A roadmap of the claimed Lean proof of the Riemann hypothesis and GRH for Dirichlet L-functions

Yan Ru Pei

yanrpei@gmail.com

November 2026

doi:10.5281/zenodo.24405360

Abstract

On 20 November 2026 L. Alpöge released a Lean 4 development, of 175,187 files, that claims a formal proof of the Riemann hypothesis and of the generalized Riemann hypothesis for all Dirichlet L-functions. At the time of writing no independent check of it had been reported. This note is a reading guide. It is based on the development's definitions and theorem statements and on the author's four-page remark; we did not build the development, and we make no claim that the proof is correct. By folder, about 81% of the development's 17.4 million lines is formalized standard background and about 5% holds the core bookkeeping of the argument; we set the background aside, except for the few inputs inside it on which the argument depends, and lay out the argument as a chain of named reductions. In our reading the argument is a version of Stepanov's auxiliary-polynomial method transplanted from function fields to $\mathbb{Z}$. The role of Frobenius is played by primes inert in an abelian CM field, where integral theta-kernel periods are forced to be highly divisible; the role of the auxiliary polynomial's degree bound is played by an averaged Bost slope inequality; and the missing multiplicity is supplied by a chain of about $\sqrt{x}$ theta lifts between totally definite unitary groups. The output is a one-sided bound of square-root quality, on windows of scales, for Riesz means of the von Mangoldt function twisted by characters, which a windowed form of Landau's theorem turns into zero-free half-planes. We describe each step with frozen Lean locators, single out the techniques that look new, list where a reader checking the argument should look first (the forced-divisibility statement, the sharpness of the slope inequality, and the Selmer and covolume bounds), and propose a plan for human readers: which statements to extract, informalize and check, and in what order.

Contents

1 Introduction **2**
1.1 The claim . . . 2
1.2 What this note does . . . 3
1.3 The architecture in brief . . . 4
1.4 Organization . . . 4

2 The analytic half: from a one-sided bound to zero-free half-planes **4**
2.1 Riesz-mean remainders . . . 4
2.2 The endgame is classical . . . 6
2.3 Level sums and positive-definite sampling . . . 6
2.4 A windowed Landau theorem . . . 6

2.5 Scale-adapted auxiliary twists . . . 7
2.6 The balance identity . . . 8
2.7 The interface . . . 8

3 The engine: Stepanov's method over $\mathbb{Z}$ 9
3.1 Stepanov's method . . . 9
3.2 Frobenius: forced divisibility at inert primes . . . 9
3.3 The degree bound: an averaged slope inequality . . . 10
3.4 Multiplicity: a chain of theta lifts . . . 11
3.5 Pigeonhole: an Iwasawa tower and a Behrend-type placement . . . 11

4 Paying for the error terms 12
4.1 Division by twisted lifts and a uniform Selmer bound . . . 12
4.2 The covolume bound . . . 13
4.3 Near-orthogonality from a relative trace formula . . . 13
4.4 Galois representations and local data . . . 13

5 What can be set aside 14
5.1 The discard map . . . 14
5.2 Abandoned designs . . . 14

6 Where to look first 15

7 A plan for human digestion and informalization 16

8 Public commentary and outlook 18

A Frozen Lean locators 18

B The author's remark and the code 21

1 Introduction

1.1 The claim

The development [1] was posted as a fourteen-part archive together with a short informal remark [2]. We read it from a byte-exact public mirror [3] at commit `a172c464` (full hash in [3]). It is pinned to Lean 4.33.1 and Mathlib, has 175,187 Lean files and about 17.4 million lines, and contains no comments or docstrings. Its main file proves five theorems. For a Dirichlet character χ we write $L(\chi, s)$ for Mathlib's `DirichletCharacter.LFunction`.

Theorem 1.1 (claimed in [1, `src/Main.lean`])**.**

(a) `RiemannHypothesis`*, Mathlib's statement: if $\zeta(s) = 0$, s is not of the form $-2(n+1)$ with $n \in \mathbb{N}$, and $s \neq 1$, then $\operatorname{Re} s = \frac{1}{2}$. The same statement is proved a second time from a definition written out in the development.*

(b) For every $N \geq 1$, every Dirichlet character χ modulo N and every $s \in \mathbb{C}$ with $\operatorname{Re} s > \frac{1}{2}$ and ($\chi \neq 1$ or $s \neq 1$), $L(\chi, s) \neq 0$.

(c) For every $N \geq 1$, every χ modulo N and every s with $0 < \operatorname{Re} s < 1$ and $L(\chi, s) = 0$, $\operatorname{Re} s = \frac{1}{2}$.

(d) Statement (c) restricted to primitive χ.

The definitions behind (b)–(d) are short and sit in `src/Statements.lean`; the mirror's maintainers read them by hand [3]. "GRH" here means Dirichlet L-functions only: Dedekind zeta functions and other L-functions are not covered. The remark says the method should extend further; this is not claimed in the code.

Status. On 20 November 2026 the mirror reported that its source is identical to the published archive, that the challenge file states the five theorems against Mathlib alone, and that a static scan finds no `sorry`, `axiom`, `native_decide`, `unsafe`, `implemented_by` or `opaque`. The `lake comparator` check, which compares the statements and checks the proof with several independent kernels, had not yet been run on the submission; the build was expected to take many hours [3]. We did not build or run anything. Everything below is a reading of definitions and theorem statements, and "proved" means only "a theorem with this statement exists in the source".

1.2 What this note does

The size of the development makes it hard to see the argument. Most of it is formalized standard mathematics: about 81% of the lines, by our directory-level inventory (Section 5), re-prove published theory such as étale cohomology, resolution of singularities in characteristic 0, p-adic Hodge theory, arithmetic intersection theory, theta correspondences and automorphy lifting. We set that material aside after naming it, and follow the chain of statements from Theorem 1.1 down to the inputs that carry the argument. For each step we give the mathematical content in ordinary notation, a Lean locator (paths are relative to `src/` in the mirror), and whether the step is classical or appears to be new. We use the author's remark [2] as a guide to intent; the author says its descriptions of the parts done by Claude are paraphrases, so where the remark and the code differ we follow the code.

We do not assess correctness. Section 6 lists the places where, in our reading, the weight of the argument rests, so that readers who check it can start there.

At a glance. Of the development's 17.4 million lines (Section 5):

- about 81% (about 14.0 million lines) is standard or published mathematics re-proved in Lean;
- about 19% (about 3.3 million lines) is candidate new argument;
- about 5% (about 850,000 lines) holds the core bookkeeping of the argument: the endgame in `Development/AutomorphicArithmetic/CharacterBoxes` and its sibling folders, the slope, divisibility and size bookkeeping in `Formalism/SharpRestriction`, `Formalism/SizeWindowBrackets`, `Formalism/RankGapBounds` and `Formalism/LocalRingCardinality`, and the Selmer bound in `.../ModularityArgument/ThresholdArguments`.

The split is by folder theme. The background can be set aside, except for a few inputs that sit inside it and on which the argument depends: the covolume bound (`VolumeTheoremX`), the Galois-representation and local inputs (`Row0b`, `RecTExists`), and parts of the archimedean budget (`meanPeriod`, `kerSide_chainS`). Conversely, the declarations mapped in this note also reach into other folders. Where to look first (Section 6): the forced-divisibility statement (`adm2` and the package `SelectionServed` it is produced from), the sharpness required of the averaged slope inequality, the uniform Selmer bound (`VWin`), the covolume bound (`VolumeTheoremX`), and two slots filled by statements weaker than their names (`RecTExists`, `Row0b`). Appendix A gives each of these at a frozen file and line.

1.3 The architecture in brief

The argument has an analytic half and an arithmetic half, joined by one inequality.

Analytic half (Section 2). For a primitive character ψ of order o, consider the first Riesz means of Λ twisted by the powers of ψ, with their explicit-formula main terms removed. A one-sided bound of size $y^{1/2+\varepsilon}$ for a fixed positive combination of them, valid only on windows $[Z, Z^{3/2}]$ along a sequence $Z \to \infty$, implies by a windowed Landau theorem that none of the $L(\psi^k, s)$ vanishes for $\operatorname{Re} s > \frac{1}{2}$. The development obtains the one-sided bound from a family of inequalities of the form "forced prime sum $\le$ archimedean main term + error", by an exact identity and finite Fourier analysis.

Arithmetic half (Sections 3–4). Those inequalities come from an Arakelov slope inequality for lattices of integral automorphic forms on totally definite unitary groups. The remark calls the method "a descendant of Stepanov's". At primes inert in a CM field, Frobenius acts anti-holomorphically, so integral theta-kernel periods of weight a are divisible by a large power of p; summed over p, the exponents form a Riesz mean of $\Lambda(n)(1-\chi(n))$. The archimedean size of the same objects is $\log a!$ corrected by the Chowla–Selberg period, which is exactly the explicit-formula main term. One application gives nothing; a chain of $m \approx \sqrt{x}$ theta lifts multiplies the signal by m while the global costs are paid once, and the relative error becomes $x/m + m$, which is where the exponent $\frac{1}{2}$ comes from.

Figure 1 shows the chain. The main inputs to the slope inequality are a forced-divisibility theorem for iterated theta lifts, an archimedean budget (Fock norms, the Chowla–Selberg formula, a covolume bound for integral cusp forms), a relative-trace-formula estimate making a family of period maps nearly a tight frame, and a uniform bound for Selmer groups of Dirichlet characters that pays for dividing out unknown constants.

1.4 Organization

Section 2 follows the analytic half down to the interface inequality. Section 3 explains the Stepanov-type engine that produces it, and Section 4 the inputs that pay for its error terms. Section 5 records what is standard background, and the abandoned designs still present in the source. Section 6 lists where to look first, Section 7 proposes a plan for human readers (what to extract, informalize and check, and in what order), and Section 8 summarizes public commentary as of the day of release. Appendix A gives frozen Lean locators, with links, for every declaration named in the note, and Appendix B relates the items of the author's remark to declarations in the code.

2 The analytic half: from a one-sided bound to zero-free half-planes

2.1 Riesz-mean remainders

For a nontrivial primitive Dirichlet character χ and $y \ge 1$ put

$$G_\chi(y) = \sum_{n \le y} \Lambda(n)\chi(n)\Big(\frac{y}{n} - 1\Big) + y\,\frac{L'}{L}(1,\chi), \qquad G_1(y) = \sum_{n \le y} \Lambda(n)\Big(\frac{y}{n} - 1\Big) - y\log y + (1+\gamma)y, \tag{2.1}$$

where γ is Euler's constant (`An.gChi` in `Formalism/SharpRestriction/RamificationFrobenius/TotientGaloisClasses/Analytic.lean`, and `Gz` in `Formalism/StatementOC.lean`). The sums are first Riesz means: for $\operatorname{Re} s > 1$,

$$\int_1^\infty \Big(\sum_{n \le y} a_n\Big(\frac{y}{n} - 1\Big)\Big) y^{-s-1}\,dy = \frac{1}{s(s-1)} \sum_{n \ge 1} \frac{a_n}{n^s},$$

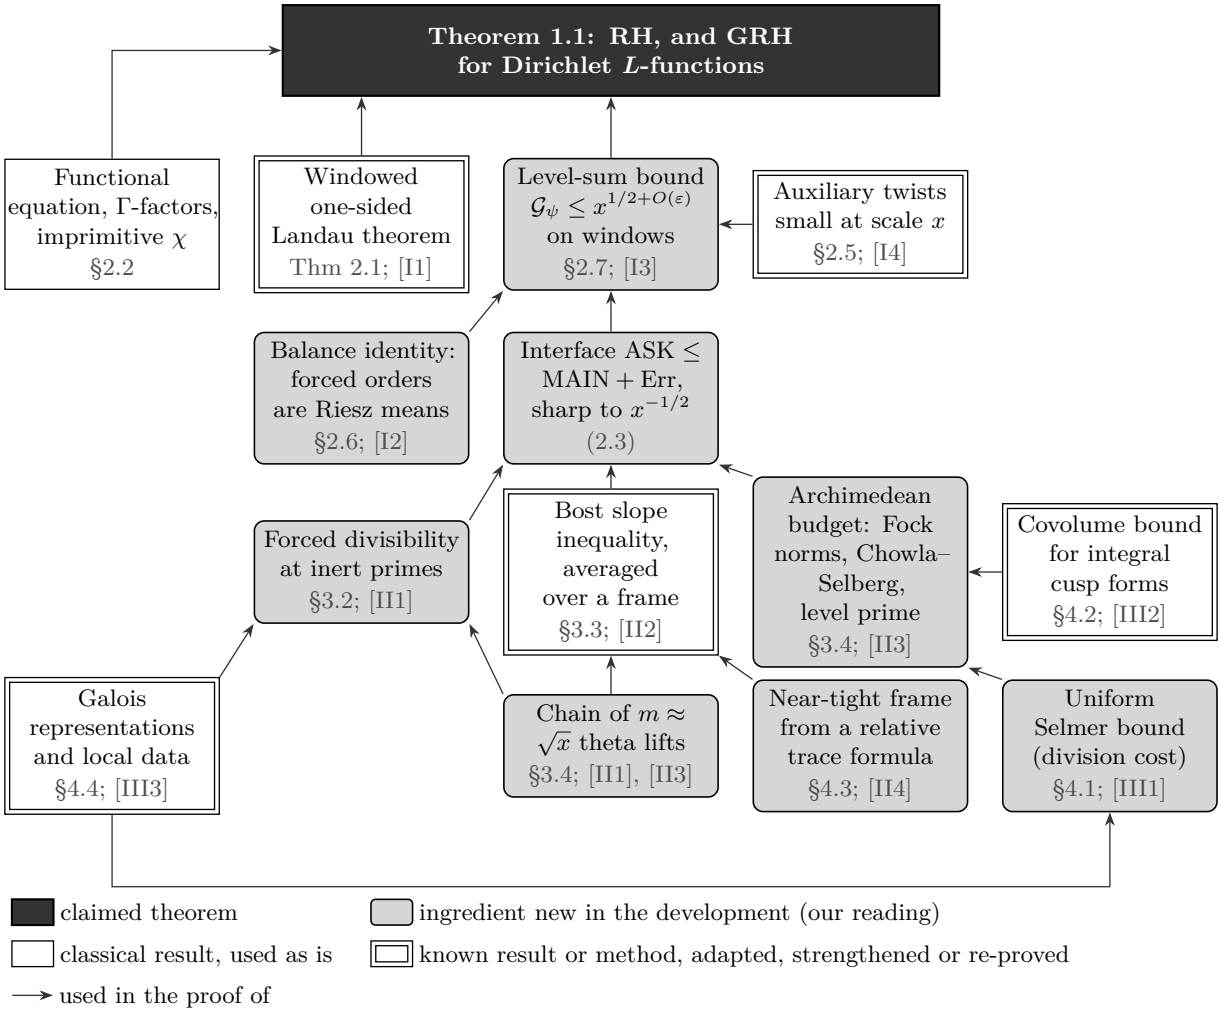


Figure 1: How Theorem 1.1 is proved, as we read the development. A solid arrow means that its source is used in the proof of its target; only the main dependencies are drawn (the lift chain also feeds the archimedean budget, which is not drawn). Box styles give provenance: a double box marks a step whose statement or method has a classical model (Landau's theorem, Bost's inequality, second-moment methods, arithmetic proportionality, Galois representations), even where the development's version is new; a shaded box marks a step without such a model. They are a reading guide, not a novelty or correctness verdict. Grey labels give the section of this note that describes each item and its tag in the reading and informalization plan of Section 7 (Table 2). Section 6 discusses the inputs on which the argument rests most heavily.

so the Mellin kernel of G_χ is $-\frac{L'}{L}(s,\chi)\,y^s/(s(s-1))$, and the terms subtracted in (2.1) are its residue at $s=1$: $y\log y-(1+\gamma)y$ for ζ, and $-y\frac{L'}{L}(1,\chi)$ for $\chi\neq 1$. Shifting the contour gives the classical explicit formula

$$G_\chi(y)=-\sum_\rho \frac{y^\rho}{\rho(\rho-1)}+O_\chi(\log 2y),$$

with the sum over nontrivial zeros, counted with multiplicity, converging absolutely [14]. In particular GRH for χ is equivalent to $G_\chi(y)\ll_{\chi,\varepsilon} y^{1/2+\varepsilon}$. Landau's theorem on Dirichlet integrals with nonnegative coefficients [14] shows that one side suffices: for ζ, a bound $G_1(y)\le Cy^\theta$ for all large y already forces $\zeta(s)\neq 0$ for $\operatorname{Re} s>\theta$. This is the normalization the remark adopts on 20 October, and the whole development is organized around it.

2.2 The endgame is classical

Statement (b) of Theorem 1.1 gives (a), (c) and (d) through the functional equation and the nonvanishing of the gamma factors (`Formalism.VW.rh_of_every` and `Formalism.VW.critLine` in `Formalism/GammaFactorNonvanishing/CritLine.lean`; for (c), `Claim.strip` in `Development/AutomorphicArithmetic/CharacterBoxes/ClaimStrip.lean`). Imprimitive characters reduce to primitive ones because the finitely many removed Euler factors do not vanish for $\operatorname{Re} s>0$ (`UpperTheory.Top.imprimitive`, `Development/AutomorphicArithmetic/CharacterBoxes/Imprimitive.lean`). So the content is: for every primitive ψ, $L(\psi,s)\neq 0$ for $\operatorname{Re} s>\frac12$.

2.3 Level sums and positive-definite sampling

The development does not bound G_ψ for one character at a time. Fix a primitive ψ of order o, let ψ_k be the primitive character inducing ψ^k, and define the *level sum*

$$\mathcal{G}_\psi(y)=G_1(y)+\frac{1}{3o+1}\sum_{k=1}^{o-1}G_{\psi_k}(y) \tag{2.2}$$

(`Level.Gm` in `.../TotientGaloisClasses/ExplicitFormulaErrorOfHostConditions.lean`, with the level `Level.cyc` ψ). The set $\{\psi_k\}$ is closed under conjugation, so $\mathcal{G}_\psi$ is real. Its Dirichlet coefficients are $\Lambda(n)$ times $1+(3o+1)^{-1}\sum_{k=1}^{o-1}\psi_k(n)$, which equals $\frac{4o}{3o+1}$ or $\frac{3o}{3o+1}$ according as $\psi(n)=1$ or not, for n prime to the modulus, and is at least $\frac{2o+2}{3o+1}$ in all cases. The Mellin transform of $\mathcal{G}_\psi$ is, up to its polar part at $s=1$, the logarithmic derivative of $P(s)=\zeta(s)^{3o+1}\prod_{k=1}^{o-1}L(\psi_k,s)$ divided by $-(3o+1)s(s-1)$. Every L-function of interest therefore occurs with a positive integer multiplicity, and at a real zero $\beta\in(0,1)$ of P the residue $-m_\beta/((3o+1)\beta(\beta-1))$ is positive. That sign is all that Landau's argument needs.

The specific weights $4/(3o+1)$ and $3/(3o+1)$ are produced by finite Fourier analysis (Section 2.6): the inequalities supplied by the arithmetic half are indexed by places of an abelian field, and summing them over a set of top places whose shares are $4/(3o+1)$ in one coset and $3/(3o+1)$ in every other coset multiplies each nontrivial character by $1/(3o+1)$ and the trivial one by 1 (`one_char`, `Frame.Shares`). Any sampling distribution with strictly positive Fourier coefficients would serve; this one is uniform plus a point mass.

2.4 A windowed Landau theorem

The arithmetic half produces bounds only for x in windows $[Z,Z^{3/2}]$ along a sequence $Z\to\infty$, because the auxiliary field it uses changes with the scale. The development therefore proves a localized Landau theorem. We transcribe its hypotheses.

Theorem 2.1 (windowed Landau theorem, `landau_window` in [1]). *Let $g : \mathbb{R} \to \mathbb{R}$ be measurable with $|g(u)| \le Ce^{2u}$ for $u \ge 0$, and let $\Phi(s) = \int_0^\infty g(u)e^{-su}\,du$ for $\operatorname{Re} s > 2$. Let $\mathcal{Z} \subset \{\operatorname{Re} s > 0\}$ have no accumulation point in the open right half-plane, and suppose Φ is holomorphic on $\{\operatorname{Re} s > 0\} \setminus \mathcal{Z}$ and has at each $\rho \in \mathcal{Z}$ a simple pole with nonzero residue c_ρ, where c_ρ is real and positive whenever ρ is real. Let $0 < \theta < 1$ and $\delta > 0$, and suppose that for every $\eta > 0$ and every U_0 there are $U \ge U_0$ and a function $R \ge \max(g, 0)$ on $[U, (1+\delta)U]$ such that $R(u)e^{-\theta u}$ is integrable there and $\int_U^{(1+\delta)U} R(u)e^{-\theta u}\,du \le e^{\eta U}$. Then $\operatorname{Re}\rho \le \theta$ for every $\rho \in \mathcal{Z}$.*

The theorem is at `Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/landauWindow.lean:83`; its hypotheses are `PoleData` and `WindowHypU` in `.../DedekindLFunctions/LandauZeroFreeCriterion.lean`. Applied to $g(u) = \mathcal{G}_\psi(e^u)$, it says that if for every $\varepsilon > 0$ there are arbitrarily large Y with $\mathcal{G}_\psi(y) \le y^{\theta+\varepsilon}$ for all $y \in [Y, Y^{1+\delta}]$, then no $L(\psi^k, s)$ vanishes for $\operatorname{Re} s > \theta$.

Idea of the proof. Take a real base point $\sigma > 2$ far to the right and look at the Taylor coefficients $T_k(z) = (-1)^k\Phi^{(k)}(z)/k! = \int g(u)u^k e^{-zu}/k!\,du$ at $z = \sigma + i\gamma$. The kernel $u^k e^{-\sigma u}/k!$ peaks at $u \approx k/\sigma$, so k can be chosen to put the peak in a window, and a Chernoff bound controls the mass outside it even against the crude growth e^{2u}. Inside the window, $|g| = 2g^+ - g \le 2R - g$ turns the one-sided bound into an L^1 bound at the price of one real coefficient: $|T_k(\sigma+i\gamma)| + \operatorname{Re} T_k(\sigma) \le Ce^{\eta U}(\sigma-\theta)^{-k}$. On the other hand the coefficients grow like d^{-k}, where d is the distance to the nearest pole. A nearest pole on the real axis contributes with positive sign and contradicts the bound at once. A complex pole is seen by moving the base point to a generic height γ near its ordinate, so that it becomes the unique nearest pole, and comparing with a Cauchy estimate at the real point. Only the length of the windows in $u = \log y$ matters, so any exponent $1 + \delta > 1$ works; the development uses $3/2$. We did not check whether this packaging of a localized Landau theorem appears in the literature.

2.5 Scale-adapted auxiliary twists

The inequalities of the arithmetic half do not see $\mathcal{G}_\psi$ directly. In the simplest case they see primes inert in an imaginary quadratic field $K = \mathbb{Q}(\sqrt{-d})$, that is, coefficients $\Lambda(n)(1 - \chi_{-d}(n))$, and Section 2.6 shows they amount to a one-sided bound for $G_1 - G_{\chi_{-d}}$. With d fixed, Landau's theorem then only shows that ζ and $L(\chi_{-d}, s)$ have the same zeros to the right of θ. The source still contains several abandoned statements of this kind (Section 5.2).

The route actually used escapes by letting the auxiliary data depend on the scale. For each window the development chooses primes $q \equiv 3 \pmod 8$ and $d \equiv 7 \pmod{16}$ with $\log q \approx Y^{\varepsilon/5}$ and every odd prime $p \le Y^{\varepsilon/5}$ dividing $d + 1$, so that small primes split in K (`WinPair`, `CharacterBoxes/ExistsPageN.lean:27`). It requires that every auxiliary twist $\theta = \psi^k\chi_q^i\chi_{-d}^j$ with $(i, j) \ne (0, 0)$ satisfy $|G_\theta(y)| \le y^{1/2+\varepsilon}$ on the window. The inequality "kept sum minus auxiliary sums $\le$ small" then becomes a bound for $\mathcal{G}_\psi$ alone (`sigA_le`, `core_gm`).

This requirement looks like GRH for the auxiliary twists, but it concerns only y up to a fixed power of Y, which is polylogarithmic in their conductors $\exp(O(Y^{\varepsilon/5}\log^2 Y))$, and only most choices of q. In that regime it is proved unconditionally (`QH.pairEx` in `Development/AutomorphicArithmetic/MeanSquareBounds/pairEx.lean:112`), by a second-moment argument with one device worth isolating. Write $G_\chi(y) = -\sum_n \Lambda(n)\chi(n)\min(1, y/n)$, truncate at $n \le Y^{O(1)}$ using a zero-free box that holds for all but $O(Q^{1/4})$ prime conductors $q \le Q$ (`count_bad`), and average over q. The truncated sums have mean square $\ll y\log y$ by the equidistribution of q in residue classes. The mean-square hypothesis is quantified over all bounded coefficient sequences κ, and the class is closed under $y\,\partial_y$, since $y\,\partial_y \sum_n \Lambda(n)\kappa(n)\min(1, y/n)$ is the same sum with κ replaced by $\kappa\mathbf{1}_{n>y}$

(`bs_mul_D_eq_S`). So the sums and their derivatives have the same mean square, a one-dimensional Sobolev inequality bounds the supremum over each dyadic block, and Markov's inequality with a geometric series over blocks shows that all but a small proportion of the q satisfy the bound for *every* y in the range (`bad_share` in `MeanSquareBounds/BadShare.lean:251`).

2.6 The balance identity

The arithmetic half produces inequalities whose left side is a sum of *forced orders*. In the simplest case, with $K = \mathbb{Q}(\sqrt{-d})$, a height $a > 0$ and a weight $w \in (0, a)$, it is

$$\mathrm{ASK} = \sum_{p \text{ inert in } K} 2\log p \sum_{i \text{ odd}} \Big(\frac{a}{p^i} - w\Big)_+ = \sum_n \Lambda(n)\big(1 - \chi_{-d}(n)\big)\Big(\frac{a}{n} - w\Big)_+,$$

up to the ramified prime, since $1 - \chi_{-d}(p^i)$ is 2 exactly when p is inert and i is odd. The right side is $\mathrm{MAIN} = a\big(\log a - 1 - \gamma + \frac{L'}{L}(1, \chi_{-d}) - \log w\big)$. With $y = a/w$ and (2.1),

$$\mathrm{ASK} = w\Big(\sum_{n \le y} \Lambda(n)\big(1 - \chi_{-d}(n)\big)\Big(\frac{y}{n} - 1\Big)\Big) = \mathrm{MAIN} + w\big(G_1(y) - G_{\chi_{-d}}(y)\big).$$

So "ASK $\le$ MAIN + Err" is exactly "$G_1(a/w) - G_{\chi_{-d}}(a/w) \le \mathrm{Err}/w$". The terms of MAIN are Stirling's $a\log a - a$, Euler's constant, and the constant $\frac{L'}{L}(1, \chi_{-d})$, which enters the arithmetic half through the Chowla–Selberg formula (Section 3.2).

The general identity (`ch_balance`; see Appendix A) replaces K by an abelian CM field E with $\mathrm{Gal}(E/\mathbb{Q}) \cong G = \{\pm 1\}^2 \times A$, the inert primes by primes whose Frobenius class lies outside a CM type $\Phi \subset G$ of index two, and the single weight by a weight $w(\rho)$ at each place $\rho \in \Phi$. It reads

$$\mathrm{ASK} - \mathrm{MAIN} = \frac{1}{|\Phi|} \sum_{g \notin \Phi} \sum_{\chi \in \widehat{G}} \overline{\chi}(g)\, w(\rho_g)\, G_\chi\big(a/w(\rho_g)\big),$$

for a place ρ_g determined by g and a top place t. The identity has no error term. Two steps of finite Fourier analysis then isolate the level sum (`sum_A`, `sum_Th`, `one_char` in `LandauTheorems/ChCore.lean`). Averaging over translates u_0 of the weight pattern by a large cyclic subgroup $A_0 \subset A$ kills every character nontrivial on A_0. Summing over the biased set of top places of Section 2.3 weights the surviving powers of ψ by $1/(3o+1)$. The remaining characters are the twists of the powers of ψ by χ_{-d}, χ_q and $\chi_q\chi_{-d}$, which Section 2.5 makes small. The weights also separate two scales. The height is $a = x \cdot b$ with $b = \exp(x^\varepsilon)$. On a sparse set of "thick" places, of density at least $x^{-\varepsilon}$, the weight is about b, so these places read G_χ at $a/w \approx x$. On all other places $w \le 2a\exp(-x^{\varepsilon/3})$, so they read G_χ at $y \ge \exp(x^{\varepsilon/3})/2$, where the prime number theorem with classical error term makes their contribution exact to within $y\exp(-cx^{\varepsilon/8})$ (`ThinScales`). The errors of the inequality live at the height a, while the signal is read at the scale x.

2.7 The interface

Putting these together (`ch_core`, `core_gm`, `level_gm_le_on_family_scales`), the analytic half needs, for each primitive ψ, each small ε and each x in windows $[Z, Z^{3/2}]$ with $Z \to \infty$:

- an abelian CM field E with $\mathrm{Gal}(E/\mathbb{Q}) \cong \{\pm 1\}^2 \times A$, $n = |A| \approx x^{2+8\varepsilon}$, containing $K = \mathbb{Q}(\sqrt{-d})$ and a 2-power cyclotomic layer, with characters as in Section 2.5;

- for every translate $u_0 \in A_0$, the inequality

$$\sum_{t\in T}\sum_{j=1}^{m}\Big[\mathrm{ASK}(a_t,t,w_{u_0,j})-\mathrm{MAIN}(a_t,w_{u_0,j})\Big] \le x^{O(\varepsilon)}\Big[|T|\,a\Big(1+\frac{m^2}{k}\Big)+m(B_0+N_c)\Big], \quad (2.3)$$

  where T is the set of top places, $a_t=(x\pm 1)b$, $m\approx\sqrt{x}$ is the number of lifts (Section 3.4), k is a rank parameter with $xm\le k\le n^2x^{1/2-\varepsilon}$, and B_0, N_c are bounded cost terms (`gapA` $\le$ the error term; `Formalism/RankGapBounds/MeanBoundSecondForm.lean`).

The left side of (2.3) has $|T|m$ terms of size about $a\log a$; the right side is about one term's worth. The inequality is therefore an equality to relative precision $x^{-1/2+O(\varepsilon)}$. Dividing by the signal, which is of order $|T|\,m\,b\,x^{-\varepsilon}$, the bound becomes $\mathcal{G}_\psi(x)\le x^{O(\varepsilon)}(x/m+m)$, and $m\approx\sqrt{x}$ gives $x^{1/2+O(\varepsilon)}$.

Remark 2.2 (sanity check). The reduction creates no strength of its own. If GRH holds for the characters of E, the left side of (2.3) is $O(\delta_2|T|\,m\,b\,x^{1/2+\varepsilon})$, where $\delta_2\ge x^{-\varepsilon}$ is the density of thick places; since the factor $x^{O(\varepsilon)}$ on the right of (2.3) has the larger exponent, this is smaller than the right side, so (2.3) holds with room. If some $L(\psi^k,s)$ had a zero with real part above $\frac12+C\varepsilon$, then by Theorem 2.1 the inequality would fail at suitable x in every long window. So (2.3) is essentially equivalent to GRH for the powers of ψ. The output is one-sided and lives only on windows, consistent with the known Ω-results for G_χ. All of the strength must come from the arithmetic half.

3 The engine: Stepanov's method over $\mathbb{Z}$

3.1 Stepanov's method

Recall the shape of Stepanov's elementary proof of the Riemann hypothesis for curves over finite fields; Bombieri's exposition [6] recasts it with p^μ-th powers and the Riemann–Roch theorem in place of derivatives. To bound the number N of $\mathbb{F}_q$-points of a curve, one builds an auxiliary function f that vanishes to order at least m at every rational point and is not identically zero, and compares mN with the degree of f. In Stepanov's version the multiplicity is cheap because $\frac{d}{dx}x^q=0$ in characteristic p, so the coefficients of f may be polynomials in x^q that behave like constants under differentiation. The degree, the global cost, grows slowly with m. Balancing the two gives $m\approx\sqrt{q}$ and $N=q+O(g\sqrt{q})$. Three things make it work: a Frobenius whose action makes high-order vanishing free at the points being counted; a global inequality (the degree bound) that turns vanishing into a count; and a multiplicity parameter traded against the global cost.

The development, in the remark's words, is a descendant of this argument. Over $\mathbb{Z}$ there is no Frobenius on a single object and no derivative to raise multiplicity. The next three subsections describe what replaces each of the three ingredients. Figure 2 sets the two arguments side by side.

3.2 Frobenius: forced divisibility at inert primes

Let K be imaginary quadratic and p a prime inert in K. An elliptic curve with CM by K has supersingular reduction at p, and on the reduction the p-th power of a holomorphic coordinate behaves anti-holomorphically. The remark records the consequence: for the invariant derivation D, $D^p\equiv 0$ at such p, and the normalized jets $D^a\theta(0)/\Omega^a$ of a theta function are divisible by p to the power $\frac{a}{p}+\frac{a}{p^3}+\cdots-O(\log a)$. Only odd powers of p appear because Frobenius at an inert prime acts on the CM type by complex conjugation, so even powers return to the same type. When the conjugate place carries weight w instead of 0, the exponent becomes $\sum_{i\ \mathrm{odd}}(a/p^i-w)_+$, the forced order of Section 2.6.

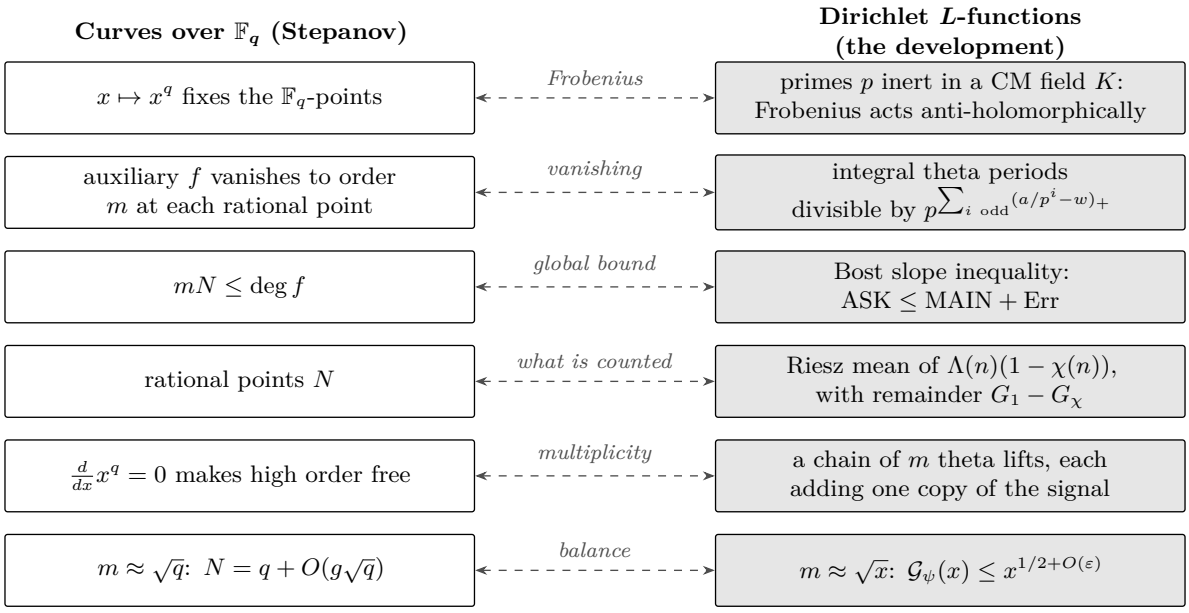


Figure 2: The dictionary between Stepanov's argument [6] and the development, as we read it. Left: the ingredients of Stepanov's proof and the balance that gives the square-root bound. Right: what plays each role in the development (Sections 3.2–3.4).

The analogy with Legendre's formula explains why the main term matches: $\sum_p \log p \cdot v_p(a!) = \sum_{n \le a} \Lambda(n) \lfloor a/n \rfloor = \log a! = a \log a - a + O(\log a)$. The p-adic and the archimedean sizes are two faces of one factorial. If the jets were integral of archimedean size about $a! \, \Omega^{-a}$, the product formula would give "inert-prime part of $\log a! \le \log a! - a \log \Omega +$ error". The Chowla–Selberg formula [10] computes the CM period, and in the development $\log \Omega = \gamma + \log(\pi/v) - \frac{L'}{L}(1, \chi_{-d})$ (`logOm` in `Formalism/RankGapBounds/MeanBoundSecondForm.lean:88`; `meanPeriod` in `Development/AbelianVarietyTheory/Proofs/DivisorProofs/MeanLogPeriodChowlaSelberg.lean:84`). This is exactly the constant in MAIN. The prime sum and its explicit-formula main term thus appear as the two sides of the product formula for one integral object. Forced divisibility selecting a Frobenius class is what lets a one-sided Chebyshev-type bound obtained from integrality, in the tradition of Gelfond's method, see a character.

In the code the divisibility is carried by integral highest-weight lattices at an unramified place of even degree, with the weight dualized at odd Frobenius positions (`lamAt` in `Formalism/SharpRestriction/GrowthBounds/UnramifiedQuadraticSetting.lean:67`), and by a height-two formal-group logarithm with Dwork's lemma (`FormalGroupSystem.lean`). The statement used downstream is `adm2` (`Formalism/SharpRestriction/Divisibility/LocalQuadraticParts.lean:206`). It says that the period of an integral column vector against an integral row vector is either 0 or has w-adic valuation at least the forced order. It is derived from properties of the theta kernel as a p-adic power series (`EvenConditionsInert`, `Cor2`, `ThetaKernelShapeAtD` in `.../Divisibility/OrbitPlaces.lean:303--355`, with `KvalGe` from `.../Divisibility/ProductZerothFirstParts.lean:99`).

3.3 The degree bound: an averaged slope inequality

The global inequality is Bost's slope inequality [7]. If $\varphi : \overline{E} \to \overline{F}$ is a morphism of Hermitian vector bundles over the ring of integers of a number field and φ is injective on generic fibres, then the normalized Arakelov degree of $\overline{E}$ is at most the maximal slope of $\overline{F}$ plus the height of φ. It plays

the role of the nonvanishing determinant in Stepanov's argument. Forced divisibility enlarges the source lattice: dividing by p^{gain} at every inert p keeps the map integral and raises the degree by $\sum_p \text{gain} \cdot \log p$, which is ASK. The archimedean side of the inequality is MAIN plus error.

No single period map is injective here, so the development averages. Its maps P_u are indexed by "words" u, binary interlacing patterns of archimedean weights, each giving one Gan–Gross–Prasad (Bessel) period [11]. If the Gram operators $P_u^* P_{u'}$ behave to second order like those of independent random isometric embeddings, the family is an approximate tight frame, and a slope inequality holds on average with explicit defects (`slope_inequality_mean` in `Formalism/HarderNarasimhan/HarderNarasimhanProofs/SlopeIneqMean.lean:249`). The required near-orthogonality,

$$\|P_u^* P_{u'}\|_{HS}^2 \le (1 + b^{-6}) \frac{d_u d_{u'}}{D_*} \quad (u \ne u'), \qquad P_u^* P_u \approx \mathrm{id}, \qquad \dim \le (1+\eta) D_*,$$

with errors at most $\exp(-x^{\varepsilon/8})$, is a second-moment statement. It comes from a relative trace formula with a Siegel–Weil main term (Section 4.3). Replacing the injective map by a frame is, in our reading, a reusable idea for slope and interpolation-determinant arguments generally.

3.4 Multiplicity: a chain of theta lifts

One slope inequality is useless: its one-time costs (heights of the end maps, slopes of the source, a level term) are of size $a\, x^{\varepsilon}$ per place, the same as the main term. Over $\mathbb{Z}$ there is no derivative, and the remark's central idea is to raise multiplicity by lifting instead. Starting from a form on a totally definite unitary group of rank k, the development lifts m times,

$$U(V_0) \to U(V_1) \to \cdots \to U(V_m), \qquad \pi \mapsto \pi \boxplus \mu_j[t], \qquad \dim V_j = k + j(t+1),$$

in the stable range, adding at each step a block of t twists of a Hecke character μ_j (`Ranks.N`, `ColumnChain` in `Formalism/LocalRingCardinality/Params.lean` and `DataKernels.lean`). Each lift contributes its own forced orders, with truncation weight w_j equal to the archimedean weight of μ_j, and its own main term. The one-time costs are paid once. By a seesaw argument (Kudla [12]; the remark credits Rallis) the period of the m-fold lift is a Bessel period on definite groups, independent of the μ_j; on definite groups such a period is a finite sum over a class set, hence algebraic with controlled denominators (`BesselId` in `Formalism/SharpRestriction/GlobalTheoremInputs.lean:65`).

Two devices complete the accounting. A level at a split prime $T_u \approx (bv/\pi)^m$, through its last branching step, turns the $m \cdot a \log a$ of m copies of Stirling into $\sum_j a \log(a/w_j)$, and so truncates each prime sum at $p \le a/w_j$. The per-lift losses add up to $O(m^2)$ copies of a unit cost, which after normalization by the signal (m copies) is $O(m)$. So the relative error is $x^{1+\varepsilon}/m$ from the one-time costs plus $m\, x^{O(\varepsilon)}$ from the lifts, optimized at $m \approx \sqrt{x}$ (the parameter model `SizesTextAt` in `Formalism/SizeWindowBrackets/SChain.lean`, a satisfiability check of the numerics, takes $x^{1/2-\varepsilon/16} \le m$). This is the analogue of Stepanov's $m \approx \sqrt{q}$, and it is the only source of the exponent $\frac{1}{2}$: the method cannot do better than $x^{1/2}$. Figure 3 draws the chain and this trade-off.

The remark adds that a first count of the cost per lift stopped at $\frac{2}{3}$, and that a "division" step was needed to reach $\frac{1}{2}$. One reading consistent with this: if the lifts' total contribution to the normalized error grew like m^2 rather than m, the bound would be $x/m + m^2$, giving the exponent $\frac{2}{3}$ at $m = x^{1/3}$; dividing out unknown constants by twisted lifts (Section 4.1) removes the extra factor.

3.5 Pigeonhole: an Iwasawa tower and a Behrend-type placement

In Stepanov's setting one passes to $\mathbb{F}_{q^n}$ to separate Frobenius orbits. Here the host field contains a 2-power cyclotomic layer of degree about $x^{2+8\varepsilon}$, so every prime $p \le x$ has Frobenius order at least

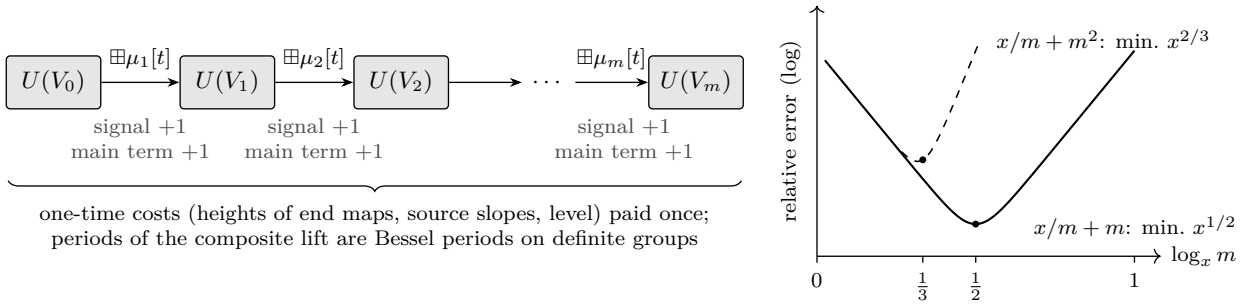


Figure 3: Left: the chain of theta lifts between totally definite unitary groups (Section 3.4). Each lift contributes its own forced orders and its own main term, while the one-time costs are paid once. Right: the resulting relative error, on log scales for $x = 10^6$ (illustrative). The solid curve $x/m + m$ is minimized at $m \approx \sqrt{x}$; the dashed curve $x/m + m^2$, one reading of the remark's first cost count, gives only $x^{2/3}$.

$2^r/(p+1)$ in it, 2^r being the order of the layer (`ShortOrdLarge` in `Formalism/BehrendSets/Statement.lean`), and the places $\tau + i \cdot \mathrm{Frob}_p$, $i \le \log a / \log p$, are distinct for every place τ. The places carrying the large weight form a translate of a set with no short three-term linear relations, built from Behrend's construction [5] (`BehrendFree`, `WindowTwo`). Along each Frobenius orbit, any window of $2L$ consecutive steps then meets it at most twice, so the partner places at odd shifts are almost always light, which is what makes the forced orders $(a/p^i - w)_+$ nonzero. The same cyclic subgroup supplies the translates used in Section 2.6.

4 Paying for the error terms

The slope inequality has four kinds of error, each paid by a large input. In our reading these inputs are where most of the difficulty lies.

4.1 Division by twisted lifts and a uniform Selmer bound

The norms of the lifts carry unknown constants of order m. The development removes them by comparing lifts twisted by different Hecke characters and dividing, which the remark calls gluing an integral structure across Hecke characters for a "gcd". The price is the congruence module of the theta-lift lattice between twist eigenspaces. By the direction of Ribet's method [15], a congruence between an eigensystem containing a block of twists and one outside it produces a class in $H^1(E, \mathbb{Z}/\ell^c(-e))$ with $0 < e < t$. The cost is therefore bounded by

$$\sum_{\ell} \log \ell \sum_{1 \le e < x^{401}} \operatorname{len} Ш^1(E, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e)),$$

and the development needs this to be at most $\exp(x^\varepsilon/2)$ for its host fields, on a window of scales (`VWin`, `VVle` in `Development/AutomorphicArithmetic/CharacterBoxes/ExistsPageN.lean:12--49`; `LinkN`). The proof (`vwin_of_row` in `.../AutomorphicVarieties/ModularityArgument/ThresholdArguments/DenseThresholdCases/VwinOfRow.lean`) works over an auxiliary imaginary quadratic field. It deforms the polarized triple $1 \oplus \Psi \oplus \Psi^{-1}$, shows every point of the deformation space is automorphic on $U(3)$ (or endoscopic, handled by Kato's Euler system), and bounds the Selmer length by the valuation of a single integer, a product of differences between cuspidal and Eisenstein Hecke traces at one Frobenius. That integer is nonzero by purity and has polynomial

height, so it bounds the Selmer groups at every ℓ at once. The remark calls this a significant weakening of Vandiver's conjecture, and its author believes it implies the vanishing of a number of K-groups of $\mathbb{Z}$.

4.2 The covolume bound

The archimedean side needs the Arakelov slope of the Petersson-metrized lattice of integral theta-lifted cusp forms on each unitary Shimura variety in the chain. The development bounds its log-covolume per rank by $(p_j - q_j)\sum_\rho w_j(\rho)\log\Omega + O(B_0 \log\log x)$ for signature (p_j, q_j) (`CovolumeBoundOfX` in `.../ThetaPairLifts/ThirdPlaceSeries/LiftLambdaBound.lean:60`). The tools are arithmetic Hilbert–Samuel, bigness theorems, Bergman-kernel sections, and an arithmetic intersection statement along towers (`VolumeTheoremX`). That statement proves proportionality only up to unknown constants per embedding, which cancel in the application.

4.3 Near-orthogonality from a relative trace formula

The frame estimate of Section 3.3 (`LinkS`; Theorem S in `Formalism/LocalToGlobal/SiegelLatticeTheory/FlagStabilizerTensors/TheoremSConclusion.lean:58,68`) comes from a relative trace formula for the Gram matrices, with the diagonal given by Siegel–Weil, and from a dimension formula equal to its identity term up to an exponentially small relative error. A separate input (`LinkIRR`) shows, from the equidistribution of depth-zero labels (Green's parametrization of cuspidal characters of $\mathrm{GL}_n(\mathbb{F}_q)$) at many auxiliary places, that all but a proportion $\exp(-x^{\varepsilon/7})$ of the forms have large image.

4.4 Galois representations and local data

Both the forced divisibility and the Selmer bound use Galois representations attached to regular algebraic polarized cuspidal representations of GL_a over CM fields, with a form of local–global compatibility (`Row0b`; `Development/AutomorphicArithmetic/Foundations/GaloisRepresentations/`). The development constructs them in-source by a hybrid of the classical route:

- in odd rank, Kuga–Sato-type families over Harris–Taylor-type unitary Shimura varieties, with Frobenius eigenvalues from an Eichler–Shimura congruence relation, followed by congruences and pseudo-representations;
- in even rank, a theta lift to two odd-rank definite groups with disjoint sets of characters, followed by subtracting the characters. This takes the place of the endoscopic transfer used in the standard proofs, and it is the same $\pi \boxplus \mu[t]$ lift that the engine iterates;
- descent from GL_a to a definite unitary group by a trace-formula comparison.

The compatibility proved is weaker than Caraiani's theorem [9]: it covers only abelian $N = 0$ places and tame places over $\mathbb{Q}_p$. The slot named for local Langlands is filled by a map (`RecTExists`) that is pinned down only on depth-zero and tame data and sends the other supercuspidals to a fixed placeholder parameter.

5 What can be set aside

5.1 The discard map

Table 1 summarizes our directory-level inventory, based on sampled file names and headers, not on reading proofs. Folder and namespace names often do not match contents; for instance `Development/WeilBound` holds Siegel–Weil, Casselman–Wallach and Nadel vanishing, not the Weil bound for curves. About 27% of the files are generated short wrapper files that re-export a sibling lemma.

Part of `src/`	Lines	Content and verdict
`Development/AbelianVarietyTheory`	8.1M	Abelian schemes, toroidal compactifications, PEL moduli, theta series, CM theory, Breuil–Kisin lattices, supplying inputs to §4.4. Standard.
Other `Development/`	4.1M	Theta correspondence and Fock models; automorphy lifting and patching; Siegel–Weil; étale cohomology; resolution of singularities; local converse theorem; the Galois-representation construction of §4.4. Mostly standard; contains the endgame (`CharacterBoxes`), the Selmer bound (`ModularityArgument`) and the covolume bound.
`GroundTheory`	1.1M	Arakelov geometry, Chowla–Selberg, class field theory, Tate's thesis, p-adic Hodge theory, local fields. Standard.
`Formalism`	3.2M	Classical analytic number theory (Landau, zero-free regions, Siegel–Walfisz) and Harder–Narasimhan slopes; in `SharpRestriction`, `SizeWindowBrackets`, `RankGapBounds`, `LocalRingCardinality`, `LocalToGlobal`, `PeriodIntegrals`, `HermitianCurvature`, the slope and divisibility arguments of §§3–4.
`ReciprocityTheory`, `CompanionTheory`, `PositivityTheory`, `Lemmata`	0.8M	Reductive groups, L^2 $\bar{\partial}$-estimates, Harish-Chandra modules, published inputs stated and proved. Standard.
Total	17.4M	About 81% standard background and 19% candidate new argument; the core mapped here is about 5% of the total.

Table 1: Where the lines are. Counts are rounded; verdicts are by theme, not exact file boundaries.

The code cites no papers: literature appears only in declaration names. An assembly file proves about 185 statements named after published results, and the three fields of the structure of "open literature statements" are filled by terms proved in the development (`CharacterBoxes/MainConclusion.lean`). So the final theorem carries no published results as hypotheses. As Section 4.4 notes, two of the slots are filled by statements weaker than their names.

5.2 Abandoned designs

The top-level files `Formalism/Statement.lean`, `StatementOC.lean`, `StatementCM.lean`, `StatementFL.lean`, `StatementU.lean` and `ZetaNonvanishing*.lean` state reductions that are never proved and never used. A reader opening the source at the top will meet them first, and they are not the route. Read in order they record a design history:

1. one real character, with a lower bound on G;
2. the difference $G_1 - G_\chi$ and equal zero orders of ζ and $L(\chi, s)$;
3. weighted families and complex combinations of characters;
4. a CM type of $\mathbb{Q}(\zeta_l)$ paired with a dual type, so that every Frobenius class $\neq \pm 1$ lies outside exactly one of the two; this would isolate ζ without making any partner small, at the price of $l \gtrsim \sqrt{x}$;

5. hypotheses on averages of G over sample points in a multiplicative window, including Chebyshev nodes;

6. the route of Section 2: an abelian CM host with an index-two CM type, scale-adapted auxiliary twists, two-scale weights, level sums and windows.

6 Where to look first

These are places where, in our reading, the weight of the argument rests. They are not findings of error.

1. **The forced divisibility.** Along one path of the code the per-prime gain is filled in by its defining formula and the lattice chain is the zero chain, so the clause "gain is a valuation" holds trivially there (`linkG_without_kernels` in `Formalism/LocalToGlobal/GroupLattices/AdmissibleFamilyNoKernel.lean:224`). The binding to an actual divisibility is `adm2` and the package it is produced from (`SelectionServed`, `hServ_of_linkR`). The claim behind it is a uniform p-adic divisibility of Bessel periods of m-fold iterated lifts by the inert-prime factorial excess, on a sublattice of small colength. We know of no published analogue. This is the mathematical heart of the method, and the first place to check that the hypotheses are met non-degenerately by the family used.

2. **Sharpness.** The slope inequality must be an equality to relative accuracy $x^{-1/2}$, uniformly in unitary groups whose rank is polynomial in x (the base rank alone is about $x^{9/2}$), over fields of degree about x^2 and with very large level. That needs dimension formulas, near-orthogonality and heights of an m-step chain far sharper than anything in the literature. The enabling device is extreme archimedean regularity: large gaps between Harish-Chandra parameters.

3. **The Selmer bound.** A bound, uniform in ℓ, for Selmer groups of Dirichlet characters twisted up to x^{401} over fields of degree about x^2 is a theorem of independent weight. Published automorphy-lifting results at residually reducible points usually assume control of the relevant Ext groups; a reader should look for circularity there.

4. **Arithmetic proportionality.** `VolumeTheoremX` is arithmetic intersection theory on unitary Shimura varieties of general signature along towers, with a rate. It is known in the literature only in special cases.

5. **Names stronger than statements.** The local Langlands slot (`RecTExists`) and the local–global compatibility slot (`RowOb`) are filled by weaker statements. This is harmless only if nothing downstream needs the full statements.

6. **In-source proofs of major theorems.** Labesse-type descent, Eichler–Shimura relations on Harris–Taylor varieties, residually reducible automorphy lifting on $U(3)$ and GSp_4, Kato's Euler system, prismatic Dieudonné theory and resolution of singularities are each a formalization milestone, and each relies on definitions whose fidelity to the standard ones would need checking.

At the level of ideas we found no inconsistency. The method cannot beat $x^{1/2}$, its output is one-sided and only on windows, and under GRH the needed inequality holds with room (Remark 2.2).

7 A plan for human digestion and informalization

The development has no human-readable proof, and its author expects the argument to simplify. Until it does, the quickest route to human understanding is to extract a small number of statements, write informal versions of them with proofs, and check those against the Lean source. Contributors on MathOverflow proposed dividing such work among several readers [13]. Table 2 lists the targets we would choose, with the Lean entry points to start from; all of them are in Appendix A, and the tags match Figure 1.

Tier I: self-contained analytic statements. These are short and need only analytic number theory. They are useful on their own, and they test whether the Lean statements mean what this note says they mean. They are good first targets.

Tier II: the arithmetic heart. These decide whether the method works. Each needs a reader at home in Arakelov geometry or in automorphic forms on unitary groups.

Tier III: inputs of independent weight. Each could be a paper of its own. Begin with a statement-level audit (what exactly is stated, and what the downstream lemmas use) before reading proofs.

Table 2: Reading and informalization targets. "Write up" is the statement to extract and prove informally; "Check first" is where, in our reading, the risk or the subtlety lies.

Tag	Write up	Lean entry points	Check first
[I1]	The windowed one-sided Landau theorem (Theorem 2.1), with a complete informal proof.	`landau_window`, `landau_window_key`, `landau_tails`, `landau_nearest`, `landau_dichotomy`, `PoleData`, `WindowHypU`	Tail control against the growth e^{2u}; the generic ordinate in the vertical shift; that `PoleData` holds for level sums. Compare with known localized Ω-theorems.
[I2]	The balance identity of §2.6 in general form: ASK − MAIN as a character sum of $w\,G_\chi(a/w)$.	`ch_balance`, `AskMainBalance`	The ramified term; the bookkeeping of the place ρ_g; that MAIN uses the normalization (2.1).
[I3]	The deduction from the family of inequalities (2.3) to the level-sum bound: translate averaging, the biased top set, thick and thin places, and the size algebra that gives $x/m + m$.	`ch_core`, `core_gm`, `sigA_le`, `one_char`, `Frame.Shares`, `level_gm_le_on_family_scales`, `ThinScales`, `RestSize`, `ErrWindows`, `SizesTextAt`	That the exponents close for small ε; the share errors $x^{-1+\varepsilon}$; the thick density $\geq x^{-\varepsilon}$; the thin-place error $\exp(-cx^{\varepsilon/8})$.
[I4]	Selection of the auxiliary primes q, d: a second-moment bound for twisted Riesz means, uniform in y through the $y\,\partial_y$ closure, and the zero-free boxes used at thin places.	`QH.pairEx`, `bad_share`, `bs_mul_D_eq_S`, `meanSq_abs`, `count_bad`, `thinH`, `WinPair`	The exponent in `count_bad`; the Siegel–Walfisz range for $\log q \approx Y^{\varepsilon/5}$; the exceptional modulus; the Linnik-range equidistribution of the d-family; the supremum over y in `bad_share`; that one pair (q, d) serves all twists of all ψ^k at once.

Tag	Write up	Lean entry points	Check first
[II1]	Forced divisibility: for integral column and row vectors, the Bessel period of the m-fold lift is 0 or divisible by the inert-prime forced order. Warm-up: the CM-curve statement of the remark, $D^a\theta(0)/\Omega^a$ divisible by $p^{a/p+a/p^3+\cdots-O(\log a)}$ at inert p.	`adm2`, `SelectionServed`, `EvenConditionsInert`, `Cor2`, `KvalGe`, `ThetaKernelShapeAtD`, `lamAt`, `hServ_of_linkR`	That the gain is bound to an actual valuation, not only along the zero-chain path `linkG_without_kernels`; non-degeneracy on the family used; uniformity in m, the rank and the place; the colength of the sublattice.
[II2]	The averaged Bost slope inequality with frame defects, as a general lemma on Hermitian lattices.	`slope_inequality_mean`, `SlopeIneqOfLat`, `BridgeFS`	The defect terms and how they enter; that near-tightness of the frame is all that is used.
[II3]	The archimedean budget: every term of size $a\log a$, a and $a\log\Omega$ per place and per lift (Fock norms, the Chowla–Selberg period, the split level prime), with a total error of about one unit per place.	`LinkDIV`, `kerSide_chainS`, `FockJet`, `BlockBound`, `meanPeriod`, `logOm`, `BesselId`, `Ranks.N`, `ColumnChain`, `DivisionPackage`	The matching of main terms to relative precision $x^{-1/2}$ (Remark 2.2; §6, item 2); that the lifts' total contribution to the normalized error is $O(m)$, not $O(m^2)$.
[II4]	Near-orthogonality of the period maps (Theorem S) from a relative trace formula with Siegel–Weil diagonal, and the big-image input.	`LinkS`, `TheoremSClosedHypotheses`, `theoremS_of_conj`, `LinkIRR`	Exponentially small relative errors, uniformly in rank, degree and level; the dimension formula; that the family-design antecedents of the frame estimate are not vacuous.
[III1]	The uniform Selmer bound: an Eisenstein congruence integer on $U(3)$ that bounds the Selmer groups at every ℓ.	`VWin`, `VVle`, `LinkN`, `vwin_of_row`, `CongruenceNumberUnitaryThree`	Residually reducible $R=T$ and possible circularity; uniformity of the constants in level and weight; as a sanity test, the K-group consequences the author expects.
[III2]	The covolume bound for integral cusp forms on $U(p,q)$, and the arithmetic proportionality statement behind it.	`CovolumeBoundOfX`, `VolumeTheoremX`	Comparison with the known cases of arithmetic proportionality; that the unknown constants per embedding cancel.
[III3]	A statement-level audit of the Galois-representation and local inputs: what is proved, and what the downstream lemmas use.	`RowOb`, `galoisLocalGlobalStrong_recParameter`, `RecTExists`	That no downstream lemma needs the full local Langlands correspondence or full local–global compatibility.

Suggested order.

1. Start with [I1] and [I2], which are short and test the reading of the Lean statements, and with the statement audit [III3], which is cheap.
2. Then do the warm-up of [II1]. If the CM-curve divisibility fails, or does not extend to iterated lifts, nothing downstream matters.
3. Then do [I3] and [II3] together. They settle whether the bookkeeping really reaches $x^{1/2}$.
4. Then the full [II1], the divisibility of Bessel periods of m-fold iterated lifts. It is the heart of the method and the statement with no published analogue.
5. Then [II2], [II4] and [I4].
6. Finally [III1] and [III2], each a research project in its own right.

Each write-up should state its result with exact hypotheses, name the Lean declarations it informalizes, and say where in the development each hypothesis is discharged. The standard background of Table 1 needs only spot-checks of statement fidelity, not re-reading.

8 Public commentary and outlook

We summarize public commentary as of 17:49 UTC on 20 November 2026; it was arriving quickly.

- T. Tao's initial reading was that the claimed proof uses the modern Langlands program and arithmetic geometry in full, and he pointed readers to the Lean Zulip discussion for verification and for extracting the strategy [16].
- On MathOverflow [13], T. Chow relayed an AI-generated summary of the code, posted on the Lean Zulip, which listed three inputs: a growth bound on Selmer-type groups of cyclotomic subfields, a covolume bound from arithmetic geometry, and theta integrals on unitary groups together with local Langlands. The summary said it could not see how the three combine. In our reading they are the error terms of one slope inequality: the division cost (Section 4.1), the archimedean budget (Section 4.2), and the multiplicity and frame structure (Sections 3.4 and 4.3). Contributors also proposed dividing the reading of such proofs among several people, each writing a partial exposition.
- K. Buzzard's checklist for large Lean proofs [8] is what verification efforts are following: read the statement and its definitions, build in a sandbox, run the comparator, and check with several kernels. A pass does not show that a prose description matches the Lean proof.
- No named expert had publicly identified a specific mathematical gap.

A precursor by the same author, with R. Furman, gives a proof, formalized in Lean, that more than two thirds of the zeros of ζ are simple and on the critical line [4]. The remark dates the first step of the present method (the Stepanov determinant) to 6 October 2026 and the completion of the GRH argument to 16 November.

Several of the techniques above look useful well beyond this problem, independently of whether the whole argument stands: the windowed one-sided Landau theorem (Theorem 2.1), which applies to any Laplace transform with isolated simple poles and positive real residues; uniform-in-parameter bounds from second moments of a family closed under the parameter derivative (Section 2.5); positive-definite sampling of group elements in positivity arguments (Section 2.3); slope inequalities with a frame in place of an injection (Section 3.3); and iterating a functorial lift to amplify a local invariant while global costs stay fixed (Section 3.4).

AI use

Inspection of the Lean code and the writing of this note were done with substantial help from the AI models Claude Opus 5.5 and Claude Fable 5.2.

A Frozen Lean locators

Every declaration named in this note, with its file and line in the mirror [3] at the pinned commit `a172c464fd4bcfc4f53e394dcba45d742b1409a6`, which the mirror reports to be byte-identical to the published archive. Each location is a permanent link to that line at that commit. The line numbers were checked against the source at that commit. Paths are relative to `src/`.

Declaration	File and line (link)
Claimed theorems (§1.1)	
riemannHypothesis	Main.lean:4
LFunction_ne_zero_of_half_lt_re	Main.lean:8
LFunction_eq_zero_re_of_isPrimitive	Main.lean:12
DirichletLNonvanishingRightOfHalf	Statements.lean:6
allChi	Supplements/Kernels/MainTheorems.lean:7
Classical endgame (§2.2)	
Formalism.VW.rh_of_every	Formalism/GammaFactorNonvanishing/CritLine.lean:29
Formalism.VW.critLine	Formalism/GammaFactorNonvanishing/CritLine.lean:11
Claim.strip	Development/AutomorphicArithmetic/CharacterBoxes/ClaimStrip.lean:59
UpperTheory.Top.imprimitive	Development/AutomorphicArithmetic/CharacterBoxes/Imprimitive.lean:5
UpperTheory.EveryCharacter	Development/AutomorphicArithmetic/CharacterBoxes/EveryCharacter.lean:13
UpperTheory.FinalTheorems.everyCharacter	Development/AutomorphicArithmetic/CharacterBoxes/MainConclusion.lean:10
UpperTheory.Top.plugWin	Development/AutomorphicArithmetic/CharacterBoxes/EveryCharacter.lean:8
UpperTheory.Top.every_character_of_green_v_win	Development/AutomorphicArithmetic/ThetaCorrespondence/ThetaPairLifts/ThirdPlaceSeries/TopPlug.lean:16
Riesz means and level sums (§§2.1–2.3)	
Gz	Formalism/StatementOC.lean:13
An.gChi	Formalism/SharpRestriction/RamificationFrobenius/TotientGaloisClasses/Analytic.lean:23
Level	Formalism/SharpRestriction/RamificationFrobenius/TotientGaloisClasses/ExplicitFormulaErrorOfHostConditions.lean:8
Level.Gm	Formalism/SharpRestriction/RamificationFrobenius/TotientGaloisClasses/ExplicitFormulaErrorOfHostConditions.lean:38
Level.cyc	Formalism/SharpRestriction/RamificationFrobenius/TotientGaloisClasses/ExplicitFormulaErrorOfHostConditions.lean:121
Frame.Shares	Formalism/SharpRestriction/RamificationFrobenius/TotientGaloisClasses/ExplicitFormulaErrorOfHostConditions.lean:139
one_char	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/ChCore.lean:202
Windowed Landau theorem (§2.4)	
landau_window	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/landauWindow.lean:83
PoleData	Development/AutomorphicArithmetic/DedekindLFunctions/LandauZeroFreeCriterion.lean:15
WindowHypU	Development/AutomorphicArithmetic/DedekindLFunctions/LandauZeroFreeCriterion.lean:25
zeroFreeAll_of_topG	Development/AutomorphicArithmetic/CharacterBoxes/ZeroFreeAllOfTopG.lean:26
zeroFreeOn_of_tiedWin_cyc	Formalism/SharpRestriction/RamificationFrobenius/TotientGaloisClasses/TopWin.lean:90
Auxiliary twists (§2.5)	
WinPair	Development/AutomorphicArithmetic/CharacterBoxes/ExistsPageN.lean:27
sigA_le	Development/AutomorphicArithmetic/PrimitiveCosets/SigALe.lean:38
core_gm	Development/AutomorphicArithmetic/PrimitiveCosets/CoreGm.lean:12
QH.pairEx	Development/AutomorphicArithmetic/MeanSquareBounds/pairEx.lean:112
count_bad	Development/AutomorphicArithmetic/MeanSquareBounds/CountBad.lean:9

Declaration	File and line (link)
bad_share	Development/AutomorphicArithmetic/MeanSquareBounds/BadShare.lean:251
bs_mul_D_eq_S	Development/AutomorphicArithmetic/MeanSquareBounds/BadShare.lean:83
Balance identity and interface (§§2.6–2.7)	
ch_balance	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/ChBalance.lean:200
ch_core	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/ChCore.lean:343
sum_A	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/ChCore.lean:28
sum_Th	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/ChCore.lean:56
ThinScales	Formalism/SizeWindowBrackets/GrowthBoundGivesRiemannHypothesis.lean:55
level_gm_le_on_family_scales	Development/AutomorphicArithmetic/PrimitiveCosets/LevelGmLeOnFamilyScales.lean:28
gapA	Development/AutomorphicArithmetic/DedekindLFunctions/SplitInertia.lean:50
MeanBoundSecondForm	Formalism/RankGapBounds/MeanBoundSecondForm.lean:187
logOm	Formalism/RankGapBounds/MeanBoundSecondForm.lean:88
Forced divisibility (§3.2)	
lamAt	Formalism/SharpRestriction/GrowthBounds/UnramifiedQuadraticSetting.lean:67
adm2	Formalism/SharpRestriction/Divisibility/LocalQuadraticParts.lean:206
EvenConditionsInert	Formalism/SharpRestriction/Divisibility/OrbitPlaces.lean:303
Cor2	Formalism/SharpRestriction/Divisibility/OrbitPlaces.lean:315
ThetaKernelShapeAtD	Formalism/SharpRestriction/Divisibility/OrbitPlaces.lean:351
KvalGe	Formalism/SharpRestriction/Divisibility/ProductZerothFirstParts.lean:99
meanPeriod	Development/AbelianVarietyTheory/Proofs/DivisorProofs/MeanLogPeriodChowlaSelberg.lean:84
Slope inequality and lift chain (§§3.3–3.5)	
slope_inequality_mean	Formalism/HarderNarasimhan/HarderNarasimhanProofs/SlopeIneqMean.lean:249
Ranks.N	Formalism/LocalRingCardinality/Params.lean:36
ColumnChain	Formalism/LocalRingCardinality/DataKernels.lean:14
BesselId	Formalism/SharpRestriction/GlobalTheoremInputs.lean:65
LinkDIV	Formalism/SharpRestriction/GlobalTheoremInputs.lean:91
SizesTextAt	Formalism/SizeWindowBrackets/SChain.lean:207
ShortOrdLarge	Formalism/BehrendSets/Statement.lean:31
BehrendFree	Formalism/BehrendSets/Statement.lean:48
WindowTwo	Formalism/BehrendSets/Statement.lean:67
Error terms (§4)	
VVle	Development/AutomorphicArithmetic/CharacterBoxes/ExistsPageN.lean:12
VWin	Development/AutomorphicArithmetic/CharacterBoxes/ExistsPageN.lean:49
LinkN	Formalism/SharpRestriction/GlobalTheoremInputs.lean:45
UpperTheory.One.vwin_of_row	Development/AutomorphicArithmetic/AutomorphicVarieties/ModularityArgument/ThresholdArguments/DenseThresholdCases/VwinOfRow.lean:4
CovolumeBoundOfX	Development/AutomorphicArithmetic/ThetaCorrespondence/ThetaPairLifts/ThirdPlaceSeries/LiftLambdaBound.lean:60
VolumeTheoremX	Development/AutomorphicSpectrum/Statements/TowerExistence/GoodTowerExistsX.lean:28
LinkS	Formalism/SharpRestriction/GlobalTheoremInputs.lean:152

Declaration	File and line (link)
`TheoremSClosedHypotheses`	Formalism/LocalToGlobal/SiegelLatticeTheory/FlagStabilizerTensors/TheoremSConclusion.lean:58
`theoremS_of_conj`	Formalism/LocalToGlobal/SiegelLatticeTheory/FlagStabilizerTensors/TheoremSConclusion.lean:68
`LinkIRR`	Formalism/SharpRestriction/GlobalTheoremInputs.lean:55
`Row0b`	Development/AutomorphicArithmetic/Foundations/GaloisRepresentations/TypeConditions/WeightsConjBis.lean:69
`RecTExists`	Development/AutomorphicArithmetic/Foundations/LocalLanglands/IsRecTOfIsRec.lean:78
Places to check first (§6)	
`linkG_without_kernels`	Formalism/LocalToGlobal/GroupLattices/AdmissibleFamilyNoKernel.lean:224
`SelectionServed`	Development/AutomorphicArithmetic/ThetaCorrespondence/ThetaPairLifts/ThirdPlaceSeries/CountsIrreducibilityConstants.lean:48
`GlobalTheory.T.hServ_of_linkR`	Development/AutomorphicArithmetic/ThetaCorrespondence/ThetaPairLifts/ThirdPlaceSeries/RowHServOfLinkStatement.lean:14
Further entry points of the plan (§7)	
`landau_window_key`	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/landauWindow.lean:40
`landau_tails`	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/landauTails.lean:80
`landau_nearest`	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/LandauNearest.lean:36
`landau_dichotomy`	Development/AutomorphicArithmetic/DedekindLFunctions/LandauTheorems/LandauDichotomy.lean:24
`AskMainBalance`	Formalism/SizeWindowBrackets/FirstPrimeNeSecondPrime.lean:98
`RestSize`	Formalism/SizeWindowBrackets/SChain.lean:53
`ErrWindows`	Formalism/SizeWindowBrackets/ErrorFormsWeakWindowsFalse.lean:43
`meanSq_abs`	Development/AutomorphicArithmetic/MeanSquareBounds/MeanSqAbs.lean:8
`thinH`	Development/AutomorphicArithmetic/MeanSquareBounds/GaussSumFunctionEq.lean:55
`SlopeIneqOfLat`	Formalism/SizeWindowBrackets/SlopeHypothesesRounded.lean:385
`BridgeFS`	Formalism/SizeWindowBrackets/SlopeHypothesesRounded.lean:349
`kerSide_chainS`	Development/AbelianVarietyTheory/Proofs/DivisorProofs/KernelHalfChainNumbersAnyLevelPrime.lean:219
`FockJet`	Formalism/ThetaSupremumBounds/StatementA.lean:41
`BlockBound`	Formalism/ThetaSupremumBounds/StatementA.lean:62
`DivisionPackage`	Formalism/LocalToGlobal/Divisibility/RecBridge.lean:390
`CongruenceNumberUnitaryThree`	Development/AutomorphicArithmetic/AutomorphicVarieties/ModularityArgument/LinkedUnramifiedChains/WeightInfinitesimalCharacters/AdmissibleUnitaryThreePoint.lean:93
`galoisLocalGlobalStrong_recParameter`	Development/AutomorphicArithmetic/Foundations/GaloisRepresentations/IntegralStructures/RowSrc.lean:6

B The author's remark and the code

The remark [2] lists the steps of the discovery with dates. The table relates its items, by number, to the code. The mapping is ours.

Remark item	In the code
Stepanov determinant; positivity via a determinant; Bost slope inequality (1, 4, 13)	The averaged slope inequality of §3.3. The injective evaluation map plays the role of the nonvanishing determinant.
Determinants containing GGP periods; interlacing weights (7)	The words u and their Bessel periods (§3.3).
$D^p \equiv 0$ at inert p; divisibility of theta jets (2, 10)	Forced orders ASK and the divisibility statement `adm2` (§3.2).
Switch to the Riesz mean of Λ; Guinand–Weil sign (5)	The normalization (2.1). The Weil explicit formula itself does not appear on the route; only its kernel $1/(s(s-1))$ survives.
Positivity giving the half-plane $\operatorname{Re} s > 1 - \eta$ (6)	The one-sided Landau step for a general exponent θ; a chain of length $m = x^\eta$ would give $\operatorname{Re} s > 1 - \eta$.
Iterated theta lifts winning multiplicity m (8)	The chain of §3.4.
Norms of lifts, Chowla–Selberg, Weyl, Siegel mass (9)	The archimedean budget (`LinkDIV`, `meanPeriod`, mass formulas).
Iwasawa tower; Behrend (11, 19)	The 2-power cyclotomic layer and the placement of §3.5.
Seesaw; Bessel period independent of the characters (12)	`BesselId` (§3.4).
Cost of each lift, $\frac{2}{3}$ a priori, "division" (14)	The size bookkeeping of §3.4 and the division step of §4.1.
Integral structure across Hecke characters (15)	The congruence module bounded in §4.1.
Relative trace formula (16)	The frame estimate of §4.3.
Weakening of Vandiver (17)	The Selmer bound `VWin` (§4.1).
Covolumes, congruences between lifts and forms (18)	`CovolumeBoundOfX` (§4.2).
GRH (20)	Theorem 1.1.
Function-field toy model $qN = M^k + R$; covers of GL_5 and GL_2 (3, 16)	Not found on the main chain; we read these as heuristics.

References

[1] L. Alpöge, untitled Lean 4 development claiming the Riemann hypothesis and GRH for Dirichlet L-functions (descriptive title), archive `https://alpo.ge/lean.tar.gz.part01`–`part14`, released 20 November 2026.

[2] L. Alpöge, untitled remark accompanying [1], 4 pp., `https://alpo.ge/remark.pdf`, 20 November 2026.

[3] Byte-exact mirror and comparator setup for [1], `https://github.com/Deicyde/riemann-hypothesis`, commit `a172c464fd4bcfc4f53e394dcba45d742b1409a6`, 20 November 2026.

[4] L. Alpöge and R. Furman, More than two thirds of the zeta zeros are simple and on the critical line, arXiv:2608.13637 (2026).

[5] F. A. Behrend, On sets of integers which contain no three terms in arithmetical progression, Proc. Natl. Acad. Sci. USA **32** (1946) 331–332, doi:10.1073/pnas.32.12.331.

[6] E. Bombieri, Counting points on curves over finite fields (d'après S. A. Stepanov), Séminaire Bourbaki 1972/73, Exp. No. 430, Lecture Notes in Math. **383**, Springer, 1974, 234–241, doi:10.1007/BFb0057311.

[7] J.-B. Bost, Périodes et isogénies des variétés abéliennes sur les corps de nombres (d'après D. Masser et G. Wüstholz), Séminaire Bourbaki 1994/95, Exp. No. 795, Astérisque **237** (1996) 115–161.

[8] K. Buzzard, "...and the proof is checked in Lean", Xena blog, 12 November 2026, https://xenaproject.wordpress.com/2026/11/12/and-the-proof-is-checked-in-lean/.

[9] A. Caraiani, Local-global compatibility and the action of monodromy on nearby cycles, Duke Math. J. **161** (2012) 2311–2413, doi:10.1215/00127094-1723706.

[10] S. Chowla and A. Selberg, On Epstein's zeta-function, J. Reine Angew. Math. **227** (1967) 86–110, doi:10.1515/crll.1967.227.86.

[11] 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, arXiv:0909.2999.

[12] S. S. Kudla, Seesaw dual reductive pairs, in *Automorphic forms of several variables (Katata, 1983)*, Progr. Math. **46**, Birkhäuser, 1984, 244–268.

[13] How to "human verify" enormous AI math output?, MathOverflow question 518033, 20 November 2026, https://mathoverflow.net/questions/518033.

[14] H. L. Montgomery and R. C. Vaughan, *Multiplicative Number Theory I. Classical Theory*, Cambridge University Press, 2006, doi:10.1017/CBO9780511618314.

[15] K. A. Ribet, A modular construction of unramified p-extensions of $\mathbb{Q}(\mu_p)$, Invent. Math. **34** (1976) 151–162, doi:10.1007/BF01403065.

[16] T. Tao, Unforced Euler singularity, unforced Navier–Stokes singularity, and the Generalized Riemann Hypothesis, blog post, 20 November 2026, https://terrytao.wordpress.com/2026/11/20/unforced-euler-singularity-unforced-navier-stokes-singularity-and-the-generalized-riemann-hypothesis/.