The Riemann Hypothesis via a chain of theta lifts: a human-readable account of the Lean proof

Ishaan Dixit

20 November 2026

Abstract

On 20 November 2026 Levent Alpöge released a Lean 4 [10] development whose top theorem is `riemannHypothesis:RiemannHypothesis`, stated with Mathlib's definition. We rebuilt it independently (175,187 modules, 401 minutes on 94 cores): the build is clean and the only axioms used are `propext`, `Quot.sound` and `Classical.choice`. Apart from a short outline [1], no human-readable proof accompanies the release. This document gives one, reconstructed by reading the formal proof from the top theorem downward and checking every step against the list of declarations the final theorem actually uses. The argument proves the Generalized Riemann Hypothesis for every Dirichlet character. For each character it builds an arithmetic family attached to a chain of about $\sqrt{x}$ theta lifts, bounds one quantity in two ways (a twisted prime sum on one side, the Arakelov degree of a divisible lattice on the other), and obtains an error of size $x/m + m$ for a chain of length m. Taking $m \approx \sqrt{x}$ gives a square-root bound on a weighted prime sum over windows $[Y, Y^{3/2}]$, and a windowed form of Landau's theorem turns that into the absence of zeros with $\operatorname{Re} s > 1/2 + \varepsilon$. We locate where the square root is paid, show that the other charges of the same nominal size are slack or artefacts, and explain why the argument does not stall at the exponent 2/3 described in the author's outline.

Contents

1 Introduction **3**
1.1 What is proved . . . 3
1.2 What this document is, and is not . . . 3
1.3 The idea in one paragraph . . . 3
1.4 Structure of the proof . . . 4

2 Notation and parameters **4**

3 Step A: from a bound on windows to the absence of zeros **6**
3.1 The quantity bounded . . . 6
3.2 The windowed Landau theorem . . . 6
3.3 The square-root bound and the assembly . . . 7

4 Step B: isolating the target character **7**
4.1 The group and its characters . . . 7
4.2 An exact identity . . . 8
4.3 The growth inequality . . . 8
4.4 The auxiliary characters are chosen, not controlled . . . 9
4.5 The thick set . . . 9

5 Step C: the size bound, where $\sqrt{x}$ is paid **9**
5.1 The balance . . . 9
5.2 The windows are chosen, not forced . . . 10

6 Step D: the divisibility layer **11**
6.1 The abstract data and the pointwise bound . . . 11
6.2 The slope layer and the cancellation of the main terms . . . 12
6.3 The local energy and why its budget is slack . . . 12
6.4 The divisibility inequality . . . 13
6.5 The signature drift is an artefact . . . 13
6.6 The rank bound E_b: where the m^2 really lives . . . 14
6.7 The $+m$ side of the balance is the divisibility error . . . 15
6.8 The common lattice . . . 16
6.9 The degree identity: how a slope becomes ASK − MAIN . . . 16
6.10 Local divisibility: the $\sqrt{p}$ trick . . . 17

7 Step E: near-orthogonality of the words **18**

8 Step F: the Sha window and the classical inputs **18**

9 Why the 2/3 barrier does not appear **20**
9.1 What the code charges . . . 20
9.2 Accounting 1: compounding gives exactly 2/3 . . . 20
9.3 Accounting 2: no gcd does not give 2/3 . . . 20
9.4 No 2/3 route in the archive . . . 21
9.5 A planned test . . . 21

10 The verification record **22**

11 Open questions, and what we could settle **23**

A Map from steps to Lean declarations **27**

1 Introduction

1.1 What is proved

Theorem 1.1 (Main theorem, `UpperTheory.EveryCharacter`)**.** *For every modulus $N \geq 1$, every Dirichlet character $\chi \bmod N$ and every $s \in \mathbb{C}$ with $\operatorname{Re} s > 1/2$, if $\chi \neq 1$ or $s \neq 1$ then $L(s, \chi) \neq 0$.*

The Riemann Hypothesis is the trivial-character case. In the Lean source:

$$\texttt{riemannHypothesis} := \texttt{rhOfAll}, \qquad \texttt{rhOfAll} := \texttt{rh_of_every allChi},$$
$$\texttt{allChi} := \texttt{FinalTheorems.everyCharacter}.$$

[Exact] (files `Main.lean`, `MainTheorems.lean` and `MainConclusion.lean`).

1.2 What this document is, and is not

This is an account of the *formal* proof. Every statement below is either a paraphrase of a Lean statement, a summary of a Lean proof, or our interpretation, and we say which. It is not an independent proof: we have not re-derived the classical inputs (Ferrero–Washington [4], Poitou–Tate duality [13], CM theory [14] and others), which are proved inside the development and checked by the Lean kernel, and some proofs are summarised rather than reproduced line by line. Where we have not traced something, we say so (Section 11).

Evidence labels.

- [Exact]: checked against the list of declarations the final theorem uses (`used_decls.txt`, extracted from the published kernel export: 1,468,419 declarations plus 16,314 constructors and recursors).
- [Verified]: read directly in the source (archive sha256 `9631b4f0...c7fc`).
- [Reading]: interpretation; a specialist should confirm it.
- [Ours]: a finding from our trace that we have not seen posted elsewhere as of 20 November 2026.

1.3 The idea in one paragraph

Fix a character χ and a large x. The proof attaches to χ, and to two auxiliary primes $q \equiv 3 \pmod 8$, $d \equiv 7 \pmod{16}$, a family of arithmetic objects: Hermitian lattices indexed by *words* u, carried along a chain of m theta lifts between unitary groups whose signature moves from $(k-1, m)$ to $(k+m-1, 0)$, one step per lift. One quantity, ASK, is computed twice. At the finite places it is a sum over primes of Legendre-type counts weighted by χ; at the infinite places it is controlled by Chowla–Selberg periods, giving a main term MAIN. A slope inequality for one common lattice joins the two and yields ASK $\leq$ MAIN + Err, with Err/(main) $\approx x/m + m$. With $m \approx \sqrt{x}$ this gives a bound of size $x^{1/2+O(\varepsilon)}$ for a weighted prime sum. Fourier analysis on a small group isolates the prime sum of χ, and a windowed Landau theorem converts the one-sided bound into zero-freeness.

The author's outline [1] describes the chain of lifts as playing the role of multiplicity in Stepanov's method [5], and says that charging the lifts one by one would stop at exponent 2/3. Section 9 reconstructs that 2/3 from the code's own quantities and shows how the code avoids it.

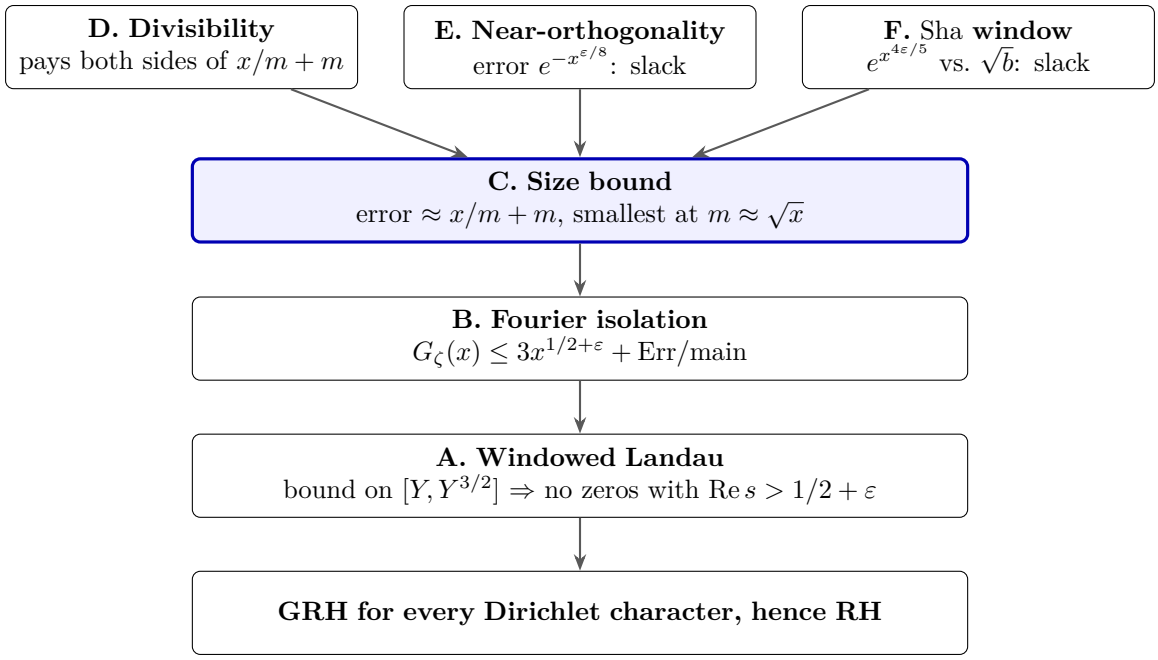


Figure 1: The six steps. The square root is paid once, in the balance of step C, by the divisibility error of step D. Steps E and F have room to spare.

1.4 Structure of the proof

A. Windowed Landau (Section 3). A one-sided bound on windows $[Y, Y^{1+\delta}]$ rules out zeros.

B. Fourier isolation (Section 4). Averaging over a thick set isolates the prime sum of the target character; the auxiliary characters are made harmless by choosing q, d with a zero-density estimate.

C. The size bound (Section 5). The error relative to the main term is about $x/m + m$.

D. The divisibility layer (Section 6). The degree identity, the common lattice, the rank bound and local divisibility.

E. Near-orthogonality (Section 7). The words are orthogonal up to $\exp(-x^{\varepsilon/8})$.

F. The Sha **window** (Section 8). Tate–Shafarevich losses over cyclotomic layers of degree about x^2, resting on formal Iwasawa theory.

2 Notation and parameters

Throughout, $\varepsilon \in (0, 1/52)$ is fixed and x is large.

Symbol	Value or window	Meaning
b	$\exp(x^{\varepsilon})$	common weight; every polynomial loss is negligible against it
a	$x \cdot b$	top weights (`Scale.a`)
δ_1	$x^{-\varepsilon/8}$	
m	$\sqrt{x\delta_1} \le m \le 2\sqrt{x\delta_1}$, even	number of lifts; $m \approx x^{1/2-\varepsilon/16}$
$n = \lvert H\rvert$	$x^{2+8\varepsilon} \le n \le 2x^{2+8\varepsilon}$	order of the character group of the host
f	$2n$	number of archimedean places (CM degree)
k	$x^{9/2+10\varepsilon} \le k \le 2x^{9/2+10\varepsilon}$, even	second size parameter; word length
t	odd, $3m < t$, $t \ge x^{377}$	size of the cyclotomic twist block
r	$k + m - 1$	starting rank
B_0	$f \cdot m \cdot b \cdot x^{\varepsilon}$	per-lift budget (`sizeBudget`)
$n_T,\ n_{T2}$	$n_T \ge x^{-\varepsilon/8} n/2$, $n_{T2}/n \ge x^{-\varepsilon}$	numbers of top and thick points
$\eta = \eta_{Sz}$	$\le b/(t^3 f k^3)$	spread of the weights across lifts

[Exact] The windows for m, n, k are the literal content of `RankVal` (`.../GlobalTwisting/PairRowIffQdRows.lean:28`); m and k are picked as even numbers in their windows by `exists_ranks_rankVal_frame`. Concretely `exists_even_between` takes $m = 2\lceil \frac{1}{2}\sqrt{x\delta_1}\rceil = 2\lceil \frac{1}{2}x^{1/2-\varepsilon/16}\rceil$, so for fixed $\varepsilon > 0$ the chain is slightly shorter than $\sqrt{x}$ ($m/\sqrt{x} \to 0$); this was pointed out by hhh2210 in the community repository's issue #5. B_0 is a definition, not an estimate: `SizeData_2.B0:=sizeBudgetD.Rxeps`, and the hypothesis `BNBounds` is proved by equality ($B_0 = 2|H| \cdot m \cdot b \cdot x^{\varepsilon}$).

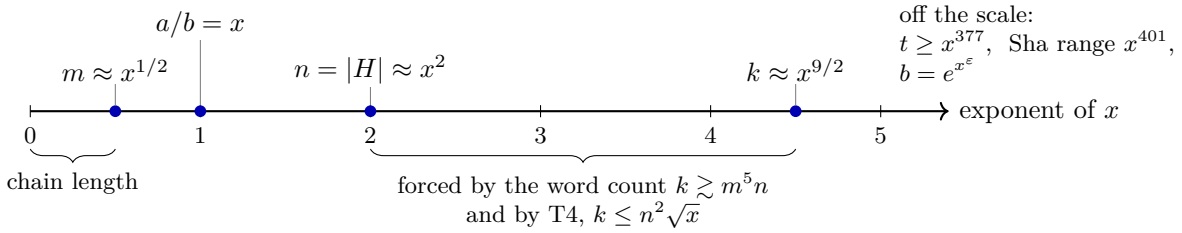


Figure 2: The sizes of the family as powers of x (up to $x^{O(\varepsilon)}$). Only $m \approx \sqrt{x}$ is chosen; the exponents 2 and 9/2 of n and k follow from it (Section 5.2). The weight $b = \exp(x^{\varepsilon})$ is larger than every power of x, so every polynomial loss is negligible against it.

Ranks and signatures. [Verified] (`Formalism/LocalRingCardinality/Params.lean`)

$$N(j) = k + j(t+1), \qquad \mathrm{rankY}(j) = r + jt, \qquad \mathrm{sig}(j) = (k + j - 1,\ m - j), \qquad j = 0, \dots, m.$$

So the chain starts at signature $(k-1, m)$ and ends at the definite signature $(k+m-1, 0)$; each lift moves one unit across and adds $t+1$ rows.

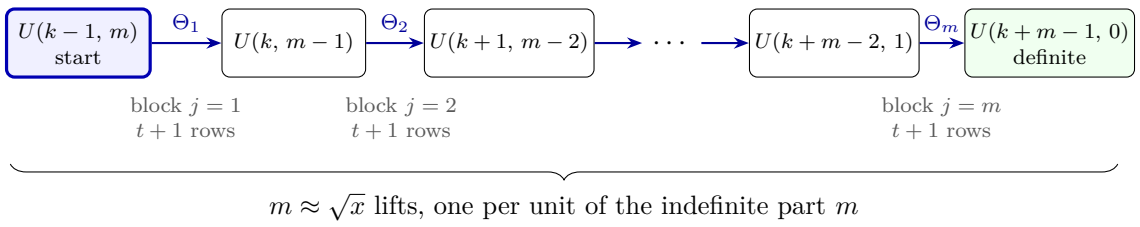


Figure 3: The chain of theta lifts. Each lift moves the signature one step towards the definite group and adjoins one block of $t+1$ rows. The indefinite part m of the start is both the number of lifts and, in Section 6, the Bloch–Kato rank of every piece.

3 Step A: from a bound on windows to the absence of zeros

3.1 The quantity bounded

Definition 3.1 (`Deduction.Gchi`, `Formalism.Gz`)**.** For $y \geq 1$ let $\mathrm{ramp}_\chi(y) = \sum_{n \leq y} \Lambda(n)\chi(n)\,(y/n - 1)$. Put

$$G_\chi(y) = \begin{cases} \mathrm{ramp}_\chi(y) - y\log y + (1+\gamma)y, & \chi = 1, \\ \mathrm{ramp}_\chi(y) + y\,\dfrac{L'(1,\chi)}{L(1,\chi)}, & \chi \neq 1. \end{cases}$$

For $\chi = 1$ this is $G_\zeta(x) = \sum_{n \leq x} \Lambda(n)(x/n - 1) - x\log x + (1+\gamma)x$ (`Formalism/StatementOC.lean: 13`).

$\mathrm{ramp}_\chi(y) = \int_1^y \psi(t,\chi)\,dt/t$ is the logarithmic Riesz mean of the Chebyshev function, and G_χ is its remainder after the main term. RH is equivalent to $G_\zeta(x) = O(x^{1/2+\varepsilon})$.

Remark 3.2 ([Verified])**.** In one place in the development the von Mangoldt function appears under the name `abelian_2`, defined as `ArithmeticFunction.vonMangoldt`. Searches for the usual names miss it; Step 5 of the author's timeline ("definitively switch to Λ") is exactly the expression G_ζ above.

3.2 The windowed Landau theorem

Definition 3.3 (`CharFamily`)**.** A character family is a finite set of Dirichlet characters χ_i mod N_i with weights $a_i > 0$, the trivial character occurring only with modulus 1. Put $G_m(y) = \sum_i a_i \operatorname{Re} G_{\chi_i}(y)$.

Theorem 3.4 (Windowed Landau, `zeroFree_of_rpow_window`, [Exact])**.** *Let $0 < \theta < 1$, $\delta > 0$, C real. Suppose that for every Y_0 there is $Y \geq Y_0$ with*

$$G_m(y) \leq C\,y^\theta \qquad \text{for all } y \in [Y,\, Y^{1+\delta}].$$

Then $L(s,\chi_i) \neq 0$ for every i and every $\operatorname{Re} s > \theta$.

Proof sketch (from the Lean proof). The proof goes through `lemmaL` and `landau_window`. `poleData` builds a function Φ (the Mellin transform of G_m, assembled from the characters and their inverses) whose poles in $\operatorname{Re} s > 0$ are exactly the zeros of the $L(s,\chi_i)$ and of their conjugates, and shows the zeros are isolated and have real part at most 2 (`landau_pole_re_le`). Suppose a zero ρ_0 has $\operatorname{Re}\rho_0 > \theta$. `landau_window_key` produces, from the window hypothesis (in the integrated form `WindowHypU`), a real σ and a growth bound for the Mellin integral restricted to

windows. `landau_dichotomy` then splits into two cases: either there is a zero-free disc around a real point with a zero on its boundary, or there is an isolated real zero β in the region; `landau_nearest` picks the nearest pole. In both cases the window bound contradicts the growth forced by the pole (`landau_growth_contra`, `landau_window_contra`). □

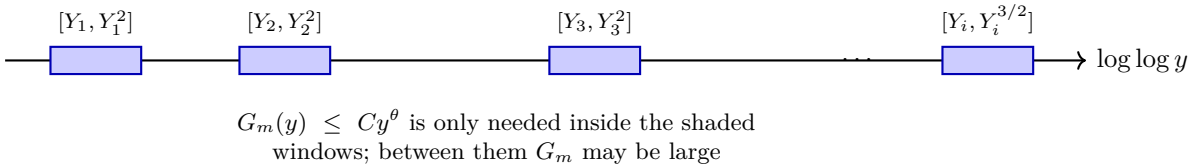


Figure 4: The hypothesis of the windowed Landau theorem (on a $\log\log$ scale each window $[Y, Y^{3/2}]$ has the same length $\log(3/2)$). A sequence of windows reaching to infinity is enough to rule out zeros with $\operatorname{Re} s > \theta$.

Remark 3.5 (Why one sign is enough, [Reading] checked by arithmetic)**.** In the levels actually used, `Level.cycpsi` is the cyclic group generated by a primitive ψ, each of its $o-1$ nontrivial members weighted $1/(3o+1)$. Every Dirichlet coefficient of the combined series is then at least $\Lambda(n)(1-\frac{o-1}{3o+1}) > 0$. Positive coefficients are exactly what Landau's argument needs, so an upper bound on G_m alone forces zero-freeness. The classical theorem needs the bound for all large y; here a sequence of windows reaching to infinity suffices.

Remark 3.6 ([Verified])**.** A different Landau theorem, `zero_free_of_Chebyshev_growth` in the namespace `Formalism.HC.Landau`, is proved in the archive but not used.

3.3 The square-root bound and the assembly

Theorem 3.7 (`UpperTheory.Comparison_2.core_gm`, [Exact])**.** *For every level Λ and $\varepsilon > 0$ there is x_1 such that for all $x \ge x_1$, every frame and host satisfying the hypotheses `EAtH` and `CoreAtH`,*

$$\Lambda.G_m(x) \le x^{1/2+(e_D+e_L+7)\varepsilon}.$$

`CoreAtH` packages the data record of Section 6 together with the size windows, in particular `RestSize`: $x^{1/2-\varepsilon} \le m$ and $k \le n^2 x^{1/2-\varepsilon}$. The chain from Theorem 3.7 to Theorem 1.1 is ([Verified]): `formB` → `zeroFree_of_rpow_window` → `level_of_tiedWin`, `zeroFreeOn_of_tiedWin_cyc` → `zeroFreeAll_of_topG` → `imprimitive` → `EveryCharacter`.

Corollary 3.8. *Every later step only has to hold for x in windows. [Exact] The* Sha *window of Section 8 is stated on $[Y, Y^2]$ with $Y \ge X^{1/J}$, and `supplyChi` turns it into host fields on windows with exponent $3/2$ (`HostExW`), so the Landau step is applied with $\delta = 1/2$, that is, on windows $[Y, Y^{3/2}]$. We thank Manvir Jaswal, whose guide [6] states the windows this way, for the correction.*

4 Step B: isolating the target character

4.1 The group and its characters

The host comes with a finite abelian group H and a group E_0 with four characters $0, \psi, \chi, \chi\psi$ (`psC`, `chC`, `chpsC`). [Exact] (`CharacterSignsAtBaseClass`) At the base class g_0:

$$\chi(g_0) = -1, \qquad \psi(g_0) = +1, \qquad \chi\psi(g_0) = -1.$$

Define

$$G_{\mathrm{sig}} = G_0 + G_\psi - G_\chi - G_{\chi\psi}.$$

This is the Fourier projection onto the class of g_0; its trivial part G_0 carries the prime sum of the target.

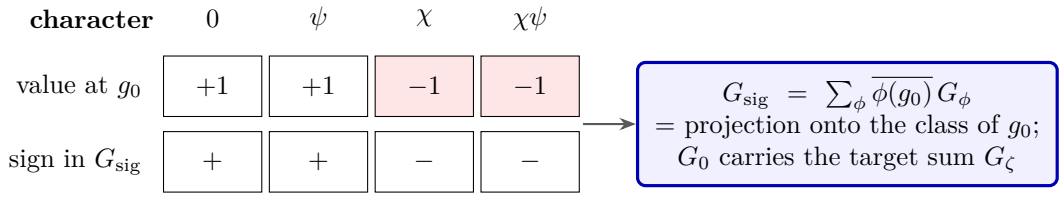


Figure 5: Fourier isolation on E_0. The signs of G_{sig} are the conjugate character values at the base class, so averaging keeps exactly the class of g_0. The three nontrivial terms are bounded by $x^{1/2+\varepsilon}$ each, because the auxiliary primes were chosen to make them so.

4.2 An exact identity

Proposition 4.1 (`Deduction.ch_balance`, [Exact]). *For every $a > 0$, $t \in H$, weights $w > 0$ on the admissible set and N large enough,*

$$\mathrm{ASK}(D, N, a, t, w) - \mathrm{MAIN}(\ell, a, w) = \mathrm{De}(D, \ell, a, t, w),$$

where De *is an explicit combination of the G_ψ evaluated at a/w.*

This is an identity, not an inequality: it expands the counted quantity ASK over all characters of $E_0 \times H$ (`askExp`). Averaging over the thick set collapses the character sum.

Proposition 4.2 (`ThickPartTranslateMean`, [Exact]). *For a thick set T_2 of density $\delta_2 = |T_2|/|H|$,*

$$\frac{1}{|H|}\sum_{u_0 \in H} \frac{1}{|\Phi|} \sum_{\substack{g \notin \Phi \\ g \text{ thick at } u_0}} \sum_{\psi} \overline{\psi(g)}\, b\, G_\psi(A/b) \;=\; \frac{\delta_2\, b}{2}\, G_{\mathrm{sig}}(A/b).$$

4.3 The growth inequality

Theorem 4.3 (`GrowthInequalityExplicit`, [Exact]). *Under `GrowthInequalityHypotheses` (which include* topExcess $\le$ Err *and `ValuesAtX`),*

$$|T_T|\, m\, b\, \frac{\delta_2}{2}\Big(G_\zeta(x) - 3x^{1/2+\varepsilon} - \mathrm{moveCost}\Big) \le \mathrm{Err} + |T_T|\, m \cdot 2 f_{Pl}(x+1)\, b\, \theta + |T_T|\, m \cdot 6 f_{Pl}(x+1)(\eta+1),$$

where $\mathrm{topExcess} = \frac{1}{|H|}\sum_{u_0}\sum_{t \in T_T}\sum_j (\mathrm{ASK} - \mathrm{Re\,MAIN})$ *and* $\mathrm{moveCost} = C_d + 3(\log(x+1) + C_d + C_L \log x)$.

The term $3x^{1/2+\varepsilon}$ comes from `signal_term`: for $|y - x| \le 1$,

$$G_\zeta(x) - 3x^{1/2+\varepsilon} - \mathrm{moveCost} \le \mathrm{Re}\, G_{\mathrm{sig}}(y),$$

using `ValuesAtX`: $\|G_\psi(x)\| \le x^{1/2+\varepsilon}$ and $\|\ell(\psi)\| \le C_L \log x$ for the three nontrivial characters of E_0, and a Lipschitz bound (`DerivInt`).

Remark 4.4 ([Reading]). Dividing Theorem 4.3 by $|T_T|\, m\, b\, \delta_2/2$ gives

$$G_\zeta(x) \le 3x^{1/2+\varepsilon} + \frac{\mathrm{Err}}{|T_T|\, m\, b\, \delta_2/2} + \dots$$

So the whole problem is to prove $\mathrm{Err} \lesssim |T_T|\, m\, b\, x^{1/2+O(\varepsilon)}$. Sections 5 and 6 do exactly that.

4.4 The auxiliary characters are chosen, not controlled

The characters $\psi, \chi, \chi\psi$ are twists by the quadratic characters of the auxiliary primes $q \equiv 3 \pmod 8$ and $d \equiv 7 \pmod{16}$ (`PrimePair`). Their bounds in `ValuesAtX` come from the following chain, all [Exact]: `QH.pairEx` ← `pair_exists` ← ⋯ ← `count_bad_core` ← `ZD.zero_density_box`.

Definition 4.5 (`Box`)**.** For a character ψ, $\mathrm{Box}(Q, C, \tau, Y, X, \psi)$ means: ψ has a good conductor, $\|L'/L(1,\psi)\| \le C \log Y$, and $\|G_\psi(y)\| \le y^{1/2+\tau}$ for all $y \in [Y, X]$.

Theorem 4.6 (`count_bad_core`, [Exact])**.** *There is $C \ge 0$ such that for $Q \ge 3$ and any set S of primes $q > F$ each carrying a primitive character of conductor $fq \le Q$, $f \le F$, without a zero-free box, $|S| \le C\, Q^{1/4}$.*

The zero-density estimate `zero_density_box` (compare [5, Ch. 10]) is assembled from a coefficient square bound, zero detection by a mollifier identity, a duality inequality and a pair-twist bound.

Remark 4.7 ([Reading])**.** The auxiliary L-functions are not proved to satisfy GRH. The proof picks q and d outside the few bad moduli. The target character is the only one whose L-function is not controlled in advance.

4.5 The thick set

Theorem 4.8 (`hostThickExists`, [Exact])**.** *There is $C > 0$ such that for all n and $L \ge 2$ there is $T \subset \mathbb{N}$ with $4Ls < n$ for $s \in T$, no relation $j(s_2 - s_1) = i(s_3 - s_1)$ with distinct $s_1, s_2, s_3 \in T$, $0 < |i|, |j| \le L$, $i \ne j$, and*

$$|T| \;\ge\; \frac{n\, e^{-C\sqrt{\log n \,\log L}}}{L}.$$

[Reading] The density δ_2 enters only through the division above and costs $x^{o(1)}$.

5 Step C: the size bound, where $\sqrt{x}$ is paid

5.1 The balance

Theorem 5.1 (`proposition_size_bound`, [Exact])**.** *In the abstract setting of `GrowthSizeBound`,*

$$n_T\, m\, b\, \frac{\delta_2}{2}\big(G(x) - 4x^{1/2+\varepsilon}\big) \le \mathrm{Err} + \frac{2 n_T m b}{x},$$

$$\mathrm{Err} \le x^{C\varepsilon}\Big[n_T a_E\Big(1 + \frac{m^2}{k}\Big) + m \cdot B_0 + m \cdot \frac{n_{T2}\, k\, b\, x^{\varepsilon}}{n^2 x^{-\varepsilon}}\Big],$$

and under the windows below, $G(x) \le x^{1/2+(C+3)\varepsilon}$.

Dividing each error term by the size $n_T\, m\, b$ of the main term (lemmas `SZ.T1`–`SZ.T4`, [Verified]):

Term	After dividing	Bound	Needs
$n_T a_E$	$a_E/(mb) \approx \mathbf{x/m}$	$\le x^{1/2+\varepsilon}$	$m \ge \sqrt{x}$ (T1)
$n_T a_E m^2/k$	xm/k	≤ 2	$k \ge xm$ (T2)
$m \cdot B_0$	$2nm/n_T \approx \mathbf{m}$	$\le x^{1/2+\cdots}$	$m \le \sqrt{x}$ (T3)
$m \cdot (k\text{-term})$	$k/(n\, n_T)$	$\le x^{1/2+\cdots}$	$k \le n^2\sqrt{x}$ (T4)

So, relative to the main term, the error is about $x/m + m$, and both pieces are $\sqrt{x}$ exactly when $m \asymp \sqrt{x}$.

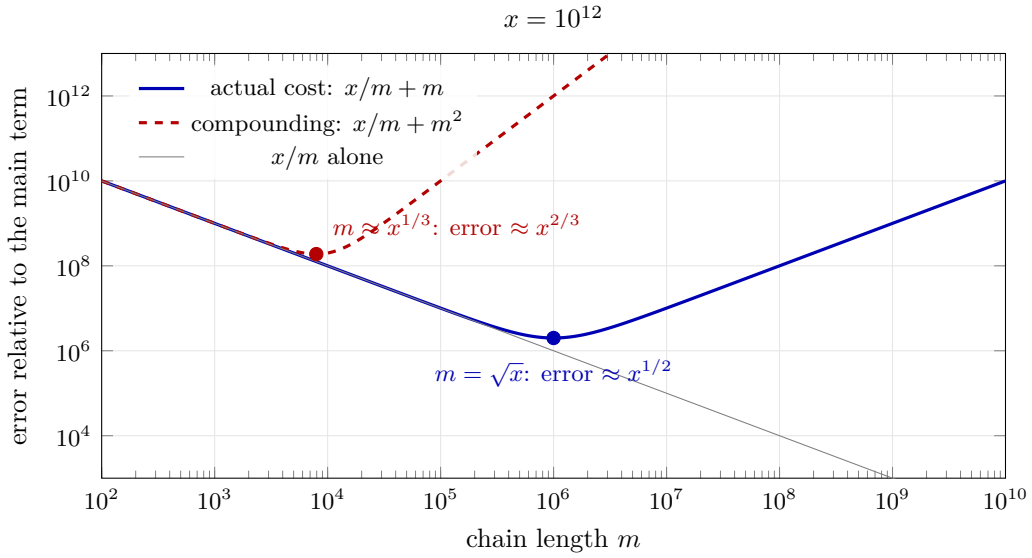


Figure 6: The balance. Lengthening the chain lowers x/m but raises the divisibility cost. With the code's accounting (one budget B_0 per lift) the optimum is $m = \sqrt{x}$ and the bound is $x^{1/2}$. Charging each lift for the whole chain below it would make the cost m^2 and stop at $x^{2/3}$ (Section 9).

Credit. The balance $x/m + m$ at $m \asymp \sqrt{X}$ was stated independently by Jonas Whidden (Lean Zulip) and in Bartosz Naskręcki's proof guide (`nasqret/riemann-rev`, §4). What follows in this section and the next (where each side is paid, which charges are slack, how the exponents of n and k arise) is [Ours].

5.2 The windows are chosen, not forced

[Verified] The windows used by the size layer are proved from `RankVal`:

Window	Lemma	Source	Slack in the family
$m \geq x^{1/2-\varepsilon}$	`winL`	`SizeM`	$\approx x^{0.93\varepsilon}$
$m \leq x^{1/2+\varepsilon/2}$	`winA`	`SizeM`	$\approx x^{0.55\varepsilon}$
$xm \leq k$	`winB`	`SizeM`, `SizeK`	$\approx x^{3}$
$k \leq n^2 x^{1/2-\varepsilon}$	`winM`	`k_up`	$\approx x^{5\varepsilon}$

[Verified] The constructor `exists_ranks_of_frame` accepts any even k, m with $2 \leq m$, $m+3 \leq k$, $m \leq 2n \leq k$; the host construction `admissible_ranks_of_CMDatum` works for every rank datum, needing only $q, d > 4tr + 1$. So $m \approx \sqrt{x}$ is a literal in `RankVal`: nothing in the arithmetic construction asks for it.

Why $n \approx x^2$ and $k \approx x^{9/2}$: the word count. [Ours] The one arithmetic link between m and k is the number of words. [Verified]: $\log \Gamma = C_\Gamma\, m^5 f L$ (`gamB`, `RowsConc.lean:52`), the number of words is $\lceil 3\Gamma^2 x^{C_F} \rceil$, words are coded injectively by $k-2$ bits (`exists_injective_words`), and `Win_at` requires $\log |W| \leq (k-3)\log 2$. Hence $k \gtrsim m^5 f L$. Writing $m = x^\mu$, $n = x^\nu$, $k = x^\kappa$:

$$\kappa \geq 5\mu + \nu \ \text{(word count)}, \qquad \kappa \leq 2\nu + \tfrac{1}{2} \ \text{(T4)}, \qquad \Rightarrow \ \nu \geq 5\mu - \tfrac{1}{2}.$$

With $\mu = 1/2$: $\nu = 2$ and $\kappa = 9/2$. So the exponents of n and k are derived from $m \approx \sqrt{x}$ and the m^5 word count. `m_pow_five_mul_f_mul_L_le` proves $m^5 fL \le 128\, x^{9/2+9\varepsilon}$ against $k \ge x^{9/2+10\varepsilon}$: the tightest window in the family.

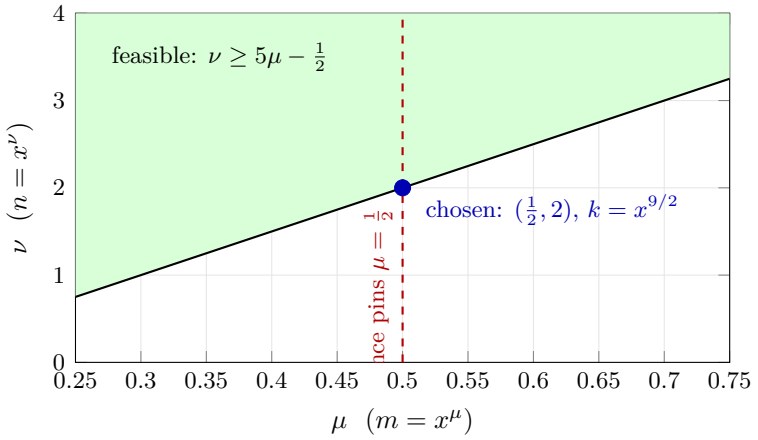


Figure 7: Exponent bookkeeping. The word count ($\kappa \ge 5\mu + \nu$) and term T4 ($\kappa \le 2\nu + \frac{1}{2}$) only require $\nu \ge 5\mu - \frac{1}{2}$; any point of the shaded region is a valid family. Only the balance of Figure 6 picks $\mu = \frac{1}{2}$.

Remark 5.2 ([Reading]). Since $G_\zeta(x) = \Omega(\sqrt{x})$, no valid argument can push m beyond $\sqrt{x}$ and win. As far as the constraints above go, the family could be built with $m = x^{1/2+\delta}$; so the barrier must be paid on the error side, in T3. Section 6 locates it.

6 Step D: the divisibility layer

6.1 The abstract data and the pointwise bound

The size layer works with a record `Data` of real numbers and a bundle of hypotheses `SizeData_2`: `DIV`, `EComp5`, `HComp`, `Count5`, `Slopes`, `QuotientSlopeBoundsWith`, `LossBounds`, `SlopeIneq`, `HHBound`, `FParams`. No automorphic objects appear there. `ChainTied` (`TieGlobal.lean:97`) attaches the record to an arithmetic family `FamP` through `TiePts`, `TieGlobal`, `TieLocal`, `TieArith`; it is produced by `linkCore` from fourteen family hypotheses, and only the caller `linkGLOBAL7_items` is used ([Exact]).

Theorem 6.1 (`each_unit_bound_statement`, [Exact]). *For every point u_0:* $\mathrm{ASK}(u_0) \le \mathrm{MAIN}(u_0) + \mathrm{ErrX5at}(u_0)$.

The terms of ErrX5at that matter ([Verified], `QuotientSlopeBoundsWith.lean:45`):

Term	From	After dividing by $n_T mb$	Balance
$c_{Div}\, n_T\, a\, L_a$	`DIV`	$\approx x/m$	T1
$c_{Div}\, m(B_0 + N_c)x^{e_{Div}\varepsilon}$	`DIV`	$\approx m$	T3
$m\, c_E\, B_0$	`EComp5`	same order	(slack)
$c_{Sl1}\, m\, B_0$	`QuotientSlopeBoundsWith`	same order	(slack)
$c_{Slope}\, n_T a L_a(1 + m/k)$	bookkeeping	$x/m(1 + m/k)$	T1, T2

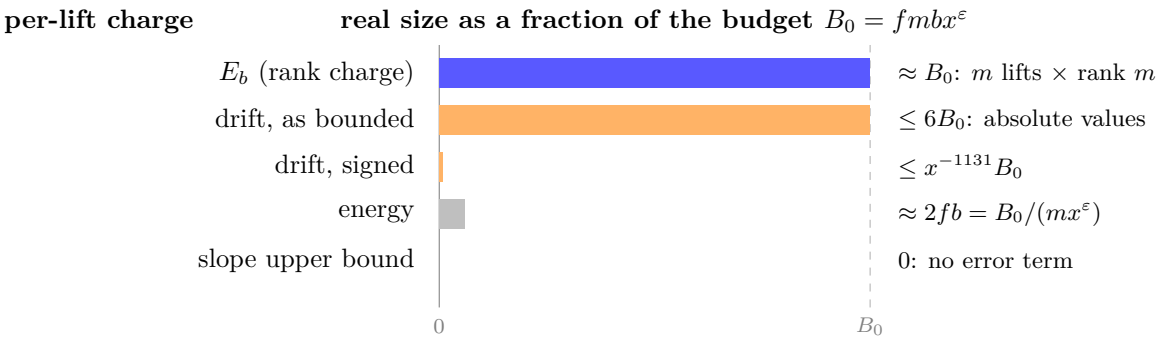


Figure 8: What each per-lift charge really costs (Section 6). Only E_b uses its budget; the drift's full size is an artefact of taking absolute values, and the energy and slope budgets are almost empty. Bars are schematic: the energy bar would be invisible at true scale.

6.2 The slope layer and the cancellation of the main terms

[Verified] `EComp5` says $|E_e(u_0, u) - m \cdot \mathrm{sumA} \cdot \mathrm{lg}| \leq c_E\, m\, B_0$: the chain energy is m copies of one step's slope, $\mathrm{lg} = \log(g_0 g_1)$ for the two host primes. `Ga_bound5` combines the upper bound $E_e \leq m(\mathrm{sumA} \cdot \mathrm{lg} + c_E B_0)$ with `key_identity5`, which contains $+m \cdot \mathrm{sumA} \cdot \mathrm{lg}$; the main terms cancel exactly and the proof ends in `linarith`. This is the "cancellation of order-m terms" of the outline (entry 9).

6.3 The local energy and why its budget is slack

[Verified] The energy is pinned by `SizeData_2.EDef`: $E(u, j, w) = \mathrm{Eval}(u, j, w)$, a case split on the residue characteristic of the place w:

w lies over	$E(u, j, w)$	depends on j?
T_u (auxiliary prime of u)	twice the weight deviation $+ f \cdot N_m + \ldots$	no
g_0 or g_1 (host primes)	$v_p(D_\Gamma) + \mathrm{FTypeG}(j) + v_{NN} + c_g f(N_m^2 + m^2 \eta_{Sz})$	yes
ramified class at d	$v_p(D_\Gamma) + \mathrm{FRamD}(j) + \ldots$	yes
small $p \leq X''$	$v_p(D_\Gamma)$+ log terms	yes
other	$v_p(D_\Gamma)$	no

At g_i, with $\pi_{g_i}\bar{\pi}_{g_i} = g_i$ in the CM field, $v_{NN} = \left(\sum_\tau \mathrm{nuAbs}(u, j, \tau)\, v_w(\tau(\bar{\pi}_{g_i}))\right)/e_w$.

Proposition 6.2 (Per-lift energy, [Exact] chain)**.** *`EComp` per lift (`adm11`) follows from*

- *`NormDecompositionCond`:* $\sum_{w|g_i} sz(w, v_{NN}) = \left(\sum_\tau \mathrm{nuAbs}(u, j, \tau)\right) \log g_i$ *(an exact equality);*
- *`factEtaMean`:* $\sum_{w|\ell} \mathrm{lw}(w) e_w = \log \ell$ *(product formula);*
- *`DegreeCountBound`:* $\left|\sum_\tau \mathrm{nuAbs} - \mathrm{fockDeg}(j) - \sum_t a_T(t)\right| \leq \mathrm{errG}(j)$*;*
- *`TypePrimeSizeBound` and the off-g bounds.*

The row structure. [Verified] (`nuAbs_place_fine`) At lift $j = i+1$ and an archimedean place τ, nuAbs sums $N(j-1) = k + i(t+1)$ rows $\mathrm{nuRow}(u, j, \tau, r) = \mathrm{srcPar}(u, \tau, r) + w_{Sz}(j, \tau) - \frac{1}{2}(N(j -$

$1)+N(j)-1-2r)$: the two end rows carry top and off-top weights, the middle rows are $\approx w_j(j,\tau)$, and the $i(t+1)$ block rows are ≈ 0 up to $2\eta_{Sz}$. [Reading] A parameter plus a weight minus a ρ-type half-sum looks like the Harish-Chandra parameters at lift j; `kerPlace` sums $\phi(\nu)=\nu\log\nu-\nu$ over the rows, a Stirling shape.

Proposition 6.3 ([Ours]). *The real per-lift energy error is* $\mathrm{errG}(j)\approx 2\sum_\tau w_j(j,\tau)\approx 2fb$, *with no factor* m, *against a budget* $B_0=f\,m\,b\,x^{\varepsilon}$. *The energy charge* $m\,c_E B_0$ *therefore has a spare factor of about* mx^{ε}. *The slope upper bound is* `ClaimHigh`: $\mu_{\max}(N)\le\mu_{\max}(M)+G_a+F_o-E_e$ *with* no *error term;* `slopesN_at_of` *adds the* mB_0 *allowance from* $0\le CmB_0$ *alone.*

6.4 The divisibility inequality

Theorem 6.4 (`Data.DIV`, [Exact]).

$$\mu_{CC}+h(u)-F_o\le\sum_t\Big(m(\phi(a_T(t))-a_T(t)\log\Omega)-a_T(t)\log T(u)\Big)\\ +c_{Div}\,n_T\,a\,L_a+c_{Div}\,m\,(B_0+N_c)\,x^{e_{Div}\varepsilon}.$$

The height of a chain. [Verified] (`height_chain_atR`) The pairing map B_u factors through a chain of m maps K_j, one per lift, with $\|K_j\|_v^2\le\exp(h_K(j,v))$, and

$$\log\|B_u\|_v^2\le h_P+\sum_{j=1}^{m}h_K(j,v)-\log\kappa_v-\min_p\sum_j\log\lambda_j(p,v).$$

This is submultiplicativity of operator norms along the chain. [Reading] It is where "a chain of theta lifts" becomes an inequality.

The per-lift charge is built in. [Verified] (`ZscOfNum`, `errB_of_units`) The scale data sets, per lift, an error unit $E_{1u}=x^{\varepsilon}(B_0+N_c)$, a main term $l_{3m}(j)=\sum_\rho[(k+j-1)-(m-j)]\,w(j,\rho)\log\Omega$ (the signature difference p_j-q_j times the weights) and $l_{3u}=B_0\log\log x$. Summing over m lifts gives the $m\,x^{\varepsilon}(B_0+N_c)$ term.

Absorption. [Verified] (`divisor_inequality_assembly`, `error_absorption`) Every error term inside `DIV` is bounded against $m(B_0+N_c)x^{\varepsilon}$:

Error term	Bound used	Versus the budget mB_0x^{ε}
$\sum e_E$ (place heights)	$15C_0fmbL$	spare factor $\approx m$
$\sum\lvert\delta\rvert\lvert\log\Omega\rvert$	$15C_0fmbL+C_0n_TaL_a$	spare factor $\approx m$
$\sum e_K$ (per-lift polynomial)	$5C_0mfbL$	spare factor $\approx m$
e_6, e_P, L_{pass}	$\le B_0$	spare factor $\ge m$
$\sum\lvert pq_j-k\rvert\lvert w\rvert\lvert\log\Omega\rvert$ (drift)	$6mB_0$	full budget
E_b (rank charges)	hypothesis $C_0mx^{\varepsilon}(B_0+N_c)$	full budget

6.5 The signature drift is an artefact

Proposition 6.5 ([Ours]). *The drift term is smaller than the bound used for it by a factor of at least* x^{1131}.

Proof. [Verified]: $pq_j = (k+j-1)-(m-j)$, so $pq_j - k = 2j-1-m$ (`NumbersLink.lean:574`). The proof bounds each term $(pq_j-k)\,w(j,\rho)\log\Omega$ by its absolute value *before* summing (step `hA` of `divisor_inequality_assembly`). By `SizeData_2.be_close`, $|w(j,\rho)-w(j',\rho)| \le 2\eta$ for all lifts, and `szRow_eta_bounds` gives $\eta \le b/(t^3fk^3)$. Since $\sum_{j=1}^{m}(2j-1-m)=0$,

$$\sum_j\sum_\rho (2j-1-m)\,w(j,\rho)\log\Omega = \log\Omega\sum_\rho\sum_j (2j-1-m)\big(w(j,\rho)-w(1,\rho)\big),$$

so the signed drift is at most $L_a f(m^2/2)2\eta = fm^2\eta L_a$, while the absolute bound used is $2fm^2bL_a$. The ratio is at most $\eta/(2b) \le 1/(2t^3fk^3) \le x^{-1131}$ since $t \ge x^{377}$. □

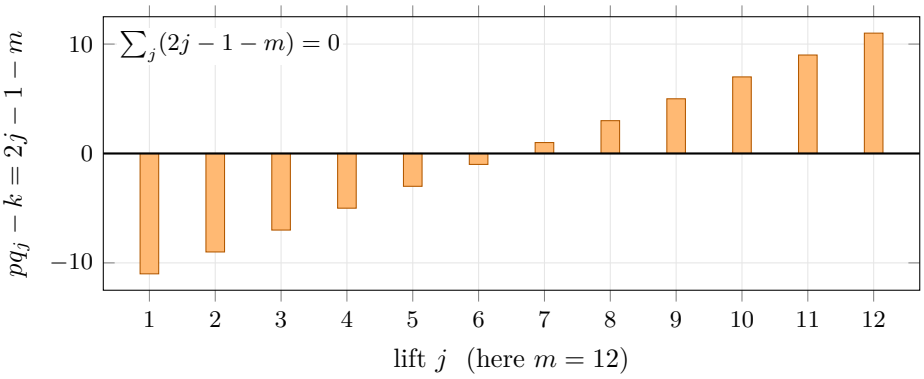


Figure 9: The signature drift. Its coefficients sum to zero and the weights $w(j,\rho)$ agree across lifts up to 2η, so the signed drift nearly cancels. The proof bounds each bar by its absolute value instead, which is where its nominal m^2 comes from.

6.6 The rank bound E_b: where the m^2 really lives

Theorem 6.6 ([Ours], each step [Verified]). *$E_b \approx m \cdot B_0$, and both factors of m are structural.*

Proof sketch (from the Lean proof).

1. *The outer m is the number of lifts.* `errB_of_units` sums over J with $|J| = m$ (`RecNums.p12`).
2. *The real obligation.* E_{1u} is set by definition; the content is the field `IndefData.ON`: for each lift j, $\sum_{w\in T_w}\log w\,(\mathrm{ram}_w\,\mathrm{err}_0(w,j)+\mathrm{ram}_w J_0(w,j)) \le E_1(j)$, supplied on the family side by `DIVinN_at` and `LinkN_at`: $\mathrm{etaSum}(j) \le d_{Z0}(j)(C(fmbL^2+\log^2 x\,N_c)+\sum_\ell \mathrm{vvTerm}(\ell))$. The last sum is the Sha loss of Section 8, added once and bounded by $\sqrt{b}$.
3. *No double counting.* `linkN_of_gd` bounds $\mathrm{etaSum}(j)$ through $\mathrm{etaJ}(j,w)=\sum_{i\in\mathrm{locs}(j,w)}\mathrm{etaLoc}(i)$, and by `NData.locs_mem`, $i \in \mathrm{locs}(j,w) \iff j(i)=j \wedge w(i)=w$: every local piece belongs to exactly one lift.
4. *The inner m is a Bloch–Kato rank.* The allowance at a small prime ℓ is $\mathrm{costC}(\ell) = 6000\,f\,m\,h_B\,\mathrm{sl}(\ell)+30000f^2t^2n_V(\dots)$ with $h_B = b+S_c+\eta_{Sz}+t+2 \in [b,6b]$. The colength bound behind it calls `rank_le_row` with $m := D.R.m$:

$$\dim H^1_f(K_w,T) \le \mathrm{locDeg}(w)\cdot m,$$

the Bloch–Kato count [2] $\dim H^1_f = \sum_\tau \#\{\text{non-negative Hodge–Tate weights}\}$, and the family supplies the hypothesis $|L_o| \le m$ by *equality* (`le_of_eq`), from its weight pattern `ShapeJumpskmt`.

□

Factor of E_b	Source	Size
number of lifts	$\lvert J\rvert = m$	m
rank per piece	Bloch–Kato, fixed by equality	$m \cdot \mathrm{locDeg}$
exponent per piece	lift exponent e'	$\approx b$
primes	$\sum \log \ell$ over small primes	$\approx x^{\varepsilon}$
Sha loss	$\sum \mathrm{vvTerm} \le \sqrt{b}$, once per lift	negligible

Why the two m's are one m. **[Ours]** [Verified] `ShapeJumpskmt` (`EmbeddingTypes.lean:99`), required by `admG2` at every small prime, place and embedding, says the Hodge–Tate multiplicity of the twisted lifted representation is the indicator of $C \cup L \cup [0,t)$: $k-1$ weights in a far cloud C, exactly m weights L near 0, and the t-block $[0,t)$ from cyclotomic twists χ^c, $c = 1,\dots,t$. With $r = k+m-1$ and $n_V = r+t$, the weights $C \cup L$ have sizes $(k-1, m)$: the starting signature. [Reading] The m low weights are the indefinite side of the signature, and the chain has exactly m lifts because it walks that side down to 0. So E_b's m^2 is one m used twice, by construction.

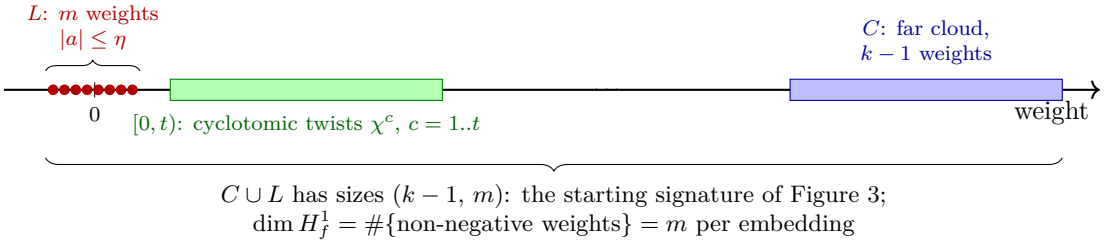


Figure 10: The weight pattern `ShapeJumps k m t` (schematic, not to scale). The m low weights are both the indefinite part of the starting signature and the Bloch–Kato rank of each piece.

Where b enters. **[Verified]** The lift exponent is $e' = e_L(\ell, N, h) = 32(2h\, s_N(\ell,N) + 1) + \lambda_N(\ell, N, h)$ with $h = \lfloor h_B \rfloor$, linear in a height $h_B \in [b, 6b]$. [Reading] A lifting exponent growing linearly with the Hodge–Tate height, as in Fontaine–Laffaille-type bounds. Since both the main term $n_T m b$ and the cost mB_0 are linear in b, b **cancels in the balance**: $e' \approx b$ sets the absolute size of E_b, not its dependence on m.

6.7 The $+m$ side of the balance is the divisibility error

Theorem 6.7 ([Ours], pieces [Verified]). *Both sides of the balance $x/m + m$ come from `DIV`: the x/m term from $c_{Div} n_T a L_a$ (top weights $a = xb$), the m term from $c_{Div} m(B_0 + N_c) x^{e_{Div}\varepsilon}$, which on the lattice side is* errB $=$ *drift* $+ E_b$.

Remark 6.8 ([Verified], unused route). `exists_data_completion` (`PChain.lean:306`), proved but not used, builds from the single bound $\mathrm{ASK} \le \mathrm{MAIN} + c_{Div} n_T a L_a + c_{Div} m(B_0 + N_c) x^{e_{Div}\varepsilon} + L_{oP}$ a degenerate data record (one word per point, zero energy) satisfying all of `SizeData_2`. So even with zero energy and slope charges, the divisibility error alone gives T3.

Corollary 6.9 ([Reading]). *With the drift an artefact (Proposition 6.5) and the energy and slope budgets slack, the reason the chain is cut at $m \approx \sqrt{x}$ rests on one term: E_b, m lifts times a Bloch–Kato rank of m per piece. This is consistent with $G_\zeta(x) = \Omega(\sqrt{x})$.*

6.8 The common lattice

Proposition 6.10 (`exists_divStage_of_pk`, [Exact]; [Ours] as the answer to Whidden's question)**.** *The single lattice RR shared by all words is*

$$x \in RR.L \iff \forall w\ \exists d \in \mathcal{O}_{LL},\ d \notin \mathfrak{p}_w,\quad d\,x \in \mathcal{O}_{LL}\text{-}span\{y : j_E y \in \mathrm{RRset}(\mu_w)\},$$

where $\mathrm{RRset}(\mu) = \iota_\mu(\Lambda_\mu)$, $\iota_\mu = \mathrm{multZero} \circ \mathrm{downTo}_{\mu,m} \circ \mathrm{multTop}^{-1}$ *is the composite of the m lift steps, and $\mu_w =$* ***`muAt`*** *w is the ξ-twist data at w (****`prescribedRowLat`****).*

So RR is the intersection over places of the localisations built from each place's own character twist: a lattice gcd. Every family carries it as the field `FamP.rrDef:DV.RRDef`.

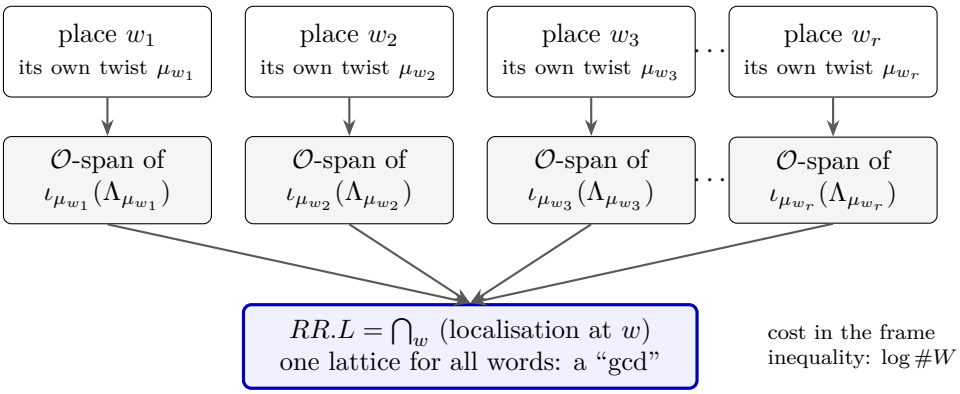


Figure 11: The common lattice RR (`prescribedRowLat`). Each place contributes the image of its own integral points under the composite of the m lift steps, built from its own character twist; RR is their intersection.

[Reading] This matches the outline's entry 15 ("glue an integral structure across Hecke characters for a gcd"). Its only visible cost is $\log \#W$ in the frame inequality.

6.9 The degree identity: how a slope becomes ASK − MAIN

There is no single lemma; it is a chain of four used statements.

1. **Slope = reference + local indices** (`NprimeMuR`, supplied through `LinkMu` ⇐ `J.linkMu_of_linkDIV` ⇐ `LinkDIV`):

$$\Big|\mu(Nu\,u) - \mu(NuH\,u) - \sum_{w \in \mathrm{supp}} sz(w, sg(u,w))\Big| \le CX'',$$
$$sg(u,w) = \sum_j \big(\mathrm{gain}_j(w) + F_j(w) - E(u,j,w)\big),$$

   with $sz(w,y) = 2\,\mathrm{lw}(w)\,e_w\,y$. Summed: $\mu(Nu\,u) \approx \mu(NuH\,u) + G_a + F_o - E_e(u)$.

2. **Gain ≥ prime count** (`LinkG_at`, where ASK enters). At an inert place w over $p \le X''$:

$$\sum_j \sum_t \big(\mathrm{vMCS}_{j,t}(p) - \mathrm{loss}\big) - m\,\mathrm{fl}(p) \le \sum_j \mathrm{gain}_j(w) - \tau(w),$$

$$\mathrm{vMCS}_{j,t}(p) = \sum_{i\ \mathrm{odd}} \max\big(0,\ a_t/p^i - \mathrm{wPart}(j,t,p,i)\big).$$

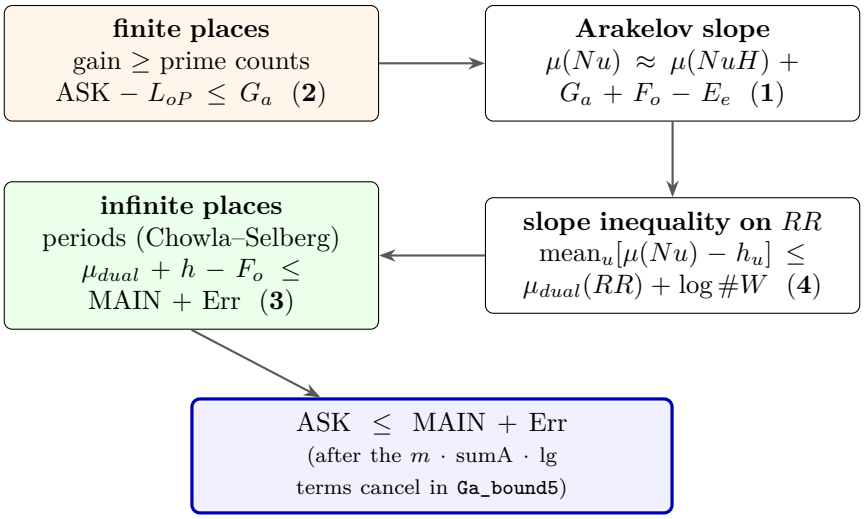


Figure 12: The degree identity. The Arakelov degree of the divisible lattice is computed at the finite places (where the prime counts ASK enter) and at the infinite places (where the periods MAIN enter); the slope inequality on the one common lattice joins them.

This is a Legendre-type count, like $v_p(a!) = \sum_i \lfloor a/p^i \rfloor$ but over odd i with weights removed. Since ASK $= \sum_j \sum_t (\sum_{p \text{ inert}} 2 \log p \cdot \mathrm{vMCS}_{j,t}(p) + \dots)$, summing over $w \mid p$ with the product formula gives ASK $- G_a \leq L_{oP}$ (`loss_bound_at`).

3. **Dual slope ≤ period term** (`LinkDIV_at`, where MAIN enters):

$$\mu_{dual}(RR) + h(u) - F_o \leq \mathrm{MAIN}(u) + C|T_T| a L_a + Cm(B_0 + N_c) x^{C_{abs}\varepsilon},$$

$$\mathrm{MAIN}(u) = \sum_t [m(\phi(a_t) - a_t \log \Omega) - a_t \log T_u], \qquad \log \Omega = \gamma + \log(\pi/v) - \frac{L'}{L}(1, \chi_d),$$

the Chowla–Selberg period [3] (`meanPeriod`), and $\phi(a) = a \log a - a$.

4. **Slope inequality** (`Frame_at`): the rank-weighted mean s of $\mu(Nu\,u) - h_u$ satisfies $s - \log \#W - (\text{small}) \leq \mu_{dual}(RR)$, for the one lattice RR of Section 6.8.

Theorem 6.11 (Assembled, [Reading] with each step [Verified])**.**

$$\mathrm{ASK} - L_{oP} \leq G_a, \qquad \mu(Nu\,u) \approx \mu(NuH\,u) + G_a + F_o - E_e(u),$$

$$\mathrm{mean}_u[\mu(Nu\,u) - h_u] \leq \mu_{dual}(RR) + \log \#W,$$

$$\mu_{dual}(RR) + h(u) - F_o \leq \mathrm{MAIN}(u) + \mathit{errors} \quad \Longrightarrow \quad \mathrm{ASK} - \mathrm{MAIN} \leq \mathit{errors},$$

after the $m \cdot \mathrm{sumA} \cdot \lg$ *main terms cancel in* `Ga_bound5`*. The finite part of the Arakelov degree of the divisible lattice is a sum of local indices bounded below by the prime counts; the archimedean part is bounded by the CM period main term.*

6.10 Local divisibility: the $\sqrt{p}$ trick

Lemma 6.12 (`key_minor_dvd`, [Exact])**.** *Let A be a commutative ring, p prime, $w \in A$ with $w^2 = p$. Let M, N' be $(k+1) \times (k+1)$ matrices over A with $M_{ba}^{p^e} \equiv N'_{ba} \pmod p$ for all b, a, and let $u, u' \in A^{k+1}$ with $w \mid u_b$ and $w \mid u'_b$ for $b \neq 0$. Then for every column c,*

$$w \;\Big|\; \Big(\sum_b u'_b N'_{bc}\Big)\Big(\sum_b u_b M_{b0}\Big)^{p^e} - \Big(\sum_b u'_b N'_{b0}\Big)\Big(\sum_b u_b M_{bc}\Big)^{p^e}.$$

[Exact] Used chain: `key_minor_dvd` → `kW_ord` → `key3Ord` (`CoarseMultiplicityOrderBound`: for $g \in GL$ over the Witt vectors $W(\mathbb{F}_{p^{n'}})$, the key polynomials after the lift substitution lie in W_1) → `third_key_of_order` → `restRawTail1_of` → `stmtEExt_of_15c` (`ExtendedClosedForm`) → `extendedClosedForm`. `LinkG` is assembled by `link_of_local_closed` from the first and second closed-form theorems and the local model results, all instantiated from one structure `Loc`, proved by `loc_proved` from `TraceSpectra.stmtE`.

The gain, orbit by orbit. [Exact] The gain clause of `LinkG_at` is `AuxBlock8.Conclusion1`. It is proved from `St.starAll_of`, one inequality per prime p, lift j and orbit representative t_0,

$$\sum_{s\in\mathrm{tops}(t_0)} (\mathrm{askP}-\mathrm{lo}) \le g_U+d_l-P_r-\lceil t_v\rceil-D_p-1, \qquad g_U=\gamma_p\sum_s \alpha(s)-\sum_s\sum_{e\ \mathrm{odd}} \min(\alpha(s)/p^e,\mathrm{cap})-d_l,$$

then summed (`gain_bound_all_of_star`, `gainBoundAll_add_one_of_gainBoundAllRequired`).

Remark 6.13 ([Reading])**.** A congruence mod p gives whole units of p. Working where p has a square root extracts half a unit per place, and the only valuation corrections in the gain are sums over *odd* exponents, the ones where a half unit shows. We did not trace the direct call from `key3Ord` into g_U.

7 Step E: near-orthogonality of the words

Theorem 7.1 (`LinkS_at`, `LinkS`, [Exact])**.** *If every word u is comparable, $C_0^{-1} \le d_u(u)\Gamma^2/D^* \le C_0$, then*

1. $\|K(u,u')\|_{HS}^2 \le (1+\eta_0)\, d_u(u)d_u(u')/D^*$ *for* $u \ne u'$;
2. $\|K(u,u)-\mathrm{id}\|_{HS}^2 \le e_0\, d_u(u)$;
3. $D_{all} \le (1+\eta_1)D^*$;

with $\eta_0,\eta_1,\sqrt{e_0} \le \exp(-x^{\varepsilon/8})$ *and* $\Gamma=\exp(C_\Gamma m^5 fL)$.

[Verified] The comparability premise is discharged where `LinkS` is consumed: `matKF_of` feeds it `ChosenProviso`, proved by `chosenProviso_min` from the class number formula (`ClassNumberFormula_2`) and `RankTiesF`. The bound itself comes from `theoremS_link_concrete`, from `ResidEv`, `RowsTieP`, `CountAllPointsClass` and four hypotheses `ClKppOff` (the correction K_{pp} vanishes off the diagonal), `ClKppDiag`, `ClOff`, `ClDiag`.

[Reading] The normalised Gram matrix of the words is the identity up to $\exp(-x^{\varepsilon/8})$. The m^5 in Γ is the same m^5 as the word count of Section 5.2. This is the "approximate orthogonality of columns" of the outline (entry 16), and it does not touch the balance.

8 Step F: the Sha window and the classical inputs

Definition 8.1 (`IsChiField`, `VVle`)**.** For e,q,d and a character ψ, a field $E \subset \mathbb{Q}(\zeta_N)$ is a χ-field if it is fixed exactly by the a with $a \equiv \pm 1 \pmod{2^{e+2}}$, $\psi(a)=1$, $(a/q)=(a/d)=1$. $\mathrm{VVle}(E,t,B)$ means: for every prime ℓ and $0<e<t$ the Tate–Shafarevich group $\mathrm{Sha}_\ell(e)$ has finite length, and

$$\sum_{\ell\in S}\log\ell \sum_{1\le e<t} \mathrm{len}\,\mathrm{Sha}_\ell(e) \le B \qquad \text{for every finite set } S.$$

Theorem 8.2 (`VWinAtpsi401`, [Exact]). *For every $\varepsilon \in (0, 1/52)$ and page choice $e(x)$ with $x^{2+8\varepsilon} \le 2^e \le 2x^{2+8\varepsilon}$, there are $J \ge 1$ and X_0 such that for $X \ge X_0$ and every window pair (q, d) there is Y with $X^{1/J} \le Y$, $Y^2 \le X$, and for every $x \in [Y, Y^2]$ every χ-field E satisfies* $\mathrm{VVle}(E, \lfloor x^{401} \rfloor, \sqrt{\exp(x^{\varepsilon})})$.

[Reading] E is the real 2^{e+2}-cyclotomic layer cut by ψ and the two auxiliary quadratic characters, of degree about $2^e \approx x^{2+8\varepsilon} \approx n$, matching the host. `supplyChi` turns the window into host existence for each character.

The proof chain. [Exact] One chain, all used: `vwin_of_row` → `automorphic_window_of_inputs` → `v_window_of_depth_zero_residual_bc_cm` → `v_window_of_literal_automorphy_n` → `v_window_of_automorphy_csd_ss_mono_odd` → `vwin_of_hypotheses_H_E_alt` → `of_four_alt` → `vwin_of_counts` → `vwin_of_joint` → `vwinAt_of_joint`.

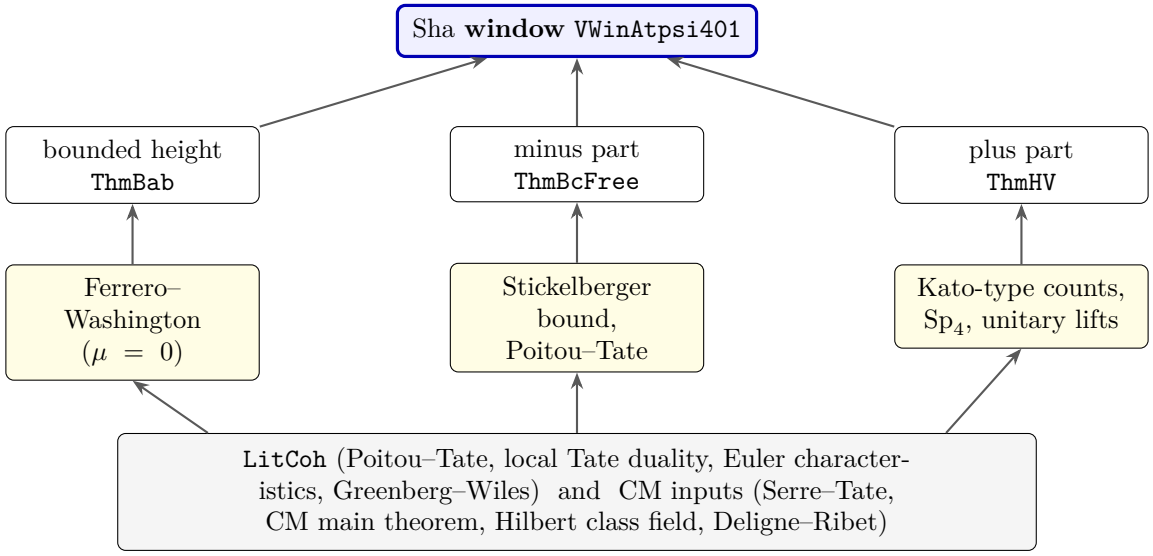


Figure 13: What the Sha window rests on. All yellow and grey inputs are proved in Lean from one `OpenImageCondition`; none is assumed.

Input	Rests on
bounded-height part (`ThmBab`)	`thmBab_of_lemmaK` with `ferreroWashington`
minus part (`ThmBcFree`)	a Stickelberger-type bound $\mathrm{len}\,\mathrm{Sha}^- \log \ell \le C\ell[E : \mathbb{Q}]^2(e+1)\log(e\ell N)^2$, via `ML.iv` (Poitou–Tate and unramified orthogonality)
plus part (`ThmHV`)	counts I, II, III, V, Z and `LemmaK`: Kato-type bounds, Sp_4 inputs, unitary lifts
duality (`LitCoh`)	Poitou–Tate, local Tate duality, local and global Euler characteristics, local and global finiteness, Greenberg–Wiles
CM inputs (`LiteratureCMHypothesis`)	Grothendieck deformation, de Rham comparison, Serre–Tate, CM main theorem, Hilbert class field via j, Siegel–Klingen and Deligne–Ribet, functional equation of partial zeta

Theorem 8.3 (Ferrero–Washington [4], `UpperTheory.V.FW.ferreroWashington`, [Exact]). *For a prime ℓ and an abelian number field K containing μ_ℓ (μ_4 if $\ell = 2$) there is r such that every layer K' of the cyclotomic $\mathbb{Z}_\ell$-tower over K satisfies $|\mathrm{Cl}(K')/\mathrm{Cl}(K')^\ell| \le \ell^r$.*

[Reading] Bounded ℓ-rank up the tower is the statement $\mu = 0$ [16]. The proof occupies a 62-file folder (`GaloisLibrary/FerreroWashingtonIndex`).

Provenance. The literature structures are filled by Lean proofs (`..._from_literature`, `..._of_open`) from one `OpenImageCondition`, which comes from `RowOb`. That in turn is proved by the theorem `galoisLocalGlobalStrong_recParameter`. The clean axiom check means none is assumed.

Budget. **[Reading]** The minus part is polynomial in $t = x^{401}$, $[E : \mathbb{Q}] \approx x^{2+8\varepsilon}$ and logarithms; the plus part is at most $\exp(x^{4\varepsilon/5})$ times a polynomial. Both are far below $\sqrt{b} = \exp(x^\varepsilon/2)$, so the window does not touch the square root.

9 Why the 2/3 barrier does not appear

The outline (entry 14) says a calculation of the cost of each lift shows one would "stop at 2/3 a priori". That calculation is not in the archive (Section 9.4). Here it is reconstructed from the code's own quantities. [Ours].

9.1 What the code charges

[Verified]

$$\mathrm{charge}(j) = \sum_{\ell \text{ small}} \log \ell \sum_{\text{pieces at } j} \text{colength} \approx f\, m\, b\, x^\varepsilon = B_0, \qquad E_b = \sum_{j=1}^{m} \mathrm{charge}(j) \approx m B_0.$$

Each piece belongs to one lift (`NData.locs_mem`), and each lift is charged on its own block: the family is `DefFam` with $L : J \to T \to$ BlockDat (`RowLat.lean:213`), and the exponents and lengths of lift j are read from $L(j,t)$ alone (`LiftLat.lean:147--151`).

9.2 Accounting 1: compounding gives exactly 2/3

If lift j were charged for the whole composite $\Theta_j \circ \cdots \circ \Theta_1$, then

$$\mathrm{charge}'(j) = \sum_{j' \le j} \mathrm{charge}(j') \approx j B_0, \qquad E'_b = \sum_{j=1}^{m} j B_0 \approx \frac{m^2}{2} B_0, \qquad \frac{E'_b}{n_T m b} \approx m^2 x^\varepsilon,$$

and the balance becomes $x/m + m^2$, minimised at $m = x^{1/3}$, giving $x^{2/3}$. This is the outline's 2/3 exactly. The two facts above remove the extra factor of m.

9.3 Accounting 2: no gcd does not give 2/3

Without the common lattice, the averaging over words in `Frame_at` would be lost. Its only visible cost is $\log \#W \le (k-3)\log 2 \approx x^{9/2+10\varepsilon}$, negligible against $n_T m b$. So in this bookkeeping the gcd does not save a factor of m; its role is to make one divisibility statement hold for all words.

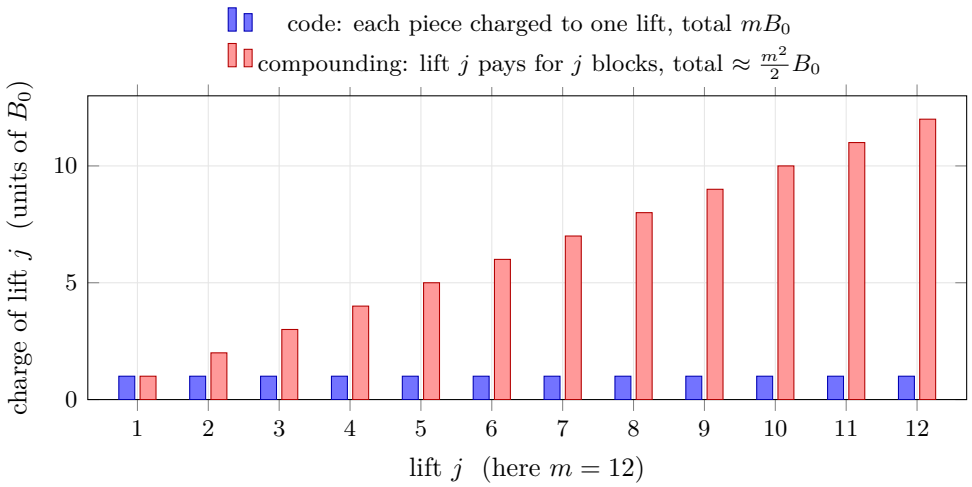


Figure 14: Two accountings of the same chain. The code assigns every local piece to exactly one lift (`NData.locs_mem`) and reads each lift's charge from its own block (`DefFam.L`); the compounding accounting adds a factor of m and gives the outline's $2/3$.

9.4 No $2/3$ route in the archive

[Exact] About twenty unused variants of the final growth bound exist, among them:

- `ChainAt`, `Chain6At`, `ChainAt7`;
- `FixedEpsilonGrowthBound`;
- `GrowthBracketBoundAtPoint`;
- `PowerBoundGivesRiemannHypothesis`.

Every one concludes $G(x) \le x^{1/2+O(\varepsilon)}$. No m^2B_0 budget, no per-word or per-character dual lattice, and no balance with exponent $2/3$ or $1/3$ appears anywhere. The only "$2/3$" constants are factors like $(2/3)^f$ in volume bounds.

9.5 A planned test

We designed a perturbation experiment (`mechanism/PERTURB.md`): loosen only the upper bound on m in `RankVal` by $x^{1/40}$, rebuild the roughly 13,000 dependent modules, and record which proofs fail. Our prediction, written before running it: the failures appear in the window lemmas (`szRow_rankSizes`, `winA`) and in the m^5 word-count lemma, not in the arithmetic layers, because m and k are locked with only about $x^{0.26\varepsilon}$ of headroom. The results will be posted as an update to this document.

10 The verification record

Check	Result (20–21 November 2026)
Source scan	0 `sorry`, `admit`, axiom declarations, `native_decide`, `reduceBool`; 5 build-time code hits (informational)
Full build	175,187 modules, 183,893 jobs, OK, 401 minutes on 94 cores
Statement	identical to Mathlib's `RiemannHypothesis` (`Mathlib.NumberTheory.LSeries.RiemannZeta`)
Axioms	`propext`, `Quot.sound`, `Classical.choice` only; `OVERALL: CLEAN`
Olean comparison	76,968 byte-identical to Sutherland's manifest, including `Main.olean`; 98,219 differ, every one embedding the absolute build path; 0 differ without a path. Recompiled at Sutherland's build path, a random sample of 300 of the 98,219 is byte-identical to the manifest, 300 of 300
Export hash	byte-identical match to the published export: 61,570,595,788 bytes, sha256 `5ea74b91...244e` (`lean4export` v4.33.0, commit `15f6055e`, 89 minutes)
Kernel replay	`Main` with imports trusted: done; full `-fresh` replay not run by us (others report full replays: Sutherland's `leanchecker`, 175,187 of 175,187 modules; Simon DeDeo, likewise)

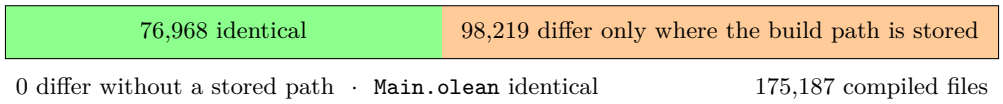


Figure 15: Our build against Sutherland's manifest of compiled files.

Remark 10.1. The differing oleans store the source path (for example `/home/.../rh/lean/src/.../OrbitMult.lean`, once per file), which differs between machines. The kernel export holds only declarations, and our re-export matches the published 61.6 GB export byte for byte, confirming that the differences cannot touch the proof.

[Verified] The path is the only difference. Sutherland's build passed sources to `lean` as `/home/ubuntu/grh/src/lean/src/<Module>.lean` (his comment in the community repository's issue #6, where patunz first reported the path dependence). We replayed, for a random sample of 300 of the 98,219 differing modules (seed 20261121), the exact `lean` command that Lake recorded for each module in its `.trace` file, changing only the output location and the source prefix to his path. All 300 outputs are byte-identical to his manifest. Before the sample, the same replay at our own path reproduced our own `.olean` files exactly, so the replay is faithful to the build (`scripts/11_path_sample.sh`).

11 Open questions, and what we could settle

Each question below is followed by what we could settle from the source and the outline. **Status:** three answered within the code (2, 3, 4), one narrowed (5), one still needs the author (1); the call-level gap in step D4, the role of the constants and a static reading of `SizesAt` are settled below.

Question 1. Is the compounding counterfactual of Section 9 the calculation behind the outline's 2/3?

Still open. The outline (entry 14, dated 15 November 2026) reports a "precise computation of the cost of each lift" and the figure 2/3, but does not show the calculation. Our reconstruction is the simplest accounting in the code's own quantities that gives exactly 2/3; only the author can say whether it is his.

Question 2. Could a family with a smaller indefinite part per lift lower the m^2 cost of E_b, or does the Bloch–Kato count force it?

Answered within this chain. [Verified] `ShapeJumpskmt` is required with the same m at every local piece, whatever lift the piece belongs to (`admG2`, `EmbeddingTypes.lean:122`), so every lift is charged rank m. [Reading] checked by arithmetic: even a family whose rank tracked the shrinking indefinite part, $m-j$ at lift j, would pay $\sum_{j=1}^{m}(m-j) = m(m-1)/2$, still of order m^2. So no refinement of the per-lift rank along this chain changes the exponent; it saves at most a factor of 2. Lowering it would need a rank of $o(m)$ per lift, that is a chain with a smaller indefinite part, hence fewer lifts, which costs on the x/m side.

Question 3. Why must the lift exponent e' be as large as $b = \exp(x^{\varepsilon})$? A polynomial e' would shrink B_0, though not its dependence on m.

Answered within this design. [Verified] e' is linear in a Hodge–Tate height $h_B \in [b, 6b]$, and the heights are of size b because the weights on thick places lie in $[b, b+\eta+1]$. [Exact] The weight scale cannot be polynomial: the Sha-window losses are themselves exponential. The bounded-height condition is stated with $\exp(x^{4\varepsilon/5})$ (`BaAt`) and the plus part with $(\mathrm{cond}\cdot t)^C + C\exp(x^{4\varepsilon/5})$ (`HVAt`), both to be absorbed by $\sqrt{b} = \exp(x^{\varepsilon}/2)$. [Reading] So b is exponential to swallow those losses, $e' \approx b$ follows, and since b cancels in the balance this sets the absolute size of E_b but not the exponent $1/2$.

Question 4. Do the formal versions of Ferrero–Washington, Poitou–Tate, the Kato-type bounds and the CM inputs match the published theorems at full strength?

Answered: statements read for all four.

- *Ferrero–Washington.* [Exact] The statement bounds $|\mathrm{Cl}(K')/\mathrm{Cl}(K')^{\ell}|$ by ℓ^r for every layer $K' = K(\mu_{\ell^n})$ over an abelian $K \supseteq \mu_{\ell}$ (μ_4 if $\ell = 2$), with Mathlib's class group. [Reading] $\mathrm{Cl}\,/\,\mathrm{Cl}^{\ell}$ is elementary abelian of order $\ell^{\mathrm{rk}_{\ell}}$, so this is bounded ℓ-rank up the cyclotomic $\mathbb{Z}_{\ell}$-tower, which is equivalent to $\mu = 0$: the published theorem, in the form for K containing μ_{ℓ}.
- *Poitou–Tate.* [Exact] `PoitouTateNormalised`: for a number field F, a prime ℓ with $\ell \neq 2$ or F totally complex, a finite set S containing the places above ℓ, and every finite $\mathbb{Z}/\ell^r$-module with a continuous G_S-action: normalised local invariant maps, the nine-term exact sequence, Sha-duality, and $H^n = 0$ for $n \geq 3$. [Reading] This is the standard theorem; the

condition on ℓ is exactly the case where the real places contribute nothing in degree ≥ 3. It is proved unconditionally by `poitouTateNormalised_holds`; the literature wrapper ignores its argument.

- *Kato-type input.* [Exact] `KatoTwoW`: for an odd prime ℓ, an imaginary quadratic K with $\ell \nmid \mathrm{disc}(K)\, h_K$, and a two-dimensional ℓ-adic representation W with determinant the quadratic character θ of K, ordinary at ℓ, residually Eisenstein (trace $\equiv \theta \cdot \chi^e\psi + (\chi^e\psi)^{-1}$), irreducible over the cyclotomic extension of K and attached to a newform of weight $2e+1$: the Bloch–Kato Selmer groups of W and of $W \otimes \theta$ vanish. [Reading] As David Loeffler pointed out on Zulip (21 November 2026), this matches Kato's theorem [8] directly without needing an explicit L-value hypothesis: in this setting the relevant critical L-value sits at $s = 1$ in the analytic normalisation (to the right of the central point $s = 1/2$), where non-vanishing is automatic. It is proved in Lean (`katoTwoW_of_printsH_of_rT_of_quadBC2`).
- *CM inputs.* [Exact] Read statement by statement: Shimura's main theorem of complex multiplication for imaginary quadratic K (`EllipticMainTheoremCM`); $K(j(\mathcal{O}_K))$ is the Hilbert class field, of degree h_K, with $j(\mathfrak{a}) = j(\mathfrak{b})$ iff $\mathfrak{a}, \mathfrak{b}$ are in the same class (`HilbertClassFieldJ`); Siegel–Klingen and Deligne–Ribet; the functional equation of partial zeta functions of totally real fields and of Hecke characters of CM fields; Grothendieck–Messing deformation theory over nilpotent divided powers; de Rham cohomology of abelian schemes (free, Hodge filtration, base change, rigidity, duality); the period relation and Gauss–Manin periods; and Serre–Tate (Néron–Ogg–Shafarevich, potential good reduction, good reduction of CM abelian varieties). [Reading] Each matches the standard theorem of that name in shape and hypotheses. We did not check the internal definitions they rely on (for example the global Artin map) line by line.

Since every one of these is proved in Lean and the axiom check is clean, Question 4 is about exposition: whether the names are honest descriptions. For the ones we read they are, with the caveat on the Kato-type statement above.

Comparison with Jaswal's guide. Manvir Jaswal's reader's guide [6] compares every classical input on the chain with its archive form (his Table 2) and reaches the same verdicts on Ferrero–Washington and Poitou–Tate. It adds two inputs we did not read, which it classes as weaker than the published theorems their names recall, both proved in the archive.

- *Galois representations* (`Row0b`, proved by `galoisLocalGlobalStrong_recParameter`): local–global compatibility is asserted only when the local parameter is abelian or tame with zero monodromy, and at $w \mid \ell$ only under extra data.
- *Local parameters* (`RecTExists`): the map exhibited by `recTExists` satisfies the required laws but sends every supercuspidal outside a class `GoodCusp` to an unramified twist of one fixed parameter $J_{w,n}$ (`cuspEx`, `JunkEx`), so it is not the local Langlands correspondence. The full correspondence is stated as `LocalLanglandsGL` and proved only conditionally.

[Verified] by name against the export of `Main`: `Row0b`, `galoisLocalGlobalStrong_recParameter`, `RecTExists`, `recTExists`, `GoodCusp`, `cuspEx` and `JunkEx` are used; `LocalLanglandsGL` is not. So the proof does not depend on the local Langlands correspondence. Jaswal also classes the Kato-type input as "Different" in shape, since it is stated with its own hypotheses (`Used`, a newform attachment, `WitAQ`), which is compatible with the reading above. His Table 1 lists the statements on the chain for which he knows no reference; several of them (`ch_balance`, `TheoremSSelected`, `LinkN`, `lemmaK`) are the steps we trace in Sections 4–8.

Question 5. Is the windowed Landau theorem (Theorem 3.4) new as a stated result?

Narrowed. Deducing zero-free regions from one-sided or localised information on prime sums is a classical line going back to Landau [9] and developed by Pintz (see the recent survey [7]), and effective versions of Landau's theorem already give quantitative lower bounds for averages over ranges $[1, Y]$. The windowed statement is therefore close to known methods in substance; whether this exact form appears in the literature is for specialists to say.

The call-level gap in step D4 is closed. [Exact] `extendedClosedForm` is the field `p15` of `local_flatness_results`, which feeds `loc8All_of`, `first_theorem_threshold_functions` and `first_theorem_closed_form`; the family bound from it is `Adm8F`, a hypothesis of `admissible_family_no_kernel`, whose output is what the gain statement (`AuxBlock8.Conclusion1`, via `starAll_of`) is proved from. So the $\sqrt{p}$ lemma reaches the gain inequality through the first closed-form theorem. Not checked: which lines of the gain proof consume `p15` specifically.

The constants do not matter for the theorem. [Exact] `ChebyshevInequalityAtWindowScale` asks for *some* constants e_D, e_L and $\varepsilon_1 > 0$ such that for every $\varepsilon < \varepsilon_1$ there is a family with $G_m(x) \le x^{1/2+(e_D+e_L+7)\varepsilon}$ on its windows. Since e_D, e_L are fixed while $\varepsilon \to 0$, the Landau step gives zero-freeness for every $\operatorname{Re} s > 1/2$ whatever their values. Explicit constants would only matter for an effective (numerical) version, which the development does not attempt.

`SizesAt` under other exponents, read statically. [Exact] Write $m = x^\mu$, $n = x^\nu$, $k = x^\kappa$. In the abstract size layer the exponents enter only through

$$\tfrac{1}{2} - \varepsilon \le \mu \le \tfrac{1}{2} + \tfrac{\varepsilon}{2}, \qquad \kappa \ge 1 + \mu, \qquad \kappa \le 2\nu + \tfrac{1}{2} - \varepsilon, \qquad \kappa \le e_K$$

(the bounds on μ are in `RestSize` and `ErrWindows0`), together with the family's word count $\kappa \ge 5\mu + \nu$. Every other window (`CostWindows`, `GrowthInequalitySizes`, the bounds on L, L_a, b') compares polynomials in x with $b = \exp(x^\varepsilon)$ and holds for any fixed exponents. So statically, μ is pinned only by the two windows that encode the balance, exactly as predicted. What reading cannot rule out is a hidden dependence elsewhere in the 13,000 modules downstream of the size choices; the perturbation run (Section 9) will check that mechanically.

Still not traced. The Lean proofs of the classical inputs of Section 8, as opposed to their statements; numerical values of the constants; the perturbation run.

Acknowledgements

The formal proof is Levent Alpöge's. This account builds on public work by others on 20 and 21 November 2026, and overlaps with some of it:

- Jonas Whidden's Zulip readings (the balance, the frame lemma, the common-lattice question);
- Bartosz Naskręcki's proof guide and book [12], which also state the $x/m + m$ balance, the localized Landau step, the choice of auxiliary primes and the Ferrero–Washington and Poitou–Tate premises;
- JoseClaudioSJr's reader's map (community issue #5), which independently traces the windowed Landau route and the gain chain through `AuxBlock8.Conclusion1` and `starAll_of`;
- Andrew Sutherland's verification record and build manifest [15]; the full kernel replays and independent rebuilds reported by Simon DeDeo and Elliot Arledge;

- the dependency and dead-code analyses of James Shore (`james-shore/rh-proof-dependencies`, 256,620 of 1,061,531 declarations unused), Sonoda, Drew Carlson and Keith Adler (`rh-trim`); patunz's report on path-dependent olean hashes;
- the expositions of Yao Xu (`cadamcat/riemann-hypothesis-exposition`) and jae-63 (`rh-alpoge-blueprint`), and the declaration index at `riemann.tamnd.com`.

Errors of interpretation are ours.

Use of AI tools. Claude (Anthropic) was used to help read the Lean source, run searches over the declaration list, and draft parts of this text and its figures. Statements marked [Exact] or [Verified] were checked against the Lean source or the build outputs in the accompanying repository, and the references against their published sources. The author has read the whole text and takes full responsibility for its content.

References

[1] L. Alpöge, note accompanying the Lean development, 20 November 2026, `https://alpo.ge/remark.pdf`.

[2] S. Bloch and K. Kato, *L-functions and Tamagawa numbers of motives*, in: The Grothendieck Festschrift, Vol. I, Progr. Math. 86, Birkhäuser, 1990, 333–400.

[3] S. Chowla and A. Selberg, *On Epstein's zeta-function*, J. reine angew. Math. 227 (1967), 86–110.

[4] B. Ferrero and L. C. Washington, *The Iwasawa invariant μ_p vanishes for abelian number fields*, Ann. of Math. (2) 109 (1979), 377–395.

[5] H. Iwaniec and E. Kowalski, *Analytic Number Theory*, AMS Colloquium Publications 53, American Mathematical Society, 2004.

[6] M. Jaswal, *The Riemann hypothesis for Dirichlet L-functions [after Levent Alpöge]*, preprint, 21 November 2026, `https://doi.org/10.5281/zenodo.24410057`.

[7] D. R. Johnston and T. Trudgian, *A round of Pintz to celebrate oscillations in sums*, arXiv:2511.09978.

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

[9] E. Landau, *Handbuch der Lehre von der Verteilung der Primzahlen*, Teubner, Leipzig, 1909.

[10] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction (CADE-28), LNCS 12699, Springer, 2021, 625–635.

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

[12] B. Naskręcki, `riemann-rev` (status dashboard, proof guide and book), `https://github.com/nasqret/riemann-rev`, November 2026.

[13] J. Neukirch, A. Schmidt and K. Wingberg, *Cohomology of Number Fields*, 2nd ed., Grundlehren der mathematischen Wissenschaften 323, Springer, 2008.

[14] G. Shimura, *Introduction to the Arithmetic Theory of Automorphic Functions*, Princeton University Press, 1971.

[15] A. V. Sutherland, build and verification page for the Lean development, `https://math.mit.edu/~drew/grh/`, November 2026.

[16] L. C. Washington, *Introduction to Cyclotomic Fields*, 2nd ed., Graduate Texts in Mathematics 83, Springer, 1997.

A Map from steps to Lean declarations

Step	TRACE.md	Key declarations
A. Windowed Landau	1, 25.1	`zeroFree_of_rpow_window`, `lemmaL`, `landau_window`, `poleData`, `core_gm`, `rh_of_every`
B. Isolation	23	`ch_balance`, `ThickPartTranslateMean`, `GrowthInequalityExplicit`, `signal_term`, `QH.pairEx`, `zero_density_box`, `hostThickExists`
C. Size bound	2, 5, 12	`proposition_size_bound`, `SZ.T1–T4`, `RankVal`, `exists_ranks_rankVal_frame`, `m_pow_five_mul_f_mul_L_le`
D. Pointwise bound	6, 7, 14	`each_unit_bound_statement`, `Ga_bound5`, `ChainTied`, `linkCore`, `linkGLOBAL7_items`
D. Energy	8, 9, 15	`SizeData_2.EDef`, `NormDecompositionCond`, `factEtaMean`, `nuAbs_place_fine`, `ClaimHigh`
D. DIV	10, 11, 16	`Data.DIV`, `height_chain_atR`, `errB_of_units`, `divisor_inequality_assembly`, `error_absorption`
D. Rank E_b	13, 18	`rank_le_row`, `ShapeJumps`, `NData.locs_mem`, `linkN_of_gd`
D. Common lattice	17	`prescribedRowLat`, `exists_divStage_of_pk`, `FamP.rrDef`
D. Degree identity	20	`NprimeMuR`, `LinkG_at`, `LinkDIV_at`, `Frame_at`, `loss_bound_at`
D. Local divisibility	21, 25.2	`key_minor_dvd`, `key3Ord`, `link_of_local_closed`, `starAll_of`, `gain_bound_all_of_star`
E. Orthogonality	22	`LinkS_at`, `theoremS_link_concrete`, `matKF_of`, `chosenProviso_min`
F. Sha window	24, 25.3	`VWinAt`, `vwinAt_of_joint`, `of_four_alt`, `ferreroWashington`, `ML.iv`
2/3 barrier	15, 18, 19	`NData.locs_mem`, `DefFam`, `Frame_at`