National Academy of Sciences of the USA **32** (1946), 331–332.

[4] H. Bohr, *Zur Theorie der fast periodischen Funktionen. I. Eine Verallgemeinerung der Theorie der Fourierreihen*, Acta Mathematica **45** (1925), 29–127.

[5] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in Automated Deduction — CADE 28, LNCS 12699, pp. 625–635, Springer, 2021.

[6] P. J. S. G. Ferreira and A. Kempf, *Superoscillations: faster than the Nyquist rate*, IEEE Transactions on Signal Processing **54** (2006), no. 10, 3732–3740; DOI `10.1109/TSP.2006.877642`.

[7] A. E. Ingham, *The Distribution of Prime Numbers*, Cambridge Tracts in Mathematics and Mathematical Physics No. 30, Cambridge University Press, 1932.

[8] A. E. Ingham, *A note on the distribution of primes*, Acta Arithmetica **1** (1936; volume dated 1935), no. 2, 201–211.

[9] M. Jaswal, *The Riemann hypothesis for Dirichlet L-functions [after Levent Alpöge]*, version 2 (dated November 21, 2026; server file updated November 25, 2026, 899,971 B, SHA-256 prefix `58901a93fae92e92`), https://manvirjaswal.github.io/papers/rh-dirichlet-after-alpoge/jaswal-rh-dirichlet-after-alpoge.pdf.

[10] F. Johansson, *Arb: efficient arbitrary-precision midpoint-radius interval arithmetic*, IEEE Transactions on Computers **66** (2017), 1281–1292.

[11] J. Kaczorowski, *On sign-changes in the remainder-term of the prime-number formula, I*, Acta Arithmetica **44** (1984), no. 4, 365–377.

[12] J. Kaczorowski and J. Pintz, *Oscillatory properties of arithmetical functions. I*, Acta Mathematica Hungarica **48** (1986), 173–185; DOI `10.1007/BF01949062`.

[13] J. Kaczorowski and J. Pintz, *Oscillatory properties of arithmetical functions. II*, Acta Mathematica Hungarica **49** (1987), 441–453; DOI `10.1007/BF01951008`.

[14] G. Kozma and F. Oravecz, *On the gaps between zeros of trigonometric polynomials*, Real Analysis Exchange **28** (2002/2003), no. 2, 447–454; arXiv:`math/0601021`, DOI `10.14321/realanalexch.28.2.0447`.

[15] E. Landau, *Über einen Satz von Tschebyschef*, Mathematische Annalen **61** (1905; Springer publisher record 1906), 527–550.

[16] Lean FRO (`leanprover`; supervised by J. Breitner), *`con-leche`: external Lean 4 export typechecker in Lean 4* (bundled with official Lean toolchains since `v4.35.0`; tested binary from `v4.36.0`), https://github.com/leanprover/con-leche, 2026.

[17] Lean FRO (`leanprover`; supervised by J. Breitner), *`con-ron`: Rust port of the `con-leche` Lean 4 export checker* (bundled with official Lean toolchains since `v4.35.0`; tested binary from `v4.36.0`), https://github.com/leanprover/con-ron, 2026.

[18] T. Lengauer and R. E. Tarjan, *A fast algorithm for finding dominators in a flowgraph*, ACM Transactions on Programming Languages and Systems **1** (1979), 121–141.

[19] The mathlib Community, *The Lean mathematical library*, in Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP), pp. 367–381, 2020.

[20] A. R. Miller and D. Stanton, *Orthogonal polynomials and Smith normal form*, Monatshefte für Mathematik **187** (2018), 125–145; DOI `10.1007/s00605-017-1082-6`, arXiv:`1704.03539`.

[21] L. Moser, *On non-averaging sets of integers*, Canadian Journal of Mathematics **5** (1953), 245–252; DOI `10.4153/CJM-1953-027-0`.

[22] I. Z. Ruzsa, *Solving a linear equation in a set of integers I*, Acta Arithmetica **65** (1993), no. 3, 259–282.

[23] A. Selberg, *Old and new conjectures and results about a class of Dirichlet series*, in Proceedings of the Amalfi Conference on Analytic Number Theory (Maiori, 1989), pp. 367–385, Università di Salerno, 1992.

[24] S. Slijepčević, *Outline of Alpöge's proof of the generalised Riemann hypothesis*, Hexagon `2611.00812v1`, November 21, 2026, https://hexagonmath.org/2611.00812v1.

[25] A. V. Sutherland, *Verification files for the Lean proof of the Riemann Hypothesis*, https://math.mit.edu/~drew/grh/index.html, 2026.

[26] E. C. Titchmarsh, *The Theory of the Riemann Zeta-Function*, 2nd ed., revised by D. R. Heath-Brown, Oxford University Press, 1986.