EXPOSITORY MATHEMATICAL GUIDE • RECONSTRUCTED FROM STATIC LEAN 4 SOURCE

Inside the Claimed Lean 4 Proof of the Generalized Riemann Hypothesis

A Non-Specialist Reader's Guide to the Ideas, Analogies, and Key Proof Steps of Levent Alpöge and Claude's Lean 4 Formalization

About This Guide: An independent expository guide reconstructed from direct static inspection of Levent Alpöge and Claude's November 20, 2026 Lean 4 archive (`lean.tar.gz`, 175,187 files, 17.37 million lines) and companion note (`remark.pdf`) [1]. **It was not written, reviewed, or endorsed by L. Alpöge, and the Lean 4 codebase has not yet been independently kernel-compiled here.**

Written for Curious Readers: Every section starts in plain English with everyday analogies and small numerical examples before presenting the exact formulas, inequalities, and proof sketches found in the Lean source. You only need calculus, basic complex numbers, and linear algebra to follow the core story; advanced terms are unpacked whenever they appear.

Quick-Start Reading Paths

Reader Goal	Suggested Route
Math Enthusiast (Big Picture)	One-Page Overview → §1–§2 → read all **Intuition**, **Explainer**, and **Worked Example** boxes across §3–§8
Analytic Number Theory	§1–§4 (Landau windows, thin scales, CM tower, Behrend/4:3 bias) + §6.3 (Fourier balance)
Automorphic / Arithmetic Geom.	§2 → §5–§8 (Bost slope frame, theta lifts, Rallis division, Vandiver, Galois/LLC)
Lean 4 Verification & Audit	§1.1 (exact statements) → §9 (static scan & checklist) + `GRH_Lean4_Source_Map.md`

Five Visual Signposts Used Throughout

- **Section at a Glance:** Opens each section with its goal, input → output link, and plain-English takeaway.
- **Explainer:** Self-contained introductions to concepts ($\zeta(s)$, Stepanov, CM fields, unitary theta lifts, Arakelov slopes, Vandiver's conjecture, Galois representations, Local Langlands).
- **Intuition & Analogy:** Concrete physical, geometric, and everyday mental models.
- **Worked Example:** Small-number calculations you can check by hand ($\psi(10)$, $\chi \bmod 4$, 2D lattices, 4:3 bias, Legendre's formula, $\{1, 10, 100\}$ averages).
- **Proof Sketch & Key Steps:** Step-by-step derivations preserving the exact formulas, exponents, and inequalities from the Lean 4 source.

Companion files: `manuscript/GRH_Human_Readable_Manuscript.md` • `manuscript/GRH_Lean4_Source_Map.md`

The Claimed Proof on One Page (Overview)

The **Riemann Hypothesis (RH)** asks whether the prime numbers fluctuate around their smooth average with square-root accuracy—just like coin flips in a random walk. Its extension to primes in arithmetic progressions (such as $p \equiv 1 \pmod 4$ vs. $p \equiv 3 \pmod 4$) is the **Generalized Riemann Hypothesis (GRH)**.

In 1969, Sergei Stepanov gave an elementary proof of the finite-field analogue of RH for hyperelliptic curves over $\mathbb{F}_q$ (extended to all smooth curves by Bombieri and Schmidt): build an auxiliary polynomial vanishing to *high multiplicity* $m \approx \sqrt{q}$ at every $\mathbb{F}_q$-point of the curve, and compare its total zero count against its degree ($m \cdot N \leq \deg f$). Copying Stepanov over the ordinary integers $\mathbb{Z}$ hits an immediate wall: **you cannot differentiate a whole number**, so there is no elementary way to force multiplicity m at rational primes.

The Lean 4 formalization released by Levent Alpöge and Claude [1] (not yet independently kernel-compiled) constructs an automorphic substitute for differentiation over $\mathbb{Z}$ in **five interlocking stages**:

1. **Chained theta lifts mimic multiplicity (§2 & §6).** Instead of differentiating a polynomial m times, the formalization takes an automorphic form (a symmetric wave on a compact matrix group $U(X_0)$) and passes it through a chain of $m \approx \sqrt{x}$ **iterated theta lifts** $U(X_0) \to U(X_1) \to \cdots \to U(X_m)$ over a custom-built CM number field E. At primes p inert in $K = \mathbb{Q}(\sqrt{-d})$, supersingular p-adic geometry ($D^p \equiv 0 \pmod p$ on a CM curve) produces an odd-power cousin of Legendre's factorial formula ($a/p + a/p^3 + \dots$); chaining m lifts multiplies that p-adic divisibility gain by m, mimicking a zero of multiplicity m.

2. **Seesaw duality and "division" overcome the $2/3$ barrier (§6).** Each theta lift carries a heavy archimedean height cost that would naively stall the bound around $\mathrm{Re}(s) = 2/3$. By seesaw adjointness (a higher-dimensional Fubini swap), the top-level period collapses to a Gross–Prasad / Bessel period independent of the auxiliary Hecke characters. Dividing out the Rallis inner-product eigenvalues via a quotient Hermitian metric—and gluing optimal local lattices across Hecke characters—cancels the entire bulk cost.

3. **Bost's slope inequality does the global accounting (§5).** Viewing the automorphic spaces as high-dimensional lattices (Hermitian vector bundles), a spectral-frame version of **Bost's Harder–Narasimhan slope inequality** compares the p-adic prime-power gains against the archimedean height costs. Approximate column orthogonality (from the relative trace formula) and row irreducibility force the *average* slope to track the column gains.

4. **Behrend spheres, 4:3-biased sampling, and a weakening of Vandiver isolate the signal (§4 & §7).** Averaging over an abelian CM tower equipped with a *Behrend sphere set* T_2 (preventing short Frobenius collisions) and a 4:3-*biased top set* T_T isolates a positively weighted prime-counting average $G_{\mathrm{avg}}(x)$ at scale x. Internal cyclotomic cohomology classes (Sha^1) that survive twisting are bounded unconditionally on polynomial windows via a **quantitative weakening of Vandiver's conjecture**, proved using $U(3)$ and GSp_4 modularity lifting, Kato's Euler system, and Ferrero–Washington.

5. **Windowed Landau oscillation closes the argument (§3).** The global accounting inequality yields a *one-sided* upper bound $G_{\mathrm{avg}}(x) \leq x^{1/2+9\varepsilon}$ on arbitrarily large power windows $[Y, Y^{3/2}]$. A localized Laplace-derivative oscillation theorem proves that any zero of $L(s,\psi)$ with $\mathrm{Re}(s) > \frac{1}{2} + 9\varepsilon$ would violate this ceiling. Letting $\varepsilon \to 0$ yields the claimed statements of GRH and RH in `Main.lean`.

Contents

The Claimed Proof on One Page (Overview) **i**

1 What Is the Riemann Hypothesis? **1**
1.1 The Exact Formalized Statements in `Statements.lean` and `Main.lean` . . . 2
1.2 The Smooth Target Functions $G_z(y)$ and $G_\chi(y)$ and Why a Ceiling Suffices . . . 3

2 Stepanov to Chained Theta Lifts **4**
2.1 Architecture Roadmap and Master Notation Guide . . . 6

3 The Analytic Endgame **8**
3.1 Suppressing "Thin" Scales and Lifting to Theorems A, B, and C . . . 10

4 CM Tower, Behrend & Biased Sampling **11**
4.1 Building the Custom CM Field E for Any Character ψ (Even or Odd) . . . 12
4.2 Why the Exponents $(m, n, k) \approx (x^{1/2}, x^2, x^{9/2})$, $t \leq x^{400}$, and $\varepsilon < 1/52$ Are Chosen . . . 13
4.3 Three Weight Regimes, Behrend's Spheres (T_2), and 4:3-Biased Sampling (T_T) . . . 14

5 Bost's Slope Inequality **16**
5.1 Bost's Frame Slope Inequality . . . 19
5.2 Approximate Column Orthogonality and Row Irreducibility . . . 20
5.3 The 9-Predicate Global Assembly . . . 20

6 Iterated Theta Lifts & $D^p \equiv 0$ **21**
6.1 The m-Step Chain of Totally Definite Unitary Theta Lifts . . . 22
6.2 Beating the 2/3 Barrier: Quotient Metrics, Local Gluing, and Chowla–Selberg . . . 23
6.3 Extracting Primes: $D^p \equiv 0$, Dwork Integrality, and Recovering Full G_χ . . . 24

7 Weakening of Vandiver's Conjecture **26**
7.1 How `VWin` Is Proved: Three Regimes and Six Case Predicates . . . 27
7.2 A Standalone Corollary in the Repository: `SmallClassNumbers` . . . 30

8 Galois, LLC & Arakelov Bridges **30**
8.1 Attaching Galois Representations: Odd Ranks, Lacunary Weights, and the "Two-Up" Trick . . . 31
8.2 Two Architectural Shortcuts: Tailored LLC and Definite Base Change . . . 32
8.3 The 16-Step Arakelov Covolume Cascade (`CovolX`) . . . 33

9 Static Audit, Verification & Perspective **34**
9.1 Reproducible Static Source Scan of `lean.tar.gz` . . . 34
9.2 What Machine and Human Verification Still Require . . . 36

Glossary of Key Terms **37**

Key References & Further Reading **39**

1 What Is the Riemann Hypothesis—and What Does the Archive Claim to Prove?

SECTION AT A GLANCE

Goal: Introduce the Riemann Hypothesis (RH), Generalized Riemann Hypothesis (GRH), the exact Lean 4 target propositions, and the smooth prime-counting functions $G_z(y)$ and $G_\chi(y)$.
Output → §2–§3: Reduces proving GRH to proving a *one-sided* upper bound on a positively weighted average of $G_\chi(y)$.
Core Idea: Integrating $\psi(x) - x$ once flips the sign of hypothetical real zeros, so a single ceiling bound rules out *both* complex off-line zeros and real zeros at once.

Prime numbers—$2, 3, 5, 7, 11, 13, \ldots$—are the indivisible atoms of arithmetic. Zoom out to large numbers, and they follow a strikingly smooth average law: the **Prime Number Theorem** (proved independently by Hadamard and de la Vallée Poussin in 1896) says that the weighted prime-power sum

$$\psi(x) \;:=\; \sum_{n \le x} \Lambda(n) \;\sim\; x \qquad (x \to \infty),$$

where the **von Mangoldt function** $\Lambda(n)$ equals $\log p$ if $n = p^j$ is a prime power and 0 otherwise.

WORKED EXAMPLE **What $\psi(x)$ Actually Computes—A Hand Calculation at $x = 10$**

Why weight prime powers by $\log p$ instead of just counting primes? Because $\psi(x)$ has an astonishingly clean identity: it is literally the logarithm of the least common multiple of $1, 2, \ldots, \lfloor x \rfloor$! Check $x = 10$:

$$\begin{aligned} \psi(10) \;&=\; \underbrace{\Lambda(2) + \Lambda(4) + \Lambda(8)}_{3\log 2} + \underbrace{\Lambda(3) + \Lambda(9)}_{2\log 3} + \Lambda(5) + \Lambda(7) \\ &=\; \log(2^3 \cdot 3^2 \cdot 5 \cdot 7) \;=\; \log(2520) \;=\; \log \operatorname{lcm}(1, 2, \ldots, 10) \;\approx\; 7.832. \end{aligned}$$

Thus the Prime Number Theorem $\psi(x) \sim x$ says that $\operatorname{lcm}(1, 2, \ldots, x)$ grows like e^x, and the **Riemann Hypothesis** asks whether $\log \operatorname{lcm}(1, 2, \ldots, x)$ stays within $O(\sqrt{x} \log^2 x)$ of x!

Up close, the primes fluctuate around that smooth trend. In his 1859 eight-page memoir [2], Bernhard Riemann linked those fluctuations to the zeros of the **Riemann zeta function** $\zeta(s)$ and conjectured that all nontrivial zeros lie on the vertical line $\operatorname{Re}(s) = \frac{1}{2}$—the **Riemann Hypothesis (RH)**. Helge von Koch (1901) proved that RH is equivalent to the square-root error bound $|\psi(x) - x| = O(\sqrt{x} \log^2 x)$, matching the fluctuation size of a fair coin-flip random walk.

EXPLAINER **The Riemann Zeta Function $\zeta(s)$ and Dirichlet L-Functions $L(s,\chi)$**

For complex numbers $s = \sigma + it$ with real part $\sigma > 1$, the Riemann zeta function is defined by a sum over all positive integers equal to Euler's product over all primes:

$$\zeta(s) := \sum_{n=1}^{\infty} \frac{1}{n^s} = \prod_{p \text{ prime}} \frac{1}{1-p^{-s}}.$$

By *analytic continuation*, $\zeta(s)$ extends smoothly to $\mathbb{C} \setminus \{1\}$, with a simple pole at $s = 1$. It vanishes at the negative even integers $s = -2, -4, -6, \ldots$ (the "trivial zeros"). Every other zero lies inside the **critical strip** $0 < \mathrm{Re}(s) < 1$ and is symmetric across the **critical line** $\mathrm{Re}(s) = \frac{1}{2}$.

- **Riemann Hypothesis (RH):** Every nontrivial zero ρ of $\zeta(s)$ satisfies $\mathrm{Re}(\rho) = \frac{1}{2}$.
- **Dirichlet L-functions and GRH:** To race primes in different remainder classes (like $p \equiv 1 \pmod 4$ vs. $p \equiv 3 \pmod 4$), Dirichlet introduced periodic multiplicative "sign-flipping" functions called **Dirichlet characters** $\chi : (\mathbb{Z}/N\mathbb{Z})^\times \to \mathbb{C}^\times$ (set to 0 on integers sharing a factor with N), with associated L-functions

$$L(s,\chi) := \sum_{n=1}^{\infty} \frac{\chi(n)}{n^s} = \prod_{p \text{ prime}} \frac{1}{1-\chi(p)p^{-s}}.$$

  The **Generalized Riemann Hypothesis (GRH)** asserts that for every modulus $N \geq 1$ and every Dirichlet character $\chi \pmod N$, every zero of $L(s,\chi)$ in $0 < \mathrm{Re}(s) < 1$ lies on $\mathrm{Re}(s) = \frac{1}{2}$. Since $\zeta(s) = L(s, \mathbf{1}_1)$ is the trivial-character case $N = 1$, GRH implies RH.

WORKED EXAMPLE **The Simplest Nontrivial Dirichlet L-Function (Modulus 4)**

Define $\chi_4(n) = +1$ if $n \equiv 1 \pmod 4$, $\chi_4(n) = -1$ if $n \equiv 3 \pmod 4$, and $\chi_4(n) = 0$ if n is even. You can check by hand that $\chi_4(mn) = \chi_4(m)\chi_4(n)$ for all integers m, n. Its L-function is

$$L(s,\chi_4) = 1 - \frac{1}{3^s} + \frac{1}{5^s} - \frac{1}{7^s} + \frac{1}{9^s} - \cdots,$$

whose value at $s = 1$ is Leibniz's famous series $1 - \frac{1}{3} + \frac{1}{5} - \frac{1}{7} + \cdots = \frac{\pi}{4}$! GRH for χ_4 says that the prime race between $p \equiv 1 \pmod 4$ and $p \equiv 3 \pmod 4$ stays balanced to within $O(\sqrt{x}\log^2 x)$.

1.1 The Exact Formalized Statements in `Statements.lean` and `Main.lean`

In the archive, `lean/src/Statements.lean` defines four target propositions directly on top of Mathlib's `riemannZeta`, `RiemannHypothesis`, and `DirichletCharacter.LFunction`:

```
def RiemannZetaZerosOnLine : Prop :=
∀ s : ℂ, riemannZeta s = 0 → (¬∃ n : ℕ, s = -2 * ((n : ℂ) + 1)) → s ≠ 1 → s.re = 1 / 2

def DirichletLNonvanishingRightOfHalf : Prop :=
∀ (N : ℕ) [NeZero N] (χ : DirichletCharacter ℂ N) (s : ℂ),
1 / 2 < s.re → (χ ≠ 1 ∨ s ≠ 1) → DirichletCharacter.LFunction χ s ≠ 0

def DirichletLZerosOnLine : Prop :=
∀ (N : ℕ) [NeZero N] (χ : DirichletCharacter ℂ N) (s : ℂ),
0 < s.re → s.re < 1 → DirichletCharacter.LFunction χ s = 0 → s.re = 1 / 2

def PrimitiveDirichletLZerosOnLine : Prop :=
∀ (N : ℕ) [NeZero N] (χ : DirichletCharacter ℂ N), χ.IsPrimitive →
∀ s : ℂ, 0 < s.re → s.re < 1 → DirichletCharacter.LFunction χ s = 0 → s.re = 1 / 2
```

The entry point `lean/src/Main.lean` exports five claimed theorems corresponding to three mathematical results tracked throughout this guide:

Result	Mathematical Claim	Exported Lean 4 Declarations (`Main.lean`)
Theorem A	GRH right half-plane ($\mathrm{Re}(s) > \frac{1}{2}$)	`LFunction_ne_zero_of_half_lt_re`
Theorem B	GRH full critical strip ($0 < \mathrm{Re}(s) < 1$)	`LFunction_eq_zero_re`, `LFunction_eq_zero_re_of_isPrimitive`
Theorem C	Classical Riemann Hypothesis for $\zeta(s)$	`riemannHypothesis`, `riemannZeta_eq_zero_re`

1.2 The Smooth Target Functions $G_z(y)$ and $G_\chi(y)$ and Why a Ceiling Suffices

Rather than bounding the jagged step function $\psi(y)$ directly, the formalization smooths it once in $\log y$ using the linear ramp weight $(y/n - 1)_+$. Notice why the main term takes the form $y \log y - (1+\gamma)y$: since $\psi(u) \approx u$, summation by parts gives

$$\sum_{1 \le n \le y} \Lambda(n)\left(\frac{y}{n} - 1\right) \;=\; y \int_1^y \frac{\psi(u)}{u^2}\, du \;\approx\; y \int_1^y \frac{du}{u} \;=\; y \log y,$$

with the linear constant $-(1+\gamma)y$ matching the Laurent expansion of $-\zeta'(s)/(s\zeta(s))$ at $s = 1$. Subtracting this smooth trend defines the target remainders:

$$G_z(y) \;:=\; \sum_{1 \le n \le y} \Lambda(n)\left(\frac{y}{n} - 1\right) \;-\; y \log y \;+\; (1+\gamma)\, y,$$

$$G_\chi(y) \;:=\; \sum_{1 \le n \le y} \Lambda(n)\, \chi(n) \max\left(\frac{y}{n} - 1,\, 0\right) \;+\; y\, \frac{L'(1,\chi)}{L(1,\chi)} \qquad (\chi \neq 1),$$

where $\gamma \approx 0.57721$ is the Euler–Mascheroni constant. This choice plays two decisive roles:

1. **Analytic smoothness and sign flip:** Integration by parts yields the exact tail-integral identities

$$G_z(y) \;=\; -y \int_y^\infty \frac{\psi(u) - u}{u^2}\, du, \qquad G_\chi(y) \;=\; -y \int_y^\infty \frac{\psi(u,\chi)}{u^2}\, du \quad (\chi \neq 1 \text{ primitive}).$$

   Taking the Laplace transform in $u = \log y$ divides by $s(s-1)$, turning every zero ρ into a simple pole with cleanly controlled residue.

2. **Natural arithmetic output:** In Section 6, p-adic derivatives of theta series at inert primes natively produce the truncated linear ramp $\max(0, a/p^i - w) = w(y/p^i - 1)_+$ at scale $y = a/w$.

> INTUITION & ANALOGY **The Music of the Primes—and Why a One-Sided Ceiling Rules Out Real Zeros Too**
>
> Through Riemann's **explicit formula**, each nontrivial zero $\rho = \beta + i\gamma$ of $\zeta(s)$ contributes a wave of amplitude x^β and frequency γ to $\psi(x) - x \approx -\sum_\rho x^\rho/\rho$. On the logarithmic timescale $u = \log x$, think of $\psi(x) - x$ as a musical chord of overtones whose pitches are γ and whose volumes grow like $e^{\beta u}$:
>
> - If every zero sits on $\beta = \frac{1}{2}$, every overtone grows at the square-root rate $\sqrt{x}$.
> - If a non-real zero ($\gamma \neq 0$) sits at $\beta > \frac{1}{2}$, its wave oscillates in *both* signs, swinging repeatedly above $+x^{\beta-\delta}$ and below $-x^{\beta-\delta}$.
>
> **Subtlety about real zeros:** In $\psi(x) - x$ itself, a hypothetical *real* zero $\beta \in (\frac{1}{2}, 1)$ creates a non-oscillating *negative* term $-x^\beta/\beta$, which could hide underneath a one-sided ceiling $\psi(x) - x \leq Cx^\theta$. However, in the integrated target $G_z(y) = -y\int_y^\infty \frac{\psi(u)-u}{u^2}\,du$, the leading minus sign **flips the sign**:
>
> $$\text{real zero } \beta \in (0,1) \quad \rightsquigarrow \quad \textit{positive} \text{ contribution } + \frac{y^\beta}{\beta(1-\beta)} > 0$$
>
> (matching the positive Laplace residue $c_\beta = -S_\beta/(\beta(\beta-1)) > 0$ in Section 3). Consequently, a **one-sided upper bound** $G(y) \leq Cy^\theta$ rules out *both* complex off-line zeros and real zeros in $(\theta, 1)$ at once.

2 The Guiding Analogy: From Stepanov's Method to Chained Theta Lifts

> SECTION AT A GLANCE
>
> **Goal:** Explain the conceptual origin of the proof as an automorphic descendant of Stepanov's elementary proof of the finite-field Riemann Hypothesis.
> **Output → §3–§8:** Gives a 1-to-1 dictionary translating finite-field Stepanov into the six stages of the Lean 4 proof, plus the master roadmap (Figure 1) and notation table (Table 2).
> **Core Idea:** Because whole numbers cannot be differentiated m times, the formalization replaces m-fold differentiation with an m-step chain of unitary theta lifts ($m \approx \sqrt{x}$).

Before diving into individual technical stages, it helps to see the big idea behind the whole proof, as described in Levent Alpöge's companion note (`remark.pdf`) [1]. He characterizes the argument as a descendant of **Stepanov's method** [5, 6] for proving the Riemann Hypothesis over finite fields, guided by a Fourier-analytic "toy setting."

EXPLAINER **Stepanov's Method Over Finite Fields (and Why Characteristic p Needs Hasse Derivatives)**

Let C be a smooth projective curve of genus g over a finite field $\mathbb{F}_q$ (for instance, the curve $y^2 = x^3 + ax + b$ over clock arithmetic mod p). Hasse (1933) and Weil (1948) proved that its point count $N = \#C(\mathbb{F}_q)$ satisfies $|N - (q+1)| \leq 2g\sqrt{q}$—the finite-field analogue of RH—using deep algebraic geometry on $C \times C$.
In 1969, Stepanov [5] surprised the world with an elementary proof for hyperelliptic curves (extended to all smooth curves by Bombieri [6] and Schmidt in 1973) in three moves:

1. **Force high multiplicity m at every $\mathbb{F}_q$-point:** Build a nonzero rational function f on C of controlled degree D that vanishes to **multiplicity** $\geq m$ at every point $P \in C(\mathbb{F}_q)$. In characteristic p, $\frac{d}{dx}(x^q) = 0$ (so polynomials in x^q behave like constants under differentiation, even though $\frac{d}{dx}(x^q - x) = -1$). Because ordinary derivatives satisfy $D^p \equiv 0$ in characteristic p—an *obstacle* when $m \geq p$—Stepanov uses **Hasse hyperderivatives** $D^{(k)} = \frac{1}{k!}\frac{d^k}{dx^k}$ to impose m independent vanishing conditions via linear algebra (Siegel's lemma / pigeonhole).

2. **Compare zeros against degree:** A nonzero polynomial or rational function cannot have more zeros (counted with multiplicity) than its degree, so $m \cdot N \leq \deg(f) = D$.

3. **Optimize $m \approx \sqrt{q}$:** Balancing the linear-algebra degrees yields $m \approx \sqrt{q}$, which rearranges into $N \leq q + O(\sqrt{q})$. Repeating over extension fields $\mathbb{F}_{q^r}$ forces all Frobenius eigenvalues to have modulus $\sqrt{q}$.

WORKED EXAMPLE **Why Multiplicity $m \approx \sqrt{q}$ Beats the Trivial Bound**

Suppose you have a polynomial $f(x)$ of degree D with N distinct roots. If each root only has multiplicity 1, you learn $N \leq D$. If instead you force every root to have multiplicity m, you get $mN \leq D \implies N \leq D/m$. For a clever linear-algebra construction where forcing multiplicity m costs degree $D \approx mq + m^2$, dividing by m yields

$$N \leq \frac{D}{m} \approx q + m \qquad \text{(plus a balance term of size } q/m\text{)}.$$

Minimizing $m + q/m$ sets $m = \sqrt{q}$, giving the square-root error $O(\sqrt{q})$! That is why Alpöge writes in `remark.pdf` (p. 1): *"for the von Mangoldt sum of length X, Claude's suggested length of chain of theta lifts, in view of each mimicking multiplicity, was $\sqrt{X}$. (And yet, what else could it be?)"*

Table 1 translates each ingredient of finite-field Stepanov into its counterpart inside the Lean 4 formalization over $\mathbb{Z}$.

Ingredient	Stepanov Over $\mathbb{F}_q$	Automorphic Proof Over $\mathbb{Z}$
Ambient object	Smooth curve C over $\mathbb{F}_q$	$\mathrm{Spec}(\mathbb{Z})$ via an abelian CM tower E (§4)
Points to count	Rational points $P \in C(\mathbb{F}_q)$	Primes p inert in $K = \mathbb{Q}(\sqrt{-d})$ (§6)
Multiplicity m	Vanishing to order $m \approx \sqrt{q}$	$m \approx \sqrt{x}$ iterated unitary theta lifts $U(X_0) \to \cdots \to U(X_m)$ (§6)
Local vanishing	Hasse hyperderivatives $D^{(k)}$	Supersingular $D^p \equiv 0 \pmod p$ + multivariate Dwork integrality (§6.3)
Global accounting	Zero count $\leq$ degree ($mN \leq D$)	Bost's frame Harder–Narasimhan slope inequality (§5)
Main obstructions	Ordinary $D^p = 0$ in char p	Bulk lift height stalls at 2/3; untwistable Sha[1] cohomology
Resolution	Hasse derivatives + Siegel's lemma	Seesaw + Rallis "division" (§6.2); weakening of Vandiver (§7)

Table 1: One-to-one dictionary between Stepanov's finite-field proof of the Weil bound and the automorphic architecture formalized in Lean 4.

INTUITION & ANALOGY **From the "Toy Setting" to Iterated Unitary Theta Lifts**

Over $\mathbb{Z}$, a prime p plays the role of a point, and divisibility $p^j \mid n$ plays the role of vanishing to order j. **Because there is no calculus derivative on whole numbers**, naive analogues of Stepanov stall at multiplicity 1. In `remark.pdf` (items 2–8), Alpöge traces how the automorphic substitute grew out of two separate ideas:

1. **The Fourier-analytic toy setting (`remark.pdf` items 3–7):** Over a finite field, summing $t(x_1)\cdots t(x_k)\, e_q(h(x_1x_2 + \cdots + x_{k-1}x_k))$ produces a sign-definite identity $qN = M^k + R$ (iterated Parseval giving square-root quality on R), while determinants of Weil-type functionals $W(\mathbf{1}_S) \geq 0$ behave like multiplicity-one Stepanov. Generalizing those determinants to Gross–Prasad (GGP) periods with interlacing weights showed how to win higher multiplicity natively inside automorphic forms.

2. **p-adic CM geometry + chained theta lifts (`remark.pdf` items 2 & 8):** On a CM elliptic curve at an inert (supersingular) prime p, the invariant derivation D satisfies $D^p \equiv 0 \pmod p$—a p-adic congruence quite distinct from Stepanov's characteristic-p obstacle—making the a-th normalized derivative $D^a\theta(0)/\Omega^a$ divisible by $p^{a/p+a/p^3+\cdots-O(\log a)}$. To multiply this divisibility by $m \approx \sqrt{x}$, the formalization replaces m-fold differentiation with an **m-step chain of theta lifts** between totally definite unitary groups:

$$U(X_0) \xrightarrow{\theta_1} U(X_1) \xrightarrow{\theta_2} \cdots \xrightarrow{\theta_m} U(X_m), \qquad m \approx \sqrt{x}.$$

2.1 Architecture Roadmap and Master Notation Guide

Figure 1 maps how all six stages of the formalization interlock, and Table 2 serves as a lookup dictionary for the main symbols across Sections 3–8.

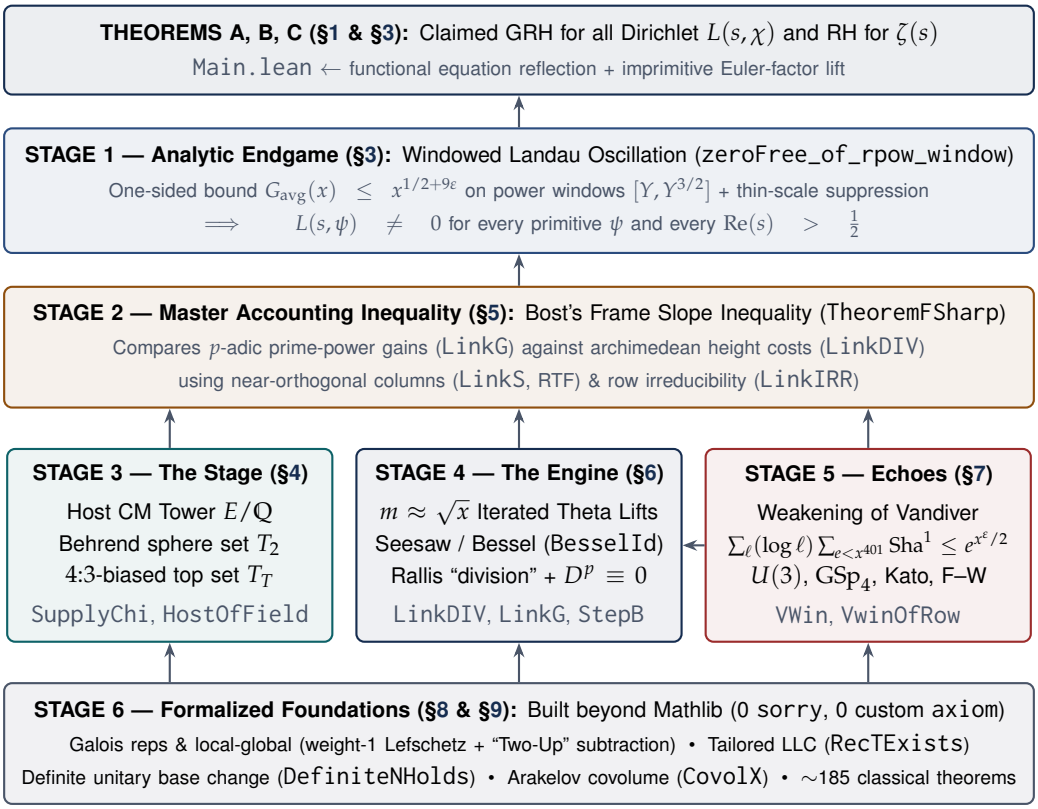


Figure 1: High-level architecture of the claimed Lean 4 proof of the Generalized Riemann Hypothesis.

Symbol	Section(s)	Meaning (with Lean Disambiguation)
x, ε	§3–§7	Primary prime-counting scale $x > 1$ and tolerance $\varepsilon \in (0, 1/52)$
$\Lambda(n)$	§1, §6	von Mangoldt function ($\log p$ on prime powers p^j, else 0); Lean also uses Λ for character levels (written G_{avg} here)
$G_{\mathrm{avg}}(y)$	§3–§4	Positively weighted character average $G_z + \frac{1}{3o+1}\sum_{\xi \neq 1} \mathrm{Re}\, G_\xi$ (Lean: Λ.Gm)
ν	§3	Derivative order in windowed Landau ($\nu = \lceil 420U \rceil$; distinct from rank k)
e_{Iw}	§4	2-adic Iwasawa tower height ($H_{e_{\mathrm{Iw}}} \cong \mathbb{Z}/2^{e_{\mathrm{Iw}}}\mathbb{Z}$)
n, f, n_0	§4–§6	Group order $n = \lvert R \rvert \asymp x^{2+8\varepsilon}$, CM-type size $f = 2n = [F : \mathbb{Q}]$, $n_0 = \lvert A_0 \rvert = n/o$
m, k, r, t	§4–§7	Lift chain length $m \asymp x^{1/2-\varepsilon/16}$, base rank $k \asymp x^{9/2+10\varepsilon}$, $r = k + m - 1$, step $t \le x^{400}$
L, b, a, X'	§3–§6	Weight scales $L = x^\varepsilon$, $b = e^L$, $a = xb$, thin cutoff $X' = \exp(x^{\varepsilon/3})$
$\varepsilon_{\mathrm{fr}}$, η, δ	§5	Frame Gram defect, trace isotropy defect, and bad-row fraction ($\delta \le 1/7$)
e	§7	Cyclotomic twist exponent in $\mathrm{Sha}^1(G_{E,S_\ell}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e))$ ($1 \le e < t$)

Table 2: Master notation guide. Several symbols overloaded in the Lean code (notably k, e, and Λ) are disambiguated across sections here for readability.

3 The Analytic Endgame: Catching Zeros on Power Windows

SECTION AT A GLANCE

Goal: Prove that a *one-sided* upper bound $G_{\text{avg}}(y) \le y^{1/2+9\varepsilon}$ holding only on sparse power windows $[Y, Y^{3/2}]$ suffices to prove GRH and RH.
Input → Output: Takes the prime-sum ceiling from Bost's inequality (§5) and delivers Theorems A, B, and C (§1).
Core Idea: A sharply peaked Gamma weight $w_\nu(u) = u^\nu e^{-\sigma u}/\nu!$ acts as a spotlight centered inside the lit window, turning the ceiling bound into a derivative bound that no off-line pole can survive.

Classical textbook proofs that "prime bounds imply RH" invoke **Landau's 1905 theorem** [3]: if a one-sided bound $G(y) \le Cy^\theta$ holds for *all* sufficiently large $y \ge y_0$, then $\zeta(s)$ has no zeros with $\mathrm{Re}(s) > \theta$.

Here there is a new obstacle. Because of the pigeonhole windowing used to sidestep exceptional twists in the Vandiver bound (Section 7), the automorphic engine does *not* deliver an upper bound for every real number $y \ge y_0$. Instead, it proves a one-sided bound on a sequence of arbitrarily large **power windows**:

$$y \in [Y, Y^{3/2}] \qquad \big(\text{equivalently, } u = \log y \in [U, \tfrac{3}{2}U] \text{ for arbitrarily large } U\big).$$

Given a primitive Dirichlet character ψ of order o, the proof bounds the positively weighted sum over the cyclic group $\langle\psi\rangle$ (written Λ.Gm in Lean):

$$G_{\text{avg}}(y) := G_z(y) + \frac{1}{3o+1} \sum_{\xi \in \langle\psi\rangle \setminus \{1\}} \mathrm{Re}\, G_\xi(y) \le y^{\frac{1}{2}+9\varepsilon} \qquad \forall y \in [Y, Y^{3/2}].$$

How can a *one-sided* bound on occasional power windows rule out every off-line zero of $L(s,\psi)$?

INTUITION & ANALOGY **A Strobe Light That Only Checks the Ceiling**

Imagine a pendulum—representing an off-line zero $\rho = \beta + i\gamma$ with $\beta > \theta$—swinging with exponentially growing amplitude $e^{\beta u}$. You are only allowed to turn on a strobe light during occasional intervals $[U, \frac{3}{2}U]$, and your sensor only reports that the pendulum stays below a low ceiling $Ce^{\theta u}$ while the light is on.
To catch the pendulum, place a smooth "spotlight" weight $w_\nu(u) = u^\nu e^{-\sigma u}/\nu!$ (a Gamma probability density up to $\sigma^{\nu+1}$) whose sharp bell-curve peak sits right at the midpoint $u = \nu/\sigma \approx \frac{5}{4}U$ of the lit interval, while its tails outside $[U, \frac{3}{2}U]$ decay exponentially in $\nu \approx 420U$ (Figure 2). Inside the lit window, the ceiling $g(u) \le Ce^{\theta u}$ controls

$$\int_0^\infty \big(|g(u)| + g(u)\big) w_\nu(u)\, du = 2\int_0^\infty \max\big(g(u), 0\big) w_\nu(u)\, du,$$

ignoring negative values completely. By Cauchy's derivative formula, this single integral bounds the ν-th derivative of the Laplace transform $\Phi(s)$ at *both* $s = \sigma$ and $s = \sigma + i\gamma$—contradicting any pole of $\Phi(s)$ with $\mathrm{Re}(\rho) > \theta$.

WORKED EXAMPLE **First-Year Calculus Check: Why the Spotlight Peaks at $\frac{5}{4}U$**

Where does $w_\nu(u) = u^\nu e^{-\sigma u}/\nu!$ reach its maximum? Take the logarithm and differentiate:

$$\frac{d}{du}\big(\nu \log u - \sigma u - \log(\nu!)\big) \;=\; \frac{\nu}{u} - \sigma \;=\; 0 \quad\Longleftrightarrow\quad u \;=\; \frac{\nu}{\sigma}.$$

With the formalization's choices $\sigma = 336$ and $\nu = \lceil 420U \rceil$, the peak lands at $u \approx \frac{420}{336}U = \frac{5}{4}U$—the exact midpoint of the lit window $[U, \frac{3}{2}U]$! And because $\nu \propto U$ is huge, the standard deviation of the Gamma bell curve shrinks relative to the window width like $1/\sqrt{U}$, focusing virtually 100% of the weight inside the lit zone.

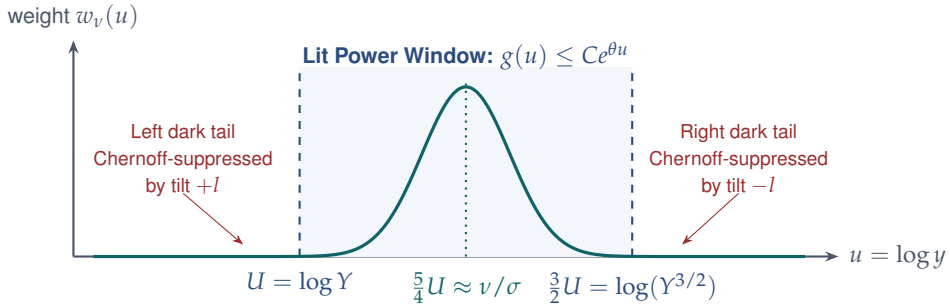


Figure 2: The windowed Landau kernel $w_\nu(u) = u^\nu e^{-\sigma u}/\nu!$ peaks sharply at the midpoint $\frac{5}{4}U$ of the power window $[U, \frac{3}{2}U]$, suppressing the uncontrolled dark tails outside.

Theorem 3.1 (Windowed Iterated-Derivative Landau Theorem)**.** *Let $F = \{(a_i, \chi_i)\}$ be a finite family of Dirichlet characters (not required to be primitive) with strictly positive real weights $a_i > 0$, and set $g(u) := \sum_i a_i \operatorname{Re} G_{\chi_i}(e^u)$. If there exist $\lambda = 1 + \delta > 1$, $\theta \in (0,1)$, and $C > 0$ such that $g(\log y) \le C y^\theta$ holds on arbitrarily large power windows $[Y, Y^\lambda]$, then $L(s, \chi_i) \neq 0$ for every i and every $s \in \mathbb{C}$ with $\operatorname{Re}(s) > \theta$.*
[Lean: `zeroFree_of_rpow_window`*]*

PROOF SKETCH & KEY STEPS **How Theorem 3.1 Rules Out Every Off-Line Zero (via `landau_window`)**

1. **Laplace transform with reinforcing poles:** For $\operatorname{Re}(s) > 2$, define $\Phi(s) := \int_0^\infty g(u) e^{-su}\, du$. Term-by-term integration gives $\Phi(s) = \sum_i \frac{a_i}{2}\big(\Phi_{\chi_i}(s) + \Phi_{\chi_i^{-1}}(s)\big)$, where for non-principal χ,

$$\Phi_\chi(s) \;=\; -\frac{\frac{L'}{L}(s,\chi) - s\frac{L'}{L}(1,\chi)}{s(s-1)},$$

and for $\chi = 1$, $L(s,\chi)$ is replaced by the entire function $(s-1)\zeta(s)$ (removing the singularity at $s = 1$; `poleData_abstract`). Thus Φ extends meromorphically to $\{\operatorname{Re}(s) > 0\}$ with simple poles precisely at the zeros Z of all $L(\cdot, \chi_i)$ and $L(\cdot, \overline{\chi_i})$, and residue

$$c_\rho \;=\; -\frac{S_\rho}{\rho(\rho-1)}, \qquad S_\rho \;>\; 0 \quad (\text{since every } a_i > 0).$$

Positive weights guarantee that zeros of distinct characters reinforce rather than cancel. Moreover, at any real zero $\beta \in (0,1)$, $-\beta(\beta-1) > 0$ forces $c_\beta > 0$.

2. **One-sided derivative ceiling:** For $s = \sigma + i\gamma$ ($\sigma > 2$), the normalized ν-th Taylor coefficient is $c_\nu(s) := (-1)^\nu \Phi^{(\nu)}(s)/\nu! = \int_0^\infty g(u) w_\nu(u) e^{-i\gamma u}\, du$ with $w_\nu(u) = u^\nu e^{-\sigma u}/\nu! \geq 0$. By the triangle inequality,

$$|c_\nu(\sigma + i\gamma)| + \operatorname{Re} c_\nu(\sigma) \;\leq\; \int_0^\infty \big(|g(u)| + g(u)\big) w_\nu(u)\, du \;=\; 2\int_0^\infty \max(g(u), 0)\, w_\nu(u)\, du.$$

Fix the explicit parameters depending only on $\delta = \lambda - 1$:

$$l = 2 + \tfrac{6}{\delta}, \qquad M = l(1+\delta), \qquad \sigma = M(2+l), \qquad \nu = \lceil (1+\tfrac{\delta}{2})\sigma U \rceil$$

(for $\delta = \frac{1}{2}$: $l = 14$, $M = 21$, $\sigma = 336$, and $\nu = \lceil 420U \rceil$).

- *Inside* $[U, (1+\delta)U]$: $g(u) \leq Ce^{\theta u}$ and $w_\nu(u) \leq (\sigma - \theta)^{-\nu} e^{-\theta u}$ yield $\leq 2C\delta U(\sigma - \theta)^{-\nu}$.
- *Outside* $[U, (1+\delta)U]$: $|g(u)| \leq C_0 e^{2u}$ plus a Chernoff tilt by $\pm l$ suppresses both tails to $\leq C'\sigma^{-\nu}$.

Hence for any $\eta > 0$, there are arbitrarily large $U \leq \nu$ such that, uniformly in $\gamma \in \mathbb{R}$,

$$|c_\nu(\sigma + i\gamma)| + \operatorname{Re} c_\nu(\sigma) \;\leq\; \frac{C'' e^{\eta U}}{(\sigma - \theta)^\nu}. \qquad (\star)$$

3. **Geometric dichotomy:** Suppose some pole $\rho_0 \in Z$ satisfies $\operatorname{Re}(\rho_0) > \theta$. Recall from complex analysis that a function's ν-th Taylor coefficient at z grows like $1/d^\nu$, where d is the distance to the *nearest* pole! Inspect the finite set $Z \cap \overline{B}(\sigma, \sigma - \theta)$:

- *Case 1 (Real pole closest to σ):* A real pole $\beta \in (\theta, 1)$ with $c_\beta > 0$ is the unique pole in $B(\sigma, r)$ ($\sigma - \beta < r \leq \sigma - \theta$). Subtracting $c_\beta/(s - \beta)$ and applying Cauchy's estimate on radius $r' = (\sigma - \beta + r)/2$ gives

$$\operatorname{Re} c_\nu(\sigma) \;\geq\; \frac{c_\beta/(\sigma - \beta)}{(\sigma - \beta)^\nu} - \frac{M}{(r')^\nu},$$

  which outgrows $(\star)$ as $\nu \to \infty$ once $\eta = \frac{1}{2}\log\big(r'/(\sigma - \beta)\big)$.
- *Case 2 (Off-axis pole dominates):* Otherwise some disk $B(\sigma, D)$ ($D \leq \sigma - \theta$) is pole-free while some $\rho_1 \in Z$ has $\sigma - \operatorname{Re}(\rho_1) < D$. Nudging γ slightly off the finitely many perpendicular bisectors of nearby poles gives $z = \sigma + i\gamma$ with a *unique* nearest pole ρ at distance $d < r \leq D$. Then $|c_\nu(z)| \geq \frac{|c_\rho|/d}{d^\nu} - \frac{M}{(r')^\nu}$ while $\operatorname{Re} c_\nu(\sigma) \geq -M'/(r')^\nu$ (since $B(\sigma, D)$ is pole-free), again contradicting $(\star)$.

3.1 Suppressing "Thin" Scales and Lifting to Theorems A, B, and C

When the global slope inequality (Section 5) extracts $G_{\text{avg}}(x)$ at scale x, unwanted echoes also enter from much larger "thin" scales $y \geq \frac{1}{3}X' = \frac{1}{3}\exp(x^{\varepsilon/3})$. Because the host conductor $N = Q_S q d$ has $\log N \leq 3x^{\varepsilon/5}(\log x)^2$, classical textbook prime bounds (Siegel–Walfisz) alone cannot control them: Siegel's theorem cannot rule out a hypothetical real "Siegel zero" very close to $s = 1$ at conductor $\approx \exp(x^{\varepsilon/5})$.

The formalization crushes these thin scales in two steps:

1. **Zero-free box + classical zero-free region:** By the choice of the prime pair (q, d) in Section 4.2, every character divisible by q or d has no real zero in $(1 - \frac{1}{400}, 1)$, while the remaining finitely many small conductors are handled by continuity at $s = 1$. Combined with the classical zero-free region (`zFRDirichlet`) and the truncated explicit formula, this gives
$$\|\psi_0(\chi, u)\| \le Cu(\log u + \lambda)^2 \exp\left(-\frac{c \log u}{\lambda + \sqrt{\log u}}\right), \qquad \lambda = 3x^{\varepsilon/5}(\log x)^2.$$
2. **The $\frac{61}{100} + \frac{39}{100} = 1$ exponent split:** Integrating over $v = \log u \ge \frac{1}{2}x^{\varepsilon/3}$, factor $v = v^{61/100} \cdot v^{39/100}$:
$$\frac{61}{100} \cdot \frac{\varepsilon}{3} = \frac{\varepsilon}{5} + \frac{\varepsilon}{300} \text{ (absorbs } \lambda), \qquad \frac{v}{\lambda} \gtrsim x^{13\varepsilon/100}, \qquad \frac{v}{\sqrt{v}} = \sqrt{v} \gtrsim x^{\varepsilon/6} \gg x^{13\varepsilon/100}.$$
Hence $v/(\lambda + \sqrt{v}) \gtrsim x^{13\varepsilon/100} \gg x^{\varepsilon/8}$, crushing every thin-scale sum to $\|G_0(\chi, y)\| \le y \exp(-\frac{1}{2}x^{\varepsilon/8})$.

Once Theorem 3.1 (letting $\varepsilon \to 0$) yields $L(s, \psi) \neq 0$ for every primitive ψ and every $\mathrm{Re}(s) > \frac{1}{2}$, three classical deductions finish Theorems A, B, and C:

1. **Imprimitive characters (Theorem A):** Any character $\chi \pmod N$ comes from a primitive character χ^* via $L(s, \chi) = L(s, \chi^*) \prod_{p|N}(1 - \chi^*(p)p^{-s})$. Every finite Euler factor is nonzero for $\mathrm{Re}(s) > 0$ since $|\chi^*(p)p^{-s}| \le p^{-\mathrm{Re}(s)} < 1$.
2. **Critical strip reflection (Theorem B):** If $L(s, \chi) = 0$ with $0 < \mathrm{Re}(s) < \frac{1}{2}$, stripping Euler factors gives $L(s, \chi^*) = 0$, and the functional equation mirrors this across the critical line to $L(1 - s, (\chi^*)^{-1}) = 0$ with $\mathrm{Re}(1 - s) > \frac{1}{2}$ (using $\Gamma_{\mathbb{R}}(s + a) \neq 0$ for $\mathrm{Re}(s) > 0$), contradicting Theorem A.
3. **Riemann Hypothesis (Theorem C):** Setting $N = 1, \chi = 1$ gives $\zeta(s) \neq 0$ for $\mathrm{Re}(s) > \frac{1}{2}$ and mirrors $\mathrm{Re}(s) < \frac{1}{2}$ via $\zeta(s) = 2(2\pi)^{-(1-s)}\Gamma(1 - s)\cos(\frac{\pi(1-s)}{2})\zeta(1 - s)$.

4 Engineering the Stage: The CM Field Tower, Behrend's Spheres, and Biased Sampling

SECTION AT A GLANCE

Goal: Build the custom abelian CM number field E, tune all asymptotic exponents, and design the three infinite-place weight bands (Top, Thick, Thin) that isolate $G_{\mathrm{avg}}(x)$.
Input → Output: Takes a primitive character ψ and Vandiver window (§7), and outputs the host stage fed to Bost's inequality (§5) and the theta-lift engine (§6).
Core Idea: High-dimensional Euclidean spheres (Behrend's set T_2) prevent short Frobenius collisions, while a 4:3 sampling bias on T_T Fourier-transforms into strictly positive weights on $\langle\psi\rangle$.

To prove GRH for a given primitive Dirichlet character $\psi \pmod{N_\psi}$, the formalization does not work directly over the rational numbers $\mathbb{Q}$. Instead, it builds a carefully engineered larger number system—an abelian **CM number field** E—whose symmetries encode ψ, splits its complex

embeddings into three weight regimes, and tunes the ranks (m, n, k, t) to balance competing error budgets.

EXPLAINER What Is a CM Field—and What Are Split vs. Inert Primes?

A **CM field** (short for *Complex Multiplication*) is a totally imaginary quadratic extension E/F of a totally real number field F—think of the Gaussian numbers $\mathbb{Q}(i)$ over $\mathbb{Q}$, or $\mathbb{Q}(\sqrt{-d})$ over $\mathbb{Q}$. A CM field comes with a canonical **complex conjugation** $\mathrm{cc} \in \mathrm{Gal}(E/F)$ that acts identically under every embedding $E \hookrightarrow \mathbb{C}$, so its $2f = [E : \mathbb{Q}]$ embeddings pair up into f conjugate pairs; picking one from each pair defines a **CM type** Φ of size f.
When you pass from $\mathbb{Z}$ to a quadratic number system like $\mathbb{Z}[i]$ inside $\mathbb{Q}(i)$, ordinary primes behave in three ways:

- **Split:** $5 = (2 + i)(2 - i)$ breaks into two distinct prime factors.
- **Inert:** 3 stays prime—it refuses to factor!
- **Ramified:** $2 = -i(1 + i)^2$ collapses into a repeated square (this happens only for finitely many primes dividing the discriminant).

Why build a huge CM field E of degree $2f = 4n \approx 4x^{2+8\varepsilon}$? First, over E/F, totally definite unitary groups $U(X_j)$ are compact at every archimedean place, collapsing infinite-dimensional wave spaces down to *finite-dimensional* algebraic spaces. Second, having $f = 2n \approx 2x^{2+8\varepsilon}$ embeddings gives room to separate Top, Thick, and Thin weights and average over symmetries to isolate $G_{\mathrm{avg}}(x)$.

4.1 Building the Custom CM Field E for Any Character ψ (Even or Odd)

The host CM field E is assembled by combining **four purposeful building blocks** inside a cyclotomic field $\mathbb{Q}(\zeta_{Q_S qd})$ (Figure 3):

$$E = \underbrace{\mathbb{Q}(\sqrt{-d})}_{\text{inert primes } K} \cdot \underbrace{\mathbb{Q}(\sqrt{-q})}_{\text{unramified } E/F \text{ \& odd } \psi} \cdot \underbrace{\mathbb{Q}(\zeta_{2^{e_{\mathrm{Iw}}+2}})^{+}}_{\text{2-adic Iwasawa layer } H_{e_{\mathrm{Iw}}}} \cdot \underbrace{K_\psi^{\mathrm{tw}}}_{\text{encodes } \langle\psi\rangle} .$$

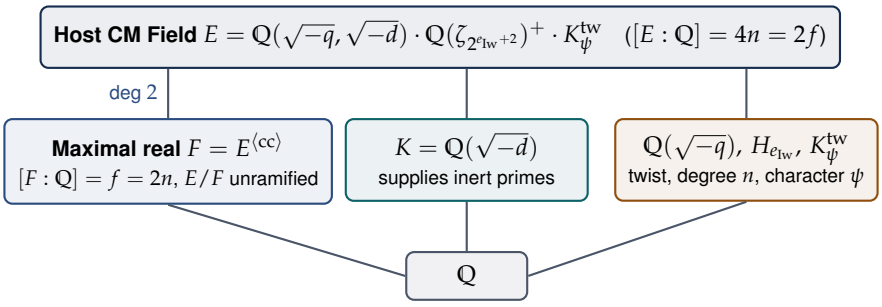


Figure 3: Structure of the custom abelian CM field E. Every finite prime is unramified in the quadratic extension E/F, and the relative different $\mathfrak{D}_{E/K}$ splits as $\mathfrak{b}\overline{\mathfrak{b}}$ with $\mathfrak{b} + \overline{\mathfrak{b}} = (1)$.

Concretely, think of the 2-adic **Iwasawa layer** $\mathbb{Q}(\zeta_{2^{e_{\mathrm{Iw}}+2}})^+$ as a clean "volume knob": dialing the integer $e_{\mathrm{Iw}} \geq 2$ doubles the field degree step by step via $H_{e_{\mathrm{Iw}}} := (\mathbb{Z}/2^{e_{\mathrm{Iw}}+2}\mathbb{Z})^\times / \langle -1 \rangle \cong \mathbb{Z}/2^{e_{\mathrm{Iw}}}\mathbb{Z}$.

Fix primes $q \equiv 3 \pmod 8$ and $d \equiv 7 \pmod{16}$ (selected below), let $Q_S := \mathrm{lcm}(2^{e_{\mathrm{Iw}}+2}, N_\psi)$, let $\Xi = \langle\psi\rangle$ (order $o = \mathrm{ord}(\psi)$), let $R \le H_{e_{\mathrm{Iw}}} \times \widehat{\Xi}$ be the image of $(\mathbb{Z}/Q_S\mathbb{Z})^\times$, and set $G_\varphi := (\{\pm1\} \times \{\pm1\}) \times R$. Then E is the fixed field of the kernel of the Artin surjection $\mathrm{art} : (\mathbb{Z}/Q_S q d\mathbb{Z})^\times \twoheadrightarrow G_\varphi$:

$$\mathrm{art}(u) := \left(\left(\left(\frac{u}{q}\right), \left(\frac{u}{d}\right)\right), \mathrm{toA}(u \bmod Q_S) \cdot w\left(\left(\frac{u}{q}\right)\right)\right), \qquad w(-1) := \mathrm{toA}(-1).$$

Why the $w((\frac{u}{q}))$ sign twist is clever: If ψ is an *odd* character ($\psi(-1) = -1$, like χ_4 above), naive evaluation sends -1 to a nontrivial element of R, breaking the clean product structure of complex conjugation. Because $q \equiv 3 \pmod 4$ guarantees the Legendre symbol $(\frac{-1}{q}) = -1$, multiplying by $w((\frac{u}{q}))$ cancels $\mathrm{toA}(-1)^2 = 1$ and forces $\mathrm{art}(-1) = ((-1,-1),1) = \mathrm{cc}$ for **both even and odd characters**.

This engineering locks in three structural properties at once:

- E/F **is unramified at every finite prime** (`unramEF`): complex conjugation cc never lies in any finite inertia group—indeed cc stays outside the decomposition group at every ramified prime (while it *does* lie in the decomposition group at primes inert in E/F, which Section 6 relies on).
- **Self-dual lattice structure:** The relative different $\mathfrak{D}_{E/K}$ over $K = \mathbb{Q}(\sqrt{-d})$ splits cleanly as $\mathfrak{D}_{E/K} = \mathfrak{b}\overline{\mathfrak{b}}$ with $\mathfrak{b} + \overline{\mathfrak{b}} = (1)$, yielding a unimodular Hermitian lattice $\mathfrak{b}^{-1}$ (meaning it equals its own dual lattice, avoiding messy local discriminant factors).
- **Clean character spectrum:** Every Galois character trivial on the cyclic subgroup $A_0 \le R$ (of index o, order $n_0 = |A_0| = n/o$) takes the form $\xi(\cdot)(\frac{\cdot}{q})^i(\frac{\cdot}{d})^j$ with $\xi \in \langle\psi\rangle$ and $i,j \in \{0,1\}$. Here $(i,j) = (0,0)$ gives our target characters, while $(i,j) \neq (0,0)$ are twisted characters pre-bounded by our choice of (q,d).

4.2 Why the Exponents $(m,n,k) \approx (x^{1/2}, x^2, x^{9/2})$, $t \le x^{400}$, and $\varepsilon < 1/52$ Are Chosen

At first glance, exponents like $k \approx x^{9/2}$, $t \le x^{400}$, and $\varepsilon < 1/52$ look arbitrary. In fact, every one of them drops out of four elementary algebra constraints (Figure 4):

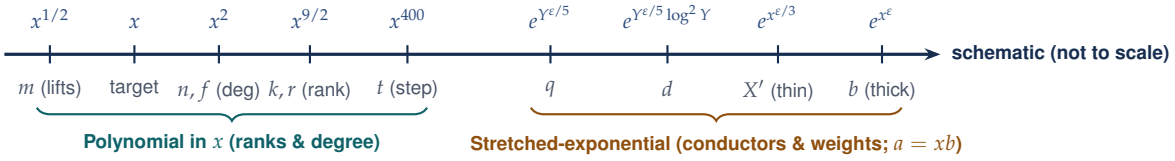


Figure 4: Schematic hierarchy of scales in the formalization ($x \in [Y, Y^{3/2}]$, not drawn to scale), separating polynomial rank parameters on the left from stretched-exponential analytic weights on the right.

1. **Error balancing uniquely pins $(h_m, h_n, h_k) = (\frac{1}{2}, 2, \frac{9}{2})$:** Write $m \asymp x^{h_m}$, $n \asymp x^{h_n}$, $k \asymp x^{h_k}$. The global error budget is

$$\mathrm{Err} = x^{C\varepsilon}\left(|T_T|\, a\left(1 + \frac{m^2}{k}\right) + m(B_0 + N_c)\right), \qquad B_0 = fmbx^{\varepsilon}, \quad N_c = \frac{|T_2|kbx^{2\varepsilon}}{n^2}.$$

Three competing inequalities lock the exponents into place:

$$\begin{aligned} mB_0 \asymp |T_T|a &\implies m^2 \asymp x\delta_1 &&\implies h_m = \tfrac{1}{2}, \\ k \gtrsim m^5 f L \ \text{(RTF proviso, §5)} &\implies h_k \geq 5h_m + h_n &&= \tfrac{5}{2} + h_n, \\ mN_c \lesssim |T_T|a &\implies k \lesssim \sqrt{x}\, n^2 &&\implies h_k \leq 2h_n + \tfrac{1}{2}. \end{aligned}$$

Comparing $\frac{5}{2} + h_n \leq h_k \leq 2h_n + \frac{1}{2}$ forces $h_n \geq 2$! Picking the smallest permitted degree exponent $h_n = 2$ yields $(h_m, h_n, h_k) = (\frac{1}{2}, 2, \frac{9}{2})$, padded with ε-margins to

$$m \in [x^{1/2-\varepsilon/16}, 2x^{1/2-\varepsilon/16}], \qquad n \in [x^{2+8\varepsilon}, 2x^{2+8\varepsilon}], \qquad k \in [x^{9/2+10\varepsilon}, 2x^{9/2+10\varepsilon}].$$

2. **Where $t \leq x^{400}$ comes from:** The auxiliary lift step t satisfies $r + t = 8fr + 4 + \mathrm{TOv}(s) \cdot \mathrm{nWv}(R, s)$, where $\mathrm{TOv}(s) = \lceil x^{377} \rceil$ and $\mathrm{nWv}(R, s) \lesssim r(f + f^2k^3L^2/m) \lesssim x^{22.67}$ for $\varepsilon < 1/52$ (summing exponents $2(2.16) + 3(4.70) + 0.04 - 0.49 + 4.70 = 22.67$). Thus

$$r + t \lesssim x^{377+22.67} = x^{399.67} \leq x^{400}.$$

3. **Choosing (q, d) via congruences on $d + 1$:** Using large-sieve mean-square estimates (`QH.pairEx`), the formalization chooses primes $q \equiv 3 \pmod 8$ ($\log q \approx Y^{\varepsilon/5}$) and $d \equiv 7 \pmod{16}$ ($\log d \approx Y^{\varepsilon/5}(\log Y)^2$) such that $p \mid (d+1)$ for $p = q$ and every odd prime $p \leq Y^{\varepsilon/5}$. These congruences force $-d \equiv 1 \pmod{8p}$, ensuring that 2, q, and all small primes split in $K = \mathbb{Q}(\sqrt{-d})$. Every (q, d)-twisted character χ of conductor $\leq Q$ ($\log Q \lesssim Y^{\varepsilon/5}(\log Y)^2$) then satisfies $\|G_\chi(y)\| \leq y^{1/2+\varepsilon/2}$ and has no zero in the box $\mathrm{Re}(s) > 1 - \frac{1}{400}$, $|\mathrm{Im}(s)| \leq Q$.

4. **Where $\varepsilon < 1/52$ comes from:** From the zero-free box of width $\frac{1}{400}$, contour integration at $z = (\log Q)^{1000}$ bounds $|L'(1, \chi)/L(1, \chi)| \leq 1000 \log\log Q + C$. Substituting $\log Q \lesssim Y^{\varepsilon/5}(\log Y)^2$ gives a leading term of

$$1000 \cdot \frac{\varepsilon}{5} \log Y = 200\varepsilon \log Y < 4 \log Y \iff \varepsilon < \frac{1}{50}.$$

Tightening slightly to $\varepsilon < \frac{1}{52}$ reserves a clean margin of $\frac{2}{13} \log Y$ for lower-order terms.

4.3 Three Weight Regimes, Behrend's Spheres (T_2), and 4:3-Biased Sampling (T_T)

The $f = 2n$ infinite places $\Psi = \tau_1(A) \sqcup \tau_2(A)$ of E are partitioned into three sharply separated weight bands (Table 3):

Regime	Infinite Places	Weight	Scale a_τ / w_j	Role in the Proof
Top	$\tau_1(\tau), \tau \in T_T \subset A$	$a_\tau \approx xb$	—	Primary readout sites for primes up to x
Thick	$\tau_2(\mathrm{dig}(s) + u_0), s \in T_2$	$w_j \approx b$	$a_\tau / w_j \approx x$	Target evaluation scale for $G_{\mathrm{avg}}(x)$
Thin	All remaining places	$w_j \approx w'' \ll b'$	$a_\tau / w_j \geq X'/3$	Crushed to $\exp(-x^{\varepsilon/8})$ by §3.1

Table 3: Partition of the $f = 2n$ infinite places of the CM field E into Top, Thick, and Thin weight bands.

When an inert prime p permutes embeddings via its Frobenius shift $h_p \in A$, two combinatorial properties are required: (i) no short arithmetic progression hits the thick set T_2 more than twice,

and (ii) the top set T_T hits cosets of $A_0 \leq A$ with prescribed densities while avoiding bad Frobenius orbits.

INTUITION & ANALOGY **High-Dimensional Spheres and 4:3 Loaded Coins**

- **Behrend's spheres prevent collisions (T_2):** Picture integer grid points lying on a circle in $\mathbb{R}^2$—say $x^2 + y^2 = 25$, which holds $(0,5), (3,4), (4,3), (5,0)$. By strict convexity, **no straight line can hit a circle or sphere in three points**, so no three points can form an equally spaced arithmetic progression! Felix Behrend (1946) [7] lifted this idea to spheres in $\mathbb{R}^S$ and read the coordinates as base-B digits of integers. Here this constructs $T_2 \subset \{1, \ldots, \lfloor n_0/(4L_w) \rfloor\}$ (where $n_0 = |A_0| = n/o$) of density $|T_2| \geq x^{-\varepsilon} n$ (explicit constant $C = 32$ before deleting 0, doubled to 64 after) satisfying `NoRel`:

$$j(s_2 - s_1) \neq i(s_3 - s_1) \qquad (\text{distinct } s_1, s_2, s_3 \in T_2,\ 0 < |i|, |j| \leq L_w,\ i \neq j).$$

  Hence any Frobenius arithmetic progression of length $L_w \asymp \log x$ meets $T_2 + u_0$ in **at most 2 points**.

- **Why the 4:3 sampling bias produces positive character weights (T_T):** Flip an independent loaded coin for each $a \in A$: pick a with probability $\delta_1 = x^{-\varepsilon/8}$ if $a \in A_0$, and $\frac{3}{4}\delta_1$ if $a \notin A_0$. Normalize the coset density on A/A_0 (order o) as $\lambda(\gamma) = \frac{3+\mathbf{1}_{\gamma \in A_0}}{3o+1}$. Summing against any character ξ of A/A_0 gives exact Fourier inversion:

$$\sum_{\gamma \in A/A_0} \lambda(\gamma)\, \xi(\gamma) = \frac{3\sum_\gamma \xi(\gamma) + 1}{3o+1} = \begin{cases} 1, & \xi = 1, \\ (3o+1)^{-1}, & \xi \neq 1. \end{cases}$$

  A second-moment and Markov argument then extracts a deterministic set T_T with coset-share error $\leq x^{-1+\varepsilon}$ and negligible bad Frobenius orbits.

WORKED EXAMPLE **Checking the 4:3 Fourier Bias by Hand ($o = 2$)**

Take a quadratic character ψ (order $o = 2$, so $A/A_0 = \{0,1\}$ has two cosets). The 4:3 normalized densities are $\lambda(0) = \frac{4}{7}$ and $\lambda(1) = \frac{3}{7}$. Now test both characters on $\{0,1\}$:

- Trivial character ($\xi \equiv 1$): $\frac{4}{7}(1) + \frac{3}{7}(1) = \mathbf{1}$.
- Nontrivial character ($\xi(0) = 1$, $\xi(1) = -1$): $\frac{4}{7}(1) + \frac{3}{7}(-1) = \frac{1}{7} = \frac{1}{3(2)+1}$.

Both weights are strictly positive—exactly what Theorem 3.1 needs so zeros never cancel!

Finally, because $n = |R| = c \cdot 2^{e_{\mathrm{Iw}}}$ carries a character-dependent factor $c \in [1, o]$, rescaling $x = \alpha x'$ (with $\alpha = c^{1/(2+8\varepsilon)}$) shrinks the window $[W, W^2]$ supplied by the Vandiver bound (Section 7) down to $[Z, Z^{3/2}]$—which is precisely where the power-window exponent $\lambda = 3/2$ in Section 3 comes from.

5 The Master Accounting Inequality: Bost's Slope Inequality & Near-Orthogonal Columns

SECTION AT A GLANCE

Goal: Replace Stepanov's elementary inequality "zeros $\leq$ degree" with an Arakelov balance sheet bounding the prime-sum excess ASK $-$ MAIN $\leq$ Err.
Input $\rightarrow$ Output: Combines prime gains (§6), quotient-metric heights (§6), Vandiver bounds (§7), and RTF column orthogonality into the ceiling bound fed to §3.
Core Idea: Viewing the automorphic spaces as high-dimensional lattices, nearly orthogonal columns force the *average* slope of the dual row space $\mu(\overline{\mathcal{R}}^{\vee})$ to stay above the mean prime-power gain.

We now reach the central balance sheet of the entire proof. Everything in Sections 4, 6, and 7 feeds into a single two-sided sandwich on the average lattice slope $\mu(\overline{\mathcal{R}}^{\vee})$ of the dual row space $\mathrm{cc} := \overline{\mathcal{R}}^{\vee}$:

THE MASTER BALANCE SHEET ON $\mu(\overline{\mathcal{R}}^{\vee})$

Lower bound (§5, Bost + RTF): Inert-Prime Gains (§6.3) $-$ Column Heights $\leq \mu(\overline{\mathcal{R}}^{\vee})$

Upper bound (§6, Quotient Metric): $\mu(\overline{\mathcal{R}}^{\vee}) \leq$ Main Term (§6.2) $+$ Vandiver & Gluing Errors (§7)

Comparing the left and right ends of this sandwich cancels $\mu(\overline{\mathcal{R}}^{\vee})$ and immediately produces the prime-sum ceiling ASK $-$ MAIN $\leq$ Err!

EXPLAINER **What Is a Hermitian Vector Bundle and Its Harder–Narasimhan Slope?**

Start with an ordinary grid of points (a **lattice**) $L \subset \mathbb{R}^d$. Its **covolume** $\mathrm{covol}(L)$ is the volume of one fundamental box: tightly packed lattices (small covolume) have a *large positive* slope $-\frac{1}{d}\log \mathrm{covol}(L)$, while stretched-out sparse lattices have a *negative* slope.

Now replace $\mathbb{Z}$ by the ring of integers $\mathcal{O}_K$ of a totally complex CM coefficient field K with $n_{\mathrm{pl}}(K) = \frac{1}{2}[K : \mathbb{Q}]$ complex places, and equip an $\mathcal{O}_K$-lattice L of K-rank $\mathrm{rk}_K(L)$ with a Hermitian inner product (framing X_v) at every complex embedding $v : K \hookrightarrow \mathbb{C}$. That pair $\overline{L} = (L, \{X_v\})$ is a **Hermitian vector bundle** [8]. Its normalized Arakelov slope is

$$\mu(\overline{L}) \;:=\; -\frac{\mathrm{lcv}(\overline{L})}{\mathrm{rk}_K(L)}, \qquad \mathrm{lcv}(\overline{L}) \;:=\; \frac{2}{[K:\mathbb{Q}]}\log \mathrm{ht}(\overline{L}),$$

where $\mathrm{ht}(\overline{L}) = 2^{n_{\mathrm{pl}}(K)\,\mathrm{rk}_K(L)}|\Delta_K|^{-\mathrm{rk}_K(L)/2}\,\mathrm{covol}(\mathrm{zl}(L))$. Because each complex dimension carries *two* real dimensions, scaling $\overline{L}$ by $c \in \mathbb{C}^\times$ shifts $\mu(\overline{L})$ by $-2\log|c|$, matching the squared operator norm in $h(\varphi)$ below.

- **Harder–Narasimhan (HN) filtration:** The densest nonzero sublattice defines $\mu_{\max}(\overline{L})$, and the sparsest quotient defines $\mu_{\min}(\overline{L})$. The canonical HN filtration sorts $\overline{L}$ from densest ($\mu_{\max}$) down to sparsest ($\mu_{\min}$), with rank-weighted average $\mu(\overline{L})$.
- **Duality:** The dual bundle $\overline{L}^\vee$ flips every slope's sign: $\mu(\overline{L}^\vee) = -\mu(\overline{L})$ and $\mu_{\min}(\overline{L}^\vee) = -\mu_{\max}(\overline{L})$.
- **Bost's Slope Inequality:** If a K-linear map $\varphi : \overline{A} \to \overline{B}$ sends A integrally into B ($\varphi(A) \subseteq B$, $\varphi(A) \neq 0$) with archimedean height $h(\varphi) = \frac{1}{n_{\mathrm{pl}}(K)}\sum_{v|\infty}\log\|\varphi_v\|_{\mathrm{op}}^2$, then

$$\mu_{\min}(\overline{A}) \;\le\; \mu_{\max}(\overline{B}) + h(\varphi).$$

  (Intuitively: a bounded-stretch linear map cannot send a dense lattice into a sparse one without enlarging its norm $\|\varphi\|_{\mathrm{op}}$!)

WORKED EXAMPLE **Visualizing $\mu_{\max}$, μ, and $\mu_{\min}$ on a 2D Rectangular Lattice**

Consider the rectangular grid $L = 0.1\,\mathbb{Z} \oplus 100\,\mathbb{Z} \subset \mathbb{R}^2$ (Figure 5): tightly packed along the x-axis (spacing 0.1) and hugely stretched along the y-axis (spacing 100). Its fundamental rectangle has area $\mathrm{covol}(L) = 0.1 \times 100 = 10$.

- **Densest direction ($\mu_{\max}$):** Along the x-axis ($0.1\,\mathbb{Z}$), the 1D real slope is $-\log(0.1) = +\log 10 \approx +2.30$.
- **Sparsest quotient ($\mu_{\min}$):** Projecting onto the y-axis ($100\,\mathbb{Z}$), the 1D real slope is $-\log(100) = -2\log 10 \approx -4.61$.
- **Average slope (μ):** $-\frac{1}{2}\log(10) = \frac{1}{2}\big(\log 10 - 2\log 10\big) \approx -1.15$, sitting strictly between $\mu_{\min}$ and $\mu_{\max}$.

Notice the catch: knowing $\mu_{\max} = +2.30$ tells you *nothing* about whether the y-direction is secretly stretched to 100 or 10^{100}! That is why a single column map (which only bounds $\mu_{\max}$) is not enough—you need beams spanning *every* direction to pin down the average slope μ.

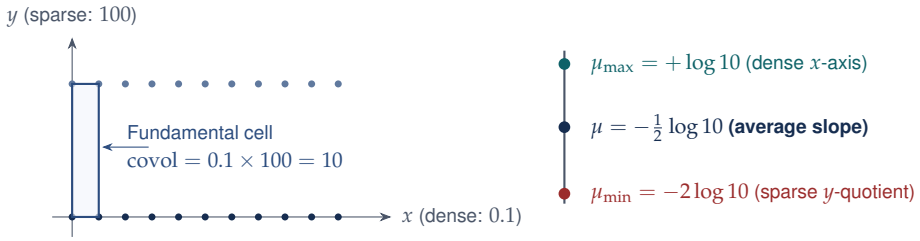


Figure 5: A skewed 2D lattice illustrates why $\mu_{\max}$ alone cannot control the average slope μ: one dense axis can hide an arbitrarily sparse orthogonal direction.

INTUITION & ANALOGY **Bounding an Average Slope via Nearly Orthogonal Beams**

Our global datum consists of:

- one target bundle $\mathrm{cc} := \overline{\mathcal{R}}^{\vee}$ (the dual of the space $\overline{\mathcal{R}}$ of irreducible cuspidal row forms, of rank $\kappa = D_{\mathrm{rows}}$);
- a family of column bundles $\overline{\mathcal{N}}_u$, indexed by words $u \in \mathcal{W}$ (each carrying an admissible prime level T_u), mapping integrally into cc via period maps $B_k(u) : \mathcal{N}_u \to \mathrm{cc}$ of height $h(B_k(u)) \le h(u)$.

Think of the column images $B_k(u)(\mathcal{N}_u)$ as beams of light shining into $\mathbb{C}^\kappa$. A single beam only tells you $\mu_{\max}(\mathrm{cc})$. But if the collection of beams points evenly in *every* direction (an approximately isotropic *tight frame*—like a spherical burst of floodlights), then at every slope cutoff τ the high-slope subspaces of the columns must illuminate almost the entire κ-dimensional room, forcing the average slope $\mu(\mathrm{cc})$ to stay close to the mean shifted column slope $s(\mathcal{N}, h)$.

5.1 Bost's Frame Slope Inequality

Define the rank-weighted mean shifted column slope and the slope excess:

$$s(\mathcal{N},h) \;:=\; \frac{\sum_u d_u\big(\mu(\mathcal{N}_u)-h(u)\big)}{\sum_u d_u}, \qquad u_F(\mathcal{N},h) \;:=\; \max_u\big(\mu_{\max}(\mathcal{N}_u)-h(u)\big)-s(\mathcal{N},h).$$

Theorem 5.1 (Bost's Frame Slope Inequality). *Suppose $B_i(\mathcal{N}_i) \subseteq \mathrm{cc}\,(\mathrm{rk}_K\,\mathrm{cc}=\kappa)$, $h(B_i)\le h_i$, $s(\mathcal{N},h)\ge s$, $u_F(\mathcal{N},h)\le u$, and $s-\mu_{\min}(\mathrm{cc})\le H_{\max}$. Suppose that at some infinite place v, the normalized matrices B_i sit as the κ-block of ambient $D\times d_i$ matrices $B_{\mathrm{all},i}$ with $(1-\delta)D\le\kappa$, and the frame operator $Q:=\sum_i \gamma_i B_{\mathrm{all},i}B_{\mathrm{all},i}^*$ is nearly isotropic:*

$$\sum_i \big\|\gamma_i B_{\mathrm{all},i}^* B_{\mathrm{all},i}-I\big\|_{\mathrm{HS}}^2 \;\le\; \varepsilon_{\mathrm{fr}}\sum_i d_i \quad (\varepsilon_{\mathrm{fr}}\le\tfrac14), \qquad \mathrm{Re}\,\mathrm{Tr}(Q^2) \;\le\; (1+\eta)\,\frac{(\mathrm{Re}\,\mathrm{Tr}\,Q)^2}{D}.$$

*Then (with **no** $-\log|\mathcal{W}|$ loss):*

$$s \;-\; (1-\delta)^{-1}\left(4u\log\left(1+\frac{H_{\max}}{u}\right)+(\eta+4\sqrt{\varepsilon_{\mathrm{fr}}})\,H_{\max}\right) \;\le\; \mu(\mathrm{cc}). \qquad \text{[Lean: TheoremFSharp]}$$

PROOF SKETCH & KEY STEPS **Four-Step Proof of Theorem 5.1**

1. **HN truncation + Markov's inequality:** For any threshold $\tau\in[0,H_{\max}]$, let $N_\tau(i)\subseteq\mathcal{N}_i$ be the HN step collecting all slopes $\ge s-\tau+h_i$. Since the weighted mean shifted slope is $\ge s$ and every slope is $\le s+u$, Markov's inequality bounds the missing rank across all columns:
$$\sum_i\big(d_i-\mathrm{rk}_K\,N_\tau(i)\big) \;\le\; \frac{u}{u+\tau}\sum_i d_i.$$

2. **Bost's slope inequality:** By `SlopeIneqMean`, every nonzero image $B_i(N_\tau(i))$ has $\mu_{\min}\ge s-\tau$, so their span $M_\tau := \bigvee_i B_i(N_\tau(i))$ lies inside the target HN step $S_\tau := \mathrm{hnStep}(\mathrm{cc},s-\tau)$.

3. **Hilbert–Schmidt frame corank bound:** Let $W := \big(\bigvee_i B_{\mathrm{all},i}(N_\tau(i))\big)^\perp \subset \mathbb{C}^D$ have dimension w, put $R=\sum_i d_i$, $t_Q=\mathrm{Re}\,\mathrm{Tr}(Q)\ge(1-\sqrt{\varepsilon_{\mathrm{fr}}})R$, and $X=\mathrm{Tr}\big(Q-\frac{t_Q}{D}I\big)^2\le \eta t_Q^2/D$. Cauchy–Schwarz gives a lower bound on $\mathrm{Tr}(P_WQP_W)$, while orthogonality $W\perp B_{\mathrm{all},i}(N_\tau(i))$ gives an upper bound:
$$\frac{t_Q}{D}\,w-\sqrt{wX} \;\le\; \mathrm{Tr}(P_WQP_W) \;\le\; \left(\frac{u}{u+\tau}+\sqrt{\frac{u}{u+\tau}\,\varepsilon_{\mathrm{fr}}}\right)R.$$
Solving for w gives $w\le\big(\eta+\frac{4u}{u+\tau}+4\sqrt{\varepsilon_{\mathrm{fr}}}\big)D$. Since projecting $\mathbb{C}^D\twoheadrightarrow\mathbb{C}^\kappa$ maps $\bigvee_i B_{\mathrm{all},i}(N_\tau(i))$ onto $M_\tau\subseteq S_\tau$, we have $\kappa-\mathrm{rk}_K(S_\tau)\le w$, and $(1-\delta)D\le\kappa$ yields
$$\kappa-\mathrm{rk}_K(S_\tau) \;\le\; (1-\delta)^{-1}\left(\eta+\frac{4u}{u+\tau}+4\sqrt{\varepsilon_{\mathrm{fr}}}\right)\kappa.$$

4. **Layer-cake integration:** Integrating this corank tail bound over $\tau\in[0,H_{\max}]$ against $\int_0^{H_{\max}}\frac{4u}{u+\tau}\,d\tau = 4u\log(1+H_{\max}/u)$ yields Theorem 5.1.

5.2 Approximate Column Orthogonality and Row Irreducibility

To check the hypotheses of Theorem 5.1, the formalization certifies that our "floodlight beams" are nearly isotropic (η, $\sqrt{\varepsilon_{\mathrm{fr}}}$ tiny) and that almost all rows are irreducible (δ tiny):

- **Approximate Column Orthogonality (via the Relative Trace Formula; `LinkS`):** Let P_u be the Poincaré-series embedding of the unramified column space (dimension d_u^{unr}) into the full automorphic space (dimension D_{all}, Weyl-law main term D^*), and define the normalized Gram matrix $K_D(u,u') = (c_0(u)c_0(u'))^{-1/2}P_u^*P_{u'}$ (measuring the inner-product overlap between beams u and u'). Expanding $K_D(u,u')$ over double cosets via the **relative trace formula** proves that distinct column spaces behave like nearly orthogonal subspaces:

$$\begin{aligned} \|K_D(u,u')\|_{\mathrm{HS}}^2 &\le (1+b^{-6})\,\frac{d_u^{\mathrm{unr}}\,d_{u'}^{\mathrm{unr}}}{D^*} \qquad (u \neq u'), \\ \|K_D(u,u) - I\|_{\mathrm{HS}}^2 &\le \left(\frac{2C_0}{\Gamma} + 2e^{-kf\log k}\right) d_u^{\mathrm{unr}}, \end{aligned}$$

  where the **proviso** $C_0^{-1} \le \frac{d_u^{\mathrm{unr}}\Gamma^2}{D^*} \le C_0$ (with $C_0 = 3$ and $\Gamma = \exp(C_\Gamma m^5 f L)$) cancels one power of Γ in the diagonal error. All three frame defects are crushed below $\exp(-x^{\varepsilon/8})$.

- **Row Irreducibility (`LinkIRR`):** Using Green's parametrization of cuspidal characters of $\mathrm{GL}_n(\mathbb{F}_g)$ at two auxiliary split primes g_0, g_1, the fraction $\delta = 1 - D_{\mathrm{rows}}/D_{\mathrm{all}}$ of reducible ("bad") rows satisfies

$$\delta \le \exp\left(-x^{\varepsilon/(7J)}\right) \ll \frac{1}{7}$$

  (where J is the Vandiver window exponent from Section 7, bottlenecked by $1/(4\min(q,d)) \le \exp(-x^{\varepsilon/(6J)})$), easily clearing the sharp threshold $\delta \le 1/7$ needed in the slope-excess bound.

5.3 The 9-Predicate Global Assembly

The master assembly theorem (`topDesWin`) fixes $(C_{\mathrm{abs}}, C_0, e_D, e_L) = (2,3,2,0)$ and intersects **nine closed family predicates** over a single stratified threshold:

Combining the lower bound on $\mu(\mathrm{cc}) = \mu(\overline{\mathcal{R}}^{\vee})$ from Theorem 5.1 with the upper bound from `LinkDIV` (Section 6) and rewriting $\mu(\mathcal{N}_u)$ in terms of prime sums via `LinkG` and `LinkMu` yields ASK − MAIN ≤ Err, feeding directly into the windowed Landau engine of Section 3.

#	Lean Predicate	Mathematical Role in the Global Assembly
1	`LinkN`	Congruence-index bound containing the Vandiver cohomology sum (paired with `HypV` $\Rightarrow$ `DIVinN`, §7)
2	`LinkEta`	Twisted places covered by normalized Hecke data (first input to `LinkDIVt`, §6)
3	`LinkIRR`	Bad-row fraction $\delta \leq \exp(-x^{\epsilon/(7J)}) \leq 1/7$ (§5.2)
4	`SmallLevelRankEq`	Surjective level projection matching local ranks $d_{Z_0}(j) = D_0$ (links `LinkN` to `DIVinN`)
5	`LinkDIVt`	Produces `LinkDIV`: quotient-metric upper bound on $\mu(\mathcal{R}^\vee) + h(u)$ minus baseline column slope (§6)
6	`NuFullF` $\wedge$ `MuRank`	Synthesizes `LinkMu` (ideal-scaling rounding error $\leq O(xnX')$) on the fly
7	`LinkG`	Inert-prime p-adic divisibility gains $\geq$ odd-power prime ramps (§6.3)
8	`LinkS`	Relative-trace-formula column near-orthogonality & Weyl law (§5.2)
9	`LinkGLOBAL`	Combines Theorem 5.1 with `LinkG, LinkS, LinkIRR, LinkDIV, LinkMu` to bound ASK $-$ MAIN $\leq$ Err

Table 4: The nine closed family predicates assembled in `topDesWin` to derive the prime-sum ceiling bound.

6 Inside the Engine: Iterated Theta Lifts, Seesaw Independence, Overcoming the $2/3$ Barrier, and $D^p \equiv 0$

SECTION AT A GLANCE

Goal: Build the m-step chain of unitary theta lifts, cancel its bulk archimedean cost via Rallis "division," extract inert odd-power prime gains via $D^p \equiv 0$, and recover the full character sums $G_\chi(x)$.
Input $\rightarrow$ Output: Supplies both sides of Bost's accounting inequality (§5): the archimedean upper bound (`LinkDIV`) and the p-adic prime-power lower bound (`LinkG`).
Core Idea: Seesaw duality (Fubini's theorem!) makes the period independent of auxiliary twists; dividing by Rallis eigenvalues beats the $2/3$ barrier; $D^p \equiv 0$ yields an odd-power cousin of Legendre's factorial formula.

This section opens the automorphic and p-adic engine at the core of the formalization: how $m \approx \sqrt{x}$ unitary theta lifts manufacture multiplicity m, how seesaw duality and Rallis "division" cancel the bulk archimedean cost, and how inert-prime odd-power ramps assemble into the full character sums $G_\chi(x)$.

EXPLAINER **Unitary Groups $U(X)$, Automorphic Forms, Theta Lifts, and the Seesaw Identity**

- **Automorphic forms on totally definite unitary groups $U(X)$:** Think of a classical periodic wave like $\sin(nx)$ on a circle. An **automorphic form** generalizes this to symmetric functions on higher-dimensional matrix groups. Over our CM field E/F, a Hermitian space (X, h) of dimension N has a symmetry group $U(X)$ preserving h. When h is positive (or negative) definite at every infinite place of F, $U(X)$ is **totally definite** (compact at ∞), and spaces of automorphic forms $\mathcal{A}(U(X))$ collapse to finite-dimensional tables of values on a finite set of lattice classes weighted by $|\operatorname{Aut}(L)|^{-1}$.
- **Theta lifts (an integral translator):** Inside a larger symplectic group, a Hermitian space X and skew-Hermitian space Y form a *dual pair* $(U(X), U(Y))$ (multiplying by $\sqrt{-d}$ converts skew-Hermitian forms back to Hermitian ones). The **Weil representation** supplies a coupling kernel $\theta(g, h)$, and integrating an automorphic form φ on G against $\theta(g, h)$ translates it into a form $\Theta_{G\to H}(\varphi)(h) = \int_{[G]} \theta(g, h)\overline{\varphi(g)}\, dg$ on H—just like multiplying a vector by a matrix to carry a melody from a smaller orchestra to a larger one!
- **The Seesaw Identity (higher-dimensional Fubini / matrix transpose):** Remember from linear algebra that $\langle Av, w\rangle = \langle v, A^* w\rangle$. Whenever two dual pairs (G_1, H_1) and (G_2, H_2) share an ambient symplectic group with $G_2 \subset G_1$ and $H_1 \subset H_2$, swapping the order of integration (Fubini's theorem!) gives

$$\left\langle \Theta_{G_2\to H_2}(\varphi)\big|_{H_1},\, F\right\rangle_{H_1} \;=\; \left\langle \varphi,\, \Theta_{H_1\to G_1}(F)\big|_{G_2}\right\rangle_{G_2}.$$

  Applying this swap m times trades the upward lift chain $U(X_0) \to \cdots \to U(X_m)$ paired on $U(Y_m) \cong U(X') \subset U(X_m)$ for a downward row tower $\iota_\mu : U(Y_m) \to \cdots \to U(Y_0)$ paired on $U(X_0) \subset U(Y_0)$.

6.1 The m-Step Chain of Totally Definite Unitary Theta Lifts

Starting from a totally positive definite space X_0 of even rank $N(0) = k$, the formalization builds a chain of $m \approx x^{1/2-\varepsilon/16}$ totally definite Hermitian spaces

$$X_0,\ X_1,\ \ldots,\ X_m, \qquad \dim_E X_j = N(j) := k + j(t+1).$$

Each step K_i lifts a form $\varphi \in M_{\text{ram}} \subset \mathcal{A}(U(X_0))$ (at admissible prime level T_u) to $U(X_{i+1})$ using an auxiliary Hecke character μ_{i+1}, appending a block of $t+1$ coordinates to the highest weight. Composing all m lifts gives $\operatorname{col}(\varphi) = (K_{m-1} \circ \cdots \circ K_0)(\varphi)$ on $U(X_m)$, whose local p-adic valuation shift accumulates additively across all m steps.

In parallel, the seesaw partner spaces $Y_0, Y_1, \ldots, Y_m$ (of rank $\dim Y_j = r + jt$, $r = k + m - 1$, with isometric embeddings $X_j \hookrightarrow Y_j$ for $j < m$) begin indefinite at $Y_0 = X_0 \perp \langle z_m\rangle \perp \mathbb{H}^{m/2-1}$ (signature $(k + \frac{m}{2} - 1, \frac{m}{2})$) and become *totally definite* at $Y_m \cong X' \hookrightarrow X_m$ (codimension 1). Seesaw factorization (`adm1`) collapses the top-level period to

$$\operatorname{perCol}(u, \mu, \varphi, F) \;=\; \zeta(u, \mu) \cdot B_u\big(\varphi,\, \iota_\mu(F)\big),$$

where B_u is a **Gross–Prasad / Bessel period** on $U(X_0) \times U(Y_0)$ *independent* of the auxiliary Hecke characters $\mu = (\mu_1, \ldots, \mu_m)$ (`BesselId`). Meanwhile, the **Rallis inner-product formula** ensures that $\iota_\mu \circ \iota_\mu^*$ acts on each isotypic block π by the positive scalar $\kappa_R \rho(\mu) \prod_{j=1}^m \lambda_P(\mu, j, \pi)$ of central L-values (Figure 6).

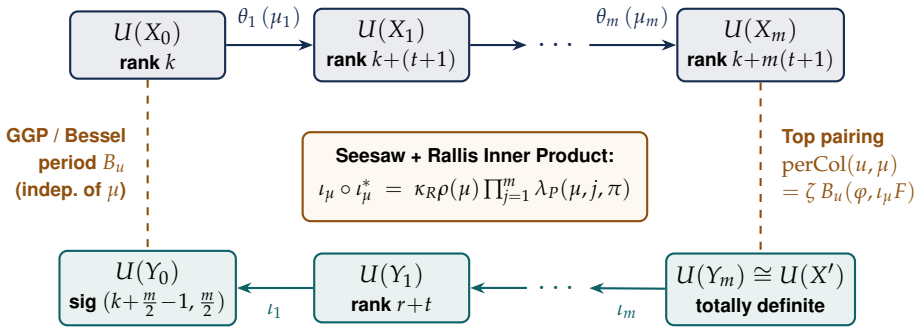


Figure 6: The m-step iterated theta lift chain ($m \approx \sqrt{x}$) paired with its seesaw row tower. Seesaw duality reduces the top pairing on $U(X_m)$ to a Bessel period B_u on $U(X_0) \times U(Y_0)$ independent of $\mu_1, \ldots, \mu_m$.

6.2 Beating the $2/3$ Barrier: Quotient Metrics, Local Gluing, and Chowla–Selberg

> **INTUITION & ANALOGY** **Weighing a Diamond Inside a Steel Safe (Why Naive Lifts Stall at $2/3$)**
>
> Each theta lift carries a massive archimedean height cost of size $O(kw_j \log w_j)$. Estimating the height of the lift directly is like weighing a tiny diamond bolted inside a two-ton steel safe: the safe's mass drowns out the diamond, stalling the zero-free bound near $\mathrm{Re}(s) = 2/3$ (`remark.pdf`, items 9, 14–15).
>
> Because seesaw duality proves that the underlying Bessel period B_u is *independent* of the auxiliary Hecke characters μ, we can "tare the scale" in two ways at once: **divide out** the known Rallis spectral factor (the safe's archimedean mass) via a quotient Hermitian metric, and **glue** the cleanest local integral lattice across different choices of μ at each prime w (like taking the greatest common divisor across a family of equations to strip away accidental prime factors!).

Table 5 tracks the four-stage cancellation ("division") that removes every bulk term and isolates the exact smoothed prime main term.

Specifically, the CM period Ω entering the integral normalization satisfies $\log \Omega = \gamma + \log(\pi/v_K) - L'(1, \chi_K)/L(1, \chi_K)$ by the **Chowla–Selberg formula** (where v_K is the archimedean discriminant scale of $K = \mathbb{Q}(\sqrt{-d})$). Telescoping the **Siegel–Weil unitary mass formula** bounds $\log T_u$ from below by $\sum_{j=1}^m \frac{1}{f} \sum_\rho \log(w_{j,\rho} v_K/\pi) + m \log(g_0 g_1) - \mathrm{err}$, canceling both $\log(g_0 g_1)$ and $\log(\pi/v_K)$ to leave the exact target main term:

$$\mathrm{MainC} = \sum_{j=1}^m \sum_{\tau \in T_T} a_\tau \left(\log a_\tau - 1 - \gamma + \frac{L'(1, \chi_K)}{L(1, \chi_K)} - \frac{1}{f} \sum_{\rho \in \Psi} \log w_{j,\rho} \right).$$

Underpinning the integral normalization is `GradedWeylDivisibility`, which builds standard monomial theory over $\mathbb{Z}$ for integral dual Weyl modules $\nabla_{\mathbb{Z}}(\nu) \subset \mathbb{Z}[X_{a,s}]$ (proving semistandard

#	Threatening Cost	Cancellation Mechanism	Surviving Term
1	Spectral factor $\prod_{j=1}^{m} \lambda_P(\mu, j, \pi)$	Quotient Hermitian metric $(PH^{-1}P^*)^{-1}$ on row space EE	Divides out Rallis eigenvalues before taking heights (`StepB`)
2	Prime-by-prime congruence noise	Glue optimal local lattices $\mathcal{R}_w = \iota_{\mu_{\mathrm{At}(w)}}(\dots)_w$ across μ	Single global lattice $\mathcal{R} \subset$ EE (a "gcd", `SecondStepGluing`)
3	Bulk archimedean cost $O(kw_j \log w_j)$	Strictly interlacing lift weights and Rallis hook-length factorials	Entire $O(kw_j \log w_j)$ bulk telescopes to zero
4	$\log(g_0 g_1)$ and $\log(\pi/v_K)$ in CM period $\log \Omega$	Chowla–Selberg + telescoping Siegel–Weil unitary mass bound for $\log T_u$	Cancels $\log(g_0 g_1)$ and $\log(\pi/v_K)$, leaving exact `MainC`

Table 5: The four-stage cancellation ledger that overcomes the $2/3$ barrier (`LinkDIV`).

bideterminants form a $\mathbb{Z}$-basis) and exact Fock-space lattice duality between Weyl and dual Weyl lattices. Combined with the quotient-metric bound (`StepB`), this yields `LinkDIV`.

6.3 Extracting Primes: $D^p \equiv 0$, Dwork Integrality, and Recovering Full G_χ

Prime powers enter the accounting ledger through `LinkG` at primes $p \le X''(n) = xnX'$ that are **inert in** $K = \mathbb{Q}(\sqrt{-d})$ (together with the single ramified prime $p = d$).

INTUITION & ANALOGY **An Odd-Power Cousin of Legendre's Factorial Formula**

Every math enthusiast knows **Legendre's formula** for the exact power of a prime p dividing $a!$:

$$v_p(a!) \;=\; \left\lfloor \frac{a}{p} \right\rfloor + \left\lfloor \frac{a}{p^2} \right\rfloor + \left\lfloor \frac{a}{p^3} \right\rfloor + \cdots .$$

Why does it arise? Every p-th integer contributes one factor of p, every p^2-th contributes another, and so on.

On a CM elliptic curve at an inert (supersingular) prime p, the invariant differentiation operator D satisfies $D^p \equiv 0 \pmod{p}$, and its formal group has **height** 2 (meaning the p-adic structure repeats every *two* powers of p instead of every one!). As a result, when you apply D to a p-adic theta series a times, the p-divisibility follows the **odd-power cousin of Legendre's formula**:

$$v_p\left(\frac{D^a\theta(0)}{\Omega^a}\right) \;\approx\; \frac{a}{p} + \frac{a}{p^3} + \frac{a}{p^5} + \cdots$$

(minus the partner weights w_j from adjacent Frobenius embeddings). Multiplying by $\log p$ turns each term into $\Lambda(p^i)(a/p^i - w_j)_+$—literally the smoothed von Mangoldt ramp from Section 1.2!

PROOF SKETCH & KEY STEPS **How $D^p \equiv 0$ and Dwork Integrality Produce Odd-Power Prime Ramps (`LinkG`)**

1. **Height-2 supersingular formal logarithm:** At a prime p inert in $K = \mathbb{Q}(\sqrt{-d})$, a CM elliptic curve has supersingular reduction: its invariant derivation satisfies $D^p \equiv 0 \pmod{p}$, and its formal group has height 2 with logarithm

$$\mathrm{lgF}(X) = \sum_{j\geq 0} \frac{X^{p^{2j}}}{p^j}, \quad g_0(X) := \frac{1}{p}\mathrm{lgF}(X^p) = \sum_{j\geq 0} \frac{X^{p^{2j+1}}}{p^{j+1}}, \quad g_0(X^p) = \mathrm{lgF}(X) - X.$$

2. **Multivariate Dwork integrality:** Rescaling by ϖ ($\varpi^{p^2-1} = p$) and applying the multivariate Dwork criterion ($\exp(\Psi)$ has p-adic integer coefficients iff $p \mid \sigma(\Psi)(X^p) - p\Psi(X)$ for a Frobenius lift σ) proves that the n-th Taylor jet of the p-adic theta series $\exp(-\Phi)U_0$ acquires p-divisibility pairing Frobenius-adjacent places $(s, s+1)$ with degrees $(p, 1)$—delivering the odd-power exponent $a/p + a/p^3 + a/p^5 + \ldots$ minus partner weights.

3. **Weyl/dual-Weyl lattice sandwich:** Sandwiching the local automorphic lattice $p^e M_{\mathrm{na}} \leq M \leq M_{\mathrm{na}}$ and averaging over Frobenius orbits (using the no-kernel shortcut `zeroCh`) yields the local divisibility gain inequality:

$$\sum_{j=1}^{m} \sum_{\tau \in T_{T,w}} \Big(\mathrm{VM}(j,\tau,p) - \mathrm{err}_{\mathrm{loc}}\Big) - m \cdot \mathrm{err}_{\mathrm{fl}} \;\leq\; \sum_{j=1}^{m} \mathrm{gain}(j,w) - \tau_{\mathrm{tor}}(w),$$

where the odd-power prime ramp (written `vonMangoldtClassSum` in Lean) is

$$\mathrm{VM}(j,\tau,p) \;:=\; \sum_{i \text{ odd}} \max\left(0,\; \frac{a_\tau}{p^i} - w_j\big(\mathrm{partner}(p,\tau,i)\big)\right),$$

$T_{T,w} \subseteq T_T$ are the top places above w, $\mathrm{err}_{\mathrm{loc}}$ and $\mathrm{err}_{\mathrm{fl}}$ are local logarithmic/polynomial losses, and $\tau_{\mathrm{tor}}(w)$ is the orbit-sublattice torsion cost.

How Odd Powers at Inert Primes Recover the Full Character Sums $G_\chi(x)$

A sharp reader will immediately ask: `LinkG` only supplies *odd* prime powers p^i at primes p *inert* in $K = \mathbb{Q}(\sqrt{-d})$ (weighted by $2\log p$), plus a ramified term $\log d \cdot \mathrm{askD}$ at $p = d$. How can we possibly recover the *full* prime-power sums $G_\chi(x)$, which include split primes and even prime powers?

The answer is a remarkably clean sign-parity trick proved in `ChBalance.lean` (`askExp`, `ch_balance`) via Fourier inversion on the Galois group G_φ. Let $\chi_K : G_\varphi \to \{\pm 1\}$ be the quadratic sign character whose kernel is the index-2 CM type Φ ($|\Phi| = |G_\varphi|/2$). Then $g \notin \Phi \iff \chi_K(g) = -1$:

- If p **splits** in K, then $\chi_K(\mathrm{Frob}_p) = +1$, so $\chi_K(\mathrm{Frob}_p^i) = (+1)^i = +1$ (never in $g \notin \Phi$).
- If p is **inert** in K, then $\chi_K(\mathrm{Frob}_p) = -1$, so $\chi_K(\mathrm{Frob}_p^i) = (-1)^i = -1$ **precisely when i is odd**!

Expanding the full prime-power sum over *all* integers $1 \le n \le N$ and *all* Galois characters $\psi \in \widehat{G_\varphi}$, restricted to elements $g \notin \Phi$, gives

$$\frac{1}{|\Phi|} \sum_{g \notin \Phi} \sum_{\psi \in \widehat{G_\varphi}} \overline{\psi(g)} \sum_{n=1}^{N} \Lambda(n)\, \psi_{\text{prim}}(n)\, \max\Big(\frac{a}{n} - w\big(\text{partner}(\tau, g)\big),\, 0\Big)\,.$$

Swapping the ψ-sum inside splits every prime into three exact cases (Table 6):

Prime Type in K	When is $\text{Frob}_p^i \notin \Phi$?	Orthogonality Factor	Exact Contribution to Sum
Split (incl. $p \mid 2Q_S q$)	Never ($\text{Frob}_p \in \Phi$)	0	Vanishes identically
Inert	Precisely for **odd** i	$\lvert G_\varphi\rvert/\lvert\Phi\rvert = 2$	$2\log p \cdot \max(a/p^i - w(\text{partner}),\, 0)$
Ramified ($p = d$)	Inertia $I_d \not\subset \Phi$ ($\lvert I_d\rvert = 2$)	$2/\lvert I_d\rvert = 1$	$\log d \cdot \text{askD}$

Table 6: Why restricting Fourier inversion on G_φ to $g \notin \Phi$ reproduces the inert odd-power sum ASK identically.

Thus ASK equals the full character expansion *identically*. Rewriting each inner $\Lambda(n)\psi_{\text{prim}}(n)$ ramp in terms of $w \cdot G_\psi(a/w)$ plus its main term—and noting that $\sum_{g \notin \Phi} \overline{\psi(g)}$ equals $+|\Phi|$ for $\psi = 1$, $-|\Phi|$ for $\psi = \chi_K$, and 0 otherwise—cancels `MAIN` (matching *both* γ from $\psi = 1$ and $L'(1,\chi_K)/L(1,\chi_K)$ from $\psi = \chi_K$!) and leaves a weighted combination of $G_\psi(a/w)$. Finally, averaging over translates $u_0 \in A_0$ kills characters nontrivial on A_0, isolating our target average $G_{\text{avg}}(x)$ plus the (q,d)-twisted characters $\xi\chi_q^i\chi_d^j$ ($(i,j) \neq (0,0)$), which are already bounded by $O(x^{1/2+\varepsilon/2})$ by Section 4.2.

7 Taming Internal Cohomology: A Quantitative Weakening of Vandiver's Conjecture

SECTION AT A GLANCE

Goal: Unconditionally bound the untwistable cyclotomic cohomology groups $\text{Sha}^1(G_{E,S_\ell}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e))$ that arise during local lattice gluing.
Input → Output: Uses Galois representations and definite base change (§8) to prove the sub-exponential window bound `VWin`, which converts `LinkN` into `DIVinN` in Bost's assembly (§5).
Core Idea: Partition (ℓ, e, ψ') into three regimes and six cases; in the generic bulk, trap the ℓ-adic lengths for *all* primes simultaneously inside the prime factorization of one nonzero integer $NN \in \mathbb{Z} \setminus \{0\}$.

There is one stubborn obstruction in the local lattice gluing of Section 6.2 (`LinkN`). When selecting the best local lattice across auxiliary Hecke characters at a twisted prime of residue characteristic ℓ, the local congruence defect (`etaSum`) is governed by the cohomology of the rank-$(r+t)$ Galois representation attached to the theta lift. Modulo the congruence ideal, that representation breaks

into a rank-r cuspidal block ρ plus t one-dimensional characters $\chi_1, \dots, \chi_t$ whose pairwise ratios are powers of the cyclotomic character:

$$\chi_1^{-1}\chi_c \;=\; \chi_{\mathrm{cyc}}^{-e} \qquad (1 \le e < t).$$

INTUITION & ANALOGY **Why Twisting Cannot Silence Internal Echoes—and What Sha^1 Measures**

Imagine tuning a radio to eliminate interference. Interactions between the cuspidal block ρ and a character χ_c can be shifted away from resonance by twisting globally by an auxiliary character ψ. However, twisting shifts *both* χ_1 and χ_c by ψ, leaving their ratio completely unchanged:

$$(\chi_1\psi)^{-1}(\chi_c\psi) \;=\; \chi_1^{-1}\chi_c \;=\; \chi_{\mathrm{cyc}}^{-e}.$$

What lives in $\mathrm{Sha}^1\big(G_{E,S_\ell}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e)\big)$? Think of an **Escher staircase or Penrose triangle**: a geometric structure that looks completely flat and trivial in every small local neighborhood (at every prime above ℓ), yet fails to glue together globally! Every such locally trivial global cohomology class survives every twist, so the formalization must bound the size of these groups unconditionally across all primes ℓ and all $1 \le e < t \le x^{400}$.

EXPLAINER **Vandiver's Conjecture and the Quantitative Weakening Formalized Here**

In the 1840s, Ernst Kummer discovered that Fermat's Last Theorem for an odd prime ℓ depends on whether ℓ divides the class number $h(\mathbb{Q}(\zeta_\ell))$, which measures the failure of unique prime factorization in the cyclotomic ring $\mathbb{Z}[\zeta_\ell]$. That class number factors into a "minus part" h^- (governed by Bernoulli numbers) and a mysterious "plus part" $h^+ = h(\mathbb{Q}(\zeta_\ell)^+)$. **Vandiver's Conjecture** (first conjectured by Kummer) predicts that $\ell \nmid h^+$ for every prime ℓ. Its higher K-theory / Soulé generalization predicts vanishing of the plus-part of the higher Shafarevich–Tate group over $\mathbb{Q}(\zeta_\ell)$ (whose trivial-character eigenspace over $\mathbb{Q}$ is $\mathrm{Sha}^1(G_{\mathbb{Q},\{\ell\}}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e))$ for even $e \ge 2$).
Full Vandiver's conjecture remains wide open. What the formalization proves unconditionally (`VWin`, `VWinAt ψ 401`) is a **quantitative weakening**: on a polynomial window $[Y, Y^2]$, the logarithmic product of orders across all primes ℓ and all $1 \le e < \lfloor x^{401} \rfloor$ is sub-exponential in x:

$$\sum_{\ell\ \text{prime}} (\log \ell) \sum_{e=1}^{\lfloor x^{401} \rfloor - 1} \mathrm{len}_{\mathbb{Z}_\ell}\, \mathrm{Sha}^1\big(G_{E,S_\ell}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e)\big) \;\le\; \exp\big(\tfrac{1}{2}x^{\varepsilon}\big)\,.$$

Here $401 = 400 + 1$: Section 4.2 gives $t \le x^{400}$, and the rescaling $x = \alpha x'$ in `WinPairs.lean` bumps $x^{400} = \alpha^{400}(x')^{400} \le (x')^{401}$.

7.1 How `VWin` Is Proved: Three Regimes and Six Case Predicates

Decomposing the cohomology of the abelian CM field E into character eigenspaces $\rho(\psi', e)$ separates the sum into three regimes:

1. **The $\ell = 2$ part and fixed bad odd primes $\ell \in B_\psi$ (`ThmBab`):** The **Ferrero–Washington theorem ($\mu = 0$)** [11] (proved inside the repository) enters via a bounded ℓ-rank bound on

Iwasawa-module generators, showing there are at most R "exceptional" twists in $[1, X^{401}]$. A **pigeonhole windowing argument** with exponent $J = 2^{R+2}(802C_K/\varepsilon)^R$ finds a clean sub-window $[Y, Y^2] \subseteq [X^{1/J}, X]$ containing no large exceptional twists (Figure 7), keeping the general bound `LemmaK` below $\exp(\frac{1}{2}x^{\varepsilon})$.

2. **The odd-ℓ minus part** Sha^- **(`ThmBcFree`):** Bounded pointwise (no window needed) by a polynomial bound (assembled into $\exp(8x^{\varepsilon/4})$ at the host scale) using generalized Bernoulli numbers (when $\ell \nmid [E : \mathbb{Q}]$) and Stickelberger annihilation (when $\ell \mid [E : \mathbb{Q}]$).

3. **The odd-ℓ plus part** Sha^+ **(`PointHV`):** Partitioned exhaustively into **six case predicates** in `Cases.lean` (with `CaseV` split into two count bounds `CountV` and `CountVs`, yielding seven count labels), summarized in Table 7.

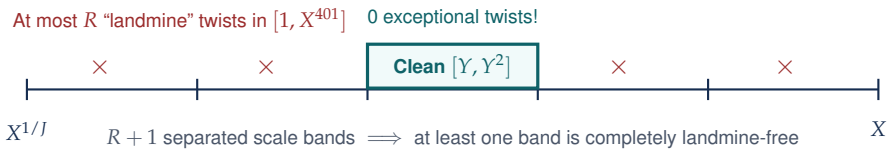


Figure 7: Pigeonhole windowing (the origin of the power windows $[Y, Y^2]$, later rescaled to $[Z, Z^{3/2}]$): dividing $[X^{1/J}, X]$ into $R+1$ well-separated scale bands when at most R exceptional twists exist guarantees a clean window $[Y, Y^2]$.

Case	Condition on (ℓ, e, ψ')	Feature / Obstruction	Proof Tool Used
`CaseI`	$\ell \notin B_\psi,\ \ell - 1 \nmid 2e$	Generic bulk	Definite $U(3)$ patching, Kato's Euler system, global integer NN (`PropGK`/`CorGK`)
`CaseZ`	$\ell \notin B_\psi,\ \ell - 1 \mid 2e,\ \psi' = 1$	Trivial character	Proves exact vanishing $\mathrm{Sha}^1 = 0$ via Bernoulli-unit coboundaries
`CaseII`	$e \geq 2$, even $\mathrm{ord}(\chi_P)$ (for $\ell \notin B_\psi$ with $\ell - 1 \mid 2e$, or $\ell \in B_\psi$)	$\chi_{\mathrm{cyc}}^{2e} \equiv 1$ collides Ψ, Ψ^{-1} in $U(3)$	Switches to 4-dimensional symplectic (GSp_4) deformations & automorphy
`CaseIII`	$\ell \notin B_\psi,\ \ell - 1 \mid 2e,\ \psi' \neq 1$, odd $\mathrm{ord}(\psi')$	Same $U(3)$ diagonal collision	Switches to GSp_4 deformations & automorphy
`CaseIV`	$\ell \in B_\psi,\ e \geq 2$, odd $\mathrm{ord}(\chi_P)$	Fixed bad odd prime	Reduces to fixed-prime Iwasawa layer bound from `ThmBab`
`CaseV`	$\ell \in B_\psi,\ e = 1$	Split into `CountV` / `CountVs`	Large 2-power conductor: $U(3)$ over $\mathbb{Q}(i)$; bounded conductor: `LemmaK`

Table 7: Exhaustive six-case partition of the odd-ℓ plus-part cohomology Sha^+ in `Cases.lean`.

INTUITION & ANALOGY **Trapping Infinitely Many Primes Inside One Concrete Integer NN**

How can you bound a sum over *all* primes ℓ at once without checking each prime individually? Use the oldest theorem in number theory—unique prime factorization! Any nonzero integer $NN = \pm\prod_\ell \ell^{v_\ell(NN)}$ satisfies

$$\sum_{\ell\,\text{prime}} (\log \ell)\, v_\ell(NN) \;=\; \log|NN|.$$

So if you can engineer **one single nonzero integer** $NN \in \mathbb{Z}\setminus\{0\}$ whose ℓ-adic valuation $v_\ell(NN)$ is at least half the cohomology length at every prime $\ell \neq p$, then $\log|NN|$ automatically bounds the *entire* infinite sum over all $\ell \neq p$ in a single stroke! Repeating at a second auxiliary prime $p' \neq p$ catches $\ell = p$ too.

PROOF SKETCH & KEY STEPS **Inside `CaseI` (`PropGK` $\to$ `CorGK`): Building the Global Congruence Integer NN**

1. **Build a 3×3 deformation ring:** Over an imaginary quadratic field $K = \mathbb{Q}(\sqrt{-p_K})$, a nontrivial class in $Ш^1$ produces a reducible extension in the universal conjugate self-dual 3-dimensional deformation ring $S.R$ (`GG 3`, Clozel–Harris–Taylor [13]) of $\Psi \oplus 1 \oplus \Psi^{-1}$ ($\Psi = \chi_{\text{cyc}}^e \psi'$).

2. **Congruences to cuspidal forms:** Commutative algebra (`LemmaM`, `LemmaP`) bounds $\operatorname{len} Ш^1$ by the ℓ-adic valuation of Frobenius trace differences $\operatorname{tr}(\rho_f(\sigma)) - \mathrm{eis}_3$ over characteristic-0 points f of $S.R$ of **Type I** (irreducible 3-dimensional) and **Type II** ($2+1$ reducible, $W \oplus 1$).

3. **Taylor–Wiles patching & Kato's Euler system:** Automorphy lifting on definite $U(3)$ proves every Type I point comes from a cuspidal form on $U(3)$ of weight $(e, 0, -e)$, and every Type II block W comes from a classical weight-$(2e+1)$ modular newform (with odd nebentypus matching K). Adjoint Selmer vanishing gives rigidity at Type I points, while **Kato's Euler system** [12] (`KatoTwoW`) kills the off-diagonal Bloch–Kato Selmer groups $H^1_f(\mathbb{Q}, W)$ and $H^1_f(\mathbb{Q}, W \otimes \theta)$ at Type II points.

4. **Global integer NN and two-prime wrap:** Multiplying $\operatorname{tr}(\rho(\mathrm{Frob}_w^{-1})) - \mathrm{eis}_3$ over all definite $U(3)$ eigensystems, weight-$(2e+1)$ newforms, and the Galois orbit of ψ' at an auxiliary unramified prime $p \geq 7$ yields $NN \in \mathbb{Z}\setminus\{0\}$ with

$$\log|NN| \;\leq\; C\varphi(H.\mathrm{index})(N(e+1))^A \log p, \qquad \sum_{\psi'} V_n(\psi') \;\leq\; 2v_\ell(NN) \quad (\forall \ell \neq p).$$

   Applying `PropGK` at a second auxiliary prime $p' \neq p$ (`CorGK`) covers every prime ℓ.

7.2 A Standalone Corollary in the Repository: `SmallClassNumbers`

The module `SmallClassNumbers/CorGKOf.lean` (compiled via `MainTheoremsAll.lean`) is **not used** in the proof of Theorems A–C. Instead, it specializes `U3.CorGK` to $K = \mathbb{Q}(\sqrt{-3})$ (class number 1), $N = 2$, $\psi = 1$, and $p = 7, p' = 11$ to obtain an independent corollary over $\mathbb{Q}$ (matching footnote 6 of `remark.pdf`):

Corollary 7.1 (Standalone `SmallClassNumbers.bound` & `large_of`)**.** *There exist absolute constants $C, A \in \mathbb{N}$ such that for every even integer $e \geq 2$, every choice of finite $\mathbb{Z}_\ell$-DVR coefficient rings $\mathcal{O}_\ell$ (with ramification index $e(\mathcal{O}_\ell/\mathbb{Z}_\ell)$), and every finite set L_s of primes $\ell \geq 5$ with $\ell - 1 \nmid 2e$,*

$$\sum_{\ell \in L_s} (\log \ell)\, \frac{\mathrm{len}_{\mathcal{O}_\ell} \mathrm{Sha}^1\big(G_{\mathbb{Q},\{2,\ell\}},\, (\mathbb{Q}_\ell/\mathbb{Z}_\ell \otimes_{\mathbb{Z}_\ell} \mathcal{O}_\ell)(-e)\big)}{e(\mathcal{O}_\ell/\mathbb{Z}_\ell)} \;\leq\; C(e+1)^A.$$

In particular (`large_of`), when $e(\mathcal{O}_\ell/\mathbb{Z}_\ell) = 1$ (e.g. $\mathcal{O}_\ell = \mathbb{Z}_\ell$) and $\log \ell > C(e+1)^A$, the group $\mathrm{Sha}^1(G_{\mathbb{Q},\{2,\ell\}}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e))$ vanishes identically.

Context: By Poitou–Tate duality and the Quillen–Lichtenbaum / Rost–Voevodsky isomorphism, for odd ℓ and even $e \geq 2$, $\mathrm{Sha}^1(G_{\mathbb{Q},\{2,\ell\}}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e))$ is dual to a subgroup of $H^2_{\text{ét}}(\mathbb{Z}[\frac{1}{2\ell}], \mathbb{Z}_\ell(e+1)) \cong K_{2e}(\mathbb{Z}) \otimes \mathbb{Z}_\ell$. Since $K_{2e}(\mathbb{Z})$ is finite (Quillen 1973; Borel 1974), qualitative vanishing for sufficiently large ℓ already holds; the new quantitative content of Corollary 7.1 is the **effective polynomial threshold** $\log \ell \leq C(e+1)^A$ and the uniform weighted sum bound across primes.

8 Building the Bridges Beyond Mathlib: Galois Representations, Tailored LLC, Base Change & Arakelov Covolumes

SECTION AT A GLANCE

Goal: Explain how the formalization builds its own deep arithmetic-geometric foundations on top of Mathlib without custom axioms—and the clever shortcuts that keep them tractable.
Output → §6–§7: Supplies Galois representations (`RowSrc`), tailored Local Langlands (`RecTExists`), definite unitary base change (`DefiniteNHolds`), and base-level covolumes (`CovolX`).
Core Idea: Lacunary weights avoid multiplicity-one, "Two-Up" avoids even-rank Shimura varieties, and a wildly ramified placeholder avoids full Rankin–Selberg ε-factors.

In a conventional paper, an author can cite landmark theorems from the last 70 years—Harris–Taylor, Clozel, Caraiani, Bernstein–Zelevinsky, Hironaka, Faltings–Chai, Gillet–Soulé, Nadel, Ohsawa–Takegoshi. A Lean 4 formalization with zero custom axioms must prove every single result it invokes from scratch on top of Mathlib, which explains why the repository spans 175,187 files. Inspecting the source reveals several remarkably clever engineering shortcuts.

EXPLAINER **Galois Representations and the Local Langlands "Rosetta Stone"**

- **Galois representations (symmetries as matrices):** Just as rotations of a cube can be written as 3×3 matrices, the hidden symmetries of algebraic numbers over E (the absolute Galois group $G_E = \mathrm{Gal}(\overline{\mathbb{Q}}/E)$) can be represented as $a \times a$ matrices $r : G_E \to \mathrm{GL}_a(\overline{\mathbb{Q}}_\ell)$. The **Langlands program** acts as a grand Rosetta Stone matching automorphic forms π on $\mathrm{GL}_a(\mathbb{A}_E)$ with a-dimensional Galois representations r_π so that at every prime w the Frobenius eigenvalues of r_π match the harmonic frequencies (Satake parameters) of π_w (**local-global compatibility**).
- **Local Langlands Correspondence (LLC):** Zooming into a single prime w (the p-adic completion E_w), LLC translates irreducible smooth representations of the p-adic matrix group $\mathrm{GL}_a(E_w)$ into a-dimensional **Weil–Deligne representations** of the local Galois/Weil group W_{E_w}, preserving dimensions, twists, central characters, and L-factors.

8.1 Attaching Galois Representations: Odd Ranks, Lacunary Weights, and the "Two-Up" Trick

The top-level file `MainConclusion.lean` reduces the entire automorphic-to-Galois engine to `RowSrc.lean` (which proves `Row0b recT` from the stronger `Row0 recT` without the $g > 16a^2$ restriction at depth-zero places $E_w \cong \mathbb{Q}_g$). It constructs Galois representations for all ranks $a \geq 1$ in two stages:

1. **Odd ranks a via Holomorphic Lefschetz & Lacunary Weights:** For odd a, unitary Shimura varieties Z of signature $(1, a-1)$ exist without endoscopic sign obstructions. Applying the **holomorphic Lefschetz fixed-point formula** to finite étale Hecke correspondences on Z produces a geometric representation ρ of some larger dimension dd carrying the Hodge–Tate weights of U. To isolate the a-dimensional piece *without knowing multiplicity one in cohomology*, the proof chooses exponentially spaced **lacunary Hodge–Tate weights** satisfying $j_{i+1} - j_0 > a(j_i - j_0)$, then passes to all odd-rank definite unitary forms via Chenevier pseudocharacter limits.
2. **Even ranks n via Theta-Parseval (`TLiftK`) and "Two-Up" Character Subtraction:** For even n, the formalization avoids even-rank unitary Shimura varieties altogether via a double-lift trick (Figure 8).

WORKED EXAMPLE **Why Lacunary Weights Isolate Exact Dimension a: The $\{1, 10, 100\}$ Trick**

Suppose a large matrix representation splits into unknown pieces, and you want to be sure that no two disjoint subsets of its weights can accidentally share the same average. Pick exponentially spaced weights like $\{1, 10, 100\}$:

- Any nonempty subset containing 100 has average $\geq \frac{100+10+1}{3} = 37$, whereas any subset without 100 has average ≤ 10!
- Among subsets without 100, any containing 10 has average $\geq \frac{10+1}{2} = 5.5$, whereas $\{1\}$ has average 1.

With the lacunary gap $j_{i+1} - j_0 > a(j_i - j_0)$, the largest weight in $A \cup B$ alone overwhelms every smaller element even when weighted by $|A| \leq a$. Hence no two disjoint subsets can share an average, locking the Satake-bearing constituent to exact dimension a without ever needing multiplicity one!

INTUITION & ANALOGY **Weighing an Elephant with Two Disjoint Counterweights (The "Two-Up" Trick)**

Suppose your scale (unitary Shimura varieties) only accepts *odd* ranks, and you need to weigh an *even*-rank representation U on $U(n)$:

1. Choose an odd jump $t = 2(d_0 + n) + 1$ so that $n + t$ is **odd**.
2. Pick **two disjoint** blocks of t algebraic 1-dimensional Hecke characters: $\psi_1 = (\psi_{1,1}, \ldots, \psi_{1,t})$ and $\psi_2 = (\psi_{2,1}, \ldots, \psi_{2,t})$ ($\psi_{1,i} \neq \psi_{2,j}$ for all i, j).
3. Apply the global definite unitary theta lift `TLiftK` (proved nonvanishing via Rallis + Siegel–Weil + compact Schur orthogonality + height-tail domination) twice, yielding two odd-rank forms U_1, U_2 on $U(n+t)$ with Satake multisets $\kappa^{-1}U + \psi_1$ and $\kappa^{-1}U + \psi_2$.
4. Attach odd-rank $(n+t)$-dimensional Galois representations r_1, r_2 to U_1, U_2. By Chebotarev density,

$$\mathrm{cp}(r_1, g) \prod_{j=1}^{t} (X - a_{2,j}(g)) \;=\; \mathrm{cp}(r_2, g) \prod_{i=1}^{t} (X - a_{1,i}(g)) \qquad (\forall g \in G_E).$$

   Because a_1 and a_2 share no characters, `PeelAll` strips all t copies of $a_{1,i}$ from r_1, leaving the exact n-dimensional Galois representation $r = s \otimes L(\kappa)$ for U.

8.2 Two Architectural Shortcuts: Tailored LLC and Definite Base Change

1. Tailored Local Langlands (`RecTExists.lean` and `JunkEx.lean`)

Formalizing Harris–Taylor's full Local Langlands correspondence for wildly ramified supercuspidals of GL_n would require full Rankin–Selberg ε-factors—like writing a complete dictionary for every dead dialect on Earth when your novel only uses three modern languages! Because the

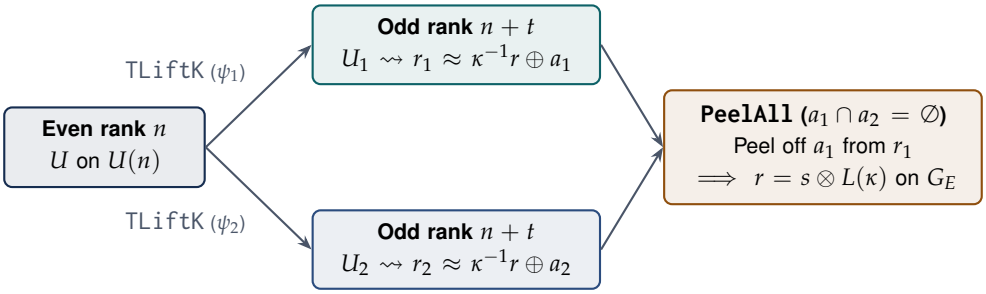


Figure 8: The "Two-Up" character subtraction trick: lifting even rank n twice to odd rank $n+t$ with disjoint character blocks ψ_1, ψ_2 and peeling off the extra t characters via Chebotarev.

GRH proof only invokes local-global compatibility at spherical, unlinked-toral, and depth-zero-over-$\mathbb{Q}_g$ places (called `Good` representations), `RecTExists.lean` builds a **tailored map** rather than a full bijection:

Supercuspidal Type	Construction in `RecTExists`	Guarantee
`Good` (rank 1, or twists of depth-zero over $\mathbb{Q}_g$)	Bushnell–Kutzko level-zero types, Green's regular characters of $\mathrm{GL}_n(\mathbb{F}_g)$, tame Weil induction	Exact local-global compatibility and bijection on `Good` types
Non-`Good` (rank $n \geq 2$)	Mapped to a single wildly ramified placeholder $J_n \otimes \lvert\cdot\rvert^{e(x)}$ (`JunkEx.lean`)	Never collides with any twist of a tame parameter; preserves injectivity whenever the second input is `Good`

Bernstein–Zelevinsky multisegments (`exists_isLExt`) then extend this map uniquely to all irreducible smooth representations of $\mathrm{GL}_n(E_w)$.

2. Definite Unitary Base Change (`DefiniteNHolds.lean` & `Narrow.lean`)

Base change from definite $U(n)$ to $\mathrm{GL}_n(\mathbb{A}_E)$ is proved directly from the simple unitary trace formula (`USideFree.lean`), Hecke-word frames, and Gelfand transforms. Then `BaseChangeIdentities/Narrow.lean` (78 lines total, with a ∼25-line core) upgrades local subquotient matching to full isomorphism using Bernstein's theorem that parabolic induction of irreducible unitarizable representations of GL_n over a characteristic-0 local field is irreducible (proved inside the repository).

8.3 The 16-Step Arakelov Covolume Cascade (`CovolX`)

Why does bounding the base-level lattice volume $\mathrm{lcv}(X_{W,j}, M_j)/D$ inside `LinkDIV` (Section 6) require so much algebraic geometry? Intuitively, measuring the arithmetic volume of a space of automorphic forms on a curved moduli space takes four geometric moves: (1) smooth out sharp singularities so calculus works, (2) control p-adic integral models at bad primes, (3) seal off the boundary at infinity cleanly, and (4) produce enough analytic sections with controlled L^2-norms to bound the volume! Table 8 summarizes the 16 formalized theorems chained across these four layers.

Layer	Formalized Theorems Chained in `CovolX`
1. Resolution & Arithmetic Intersection	Hironaka's resolution of singularities and ideal principalization in char 0 (499 files) + Poincaré–Lelong symmetry $\Longrightarrow$ Bost–Gillet–Soulé arithmetic intersection theory, arithmetic Hilbert–Samuel, and Moriwaki arithmetic bigness
2. p-adic Hodge & Deformation Theory	Zarhin's trick + Fargues's Hodge–Tate exact sequence + Grothendieck semistable reduction + prismatic Dieudonné–Hodge + Grothendieck–Messing deformation theory
3. Shimura Compactifications & Metrics	Faltings–Chai integral toroidal compactifications of Siegel and unitary Shimura varieties + Mumford's good Chern–Weil metrics
4. Complex Analytic Extension & Vanishing	Ohsawa–Takegoshi L^2 extension (explicit constant $34/\delta \cdot 2^{r+1}\pi^r/(r-1)!$) + Nadel and Kawamata–Viehweg vanishing + Angehrn–Siu effective point separation

Table 8: The four layers of the 16-step Arakelov covolume cascade (`CovolX`).

In the production wiring (`EveryCharacterOfLiteratureOfVWin.lean` $\rightarrow$ `SiegelEisensteinSecondTermHypotheses.lean`), the formalization trims the earlier prototype `EveryCharacterE.lean` by dropping roughly 20 heavy literature hypotheses (including `ScholzeWeinsteinB`, `TamagawaNumberSU`, `Labesse-Waldspurger`, `MinguezThetaGL`, and `PDivisibleGroupsOverOC`) that proved unnecessary along the main path.

9 Behind the 17.37 Million Lines: Static Audit, Verification Checklist, and Perspective

SECTION AT A GLANCE

Goal: Report reproducible static metrics on `lean.tar.gz`, visualize how the 175,187 files are distributed, and specify the exact machine and human verification steps still required.
Status: 0 `sorry`, 0 custom `axiom`, and 0 `native_decide` found in static text scanning; full kernel compilation (`lake build Main` + `lean4checker`) remains to be run on a large machine.

9.1 Reproducible Static Source Scan of `lean.tar.gz`

A direct static scan of `lean.tar.gz` (239,275,554 bytes, SHA-256 `9631b4f020199e580ba86178cb56c6e9b09f690c2369a80dc4334aa93f49c7fc`) gives the measurements in Table 9 (all counts come from static text parsing without compiling Lean):

Figure 9 shows how the 175,187 files are distributed across the codebase—making immediately visible that `Development/AbelianVarietyTheory` alone accounts for more than half the archive (99,187 files, 56.6%).

Scan methodology and metaprogramming notes:

Metric	Static-Scan Value (Reproducible Methodology)
Lean 4 toolchain & Mathlib rev	`leanprover/lean4:v4.33.1` Mathlib `0df444a360eaa60ab8c11dca51a86af692955474`
Total / Import-reachable files	**175,187 total** | **175,187 import-reachable from `Main.lean`** (100%, Unicode-aware BFS; 0 unreachable)
Total source lines / size	**17,374,215 lines** | 1,030,531,854 bytes (982.8 MiB / 1.03 GB)
Maximum import-DAG BFS depth	**43 layers** from `Main.lean`
Line-start declarations	**778,887 `theorem`** + **34,307 `lemma`** = **813,194** (plus 211,896 `def`, 8,230 `structure`, 13,202 `abbrev`, 8,199 `instance`, 121 `inductive`, 43 `class`)
Static trust scan (comment-stripped)	**0 `sorry`**, **0 custom `axiom`**, **0 `native_decide`**

Table 9: Summary metrics from static text inspection of the 175,187-file Lean 4 repository.

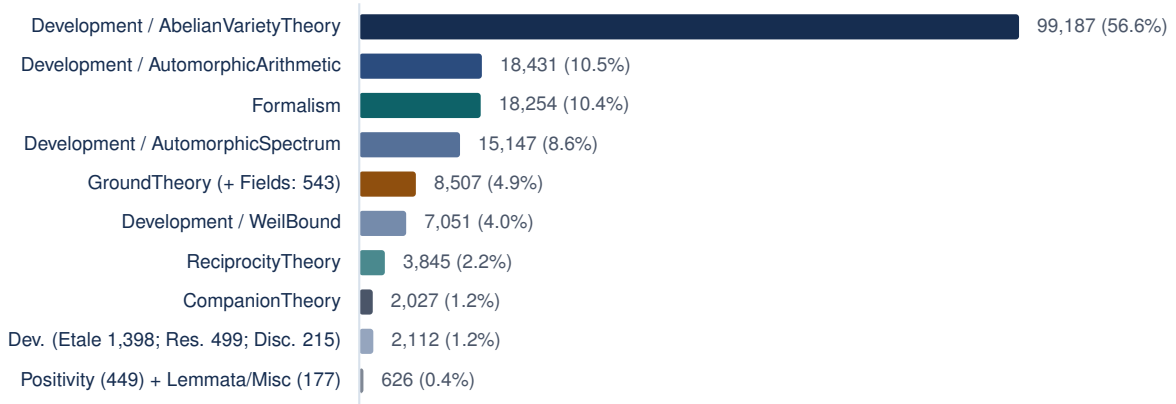


Figure 9: Distribution of all 175,187 `.lean` files across top-level libraries (and `Development` sub-libraries). Here `GroundTheory` (8,507) splits in `lakefile.toml` into `GroundTheory` (7,964) and `GroundTheoryFields` (543: `GlobalFields` 384, `LocalFields` 159); the 177 miscellaneous files are `Lemmata` (167), `BuildAid` (5), `Supplements` (3), `Statements.lean` (1), and `Main.lean` (1).

- **Declaration regex:** Counts match line-start keywords preceded optionally by modifiers (`private`, `protected`, `noncomputable`, `unsafe`, `partial`, `scoped`, `local`, `nonrec`); counts rise slightly (e.g. ∼784k `theorem`) if inline `@[...]` attribute-prefixed declarations are also included.
- **Trust-scan patterns:** After stripping `-- ...` line comments and `/- ... -/` block comments across all 175,187 files, we searched for `\bsorry\b`, line-start `axiom`, `\bnative_decide\b`, `\baddDecl\b`, and `\bmodifyEnv\b`.
- **Two benign metaprogramming locations found:** (i) `CutCptCopy.lean` (under `Formalism/HermitianCurvature/`) programmatically duplicates definitions and theorems via `addDecl` (every declaration added via `addDecl` is still type-checked by the Lean kernel); and (ii) `BuildAid/D1.lean–D5.lean` invoke `run_cmd Lean.modifyEnv` solely to append library names to `Lean.LibrarySuggestions.nameDenyListExt`.

9.2 What Machine and Human Verification Still Require

1. **Required machine-verification protocol (not yet run here):** Static inspection alone cannot establish that a Lean 4 repository compiles or is sound. Before the claimed proof can be regarded as machine-checked, three commands must be completed on a machine capable of building 175,187 files:

```
1. lake exe cache get && lake build Main    (under leanprover/lean4:v4.33.1)
2. #print axioms    (on all five theorems in Main.lean: only propext, Classical.choice, Quot.sound)
3. lean4checker    (independent kernel re-check of the compiled environment)
```

2. **What human mathematical review still checks:** Even after full kernel compilation, human reviewers must confirm: (i) that Mathlib's definitions of `riemannZeta` and `DirichletCharacter.LFunction` faithfully match classical definitions; (ii) that the four custom `Prop` definitions in `Statements.lean` (quoted verbatim in Section 1.1) accurately express RH and GRH; and (iii) that the high-level mathematical argument can be extracted, verified conceptually, and simplified. Note also that several literature-statement stubs exist in the repository without proofs (`LabesseWaldspurger`, `MinguezThetaGL`, `TamagawaNumberSU`, `PDivisibleGroupsOverOC`); these appear only as hypothesis types in off-path modules and are not invoked along the production path to `Main.lean`.

3. **Prospects for simplification:** In `remark.pdf` [1], Alpöge writes that he *"may have baked lots of strange suboptimal choices into the argument, and expect[s] short arguments to follow fairly quickly"* (p. 3), and remarks in footnote 7 on Behrend's trick that *"this is the sort of trick that intuitively to me is obviously going to be unnecessary."* Several simplifications already present in the codebase (`zeroCh`, `JunkEx`, `TheoremFSharp` without $-\log|\mathcal{W}|$, and dropping ∼20 literature inputs in `EveryCharacterOfLiteratureOfVWin.lean`) point in the same direction.

Glossary of Key Terms

Arakelov Geometry & Hermitian Vector Bundle
A framework pairing $\mathcal{O}_K$-lattices with Hermitian inner products at complex embeddings $K \hookrightarrow \mathbb{C}$, measuring arithmetic size (covolume/height) across finite and infinite places simultaneously.

Automorphic Form / Representation
A symmetric function (or representation) on an adelic group such as $\mathrm{GL}_n(\mathbb{A})$ or $U(n)(\mathbb{A})$ generalizing classical modular forms. On a *totally definite* unitary group, automorphic forms live on a finite set of lattice classes.

Behrend Set
A dense subset of $\{1, \ldots, \lfloor n_0/(4L_w)\rfloor\}$ containing no short affine relation $j(s_2 - s_1) = i(s_3 - s_1)$ $(i \neq j)$, built from lattice points on a high-dimensional Euclidean sphere encoded in base B.

Bessel / Gross–Prasad (GGP) Period
A period integral of an automorphic form over a subgroup paired against a character; here it provides the pairing B_u on $U(X_0) \times U(Y_0)$ independent of auxiliary Hecke character twists $\mu_1, \ldots, \mu_m$.

Bost's Slope Inequality
An Arakelov bound comparing Harder–Narasimhan slopes of two Hermitian vector bundles linked by a linear map φ: $\mu_{\min}(\overline{A}) \leq \mu_{\max}(\overline{B}) + h(\varphi)$.

Chowla–Selberg Formula
Relates periods of CM abelian varieties to Gamma values and $L'(1, \chi_K)/L(1, \chi_K)$, supplying the exact constants γ and $L'(1, \chi_K)/L(1, \chi_K)$ in `MainC`.

CM Field (Complex Multiplication Field)
A totally imaginary quadratic extension E of a totally real number field F, carrying a canonical complex conjugation across all embeddings into $\mathbb{C}$.

Dirichlet Character χ & L-Function $L(s, \chi)$
A periodic multiplicative homomorphism $\chi : (\mathbb{Z}/N\mathbb{Z})^\times \to \mathbb{C}^\times$ and its Dirichlet series $L(s, \chi) = \sum_{n \geq 1} \chi(n) n^{-s}$.

Dual Weyl Module & Fock Space
Integral polynomial representations $\nabla_{\mathbb{Z}}(\nu) \subset \mathbb{Z}[X_{a,s}]$ spanned by semistandard bideterminants, equipped with the Fock/Fischer inner product $\langle X^m, X^m \rangle = \prod_k (m_k)!$.

Dwork Integrality Criterion
A p-adic test characterizing when $\exp(\Psi(X))$ has integral coefficients ($p \mid \sigma(\Psi)(X^p) - p\Psi(X)$); combined with $D^p \equiv 0 \pmod{p}$, it yields the $a/p + a/p^3 + \ldots$ gains at primes inert in $\mathbb{Q}(\sqrt{-d})$.

Explicit Formula
A duality relating sums over prime powers (such as $\psi(x)$ or $G_z(x)$) to sums over nontrivial zeros ρ of $\zeta(s)$ or $L(s, \chi)$.

Harder–Narasimhan (HN) Filtration & Slopes
The canonical filtration of a Hermitian vector bundle from densest sublattice ($\mu_{\max}$) to sparsest quotient ($\mu_{\min}$), with rank-weighted mean slope $\mu = -\operatorname{lcv} / \operatorname{rk}_K$.

Iwasawa Tower
An ascending tower of number fields with Galois group $\mathbb{Z}_p$ (here the 2-adic cyclotomic layers $\mathbb{Q}(\zeta_{2^{e_{\mathrm{Iw}}+2}})^+$ of degree $2^{e_{\mathrm{Iw}}}$).

Kato's Euler System
A compatible system of Galois cohomology classes built from modular units, used in Section 7 to kill off-diagonal Bloch–Kato Selmer groups at Type II points of $U(3)$ deformation rings.

Local Langlands Correspondence (LLC)
A dictionary matching irreducible representations of $\mathrm{GL}_n(E_w)$ over a p-adic field E_w with n-dimensional Weil–Deligne representations of W_{E_w}.

Rallis Inner Product Formula
Expresses the L^2-norm of a theta lift in terms of special values of automorphic L-functions; dividing out these eigenvalues cancels the bulk archimedean lift cost.

Relative Trace Formula (RTF)
Compares automorphic period integrals against geometric orbital integrals; used here to establish approximate orthogonality (`LinkS`) of column spaces.

Seesaw Identity
An adjointness identity for two dual pairs (G_1, H_1) and (G_2, H_2) with $G_2 \subset G_1$ and $H_1 \subset H_2$, swapping a theta lift from G_2 to H_2 (restricted to H_1) with the opposite lift from H_1 to G_1 (restricted to G_2).

Shafarevich–Tate Group Sha^1 & Vandiver's Conjecture
$\mathrm{Sha}^1(G_{E,S_\ell}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e))$ consists of global Galois cohomology classes unramified outside S_ℓ and locally trivial at places above ℓ. Vandiver's conjecture predicts $\ell \nmid h(\mathbb{Q}(\zeta_\ell)^+)$; this formalization proves an unconditional sub-exponential window bound (`VWin`).

Siegel–Weil Formula
Equates an integral of a theta function with a special value of an Eisenstein series, yielding the unitary mass ratios used to cancel $\log T_u$.

Stepanov's Method
An elementary proof of the finite-field Riemann Hypothesis (initially for hyperelliptic curves, extended by Bombieri and Schmidt) using an auxiliary function vanishing to high multiplicity $m \approx \sqrt{q}$ at $\mathbb{F}_q$-points.

Theta Lift
An integral transform via the Weil representation kernel $\theta(g, h)$ transferring automorphic forms between members of a dual pair.

von Mangoldt Function $\Lambda(n)$
Equals $\log p$ if $n = p^j$ is a prime power and 0 otherwise.

Key References & Further Reading

Starting Points for Non-Specialist Readers

- **Riemann Hypothesis & Prime Waves:** B. Mazur & W. Stein, *Prime Numbers and the Riemann Hypothesis* (Cambridge, 2016) — visually rich, requiring only calculus; E. Bombieri's official Clay Mathematics Institute problem description (2000).
- **Stepanov's Method:** S. A. Stepanov's original approach is beautifully exposed in E. Bombieri [6] and in T. Tao's expository blog notes on Stepanov's method and Behrend's sphere construction.
- **Langlands Program & Galois Representations:** E. Frenkel, *Love and Math* (2013), or the accessible overview articles in T. Gowers (ed.), *The Princeton Companion to Mathematics* (2008).

[1] L. Alpöge (with Claude), *Remark* (`remark.pdf`) and Lean 4 formalization archive (`lean.tar.gz`), https://alpo.ge/, November 20, 2026.

[2] B. Riemann, *Über die Anzahl der Primzahlen unter einer gegebenen Grösse*, Monatsberichte der Berliner Akademie (1859), 671–680; and H. von Koch, *Sur la distribution des nombres premiers*, Acta Math. **24** (1901), 159–182. (Accessible guides: H. M. Edwards, *Riemann's Zeta Function*, Dover, 2001; B. Mazur & W. Stein, *Prime Numbers and the Riemann Hypothesis*, Cambridge Univ. Press, 2016.)

[3] E. Landau, *Über einen Satz von Tschebyschef*, Math. Ann. **61** (1905), 527–550.

[4] H. Davenport, *Multiplicative Number Theory*, 3rd ed., Springer GTM 74, 2000.

[5] S. A. Stepanov, *On the number of points of a hyperelliptic curve over a finite prime field*, Izv. Akad. Nauk SSSR Ser. Mat. **33** (1969), 1171–1181.

[6] E. Bombieri, *Counting points on curves over finite fields (d'après S. A. Stepanov)*, Séminaire Bourbaki, Exp. 430 (1972/73), Lecture Notes in Math. **383**, Springer, 1974; and W. M. Schmidt, *Zur Methode von Stepanov*, Acta Arith. **24** (1973), 347–367.

[7] F. A. Behrend, *On sets of integers which contain no three terms in arithmetical progression*, Proc. Nat. Acad. Sci. USA **32** (1946), 331–332.

[8] J.-B. Bost, *Algebraic leaves of algebraic foliations over number fields*, Publ. Math. IHES **93** (2001), 161–221.

[9] S. S. Kudla and S. Rallis, *A regularized Siegel–Weil formula: the first term identity*, Ann. of Math. **140** (1994), 1–80; and S. Rallis, *Injectivity properties of liftings associated to Weil representations*, Compositio Math. **52** (1984), 139–169.

[10] W. T. Gan, B. H. Gross, and D. Prasad, *Symplectic local root numbers, central critical L-values, and restriction problems in the representation theory of classical groups*, Astérisque **346** (2012), 1–109.

[11] B. Ferrero and L. C. Washington, *The Iwasawa invariant μ_p vanishes for abelian number fields*, Ann. of Math. **109** (1979), 377–395.

[12] K. Kato, *p-adic Hodge theory and values of zeta functions of modular forms*, Astérisque **295** (2004), 117–290.

[13] L. Clozel, M. Harris, and R. Taylor, *Automorphy for some ℓ-adic lifts of automorphic mod ℓ Galois representations*, Publ. Math. IHES **108** (2008), 1–181.

[14] M. Harris and R. Taylor, *The Geometry and Cohomology of Some Simple Shimura Varieties*, Annals of Mathematics Studies **151**, Princeton Univ. Press, 2001; and A. Caraiani, *Local-global compatibility and the action of monodromy on nearby cycles*, Duke Math. J. **161** (2012), 2311–2413.

[15] I. N. Bernstein and A. V. Zelevinsky, *Induced representations of reductive* $\mathfrak{p}$*-adic groups I*, Ann. Sci. École Norm. Sup. **10** (1977), 441–472; and A. V. Zelevinsky, *II*, ibid. **13** (1980), 165–210.

[16] G. Faltings and C.-L. Chai, *Degeneration of Abelian Varieties*, Ergebnisse der Mathematik **22**, Springer, 1990.

[17] H. Hironaka, *Resolution of singularities of an algebraic variety over a field of characteristic zero*, Ann. of Math. **79** (1964), 109–326.

[18] H. Gillet and C. Soulé, *Arithmetic intersection theory*, Publ. Math. IHES **72** (1990), 93–174; and A. Moriwaki, *Arakelov Geometry*, Translations of Mathematical Monographs **244**, AMS, 2014.

Credits & Colophon: Written by **Antigravity with Gemini Neon**, prompted by **Gokhan Mergen** (November 20, 2026), reconstructed from static inspection of Levent Alpöge and Claude's Lean 4 formalization (`lean.tar.gz` and `remark.pdf`). Not written, reviewed, or endorsed by L. Alpöge.

Verification Scope: A static source scan of all **175,187 import-reachable files** finds 0 `sorry`, 0 custom `axiom`, and 0 `native_decide` declarations. Because compiling 175,187 Lean files requires substantial compute, full kernel verification (`lake build Main`, `#print axioms`, `lean4checker`) was not run on this local machine; Section 9 details the scan methodology and full verification steps.