DIRICHLET L-FUNCTIONS AND THE CRITICAL LINE

Abstract. We give a mathematical companion to Alpöge's Lean development of the Riemann hypothesis and its analogue for Dirichlet L-functions. The analytic step turns upper bounds on selected logarithmic windows for a positive-weight family of compensated prime sums into nonvanishing to the right of the critical line. Finite Fourier averaging identifies the relevant prime sums in an arithmetic evaluation inequality, while an approximate frame transfers the mean slopes of many small lattices to a large target. We prove these analytic, Fourier, and slope-transfer arguments in detail and explain the exact cancellation of the finite and infinite local terms. The arithmetic construction and cohomological window are stated as explicit inputs from the formal development, with their quantifiers and principal source declarations recorded. This separates the self-contained arguments given here from the deeper construction results used by reference and from the independent verification of the formal proof.

1. The result and the organization of the argument

The Riemann hypothesis concerns the zeros of the meromorphic continuation of $\zeta(s) = \sum_{n\ge1} n^{-s}$, initially defined for $\operatorname{Re} s > 1$. For a Dirichlet character χ modulo a positive integer N, its analogue concerns $L(s,\chi) = \sum_{n\ge1} \chi(n)n^{-s}$. Alpöge's development [2] establishes nonvanishing for all these functions in $\operatorname{Re} s > 1/2$, with the principal pole at $s = 1$ excluded. The functional equation then places every zero in the open critical strip on its central line. The following is the principal mathematical conclusion of that development.

Theorem 1.1 (`riemannHypothesis`)**.** *Every nontrivial zero of ζ has real part $1/2$. More generally, for every Dirichlet character χ, the meromorphic function $L(s,\chi)$ is nonzero when $\operatorname{Re} s > 1/2$, apart from its possible pole at 1, and every zero with $0 < \operatorname{Re} s < 1$ has real part $1/2$.*

Here and throughout, a pole is not a zero. The precise translations of the five declarations in `Main.lean` appear in Appendix A. In particular, the declaration for primitive characters in `Statements.lean` retains both inequalities $0 < \operatorname{Re} s < 1$. The concordance records the exact hypotheses of each public statement.

This article provides a mathematical guide to Theorem 1.1 and its formal source. We prove the arguments concerning Mellin transforms, localization to windows, finite character sums, frames, slope distributions, and the final numerical cancellation. The integral theta construction and the cohomological window theorem in Sections 6 and 7 are used as cited inputs, with their required conclusions and quantifiers stated explicitly. This is therefore an expository proof using those construction results; it does not independently reconstruct their arithmetic proofs.

The organizing estimate concerns a smoothed remainder, not the unsmoothed prime-counting function. For a primitive character χ write $S_\chi(y) = \sum_{n\le y} \Lambda(n)\chi(n)(y/n - 1)$, where Λ is the von Mangoldt function. The principal sum is compensated by subtracting $y\log y - (1+\gamma_E)y$; a nonprincipal sum is compensated by adding $yL'(1,\chi)/L(1,\chi)$. Denote the resulting functions by G_χ. If ψ has order o, let $\mathcal{C}_\psi$ consist of the primitive characters inducing its powers. The real function

$$H_\psi(y) = G_1(y) + \frac{1}{3o+1} \sum_{\chi\in\mathcal{C}_\psi\setminus\{1\}} \operatorname{Re} G_\chi(y)$$

is the quantity the arithmetic construction controls. All coefficients are positive, so the logarithmic derivatives of different L-functions cannot cancel the singularity associated with a zero. The task is to prove, for each sufficiently small $\varepsilon > 0$, an upper bound $H_\psi(y) \le y^{1/2+9\varepsilon}$ throughout $[Y, Y^{3/2}]$, for arbitrarily large Y. The endpoint Y is allowed to depend on the character and on ε.

There are three reasons why this particular form of the estimate is useful. First, it is one-sided. The arithmetic evaluation argument produces an inequality of one sign, so imposing an absolute-value estimate would discard useful information. Second, it concerns a finite family of characters with prescribed positive coefficients. The positivity is used at the level of zero multiplicities, not as a claim that H_ψ is everywhere positive. Third, the estimate is needed only on a sequence of intervals whose lengths are proportional to their left endpoints after taking logarithms. This allows the arithmetic construction to select intervals where the necessary cohomological losses are controlled.

The analytic argument makes this flexibility precise. The k-th derivative of a Laplace transform integrates its original function against $u^k e^{-\sigma u}/k!$. Choosing k proportional to $\log Y$ concentrates this positive kernel inside the available window. A triangle inequality bounds the sum of the modulus of a derivative at a complex point and the real part of the corresponding derivative on the real axis by the positive part of the original function. A zero to the right of the desired line would create a nearer pole and force faster derivative growth. A small vertical perturbation isolates a complex pole; a real pole is handled by its positive residue. Neither a globally nonnegative remainder nor a window-independent starting point is needed.

The arithmetic side supplies the estimate through evaluation maps of automorphic forms over an auxiliary abelian CM field. The difference between the finite-place divisibility contribution and the archimedean period contribution is an exact finite Fourier transform of the compensated prime sums. A constant metric would retain only two characters. The construction instead assigns different weights to top, thick, and thin sets of embeddings. Thick places sample the prime sums at the desired scale; the other places sample much larger arguments. Averaging translates removes unwanted quotient characters, and a deliberately biased distribution of top places gives the coefficients 1 and $1/(3o+1)$ above.

The geometric difficulty is a rank mismatch. Each individual source lattice is too small to span the target, but the family of evaluation maps is approximately orthogonal and approximately fills the target in the sense of a trace inequality. The frame estimate converts a bound for the total corank of selected source subspaces into a bound for the codimension of their image span. Applied to the Harder–Narasimhan filtrations and integrated over the slope threshold, this gives a mean-slope inequality with a logarithmic loss $u\log(1+H/u)$. The local and global geometric estimates can then be combined without destroying the cancellation of the main terms.

It is useful to distinguish the classical ingredients from the features particular to this architecture. Euler products, functional equations, logarithmic-derivative identities, and Landau's sign–singularity method belong to classical analytic number theory [15, 13]. Finite character orthogonality and the trace Cauchy–Schwarz inequality are elementary. Harder–Narasimhan slopes and arithmetic evaluation inequalities belong to the established Arakelov method [6]. The new work to be understood here is their quantitative coordination: localization to selected windows, the retained positive-weight character family, the frame-based transfer across a rank gap, and the construction that makes the finite-place gains and infinite-place costs match. The distinctions below concern the role of each ingredient in this proof; no comprehensive claim of historical priority is intended.

The deeper background is also established mathematics: theta correspondence and Siegel–Weil identities, CM periods, arithmetic Hilbert–Samuel theory, integral p-adic comparison, automorphic Galois representations, and Iwasawa theory [12, 10, 8, 14, 5, 7, 9]. The argument requires versions with compatible normalizations, integral structures, and estimates uniform in varying fields and weights. The construction-specific work lies in the compatibility and uniformity statements that permit these theories to be used simultaneously. Their conclusions and points of application are stated explicitly below.

Only principal propositions carry Lean names in their headings. Ordinary notation is used for auxiliary objects, and Appendix A gives the corresponding files and declarations. The source version is fixed there, together with a short account of the independently reported verification.

2. A windowed zero-free criterion

Write $\Lambda(n)$ for the von Mangoldt function and γ_E for Euler's constant. Throughout this section, an L-function is understood as a meromorphic function, so that its pole at 1 is never counted as a zero. The primitive principal character has conductor 1, and its L-function is ζ. The basic analytic facts used below are analytic continuation, the functional equation, nonvanishing on $\operatorname{Re} s = 1$, and the logarithmic-derivative Dirichlet series on $\operatorname{Re} s > 1$. These are classical facts about Dirichlet L-functions; they are also the analytic inputs imported by the corresponding Lean modules. In particular, this section does not use the Riemann Hypothesis or any zero-free region inside the critical strip.

For a primitive character χ, put $S_\chi(y) = \sum_{n\le y} \Lambda(n)\chi(n)(y/n - 1)$ for $y \ge 1$. Define the compensated remainders by

$$G_1(y) = S_1(y) - y\log y + (1+\gamma_E)y, \qquad G_\chi(y) = S_\chi(y) + y\frac{L'(1,\chi)}{L(1,\chi)} \quad (\chi \ne 1). \tag{2.1}$$

The compensating terms remove the contribution of $s = 1$ from the Mellin transform. They are part of the normalization and cannot be dropped when deriving identities. The ramp $y/n - 1$ vanishes at $y = n$, so these functions are continuous, even though their derivatives have jumps at prime powers. The elementary estimate $\Lambda(n) \le \log n$ gives $S_\chi(y) = O(y(1+\log y)^2)$ and hence $G_\chi(y) = O(y^2)$; all constants in this section may depend on the fixed finite family of characters.

Let $\mathcal{F} = (\chi_j, a_j)_{j\in J}$ be a finite family of primitive characters with real weights $a_j > 0$. Set $H(y) = \sum_j a_j \operatorname{Re} G_{\chi_j}(y)$ and $g(u) = H(e^u)$ for $u \ge 0$. Thus g is real and measurable, and $|g(u)| \le C_0 e^{2u}$ for a suitable C_0. We first record exactly what information its Laplace transform retains about the zeros.

Proposition 2.1 (`poleData`)**.** *The function $\Phi(s) = \int_0^\infty g(u)e^{-su}\,du$, initially defined for $\operatorname{Re} s > 2$, extends meromorphically to $\operatorname{Re} s > 0$. Its poles are precisely the zeros in this half-plane of the functions $L(s,\chi_j)$ and $L(s,\overline{\chi}_j)$. Each pole ρ is simple, with residue*

$$c_\rho = -\frac{b_\rho}{\rho(\rho-1)}, \qquad b_\rho = \frac{1}{2}\sum_j a_j(\operatorname{ord}_\rho L(s,\chi_j) + \operatorname{ord}_\rho L(s,\overline{\chi}_j)) > 0. \tag{2.2}$$

There is no pole at 1. Every real pole β lies in $(0,1)$ and has strictly positive real residue $b_\beta/(\beta(1-\beta))$.

Proof. For $n \ge 1$ and $\operatorname{Re} s > 1$, direct integration gives $\int_n^\infty (y/n - 1)y^{-s-1}\,dy = n^{-s}/(s(s-1))$. Absolute convergence justifies summing this identity against $\Lambda(n)\chi(n)$: the sum of the integrals of the absolute values is bounded by a constant times $\sum_n \log n\, n^{-\operatorname{Re} s}$. Therefore, if $F_\chi(s) = \int_1^\infty G_\chi(y)y^{-s-1}\,dy$, then

$$\begin{aligned} F_\chi(s) &= -\frac{L'(s,\chi)/L(s,\chi) - sL'(1,\chi)/L(1,\chi)}{s(s-1)} \qquad (\chi \ne 1), \\ F_1(s) &= -\frac{\zeta'(s)/\zeta(s) + 1/(s-1) + 1 - (1+\gamma_E)s}{s(s-1)}. \end{aligned} \tag{2.3}$$

For the second formula, the two elementary integrals contributed by the compensation are $-1/(s-1)^2 + (1+\gamma_E)/(s-1)$. Equivalently, the second expression in (2.3) is $-[\zeta'/\zeta + 1/(s-1) - \gamma_E s]/(s(s-1)) + 1/s$. This final $1/s$ is required by the normalization (2.1).

The change of variables $y = e^u$, together with $G_{\overline{\chi}}(y) = \overline{G_\chi(y)}$, now gives $\Phi(s) = \frac{1}{2}\sum_j a_j(F_{\chi_j}(s) + F_{\overline{\chi}_j}(s))$. The formulas on the right supply the continuation. For $\chi \ne 1$, their numerator vanishes at $s = 1$ because $L(1,\chi) \ne 0$. For $\chi = 1$, use $\zeta'/\zeta = -1/(s-1) + \gamma_E + O(s-1)$: the numerator is again $O(s-1)$. Both apparent singularities at 1 are consequently removable.

If a holomorphic function has a zero of order m at ρ, it can be written $(s-\rho)^m h(s)$ with $h(\rho) \ne 0$, so its logarithmic derivative is $m/(s-\rho) + h'(s)/h(s)$. Applying this factorization to each L-function proves (2.2). Every multiplicity is nonnegative and every weight is positive. At a zero of at least one member of the family, their weighted sum is strictly positive, and no pole can disappear by cancellation. The zero sets are locally finite because none of the finitely many

L-functions is identically zero. Finally, classical nonvanishing on and to the right of $\operatorname{Re} s = 1$ places every real pole in $(0,1)$, where the sign in (2.2) is positive. □

The positivity in this proposition is a statement about residues, not a claim that g itself is nonnegative. This distinction is essential. The available arithmetic estimates control an upper bound for g only on selected intervals, and g is free to change sign both inside and outside them.

Fix $0 < \theta < 1$ and $\delta > 0$. The precise window hypothesis is that, for every $\eta > 0$ and every U_0, there exist $U \ge U_0$ and a nonnegative measurable majorant R on $[U, (1+\delta)U]$ such that $g \le R$ there and

$$\int_U^{(1+\delta)U} R(u) e^{-\theta u}\, du \le e^{\eta U}. \tag{2.4}$$

The integral is assumed finite. A pointwise estimate $g(u) \le e^{\theta u}$ on arbitrarily distant such windows implies (2.4), since one may take $R(u) = e^{\theta u}$ and then use $\delta U \le e^{\eta U}$ for large U. On the original scale, the window is $[Y, Y^{1+\delta}]$ with $Y = e^U$, and the integral is $\int_Y^{Y^{1+\delta}} R(\log y) y^{-\theta-1}\, dy$. The integral formulation is useful because it permits local averaging errors and need not arise from a pointwise power bound.

Lemma 2.2. *For every $\delta > 0$ there are $\sigma > 2$ and $C_\delta > 0$ such that, for every $U \ge 1$, some integer $k \ge U$ satisfies*

$$\int_{[0,\infty)\setminus[U,(1+\delta)U]} \frac{u^k e^{-(\sigma-2)u}}{k!}\, du \le \frac{C_\delta}{\sigma^k}. \tag{2.5}$$

The same σ and C_δ work for all U.

Proof. We give explicit parameters, so the dependence on the window length is transparent. Put $\ell = 2 + 6/\delta$, $M = \ell(1+\delta)$, $\sigma = M(\ell+2)$, and $b = 1 + \delta/2$, and take $k = \lceil b\sigma U \rceil$. These choices satisfy

$$b(\ell+2) = M - 1, \qquad b(\ell-2) = \ell + 1, \qquad b\sigma U \le k \le b\sigma U + 1. \tag{2.6}$$

In particular $\sigma > 2 + \ell$, $M - 1 > 1$, and $k \ge U$.

For $s > \ell \ge 0$ and $V \ge U$, the elementary bounds $1_{u>V} \le e^{\ell(u-V)}$ and $1_{u<U} \le e^{\ell(U-u)}$, together with $\int_0^\infty u^k e^{-au}/k!\, du = a^{-k-1}$ for $a > 0$, give

$$\int_{[0,\infty)\setminus[U,V]} \frac{u^k e^{-su}}{k!}\, du \le \frac{e^{-\ell V}}{(s-\ell)^{k+1}} + \frac{e^{\ell U}}{(s+\ell)^{k+1}}. \tag{2.7}$$

The gamma integral here follows by k integrations by parts from the case $k = 0$. Apply (2.7) with $s = \sigma - 2$ and $V = (1+\delta)U$; write $a_- = \sigma - 2 - \ell$ and $a_+ = \sigma - 2 + \ell$. For any $a > 0$, the inequality $\log(\sigma/a) \le (\sigma - a)/a$ yields

$$e^x a^{-k-1} \le a^{-1} \sigma^{-k} \exp\left(x + k \frac{\sigma - a}{a} \right).$$

For $a = a_-$, the exponent on the right is $-MU + k/(M-1) \le 1$, by (2.6). For $a = a_+$, it is $\ell U - k(\ell-2)/a_+ \le 0$: indeed $k(\ell-2) \ge (\ell+1)\sigma U \ge \ell a_+ U$. The last inequality follows from $\sigma \ge \ell(\ell-2)$. Thus (2.5) holds with $C_\delta = e/a_- + 1/a_+$. □

Lemma 2.3. *Suppose g is real and measurable, $|g(u)| \le C_0 e^{2u}$ for $u \ge 0$, and (2.4) holds. Let Φ be its Laplace transform on $\operatorname{Re} s > 2$. There exist $\sigma > 2$ and $C > 0$ such that, for every $\eta > 0$ and every U_0, there are $U \ge U_0$ and an integer $k \ge U$ for which, simultaneously for all $t \in \mathbb{R}$,*

$$\left| \frac{(-1)^k \Phi^{(k)}(\sigma + it)}{k!} \right| + \operatorname{Re} \frac{(-1)^k \Phi^{(k)}(\sigma)}{k!} \le \frac{C e^{\eta U}}{(\sigma - \theta)^k}. \tag{2.8}$$

Proof. Choose σ and C_δ from Lemma 2.2, and choose a window with $U \ge \max(U_0, 1)$. The growth bound justifies differentiation under the integral: every differentiated integrand is

dominated locally in $\operatorname{Re} s > 2$ by an integrable polynomial times a decreasing exponential. With $w_k(u) = u^k e^{-\sigma u}/k! \geq 0$ and $A_k(z) = (-1)^k \Phi^{(k)}(z)/k!$, we obtain

$$|A_k(\sigma + it)| + \operatorname{Re} A_k(\sigma) \leq \int_0^\infty (|g(u)| + g(u)) w_k(u)\, du = 2\int_0^\infty \max(g(u), 0) w_k(u)\, du. \tag{2.9}$$

This is the point at which a one-sided estimate becomes sufficient. On the chosen window, use $\max(g, 0) \leq R$ and $u^k e^{-(\sigma-\theta)u}/k! \leq (\sigma - \theta)^{-k}$, which follows from $e^v \geq v^k/k!$ for $v \geq 0$. The contribution of the window is at most $2e^{\eta U}(\sigma - \theta)^{-k}$. Off the window, use the global growth bound and Lemma 2.2; the contribution is at most $2C_0 C_\delta \sigma^{-k}$. Since $0 < \sigma - \theta < \sigma$ and $e^{\eta U} \geq 1$, their sum proves (2.8) with $C = 2 + 2C_0 C_\delta$. □

We next isolate the geometric fact needed to turn (2.8) into a contradiction. Let Z be locally finite in $\Omega = \{\operatorname{Re} s > 0\}$, with all its points strictly to the left of the vertical line $\operatorname{Re} s = \sigma$, and suppose some $\rho_0 \in Z$ has $\operatorname{Re} \rho_0 > \theta > 0$. Put $E = \sigma - \theta$. At least one of the following alternatives holds. Either there is a radius $D \leq E$ such that $B(\sigma, D)$ contains no point of Z and some pole ρ_1 has $\sigma - \operatorname{Re} \rho_1 < D$, or there is a real point $\beta \in Z$, with $\beta > \theta$, which is the only point of Z in some disk $B(\sigma, r)$ with $\sigma - \beta < r \leq E$.

Here is a proof that also explains the role of real poles. The set $F = Z \cap \overline{B}(\sigma, E)$ is finite because this closed disk is a compact subset of Ω. If it is empty, take $D = E$ and $\rho_1 = \rho_0$. Otherwise let D be the smallest distance from σ to F. The open disk of radius D is pole-free. If the first alternative fails, every $\rho \in Z$ satisfies $D \leq \sigma - \operatorname{Re} \rho$. A point attaining distance D must then have imaginary part zero, since $D \leq \sigma - \operatorname{Re} \rho \leq |\sigma - \rho| = D$. It is therefore the single point $\beta = \sigma - D$. Applying the same inequality to ρ_0 gives $\beta > \theta$, so $D < E$. Finiteness of F supplies a radius r strictly between D and the distance to every other point of F, and no greater than E. This proves the second alternative.

In the first alternative one can move the center a little vertically so as to obtain a unique nearest pole closer than D. To justify this without losing control of poles entering the disk, start with $z_0 = \sigma + i \operatorname{Im} \rho_1$ and choose $\varepsilon > 0$ so small that $D + \varepsilon < \sigma$ and $\sigma - \operatorname{Re} \rho_1 + \varepsilon < D$. All points that can enter $B(\sigma + it, D)$ while $|t - \operatorname{Im} \rho_1| < \varepsilon$ belong to the fixed finite set $Z \cap \overline{B}(z_0, D + \varepsilon)$. For two distinct points a, b in this set, equality of their distances to $\sigma + it$ is a linear equation in t after squaring. It has at most one solution if their imaginary parts differ, and no solution if their imaginary parts agree, because both real parts are smaller than σ. Avoiding these finitely many values of t gives a unique nearest point ρ at distance $d < D$. One may then choose $d < r \leq D$ so that $B(\sigma + it, r)$ contains precisely ρ.

Theorem 2.4 (`landau_window`). *Let g be real and measurable with $|g(u)| \leq C_0 e^{2u}$ for $u \geq 0$. Suppose its Laplace transform extends meromorphically to $\Omega = \{\operatorname{Re} s > 0\}$, with a locally finite set of simple poles, all having nonzero residue, and with positive real residue at every real pole. If (2.4) holds for $0 < \theta < 1$ and $\delta > 0$, there is no pole with real part greater than θ.*

Proof. The integral defines a holomorphic function on $\operatorname{Re} s > 2$, so no pole lies in that half-plane. Choose $\sigma > 2$ and C as in Lemma 2.3. Assume there is a pole to the right of θ, and apply the preceding geometric argument with $E = \sigma - \theta$.

In the first case there are a pole-free disk $B(\sigma, D)$, a point $z = \sigma + it$, and a unique pole ρ in $B(z, r)$, where $0 < d = |z - \rho| < r \leq D \leq E$. Write its principal part as $c_\rho/(s - \rho)$. Subtracting this term gives a holomorphic function on $B(z, r)$; removal at ρ follows from the assumed local Laurent expansion. Fix $d < r' < r$. Cauchy's estimate on the circle of radius r' gives constants M, M' independent of k such that

$$\left| A_k(z) - \frac{c_\rho}{(z - \rho)^{k+1}} \right| \leq M(r')^{-k}, \qquad |A_k(\sigma)| \leq M'(r')^{-k}. \tag{2.10}$$

For the second estimate use $r' < D$. Consequently (2.8), along its arbitrarily distant windows, implies $|c_\rho| d^{-k-1} \leq (M + M')(r')^{-k} + Ce^{\eta U} E^{-k}$. Choose $\eta > 0$ with $e^\eta d/E < 1$ and multiply by d^k. Because $k \geq U \to \infty$, both terms on the right tend to zero, whereas the left is the positive constant $|c_\rho|/d$. This is impossible.

In the second case the isolated pole is real, say $\beta > \theta$, and $d = \sigma - \beta < r \le E$. Choose $d < r' < r$ and apply Cauchy's estimate after subtracting $c_\beta/(s-\beta)$. Positivity of the real residue now gives $\operatorname{Re} A_k(\sigma) \ge c_\beta d^{-k-1} - M(r')^{-k}$. The first term on the left of (2.8) is nonnegative. We therefore obtain $c_\beta d^{-k-1} \le M(r')^{-k} + Ce^{\eta U}E^{-k}$, and the same choice of η gives the same contradiction. Both alternatives have been excluded. □

The proof uses the classical relation between signs and singularities of Mellin transforms, familiar from Landau's oscillation theorem [13]. Its distinctive feature in the present development is the localization to arbitrarily distant logarithmic windows: the derivative order is chosen so that the positive kernel is concentrated inside the available window. The two-center argument compares a complex derivative with a real derivative, and the positive real-pole residue treats the remaining case. Ordinary global one-sided Landau arguments can already detect nonreal singularities; the improvement needed here is the weaker hypothesis on where the upper bound is available.

For a primitive character ψ of order o, let $\mathcal{C}_\psi$ be the set of primitive characters inducing the powers of ψ. This set is closed under conjugation, contains 1 and ψ, and is the retained family in the source. The corresponding weighted remainder is

$$H_\psi(y) = G_1(y) + \frac{1}{3o+1} \sum_{\chi \in \mathcal{C}_\psi \setminus \{1\}} \operatorname{Re} G_\chi(y). \tag{2.11}$$

All the weights are positive. Their exact values are supplied by the arithmetic construction; the analytic criterion needs only positivity and the presence of the target character.

Corollary 2.5. *Suppose that, for every primitive ψ and every sufficiently small $\varepsilon > 0$, there are arbitrarily large Y such that $H_\psi(y) \le y^{1/2+A_\psi\varepsilon}$ for all $Y \le y \le Y^\lambda$, where $\lambda > 1$ and A_ψ is fixed independently of ε and Y. Then every Dirichlet L-function is nonzero on $\operatorname{Re} s > 1/2$, apart from the principal pole at $s = 1$.*

Proof. Fix ψ and a point s with $\operatorname{Re} s > 1/2$. Choose θ strictly between $1/2$ and $\min(\operatorname{Re} s, 1)$, and then choose a permitted $\varepsilon > 0$ so small that $A_\psi\varepsilon \le \theta - 1/2$. For $y \ge 1$, the assumed upper bound implies $H_\psi(y) \le y^\theta$. Set $\delta = \lambda - 1$ and apply Theorem 2.4 to the family (2.11). Proposition 2.1 then excludes any zero of $L(s,\psi)$ to the right of θ, and in particular at the chosen point s. The principal pole is excluded from the statement by convention.

If χ modulo N is induced by a primitive character χ^* of conductor $q \mid N$, then $L(s,\chi) = L(s,\chi^*)\prod_{p\mid N}(1-\chi^*(p)p^{-s})$. When $\operatorname{Re} s > 0$, every factor is nonzero, since $|\chi^*(p)p^{-s}| \le p^{-\operatorname{Re} s} < 1$. Thus primitive nonvanishing implies nonvanishing for every character. □

The quantifier order is essential. The constant A_ψ is fixed before ε is sent to zero; the character, ε, and the auxiliary arithmetic data may influence the beginning of a good window. A proof producing $A_\psi = A_\psi(\varepsilon)$ with no control as $\varepsilon \to 0$ would not justify this corollary. In the source the corresponding exponent is $1/2 + (e_D + e_L + 7)\varepsilon$, and `TopWin.lean` chooses e_D, e_L before ε. The source theorem `zeroFreeAll_of_topG` packages the passage from the geometric family statements to this analytic implication. This section proves the implication from the window estimate; establishing that estimate remains the arithmetic task.

Proposition 2.6 (`critLine`; `rh_of_every`)**.** *Assume nonvanishing of every Dirichlet L-function on $\operatorname{Re} s > 1/2$, apart from its possible pole at 1. Then every zero in $0 < \operatorname{Re} s < 1$ of any Dirichlet L-function lies on $\operatorname{Re} s = 1/2$, and every nontrivial zero of ζ lies on that line.*

Proof. For a primitive character χ of conductor q, let $a \in \{0,1\}$ satisfy $\chi(-1) = (-1)^a$. Its completed L-function is $\mathcal{L}(s,\chi) = (q/\pi)^{(s+a)/2}\Gamma((s+a)/2)L(s,\chi)$. The functional equation is $\mathcal{L}(s,\chi) = \epsilon_\chi \mathcal{L}(1-s,\overline{\chi})$ with $|\epsilon_\chi| = 1$. In the open strip $0 < \operatorname{Re} s < 1$, both gamma factors are finite and nonzero. A zero with real part below $1/2$ would therefore give a zero of $L(1-s,\overline{\chi})$ with real part above $1/2$, contrary to the hypothesis. A zero to the right is excluded directly. The finite Euler-factor identity used in the preceding proof transfers the conclusion to imprimitive characters throughout the open strip.

For ζ, the same argument handles $0 < \operatorname{Re} s < 1$. To check that no nontrivial zero has been missed on the left, use $\zeta(0) = -1/2$ and the functional equation $\zeta(s) = 2(2\pi)^{s-1}\Gamma(1-s)\cos(\pi(1-s)/2)\zeta(1-s)$. For $\operatorname{Re} s < 1/2$, $s \neq 0$, the first three factors other than the cosine are finite and nonzero, and $\zeta(1-s) \neq 0$ by the hypothesis. The cosine vanishes precisely when s is an even integer; under $\operatorname{Re} s < 1/2$ and $s \neq 0$, these are the negative even integers. They are exactly the trivial zeros. Nontrivial zeros to the right of the critical line are excluded by assumption, completing the reduction. □

3. Finite Fourier analysis and the arithmetic balance

The finite Fourier calculation is best separated from the construction of the number field. Let A be a finite abelian group, written additively, with $|A| = n$, and put $G = (\mathbb{Z}/2\mathbb{Z})^2 \times A$ and $f = 2n$. Throughout this section we use the coordinates of the character-sum modules: $\eta(\alpha,\beta,t) = (-1)^\alpha$, $\nu(\alpha,\beta,t) = (-1)^\beta$, $\Phi = \ker\eta$, and $c = (1,1,0)$. Thus the two copies of A in Φ are $\tau_1(t) = (0,0,t)$ and $\tau_2(t) = (0,1,t)$. For $t \in A$ the map $g \mapsto \rho_t(g) := c + g + \tau_1(t)$ bijects $G \setminus \Phi$ with Φ. The geometric modules sometimes interchange the two quadratic coordinates; this change of coordinates does not change any character sum, but it must be made before comparing formulas.

For each rational prime p, let $I_p \leq G$ be an inertia subgroup and let $F_p \in G$ be a lift of a Frobenius class in G/I_p. For $\chi \in \widehat{G}$ and $r \geq 1$ set $v_\chi(p^r) = |I_p|^{-1}\sum_{h\in I_p}\chi(rF_p + h)$. The definition is independent of the chosen lift. It equals $\chi(F_p)^r$ when χ is trivial on inertia and is zero otherwise. In an abelian extension this is precisely the prime-power coefficient of the logarithmic derivative of the primitive Artin L-function. In particular $v_1(p^r) = 1$, including at ramified primes.

Define $S_\chi(y) = \sum_{p^r \leq y} \log p\, v_\chi(p^r)(y/p^r - 1)$ for $y > 0$. Given numbers $\ell_\chi \in \mathbb{C}$ for $\chi \neq 1$, define the compensated sums by $G_1(y) = S_1(y) - y\log y + (1+\gamma_E)y$ and $G_\chi(y) = S_\chi(y) + y\ell_\chi$ for $\chi \neq 1$. In the application, $\ell_\chi = L'(1,\chi)/L(1,\chi)$. The plus sign is essential: the order-y contribution in S_χ is $-yL'(1,\chi)/L(1,\chi)$.

Let $a > 0$ and $w : \Phi \to \mathbb{R}_{>0}$. Define the nonnegative arithmetic functional

$$\mathcal{A}(a,t,w) = \sum_p \sum_{r\geq 1} \frac{2\log p}{|I_p|} \sum_{g \in (rF_p + I_p)\setminus\Phi} \left(\frac{a}{p^r} - w(\rho_t(g))\right)_+, \qquad z_+ := \max\{z, 0\}.$$

The sum is finite, since a nonzero term requires $p^r < a/\min_\Phi w$. Put

$$\mathcal{M}(a,w) = a\left(\log a - 1 - \gamma_E + \ell_\eta - \frac{1}{f}\sum_{\rho\in\Phi}\log w(\rho)\right).$$

Proposition 3.1 (`ch_balance`)**.** *With these definitions,*

$$\mathcal{A}(a,t,w) - \mathcal{M}(a,w) = \frac{1}{f}\sum_{g\notin\Phi}\sum_{\chi\in\widehat{G}} \overline{\chi(g)}\, w(\rho_t(g)) G_\chi\!\left(\frac{a}{w(\rho_t(g))}\right).$$

Proof. Character orthogonality gives $f^{-1}\sum_\chi \overline{\chi(g)} v_\chi(p^r) = 2|I_p|^{-1}\mathbf{1}_{rF_p+I_p}(g)$. Multiplication by $\log p\,(a/p^r - w(\rho_t(g)))_+$ and summation therefore express $\mathcal{A}$ as the right-hand side with G_χ replaced by S_χ. To replace the latter, use $S_1(y) - G_1(y) = y\log y - (1+\gamma_E)y$ and $S_\chi(y) - G_\chi(y) = -y\ell_\chi$ for $\chi \neq 1$. Since $\mathbf{1}_{G\setminus\Phi} = (1-\eta)/2$, the sum $\sum_{g\notin\Phi}\overline{\chi(g)}$ is f for $\chi = 1$, $-f$ for $\chi = \eta$, and zero otherwise. The nontrivial characters consequently contribute $a\ell_\eta$. The trivial character contributes $a\log a - (1+\gamma_E)a - (a/f)\sum_{\rho\in\Phi}\log w(\rho)$, where we used the bijectivity of ρ_t. Their sum is $\mathcal{M}(a,w)$. □

The inertia average treats ramified and unramified primes in the same identity. In the host used by the development, an unramified prime inert for the quadratic character η contributes only at odd powers and with coefficient $2\log p$. At the distinguished ramified prime d, the inertia group has order two and meets both cosets of Φ. Every coset $rF_d + I_d$ consequently contains exactly one element outside Φ, and the coefficient is $\log d$ at each integer power d^r. At the

other exceptional primes, the Frobenius coset lies in Φ, so its contribution vanishes. These contributions are the inert and ramified parts of the arithmetic functional.

A constant weight removes every character except 1 and η from the balance. The point of varying the metric is therefore already visible in finite Fourier analysis: different places must sample different arguments if one is to retain a prescribed nonquadratic character. This elementary observation explains the role of the three weight ranges, without attributing a new theorem to character orthogonality itself.

Let $A_0 \le A$, suppose $T_2 \subseteq A_0$, and write $\delta_2 = |T_2|/|A|$. For an arbitrary family of complex numbers z_χ, the exact translate average is

$$\frac{1}{|A_0|f} \sum_{u\in A_0} \sum_{s\in T_2} \sum_{\chi\in\widehat{G}} \overline{\chi(1,0,s+u-t)}\, z_\chi = \frac{\delta_2}{2} \sum_{\chi|_{A_0}=1} \chi(1,0,t)\, z_\chi.$$

Indeed, the average over u is zero unless $\chi|_{A_0} = 1$; for such a character $\chi(s) = 1$ and $\overline{\chi(1,0,-t)} = \chi(1,0,t)$. There are $4[A : A_0]$ surviving characters. This projection depends only on character orthogonality.

The distribution of top places determines their surviving coefficients. If A/A_0 has order o, give its zero coset mass $4/(3o+1)$ and each of its other cosets mass $3/(3o+1)$. The Fourier coefficient of this probability distribution is 1 at the trivial character and $1/(3o+1)$ at every nontrivial character. To check the latter statement, sum a nontrivial character over the quotient: the sum over the nonzero cosets is -1, leaving $(4-3)/(3o+1)$. The construction arranges that the empirical distribution of the top set approximates these masses. The resulting positive coefficients explain why the analytic argument can detect every character in the retained cyclic family.

The estimate for discarded characters is a separate analytic input. Auxiliary-prime selection supplies a square-root bound for all characters in the designated bounded set, together with control of their logarithmic derivatives at one and an explicit treatment of real characters. A large-argument bound is also supplied uniformly for every host character. These selection results are described at the end of the section.

For orientation, the scales used in the construction are $b = \exp(x^\varepsilon)$, $a = xb$, $X' = \exp(x^{\varepsilon/3})$, $b' = a/X'$, and $w'' = b'/(xn)$. Thus $a/w'' = xnX'$ exactly. Weights near b sample G_χ near x, whereas weights bounded above by $2b'$ sample it at arguments at least $X'/3$, provided a_t/b lies within one of x and $x \ge 3$. Thin weights are small relative to b, but their logarithms need not be small: $\log w'' = x^\varepsilon - x^{\varepsilon/3} - \log n$. The proof must therefore preserve the exact cancellations involving logarithmic weights rather than simply declaring their geometric mean negligible.

The sparse sets used to control arithmetic progressions satisfy a bounded three-point relation condition. For a set T of nonnegative integers and an integer $L \ge 2$, the required condition excludes $j(s_2 - s_1) = i(s_3 - s_1)$ whenever s_1, s_2, s_3 are distinct elements of T and i, j are distinct nonzero integers with $|i|, |j| \le L$. The relevant construction adapts Behrend's sphere method [4] to these bounded relations; it controls the number of intersections of a translate of T with short progressions.

The required continuity estimate also has a short proof. Write $\psi_0(y) = \sum_{p^r\le y} \log p$. The inequalities $\psi_0(2n) - \psi_0(n) \le \log \binom{2n}{n} \le 2n \log 2$ imply, by summing over dyadic integers and using monotonicity, that $\psi_0(y) \le 4(\log 2)y < 3y$ for $y \ge 1$; the bound is trivial for $0 < y < 1$. Since $|v_\chi(p^r)| \le 1$, the function $w \mapsto wG_\chi(a/w)$ is continuous and piecewise differentiable on compact subintervals of $(0,\infty)$. For $\chi \ne 1$ its derivative, away from the breakpoints, is $-\sum_{p^r\le a/w} \log p\, v_\chi(p^r)$; the compensating term $a\ell_\chi$ is constant. For $\chi = 1$ the derivative is $a/w - \psi_0(a/w)$. Both derivatives have absolute value at most $3a/w$. Thus

$$|wG_\chi(a/w) - w'G_\chi(a/w')| \le \frac{3a}{\min\{w,w'\}}|w - w'|.$$

The dependence on a and on the least weight is retained when this estimate is applied.

Here is a quantitative version of the extraction step. It makes clear exactly where the discarded-character estimates enter. Let $\Psi = \{\chi \in \widehat{G} : \chi|_{A_0} = 1\}$, let $K \subseteq \Psi$ contain 1, and

assume $\chi(1,0,0)=1$ for $\chi\in K$. Put $M=|\Psi|=4[A:A_0]$ and $q_\chi=1$ if $\chi|_A=1$, and $q_\chi=c_0$ otherwise, where $0\le c_0\le 1$. Define $\mathcal{G}(x)=\sum_{\chi\in K}q_\chi\operatorname{Re}G_\chi(x)$ and $\Sigma(x)=\sum_{\chi\in\Psi\setminus K}|G_\chi(x)|$. Let $b,b',X'>0$, $\theta,\rho\ge 0$, and let $T\subseteq A$ and $T_2\subseteq A_0$ be nonempty, let $m\ge 1$ be an integer, write $r=|T|$ and $\delta=|T_2|/|A|$, and suppose $x\ge 3$ and $|a_t/b-x|\le 1$ for every $t\in T$. Assume $|r^{-1}\sum_{t\in T}\chi(0,0,t)-q_\chi|\le\rho$ for every $\chi\in K$; for characters trivial on A this equality is exact with $\rho=0$. Suppose $|\ell_\chi|\le K_0$ for $\chi\in K\setminus\{1\}$ and $|\ell_\chi|\le C\log x$ for $\chi\in\Psi\setminus K$, with $K_0,C\ge 0$.

For $u\in A_0$ and $1\le j\le m$, let $w_{u,j}:\Phi\to\mathbb{R}_{>0}$ have values in $[b,b+h]$ at $\tau_2(T_2+u)$ and values in $(0,2b']$ elsewhere, with $h\ge 0$ and $b'X'\le xb$. Assume the uniform large-argument bound $|G_\chi(y)|\le\theta y$ for every $\chi\in\widehat{G}$ and $y\ge X'/3$, and the two numerical inequalities $2f(x+1)\theta\le x^{-1}$ and $6f(x+1)h\le b/x$. Finally assume, for every $u\in A_0$, that $\sum_{t\in T}\sum_{j=1}^m[\mathcal{A}(a_t,t,w_{u,j})-\operatorname{Re}\mathcal{M}(a_t,w_{u,j})]\le E$. Set

$$V_x=4(1+\log(x+1))+|\gamma_E|+K_0+C\log x,$$
$$P_x=x\{4(1+\log(x+1))+2+|\gamma_E|+K_0\},\qquad R_x=V_x+\rho P_x.$$

Proposition 3.2 (`ch_core`, quantitative form)**.** *Under these hypotheses,*

$$\mathcal{G}(x)\le\Sigma(x)+MR_x+\frac{2E}{rm\delta b}+\frac{4}{\delta x}.$$

Proof. Fix u,t,j, and freeze the thick weights at b. The preceding Lipschitz estimate bounds the error for each thick summand by $3(x+1)h$. At every other place, $a_t/w\ge(x-1)b/(2b')\ge X'/3$, so the absolute value of $wG_\chi(a_t/w)$ is at most $(x+1)b\theta$. There are f places and $2f$ characters, and the balance has prefactor $1/f$. Its difference from the frozen thick part therefore has absolute value at most $6f(x+1)h+2f(x+1)b\theta\le 2b/x$. Sum over t,j, average over u, and apply the translate identity. With $y_t=a_t/b$ this gives

$$\frac{m\delta b}{2}\sum_{t\in T}\operatorname{Re}\sum_{\chi\in\Psi}\chi(1,0,t)G_\chi(y_t)\le E+\frac{2rmb}{x}.$$

Partial summation of $\psi_0(y)\le 3y$ gives $\sum_{p^r\le y}(\log p)/p^r\le 3(1+\log y)$ for $y\ge 1$. Differentiating the compensated sums between their breakpoints consequently shows that $|G_\chi(y_t)-G_\chi(x)|\le V_x$ for $\chi\in\Psi$; the same elementary bound gives $|G_\chi(x)|\le P_x$ for $\chi\in K$. For a retained character, changing y_t to x and then using its top-place Fourier coefficient loses at most $V_x+\rho P_x=R_x$ after division by r. For a discarded character, the corresponding expression is at least $-|G_\chi(x)|-V_x$. Thus the sum on the left is at least $r(\mathcal{G}(x)-\Sigma(x)-MR_x)$. Divide by the positive quantity $rm\delta b/2$ to obtain the conclusion. □

For the cyclic family of a fixed primitive character, K reads the retained primitive characters, $c_0=(3o+1)^{-1}$, $M=4o$, and $\mathcal{G}=H_\psi$. The empirical coset estimate gives $\rho\le Mx^{-1+\varepsilon}$, so $MR_x=O_{\psi,C}(x^\varepsilon\log x)$. The host package supplies $\Sigma(x)\le 4o\,x^{1/2+\varepsilon}$. If $\delta\ge x^{-\varepsilon}$ and the normalized geometric error satisfies $2E/(rm\delta b)\le x^{1/2+8\varepsilon}$, Proposition 3.2 yields $H_\psi(x)\le x^{1/2+9\varepsilon}$ for all sufficiently large admissible scales. Indeed all terms except the displayed geometric error are $O_\psi(x^{1/2+\varepsilon})$, and x^ε eventually absorbs their fixed constants. The source theorem `core_gm` carries out the same normalization with the full geometric error bracket and gives the more general exponent $1/2+(e_D+e_L+7)\varepsilon$; the construction uses $e_D=2$ and $e_L=0$.

There is an essential analytic selection step behind the estimate for Σ. The auxiliary-prime selection chooses the auxiliary primes q,d so that, for every character ξ modulo the fixed original modulus and every nontrivial quadratic twist $\xi(\frac{\cdot}{q})^i(\frac{\cdot}{d})^j$, its associated primitive character χ satisfies $|G_\chi(y)|\le y^{1/2+\varepsilon/2}$ and $|L'(1,\chi)/L(1,\chi)|\le 4\log y$ throughout the prescribed scale interval. It also supplies a zero-free rectangle $\operatorname{Re}s>1-1/400$, $|\operatorname{Im}s|\le Q_j$ for the relevant real twists, with $Q_j=8M_0\exp(Y^{\varepsilon/5})\exp(Y^{\varepsilon/5}(\log Y)^2)^j$ and M_0 the original modulus. A mean-square selection argument constructs such pairs. Every discarded host character is then identified with one of these nontrivial twists, and its bound is summed over at most $4o$ characters. The selected family is what permits the stated square-root estimate.

For the thin-place bound, the zero-free rectangles are first transferred to all real host characters involving q or d. This information is then combined with the classical zero-free-region estimates

and the common conductor bound. Real characters involving neither auxiliary prime lie in a fixed finite family, since the odd part of the remaining conductor is bounded and a primitive quadratic character has bounded conductor at 2. Their positive distance from 1 can therefore be absorbed by increasing the starting scale. This treats the possible exceptional real zero within the host selection. The resulting estimates are available before the geometric extraction of H_ψ.

4. Frames and transfer of slopes

The role of the frame inequality is to pass from many maps of small rank to information about almost all directions of a much larger target. The following form is elementary and includes all constants that are subsequently integrated.

Proposition 4.1 (`frame_linear_algebra_first`)**.** *For a finite index set I, let $B_i : V_i \to H$ be linear maps between finite-dimensional complex Hermitian spaces, let $c_i > 0$, and put $R = \sum_i \dim V_i > 0$, $D = \dim H > 0$, and $Q = \sum_i c_i B_i B_i^*$. Suppose $0 \le e_F \le 1/4$, $\eta_F \ge 0$, and*

$$\sum_i \|c_i B_i^* B_i - I\|_{\mathrm{HS}}^2 \le e_F R, \qquad \operatorname{Tr}(Q^2) \le \frac{1+\eta_F}{D}(\operatorname{Tr} Q)^2.$$

For subspaces $N_i \subseteq V_i$ with $\sum_i(\dim V_i - \dim N_i) \le \epsilon R$, their image span $S = \sum_i B_i N_i$ satisfies $D - \dim S \le (\eta_F + 4\epsilon + 4\sqrt{e_F})D$.

Proof. Let P_i project onto $N_i^\perp$ and let P project onto $S^\perp$. Since $PB_i = PB_iP_i$, positivity and the trace identity give $\operatorname{Tr}(PQ) \le \sum_i \operatorname{Tr}(P_i c_i B_i^* B_i)$. Write $\Delta_i = c_i B_i^* B_i - I$. Cauchy–Schwarz, applied to the direct sum of all matrix spaces, yields $\operatorname{Tr} Q \ge (1-\sqrt{e_F})R > 0$ and $\operatorname{Tr}(PQ) \le (\epsilon + \sqrt{e_F})R$; for the second inequality one may use $\sum_i \|P_i\|_{\mathrm{HS}}^2 \le R$. Set $r = \operatorname{Tr}(PQ)/\operatorname{Tr} Q$. Then $0 \le r \le 1$ and $r \le 2(\epsilon + \sqrt{e_F})$. Cauchy–Schwarz for the Hilbert–Schmidt pairing with $I - P$ now gives $(1-r)^2(\operatorname{Tr} Q)^2 \le \dim S \operatorname{Tr}(Q^2)$. Consequently $\dim S/D \ge (1-r)^2/(1+\eta_F) \ge 1 - \eta_F - 2r$, proving the assertion. The last inequality follows by multiplying by $1+\eta_F$; the difference of the two sides is $r^2 + \eta_F^2 + 2\eta_F r \ge 0$. □

For the arithmetic application, fix a number field K. A Hermitian lattice is a finite projective $\mathcal{O}_K$-module with compatible Hermitian norms at its archimedean embeddings. For a lattice E of positive rank r and a nonzero rational section v of its determinant line, define

$$\widehat{\deg}_{\mathrm{Ar}}(E) = \frac{1}{[K:\mathbb{Q}]}\left(\sum_{\mathfrak{p}} \operatorname{ord}_{\mathfrak{p}}(v; \det E)\log N\mathfrak{p} - \sum_{\sigma: K\hookrightarrow\mathbb{C}} \log\|v\|_\sigma\right), \qquad \mu_{\mathrm{Ar}}(E) = \frac{\widehat{\deg}_{\mathrm{Ar}}(E)}{r}.$$

The valuation is measured relative to a local generator of $\det E$; the sum over embeddings counts a conjugate pair twice. The product formula proves independence of v. The source uses $\mu_2 = 2\mu_{\mathrm{Ar}}$ and doubles the averaged logarithmic operator height as well. Its normalized logarithmic covolume satisfies $\operatorname{lcv}(E) = -r\mu_2(E)$. The slope-transfer inequality below is homogeneous: either convention gives the same statement when μ, h_i, s, u, H_0 are scaled together.

To state the slope consequence cleanly, suppose that each lattice E has a decreasing Harder–Narasimhan filtration $E^{\ge t}$ and a multiset of slopes $\lambda_j(E)$, counted with multiplicity. Thus $\operatorname{rk} E^{\ge t}$ is the number of slopes at least t, and $\mu(E)$ is their mean. The filtration property used from arithmetic slope theory [6] is the morphism inequality: an integral map $B : E \to F$ of logarithmic height at most h satisfies $B(E^{\ge t}) \subseteq F^{\ge t-h}$. For a nonzero integral map, the standard height is $h_{\mathrm{Ar}}(B) = [K:\mathbb{Q}]^{-1}\sum_\sigma \log\|B_\sigma\|_{\mathrm{op}}$; the source height is twice this quantity. The average here uses all archimedean embeddings. A zero map satisfies the filtration inclusion without a height condition.

For clarity, the inclusion follows directly from the usual extremal properties of the Harder–Narasimhan slopes. If $B : E \to F$ is nonzero, its image I is a quotient of E algebraically. Compare the quotient metric on $E/\ker B$ with the metric on I induced from F. The determinant operator-norm inequality gives $\mu(E/\ker B) \le \mu(I) + h$. The extremal properties give $\mu_{\min}(E) \le \mu(E/\ker B)$ and $\mu(I) \le \mu_{\max}(F)$, with saturation only increasing the degree of a sublattice. Hence $\mu_{\min}(E) \le \mu_{\max}(F) + h$. Apply this inequality to the composite from $E^{\ge t}$ to $F/F^{\ge t-h}$. If it were nonzero, its source would have minimal slope at least t, whereas its target would have

maximal slope strictly below $t-h$, a contradiction. Restriction and quotient do not increase the operator norms, so the same height bound applies. This proves the required inclusion.

Proposition 4.2 (theoremFSharp)**.** *Let I be finite and nonempty, let E_i and F be nonzero full Hermitian lattices over K, with the filtration and morphism property just stated, and let $B_i : E_i \to F$ be integral maps, with heights bounded by h_i, and let $d_i = \operatorname{rk} E_i$, $R = \sum_i d_i > 0$. At one complex embedding suppose there are maps $\widetilde{B}_i : E_{i,\mathbb{C}} \to \widetilde{H}$ whose projections to $F_{\mathbb{C}}$ are B_i, where $F_{\mathbb{C}}$ is a quotient of $\widetilde{H}$ and $\dim F_{\mathbb{C}} \geq (1-\delta_F)\dim \widetilde{H}$ for $0 \leq \delta_F < 1$. Assume all the hypotheses of Proposition 4.1 for the maps $\widetilde{B}_i$. Let $u > 0$, $H_0 \geq 0$, and $s \in \mathbb{R}$ satisfy*

$$\frac{1}{R}\sum_i d_i(\mu(E_i) - h_i) \geq s, \qquad \mu_{\max}(E_i) - h_i \leq s + u, \qquad \mu_{\min}(F) \geq s - H_0.$$

Then

$$\mu(F) \geq s - \frac{4u\log(1 + H_0/u) + (\eta_F + 4\sqrt{e_F})H_0}{1-\delta_F}.$$

Proof. Consider the pooled list of shifted source slopes $\lambda_j(E_i) - h_i$. Its mean is at least s, and every entry is at most $s+u$. If $k(\tau)$ entries are below $s-\tau$, then $Rs \leq (R-k(\tau))(s+u)+k(\tau)(s-\tau)$, so $k(\tau)/R \leq u/(u+\tau)$ for $\tau \geq 0$. Notice that this is a statement about the pooled list; no lower bound on the mean of an individual source is needed.

Put $N_i(\tau) = E_i^{\geq s+h_i-\tau}$. Proposition 4.1 shows that the span of the images $\widetilde{B}_i N_i(\tau)$ has codimension at most $(\eta_F + 4u/(u+\tau) + 4\sqrt{e_F})\dim \widetilde{H}$. A quotient cannot increase codimension. By the morphism inequality, the images in $F_{\mathbb{C}}$ lie in $F_{\mathbb{C}}^{\geq s-\tau}$. Hence the fraction of target slopes below $s-\tau$ is at most $(1-\delta_F)^{-1}(\eta_F + 4u/(u+\tau) + 4\sqrt{e_F})$.

For any target slope $\lambda \geq s - H_0$, the equality $\min\{\lambda, s\} = s - \int_0^{H_0} \mathbf{1}_{\lambda < s-\tau}\, d\tau$ holds; the choice of strict inequality at a single integration endpoint is irrelevant. Average this identity over the target slopes and use $\lambda \geq \min\{\lambda, s\}$. Integrating the corank bound and evaluating $\int_0^{H_0} u/(u+\tau)\, d\tau = u\log(1 + H_0/u)$ proves the result. □

In the application one takes $s = R^{-1}\sum_i d_i(\mu(E_i) - h_i)$ itself. This gives a bound in terms of the weighted mean that actually enters the geometric bookkeeping. The same mean is used in the upper bound for the shifted source slopes and in the lower bound for the target slopes. The case $u = 0$ follows by applying the proposition with any positive u', then letting $u' \downarrow 0$; by convention $u\log(1 + H_0/u) = 0$ at $u = 0$. The sharper source theorem used here has no additional $\log|I|$ loss.

5. The cancellation of the arithmetic main terms

The slope-transfer inequality becomes useful only after the local normalizations have been matched. We give the numerical argument explicitly, since it is short, contains the decisive cancellation, and clarifies exactly what the construction must supply. This section uses the doubled slope and map-height normalization of the source. The local gains and evaluation estimates below are expressed in that same normalization.

Fix one translate of the thick set. Let T be the set of top places, let $1 \leq j \leq m$ index the transforms in a column, and let $a_t > 0$ and $w_j : \Phi \to \mathbb{R}_{>0}$ be their parameters. Put $A = \sum_{t\in T} a_t$. Let the finite nonempty set I index the column lattices, and let $p_i = d_i/\sum_{i'} d_{i'}$, where all $d_i > 0$. We write $\mathbb{E}_i Z_i = \sum_i p_i Z_i$. Thus the averaging in this section is weighted by ranks; replacing it by an unweighted average would change the statement.

Let N_i be the i-th integral column lattice, let $\mu_i = \mu(N_i)$, let μ_T be the slope of the dual row lattice, and let h_i be a bound for the height of the corresponding map. We use G for a common finite-place gain and F for a common ramified-place term. The individual columns have nonnegative level losses E_i, nonnegative rounding or comparison losses L_i, and positive degree parameters T_i. The period normalization is $\log\Omega = \gamma_E + \log(\pi/v) - \operatorname{Re}\ell_\eta$, where $v > 0$ and $\ell_\eta = L'(1,\eta)/L(1,\eta)$.

Proposition 5.1 (`key_identity`). *Let the bounds M, E_{HN}, E_D, E_C, E_T, and E_P be nonnegative, and let $g_0 > 0$. Suppose*

$$\mu_i \ge -M + G + F - E_i - L_i, \qquad \mathbb{E}_i(\mu_i - h_i) \le \mu_T + E_{\mathrm{HN}}, \tag{5.1}$$

$$\mu_T + h_i - F \le \sum_{t\in T} m a_t(\log a_t - 1 - \log\Omega) - A\log T_i + E_D, \tag{5.2}$$

$$E_i \le mA\log g_0 + E_C, \qquad \log T_i \ge \sum_{j=1}^{m}\frac{1}{f}\sum_{\rho\in\Phi}\log\frac{w_j(\rho)v}{\pi} + m\log g_0 - E_T, \tag{5.3}$$

$$\sum_{t\in T}\sum_{j=1}^{m}\mathcal{A}(a_t,t,w_j) \le G + E_P. \tag{5.4}$$

Then

$$\sum_{t\in T}\sum_{j=1}^{m}(\mathcal{A}(a_t,t,w_j) - \operatorname{Re}\mathcal{M}(a_t,w_j)) \le E_{\mathrm{tot}}, \qquad E_{\mathrm{tot}} = M + E_{\mathrm{HN}} + \mathbb{E}_i L_i + E_D + E_C + AE_T + E_P. \tag{5.5}$$

The functionals $\mathcal{A}$ and $\mathcal{M}$ are those of Proposition 3.1.

Proof. Average the first inequality in (5.1), subtract $\mathbb{E}_i h_i$, and apply the second one. Rearranging gives

$$G \le \mathbb{E}_i(\mu_T + h_i - F + E_i) + M + \mathbb{E}_i L_i + E_{\mathrm{HN}}.$$

Insert (5.2) and (5.3). The term $mA\log g_0$ from E_i cancels the same term from $-A\log T_i$. The remaining main term is

$$\sum_{t\in T} m a_t(\log a_t - 1 - \log\Omega) - A\sum_{j=1}^{m}\frac{1}{f}\sum_{\rho\in\Phi}\log\frac{w_j(\rho)v}{\pi}.$$

Since $|\Phi| = f$, expanding the logarithm in the second sum contributes $-mA\log(v/\pi)$. Expanding $-\log\Omega$ in the first sum contributes $-mA\log(\pi/v)$. They cancel exactly, leaving $\sum_{t,j} a_t(\log a_t - 1 - \gamma_E + \operatorname{Re}\ell_\eta - f^{-1}\sum_\rho \log w_j(\rho))$, which is $\sum_{t,j}\operatorname{Re}\mathcal{M}(a_t,w_j)$. The remaining error terms, followed by (5.4), are exactly those in (5.5). □

This proposition is the numerical form of the cancellation in `WeightedMeanData.lean`; its named identity is the equality of the two main-term expressions in the proof. The fuller mean inequality retains the individual quantitative losses before they are absorbed into one error bound. The extraction requires this budget separately for every translate of the thick set. The construction supplies the corresponding pointwise form; its source location is recorded in Appendix A. Three cancellations should be kept distinct. The ramified term F cancels between the source slope and the evaluation height. The level term cancels between the finite index loss and the count of integral sections. The volume factor v/π cancels between the period formula and the archimedean degree. Each cancellation uses the same normalization on both sides. None is justified by calling a main term negligible.

The first inequality in (5.1) has a simple arithmetic explanation. Suppose $N_i = J_i N_i^0$, where J_i is a fractional ideal with the metric inherited from the ambient vector space. If $J_i = \prod_{\mathfrak{p}} \mathfrak{p}^{-n_{i,\mathfrak{p}}}$, the conventional normalized Arakelov degree of J_i is $[K:\mathbb{Q}]^{-1}\sum_{\mathfrak{p}} n_{i,\mathfrak{p}}\log N\mathfrak{p}$. Taking determinants shows that $\mu_{\mathrm{Ar}}(J_i N_i^0) = \mu_{\mathrm{Ar}}(N_i^0) + \widehat{\deg}(J_i)$. In the source normalization $\mu_2 = 2\mu_{\mathrm{Ar}}$, both sides of this degree increment are doubled. Thus expanding the integral lattice by prescribed divisibility raises its slope by a weighted local sum. In the construction, the terms in that sum are separated into the favorable gain, the common ramified term, the level losses, and errors caused by integral rounding and lattice comparison. A lower bound for $\mu(N_i^0)$ gives (5.1).

The second inequality in (5.1) is precisely the role of Proposition 4.2. If $s = \mathbb{E}_i(\mu_i - h_i)$, it permits

$$E_{\mathrm{HN}} = \frac{4u\log(1 + H_0/u) + (\eta_F + 4\sqrt{e_F})H_0}{1 - \delta_F}.$$

The parameters e_F, η_F, δ_F measure different defects: the failure of the column maps to be isometric, the lack of flatness of the total frame, and the proportion of discarded row directions. They should not be identified merely because all tend to zero. The spread u is the difference between the largest shifted source slope and the rank-weighted mean. Explicitly, $u = \max_i(\mu_{\max}(N_i) - h_i) - s$. It is a pooled quantity. A lower bound for every individual source mean is neither assumed nor required.

There are two elementary observations behind the subsequent error estimates. If $u > 0$, $H_0 \le Xb$, $u \le U$, and $X, b > 0$, then

$$u \log(1 + H_0/u) \le (U + b) \log(1 + X). \tag{5.6}$$

To see this, concavity of the logarithm gives $\log(1+H_0/u) \le \log(1+X)+((1+H_0/u)/(1+X)-1)$. After multiplication by u, the correction is at most $bX/(1 + X)$. The inequality $X/(1 + X) \le \log(1 + X)$ then proves (5.6). This avoids replacing a logarithmic loss by an unnecessarily large linear loss. The limiting case $u = 0$ contributes zero, because $u \log(1 + H_0/u) \to 0$ as $u \downarrow 0$.

Second, for any fixed $A, B, c > 0$, $x^A \exp(-x^c) \le x^{-B}$ eventually. Taking logarithms reduces this to $(A + B) \log x \le x^c$, and $\log x/x^c \to 0$. The constants and exponents must be fixed before making this comparison. It is this elementary domination, together with bounds such as (5.6), that makes exponentially small frame and rank defects affordable after multiplication by polynomial factors. It does not, by itself, prove that the arithmetic construction has the required defects or uniform constants.

After dividing (5.5) by the positive thick-place normalization, the Fourier extraction gives an upper bound for $H_\psi(x)$. The full source records every loss before performing this division. In particular, a small unscaled error may cease to be small after division by the density of the thick set, and a statement with constants depending on the final scale x would be unusable. Section 8 explains how the construction fixes the constants in the necessary order.

6. The compatible arithmetic construction

The arithmetic construction must supply more than an auxiliary number field and a collection of automorphic forms. It must supply integral structures, compatible Hermitian norms, rational evaluation maps, and estimates uniform in every parameter that will later vary with the analytic scale. The role of the theta correspondence is to put these objects into a single construction. The role of the geometric estimates is to prevent the integral structures from becoming too expensive. Keeping these two tasks separate makes the final slope argument considerably easier to follow.

Write E/F for the CM extension in the host datum, and write k, m, t for its rank parameters. Here k and m are even, t is odd, $2 \le m$, $m+3 \le k$, $t \ge k+m+4$, and $t > 3m$. Put $r = k+m-1$. The construction uses definite Hermitian spaces X_j of rank $k + j(t + 1)$ and Hermitian spaces Y_j of rank $r + jt$, together with isometric embeddings $X_j \hookrightarrow Y_j$ for $j < m$. At the last stage, Y_m is isometric to a space X' of codimension one in X_m. A second collection of Hermitian spaces W_j, of fixed rank r, has signature $(k + j - 1, m - j)$ for $1 \le j \le m$. Thus W_m is positive definite. The changing-signature spaces W_j must be distinguished from the increasing-rank spaces Y_j: they play different roles in the geometric and theta constructions. The signatures of W_j are imposed at every archimedean place and sum to the fixed rank r.

The analytic form of a column is a composite of m integral transforms. At stage j, let $\mathcal{C}_j$ be the finite set indexing the relevant lattice classes and let $a_j(c)$ be the order of the stabilizer of $c \in \mathcal{C}_j$. The transform from stage j to stage $j + 1$ is a weighted sum over $\mathcal{C}_j$, with weight $a_j(c)^{-1}$. The individual summands contain the unfolding operator and the determinant powers prescribed by the archimedean character. If $U_j(c, c')$ denotes the unfolding operator, the normalized summand has the form

$$K_j(c, c') = \rho_j(c)\rho'_j(c')\, D_{j+1}^{-s_j}\, U_j(c, c')\, D_j^{s'_j}, \qquad \Phi_{j+1}(c') = \sum_{c \in \mathcal{C}_j} a_j(c)^{-1} K_j(c, c')\Phi_j(c).$$

The factors ρ_j and ρ'_j have absolute value one. The symbols $D_j^{s'_j}$ and $D_{j+1}^{-s_j}$ abbreviate multiplication by the products of the determinant functions at the archimedean embeddings, with the integral exponents specified by the character data. Their purpose is to use the same normalization for the automorphic transformation law, the period pairing, and the integral lattice. The common denominator D_Γ is divisible by every $a_j(c)$ occurring in the chain. This denominator is part of the arithmetic estimates; it is not discarded after the analytic transforms have been defined.

The finite column index set will be denoted by $\mathcal{U}$. For each $u \in \mathcal{U}$, the column transform and the terminal period pairing produce a bilinear pairing between a column space and the common row space. All these pairings are realized over a common coefficient number field K. Their integral realizations are maps

$$B_u : N_u \otimes K \longrightarrow R^\vee \otimes K,$$

where N_u and R are finitely generated full lattices. Compatibility with complex realization is essential: the algebraic pairing obtained from B_u must become the same pairing used to prove the analytic norm estimates. The corresponding Bessel comparison asserts that, after unitary changes of coordinates on the row and column spaces and multiplication by a nonzero scalar, this pairing is the Bessel pairing. It is a rationality and normalization statement, distinct from the subsequent estimates for the Gram operators.

The geometric input is best understood by examining the inequality it supplies. Use the source's slope normalization μ_2, in which a full lattice M of rank d satisfies $\mu_2(M) = -\operatorname{lcv}(M)/d$; as explained in Section 4, $\mu_2 = 2\mu_{\mathrm{Ar}}$. At a fixed stage, the construction has a core lattice, a lattice of the chosen automorphic type, and its good part. Denote their slopes by μ_{core}, μ_{type}, and μ_{good}. Let $\mathcal{P}_j$ be the prescribed period main term, let $E_0 \geq b > 0$ be the common error scale, and let $C_1, C_2 \geq 0$. The two lower bounds that matter are

$$-\mathcal{P}_j - C_2 E_0 \leq \mu_{\mathrm{core}}, \qquad \mu_{\mathrm{core}} - C_1 E_0 \leq \mu_{\mathrm{type}}.$$

The first comes from the geometric estimates on integral automorphic sections; the second compares the core lattice with its realization at the chosen level. These estimates must hold uniformly at all points of the family, after one common threshold has been chosen.

There is also an orthogonal projector onto the good part. If e is this projector, there is an integer $g \geq 1$ such that $ge(M_{\mathrm{type}}) \subseteq M_{\mathrm{type}}$ and $2\log g \leq b \leq E_0$. The rank discarded by the projection is at most $3\delta(x)$ times the rank of that lattice, where $\delta(x) = \exp(-x^{e_z/7})$ and $e_z > 0$ is fixed in the family. Increase the common threshold so that $\delta(x) \leq 1/12$; the good part then has positive rank. The maximal slope of the ambient lattice is at most $x^A b$, while $|\mathcal{P}_j| \leq x^B b$. The lattice truncation estimate then gives

$$\mu_{\mathrm{good}} \geq \mu_{\mathrm{type}} - 12\delta(x)(x^A + x^B + C_1 + C_2 + 1)E_0.$$

Because $\delta(x)$ decreases faster than every negative power of x, the factor multiplying E_0 is at most one for all sufficiently large x. Combining the three inequalities gives $\mu_{\mathrm{good}} \geq -\mathcal{P}_j - (C_1 + C_2 + 1)E_0$. The recursive comparison between the row lattice at this stage and the good automorphic lattice therefore yields

$$\frac{\operatorname{lcv}(M_j)}{D} \leq -\mu_{\mathrm{good}} \leq \mathcal{P}_j + (C_1 + C_2 + 1)E_0.$$

Here D is the common dimension appearing in the row construction. This is the covolume estimate required by the global argument. The derivation explains why a small exceptional rank is useful only together with a maximal-slope bound and control of the projector denominator: a small vector-space complement can still carry a large arithmetic degree.

The proof of the first lower bound is the substantial geometric part. Away from the terminal definite stage, the source passes to finite covering levels carrying integral automorphic bundles. It constructs maps from negatively shifted weights to the central weight, bounds their operator norms, and compares the ranks at positive and negative shifts. Geometric degrees control the dimensions of the section spaces on the individual components. Arithmetic Hilbert–Samuel estimates convert the corresponding height bounds into degree bounds for the lattices of sections.

The component estimates must then be summed with the same rank weights that occur in the slope, after which descent transfers the estimate back to the original level. At the terminal definite stage a separate bound supplies the end of the recursion. The construction also checks that the geometric norms and the automorphic L^2 norms differ by factors whose logarithms are bounded by a constant times E_0.

The theorem `topCovolSlim2_use_alt_alt` is the main source interface for this geometric input. It combines the covering and height estimates into the uniform covolume bound just described. Its hypotheses include the relevant toroidal compactification, deformation, comparison, arithmetic intersection, and section-extension results. These hypotheses are supplied by proved declarations in the development. Citing the names of those general theories is not a proof of the uniform inequality: the maps, integral structures, rank comparisons, and dependence of the constants are specific to this construction. The geometric part of the source is therefore a substantive input to the present guide, not an elementary consequence of a single classical slope inequality.

At finite places, the construction has two separate duties. First, the divisibility estimates turn the inert-prime sums into gains for the column lattices. For an inert prime and its places above it, the source compares the sum of the local gains with the appropriate von Mangoldt class sums, subtracting explicitly defined losses. It also bounds the indices between the local lattices and proves that the relevant orbit sets are contained in the sets for which the local estimates hold. The residual term is allowed to be $b^{\delta_0(x)}$ with $\delta_0(x) \to 0$, rather than being silently treated as a uniform constant. These are the contents of the main divisibility input.

Second, changing the local type can introduce an index loss. If η_j denotes the weighted sum of these local losses, the source bounds it by

$$\eta_j \le d_j \left[C\{fmbL^2 + (\log x)^2 N_c\} + \sum_{\ell} \log \ell \sum_{1 \le e < t} \mathrm{length}_{\ell,e} \right].$$

Here n is the order of the finite frame group, $f = 2n$, and T_2 is the host's finite set of second-coordinate labels. The auxiliary error budget is $N_c = |T_2| kbL^2/n^2$, where $L = x^{\varepsilon}$ and $b = \exp(x^{\varepsilon})$. The remaining quantities are $\mathrm{length}_{\ell,e} = v_{\ell,e}(E)$, the $\mathbb{Z}_\ell$-lengths of the localization kernels $\mathrm{Sha}_{\ell,e}(E)$ defined in (7.1), over the host field E. Thus the cohomological window $\mathcal{V}(E,t;\sqrt{b})$ supplies both their finiteness and the bound $\sum_\ell \log \ell \sum_{1\le e<t} \mathrm{length}_{\ell,e} \le \sqrt{b}$. The equality between these lengths and the rational bookkeeping values used in the local lattice construction is part of the compatibility data. The separate equality $d_j = D_0$ identifies the small-level rank with the common row dimension. Since $(\log x)^2 \le x^{\varepsilon}$ eventually, the displayed bound becomes the loss estimate needed by the divisibility theorem. This is the precise point at which the cohomological window is used in the geometric comparison.

The archimedean and finite-place comparisons meet in two final inequalities. The first controls the height of each map B_u and compares that height, plus the slope of the dual row lattice, with the period main term. It also controls the common amplitude $h_u + A \log T_u$ as u varies. The second compares the actual column slope with the slope of a simpler reference lattice:

$$\left| \mu_2(N_u) - \mu_2(N_u^{\mathrm{ref}}) - \sum_w \mathrm{sz}_w(s_{u,w}) \right| \le CX''.$$

Here the finite sum records the local fractional-ideal shifts, with the same normalization used for the map heights. The full-lattice and positive-rank conditions are necessary to make these slope comparisons meaningful. In particular, positive column rank is supplied separately from the covolume estimate.

Once these compatible objects and estimates are available, the remainder of the construction is a finite assembly. The frame estimates supply the dimension and near-orthogonality information needed for slope transfer. The irreducibility estimate controls the exceptional row fraction. The local index bounds and the covolume bound supply the height comparison. Finally the global comparison converts this data to the tied chain to which the explicit-formula argument applies. The two principal assembly interfaces are `shapeWin` and `topDesWin`. Their importance is the

order of the quantifiers: all estimates hold for one selected family and one common threshold, and the family satisfies the host and cohomology conditions on its windows.

The final arithmetic-to-analytic interface takes the cohomological window and the explicitly packaged representation-theoretic input, constructs compatible local theta data and evaluation grids, supplies the geometric covolume estimate, and invokes the uniform family construction. The resulting tied chains feed the one-sided analytic estimate and hence the zero-free theorem. This separates the logical roles cleanly: the window theorem chooses scales on which the arithmetic loss is affordable; the theta and geometric construction produces maps realizing the corresponding prime sums; and the analytic argument turns the resulting one-sided bounds into nonvanishing.

The established ingredients include the theta correspondence, rationality and local matching results, Hermitian lattice theory, and arithmetic intersection and section estimates. The distinctive feature that requires explanation here is their simultaneous quantitative use: the same column maps must have rational integral realizations, favorable local divisibility, controlled archimedean heights, and a sufficiently flat collective Gram operator. Their compatibility is the substantial content of the construction result invoked here. The covering, descent, integral comparison, and local matching arguments are located by the source concordance.

7. Cohomology and the choice of windows

We now describe the arithmetic estimate that allows the construction to be carried out on suitable intervals of scales. Its formulation is more informative than a statement that certain cohomology groups are finite. Both their lengths and the order of the quantifiers matter. In particular, the same auxiliary primes must work throughout a chosen interval, while the degree of the host field is allowed to change with the scale.

Let E be a number field and let ℓ be a rational prime. Write S_ℓ for the places of E above ℓ, and let G_{E,S_ℓ} be the Galois group of the maximal extension unramified outside S_ℓ and the infinite places. For an integer $e \geq 1$, the coefficient module used in the source is $\mathbb{Q}_\ell/\mathbb{Z}_\ell(-e)$: the cyclotomic character acts to the power $-e$. Define

$$\mathrm{Sha}_{\ell,e}(E) = \ker\left(H^1(G_{E,S_\ell}, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e)) \longrightarrow \prod_{v\mid\ell} H^1(E_v, \mathbb{Q}_\ell/\mathbb{Z}_\ell(-e)) \right). \tag{7.1}$$

Here H^1 denotes continuous Galois cohomology. The definition imposes both the global unramified-outside-ℓ condition and the vanishing of the local classes at the primes above ℓ. The negative twist is the one used in the formal source.

Let $v_{\ell,e}(E)$ be the $\mathbb{Z}_\ell$-module length of $\mathrm{Sha}_{\ell,e}(E)$, allowing $+\infty$. For an integer $t \geq 1$ and a real $B \geq 0$, the assertion $\mathcal{V}(E,t;B)$ means that all $v_{\ell,e}(E)$ are finite for $1 \leq e < t$ and that

$$\sum_{\ell\in P} \log \ell \sum_{e=1}^{t-1} v_{\ell,e}(E) \leq B \quad \text{for every finite set } P \text{ of rational primes.} \tag{7.2}$$

This is the content of `VVle`. The finiteness requirement must be retained: converting an extended natural number to an ordinary natural number before establishing finiteness would not express the intended assertion.

One may then write the sum over all primes without ambiguity. Indeed, if infinitely many primes had positive inner sum, choosing more than $B/\log 2$ of them would contradict (7.2). Thus only finitely many primes contribute for fixed E, t, B. Equivalently, the nonnegative infinite sum is the supremum of its finite partial sums and is at most B. The finite-partial-sum formulation in the code avoids using summability before it has been proved.

We next specify the host fields. Fix a primitive character ψ. Let F_ψ be the subfield of the cyclotomic field of conductor N_ψ fixed by $\ker\psi$, and put $F_h^+ = \mathbb{Q}(\zeta_{2^{h+2}})^+$. The field used at level h, with auxiliary primes q, d, is the compositum

$$E_{\psi,h,q,d} = F_\psi F_h^+ \mathbb{Q}(\sqrt{-q})\mathbb{Q}(\sqrt{-d}). \tag{7.3}$$

A definition independent of choices of embeddings is obtained inside a common cyclotomic field: its fixing subgroup consists of residues a satisfying $a \equiv \pm 1 \pmod{2^{h+2}}$, $\psi(a) = 1$, $(a/q) = 1$, and $(a/d) = 1$. This is the fixed-field condition called `IsChiField`. It also explains why the discrete level is a power of two, up to a factor depending only on ψ.

Fix $0 < \varepsilon < 1/52$. Choose a function $h : [1, \infty) \to \mathbb{N}$ with

$$x^{2+8\varepsilon} \le 2^{h(x)} \le 2x^{2+8\varepsilon} \qquad (x \ge 1). \tag{7.4}$$

Such a function is obtained by rounding $(2 + 8\varepsilon) \log x / \log 2$ upward; an arbitrary value may be assigned for $x < 1$. The formal window assertion is uniform over every choice satisfying (7.4). It is not a statement about one specially chosen sequence of favorable discrete levels.

For $T > 1$, call q, d an admissible window pair if they are primes, $q \equiv 3 \pmod 8$, $d \equiv 7 \pmod{16}$, and

$$\tfrac{1}{2}T^{\varepsilon/5} < \log q \le T^{\varepsilon/5}, \qquad \tfrac{1}{2}T^{\varepsilon/5}(\log T)^2 < \log d \le T^{\varepsilon/5}(\log T)^2,$$

with $q \mid d + 1$ and $p \mid d + 1$ for every odd prime $p \le T^{\varepsilon/5}$. Notice the logarithmic factor in the size of d. The two primes have different assigned sizes. Further analytic conditions on the pair, discussed in Section 3, are imposed when constructing the actual character host.

Proposition 7.1 (`vwin_of_row`)**.** *For each primitive character ψ, each $0 < \varepsilon < 1/52$, and each selector h satisfying (7.4), there exist $J \ge 1$ and $X_0 > 1$ such that the following holds. For every $X \ge X_0$ and every admissible window pair q, d at $T = X^{1/J}$, there is Y with $X^{1/J} \le Y$ and $Y^2 \le X$ such that, for every $x \in [Y, Y^2]$,*

$$\mathcal{V}\Big(E_{\psi,h(x),q,d}, \lfloor x^{401} \rfloor; \exp(\tfrac{1}{2}x^{\varepsilon})\Big) \quad \textit{holds.} \tag{7.5}$$

This is the arithmetic window theorem supplied by the formal development, specialized to its constructed Galois-representation input [2]. More precisely, `vwin_of_row` proves the assertion from that input, and `galoisLocalGlobalStrong_recParameter` supplies it unconditionally. The formulation above spells out `VWinAt` ψ `401`; it is a cited arithmetic theorem in this article. Its construction proof is not replaced by the observations about finite sums or dyadic rounding just given.

There are several points of substance in this statement. The constants J, X_0 are chosen before the large parameter X and before the pair q, d. After q, d are fixed, one interval must control every twist $1 \le e < \lfloor x^{401} \rfloor$, every prime ℓ, and every scale x in that interval. The field in (7.5) is the field at the corresponding discrete level $h(x)$; the proposition does not keep one number field fixed as x grows through the window. The window is allowed to begin at a later Y, but it cannot be chosen separately for each cohomology group. These distinctions are essential in the later assembly.

The arithmetic proof first separates the character components of the relevant cohomology and then treats their different local regimes. The source includes a uniform treatment of the components controlled by Stickelberger and p-adic L-function estimates, and a separate construction for the components requiring automorphic congruence bounds. The latter is the part for which a qualitative finiteness theorem is far too weak. The estimates have to be summed over coefficient primes and over a growing range of twists.

For the generic odd-prime regime, the source bounds sums of normalized cohomological lengths over character orbits by polynomials in the conductor, twist, and auxiliary discriminant. The orbit sum is organized by kernels of characters. This allows the elementary identity for the sum of the sizes of cyclic quotient-character orbits to control the total contribution instead of estimating each character independently. At resonant primes, where a congruence such as $\ell - 1 \mid 2e$ changes the residual representation, additional cases are required. The trivial-character case and the boundary $e = 1$ are also handled separately. The source's case-partition theorem ensures that these regimes cover the required data before their estimates are combined.

The automorphic congruence input uses ordinary deformation and automorphy arguments for unitary groups together with the two-dimensional Euler-system bounds and symplectic input present in the development. The foundational literature includes the Taylor–Wiles method and

Kato's Euler system [16, 11], but neither reference by itself is the window theorem above. In particular, the numerical bounds, the treatment of residual degeneracies, and their uniformity as the host field varies are part of the argument being explained. It would be incorrect to infer (7.5) merely from the existence of automorphic Galois representations.

The tower estimate uses the vanishing of Iwasawa's μ-invariant for abelian fields [9], local and global duality, and a bound for the relevant second cohomology. The formal interface `thmBab_of_lemmaK` produces constants that work for all twist-growth exponents $C_t \geq 1$, after its cohomology input has been supplied. This estimate is combined with the odd-prime character estimates by `jointWin_of_counts`, and then with the other character components to give the window assertion. The simultaneous interval comes from this joint statement. Finding separate intervals for several estimates and intersecting them would not suffice: intervals with arbitrarily large endpoints need not have a common intersection.

The elementary rescaling that turns the cohomological window into the window used by the prime-sum argument is worth proving. Let $c_\psi \geq 1$ be the fixed factor in the frame order, and choose $\kappa = c_\psi^{1/(2+8\varepsilon)}$. If $x = \kappa x'$, then $c_\psi (x')^{2+8\varepsilon} = x^{2+8\varepsilon}$. Thus a level satisfying (7.4) at x' has the correct size for the frame at x.

Lemma 7.2. *Let $\kappa \geq 1$ and $W \geq \max(\kappa^3, \kappa^{400}, 1)$. Put $Z = \kappa W$. For every $x \in [Z, Z^{3/2}]$, the number $x' = x/\kappa$ lies in $[W, W^2]$, and $\lfloor x^{400} \rfloor \leq \lfloor (x')^{401} \rfloor$. If $0 < \varepsilon$, then $\exp(\frac{1}{2}(x')^\varepsilon) \leq \exp(\frac{1}{2}x^\varepsilon)$.*

Proof. The lower bound $x' \geq W$ is immediate. Since $W \geq \kappa^3$, $Z^{3/2} = \kappa^{3/2} W^{3/2} \leq W^2$; hence $x' \leq x \leq W^2$. Also $x^{400} = \kappa^{400}(x')^{400} \leq (x')^{401}$ because $x' \geq W \geq \kappa^{400}$. Monotonicity of the floor function gives the asserted integer inequality. The last inequality follows from $x' \leq x$ and monotonicity of the positive power and exponential functions. □

The source's rank estimates require twists below $\lfloor x^{400} \rfloor$ at the geometric scale. Lemma 7.2 explains the extra exponent in 401: it absorbs the fixed character-dependent rescaling. It also explains the use of $[Z, Z^{3/2}]$ inside a quadratic window. Neither number is a purported natural invariant of the zeta function. They are sufficient numerical choices for this construction.

Finally, $\mathcal{V}(E, t; B)$ is monotone when t is decreased or B is increased. The finiteness conditions become weaker, and every nonnegative partial sum in (7.2) can only decrease. Combining this monotonicity with the preceding lemma transfers Proposition 7.1 to the geometric scale, with bound $\sqrt{b}$ for $b = \exp(x^\varepsilon)$. This is the precise form in which the arithmetic estimate enters the local lattice-index inequality of Section 6.

8. Uniform assembly and the conclusion

The estimates proved or stated above do not yet imply the prime-sum bound merely by being individually true for sufficiently large parameters. They must hold for the same family of fields, lattices, maps, and weights, with a compatible choice of constants. This last issue is logical as well as quantitative. An estimate valid after a threshold depending on a later choice cannot be combined with an existence assertion that was already used to make that choice.

The formal construction resolves this by making its threshold dependence explicit. A family is supplied with several tiers of constants; later tiers may depend on data already selected at earlier tiers. A threshold is a collection of lower-bound functions on these data. Two thresholds have a common upper bound obtained by taking the pointwise maximum of their requirements. Increasing the threshold preserves the relevant conclusions. The following elementary observation captures the finite assembly used by the construction.

Proposition 8.1 (`topDesWin`, logical form)**.** *For each $\varepsilon > 0$, let $\mathcal{F}_\varepsilon$ be a class of admissible families, and suppose its threshold requirements have finite joins. Write $F \succcurlyeq K$ when F satisfies a threshold K. Assume that $F \succcurlyeq K \vee K'$ implies $F \succcurlyeq K$ and $F \succcurlyeq K'$. Let $P_1, \dots, P_r$ be properties such that, for every i, there is $\varepsilon_i > 0$ with the following property: for each $0 < \varepsilon < \varepsilon_i$, some threshold $K_i(\varepsilon)$ forces P_i for every admissible family satisfying it. Suppose also that for every sufficiently small $\varepsilon > 0$ and every threshold K, there is an $F \in \mathcal{F}_\varepsilon$ with $F \succcurlyeq K$. Then, for every sufficiently small $\varepsilon > 0$, there is a single family satisfying all P_i.*

Proof. Choose ε below every ε_i and below the range allowed by the existence assertion. Form $K = \bigvee_{i=1}^{r} K_i(\varepsilon)$, and apply the existence assertion to this K. The resulting family satisfies each $K_i(\varepsilon)$, hence each P_i. All constants implicit in these conclusions are therefore chosen before the scale is allowed to tend to infinity within this family. □

The proposition remains valid when thresholds are functions of previously chosen geometric data: the join is then pointwise, and the hypotheses must include the indicated monotonicity. The geometric existence theorem in the source is formulated for precisely these threshold objects. It is the universal quantifier over the threshold in that existence theorem that makes the argument work. An existence theorem producing a family before the threshold is specified would not prove the proposition's conclusion.

The actual assembly combines nine compatible requirements. They control local lattice indices, local type-change losses, the exceptional row rank, the equality of the small-level rank with the row rank, the evaluation-height inequality, full column lattices with positive ranks, the finite-place gains, the frame estimates, and the global comparison with the prime-sum data. The cohomological window supplies the bound needed in the local index step. The frame theorem supplies the mean-slope comparison. The period and section-count estimates then give the cancellation of Proposition 5.1. This is the concrete content of the otherwise abstract finite assembly.

At the final specialization the construction fixes the error exponents $e_D = 2$ and $e_L = 0$, together with the window ratio $\lambda = 3/2$. In the general comparison the exponent of the resulting remainder bound is $1/2 + (e_D + e_L + 7)\varepsilon$. These exponents are fixed before ε, and the family and its starting scale may subsequently depend on ε. This is exactly the order required by Corollary 2.5.

The geometric construction supplies a pointwise error bound of the form $E \le x^{C\varepsilon}\mathcal{B}$, uniformly in the translate of the thick set. Here n is the order of the finite group in Section 3, $r = |T|$, and $n_2 = |T_2|$, and m is the chain length. The parameter a_E is the common amplitude parameter in the geometric estimates; with $a_t \le 2a_E$ in the source size conditions. It equals xb in the displayed scale model and is distinct from the sum of all top-place amplitudes. The following elementary calculation proves the numerical normalization once these construction bounds have been supplied.

Lemma 8.2. *Let $x \ge 1$, $\varepsilon > 0$, $C \ge 0$, and $n, r, m, b, k > 0$. Suppose $0 < n_2 \le n$, $0 \le a_E \le 2xb$, and*

$$r \ge \tfrac{1}{2} n x^{-\varepsilon/8}, \qquad x^{1/2-\varepsilon} \le m \le x^{1/2+\varepsilon/2}, \qquad xm \le k \le n^2 x^{1/2-\varepsilon}.$$

Put $\delta_2 = n_2/n$, assume $\delta_2 \ge x^{-\varepsilon}$, and define

$$\mathcal{B} = r a_E \left(1 + \frac{m^2}{k}\right) + m\left(2nmbx^{\varepsilon} + \frac{n_2 k b x^{\varepsilon}}{n(nx^{-\varepsilon})}\right).$$

If $0 \le E \le x^{C\varepsilon}\mathcal{B}$, then

$$\frac{2E}{rm\delta_2 b} \le 20x^{1/2+(C+21/8)\varepsilon}. \tag{8.1}$$

In particular, if $26 \le x^{3\varepsilon/8}$, the right-hand side is at most $\frac{20}{26}x^{1/2+(C+3)\varepsilon}$.

Proof. Every denominator below is positive. Dividing the bracket by rmb, using $k \ge xm$ in its first term, gives

$$\begin{aligned}\frac{\mathcal{B}}{rmb} &\le \frac{2x}{m} + 2 + 2\frac{n}{r}mx^{\varepsilon} + \frac{n_2 k x^{2\varepsilon}}{rn^2} \\ &\le 2x^{1/2+\varepsilon} + 2 + 4x^{1/2+13\varepsilon/8} + 2x^{1/2+9\varepsilon/8} \le 10x^{1/2+13\varepsilon/8}.\end{aligned}$$

For the third term we used $n/r \le 2x^{\varepsilon/8}$ and the upper bound for m; for the fourth we used $n_2 \le n$ and the upper bound for k. The last inequality follows from $x \ge 1$ and $\varepsilon > 0$. Multiplication by $2x^{C\varepsilon}/\delta_2$, followed by $\delta_2^{-1} \le x^{\varepsilon}$, proves (8.1). The additional threshold absorbs the constant because $C + 3 - (C + 21/8) = 3/8$. □

For a fixed retained character family, the other errors in Proposition 3.2 satisfy $MR_x + 4/(\delta_2 x) = O(x^{\varepsilon}\log x) + O(x^{\varepsilon-1})$, with constants fixed before x increases. Since $C \ge 0$, increasing the common starting threshold makes their sum at most $\frac{6}{26}x^{1/2+(C+3)\varepsilon}$. The proposition and the lemma consequently give $\mathcal{G}(x) \le \Sigma(x) + x^{1/2+(C+3)\varepsilon}$. The construction has $C = e_D + e_L + 3$, so the exponent here is $1/2 + (e_D + e_L + 6)\varepsilon$; at $e_D = 2, e_L = 0$ it is $1/2 + 8\varepsilon$.

For completeness, the last absorption can be seen without any geometric notation. The one-scale comparison first gives, with $a = e_D + e_L \ge 0$,

$$H_\psi(x) \le x^{1/2+(a+6)\varepsilon} + 4o\, x^{1/2+\varepsilon}.$$

For $x \ge 1$, the second power is no greater than $x^{1/2+(a+6)\varepsilon}$. Increasing the starting scale until $1 + 4o \le x^{\varepsilon}$ gives $H_\psi(x) \le x^{1/2+(a+7)\varepsilon}$. With $a = 2$, this is the required exponent $1/2 + 9\varepsilon$.

Proposition 8.3 (`zeroFreeAll_of_topG`)**.** *The host-existence, arithmetic-window, and compatible construction results of the formal development give, for every primitive character ψ and every sufficiently small $\varepsilon > 0$, arbitrarily large Y such that*

$$H_\psi(y) \le y^{1/2+9\varepsilon} \qquad (Y \le y \le Y^{3/2}). \tag{8.2}$$

Consequently every primitive Dirichlet L-function is nonzero in $\operatorname{Re} s > 1/2$, away from the principal pole.

Proof from the cited construction inputs. Fix ψ, a permitted $\varepsilon > 0$, and a selector h satisfying (7.4). Proposition 7.1 first supplies $J \ge 1$ and X_0. Apply the auxiliary-prime selection theorem with interval exponent J: for every sufficiently large T, it supplies one admissible pair q, d satisfying all the required analytic estimates throughout $[T, T^J]$. Let $c_\psi \ge 1$ be the fixed factor in the frame order and put $\kappa = c_\psi^{1/(2+8\varepsilon)}$. Given any desired lower bound for the final window, choose M above that bound, the auxiliary-prime threshold, and $\max(1, \kappa^3, \kappa^{400})$. Choose $X \ge \max(X_0, M^J)$, put $T = X^{1/J}$, and select q, d by the analytic theorem at T. For this same pair the arithmetic window theorem supplies W with $T \le W$, $W^2 \le X$, and the cohomological estimate at every point of $[W, W^2]$.

Put $Z = \kappa W$. For $Z \le x \le Z^{3/2}$, set $x' = x/\kappa$. Lemma 7.2 gives $T \le W \le x' \le x \le W^2 \le X$; hence the analytic estimates hold at the geometric scale x, while the cohomological estimate holds at x'. Moreover $x^{1/J} \le T \le x$, and the level $h(x')$ satisfies $x^{2+8\varepsilon} \le c_\psi 2^{h(x')} \le 2x^{2+8\varepsilon}$. Monotonicity of $\mathcal{V}$, together with $\lfloor x^{400} \rfloor \le \lfloor (x')^{401} \rfloor$ and $\exp((x')^{\varepsilon}/2) \le \exp(x^{\varepsilon}/2)$, therefore supplies the required cohomological bound for the host at that level. Since M was arbitrary, these common analytic and arithmetic windows begin arbitrarily far out. The host-existence theorem incorporates the retained character family and the prescribed top-place distribution.

The compatible-family construction described in Section 6 and the finite threshold assembly above supply all nine inequalities for one family on these windows. Apply the slope transfer, Proposition 5.1, and Proposition 3.2, then the preceding exponent absorption. This proves (8.2). The analytic conclusion is Corollary 2.5.

The exact source interfaces for the steps invoked here are the character-host existence theorem, the compatible-family result assembled by `shapeWin`, and its application in `zeroFreeAll_of_topG`; the concordance in Appendix A gives their locations. The paragraph is a proof using those construction results, not an independent proof of their geometric and arithmetic hypotheses. □

Deduction of Theorem 1.1. The preceding proposition excludes zeros of primitive Dirichlet L-functions to the right of the critical line. Removing the finitely many Euler factors transfers this assertion to every Dirichlet character. Proposition 2.6 then applies the functional equation to rule out zeros to the left of the line in the open strip. For ζ, the same proposition checks the remaining half-plane and identifies the negative even integers as the trivial zeros. This proves the stated conclusion. □

The mathematical contribution of the arrangement can now be stated precisely. The analytic end accepts a weaker input than a global error estimate; the arithmetic end is designed to

deliver exactly that weaker input. Positive character weights retain the relevant singularities. The nonconstant metrics prevent the finite Fourier transform from collapsing to two characters. Collective frame estimates overcome the rank mismatch of the individual maps. The period, level, and ramification main terms cancel before any errors are estimated. Finally, the joint window theorem and the ordered choice of constants ensure that all these advantages occur simultaneously.

The author's announcement identifies chains of theta lifts of approximately square-root length as the arithmetic analogue of multiplicity in a Stepanov argument [1]. This is a useful account of the idea behind the construction. The actual proof requires the integral, norm, and congruence estimates described above, and the present exposition does not replace them by the analogy. Nor does it claim that each auxiliary inequality is historically new. The source-backed conclusions and the fully written arguments have been separated so that further simplifications can be judged at the precise point where they would shorten the proof.

Appendix A. Guide to the formal development

All source links in this article point to the fixed revision `7cbf7a0c32fc` of the repository in [2]; their paths are relative to its `src/` directory. The starting points are `Main.lean`, containing the five public theorems, and `Statements.lean`, which imports only Mathlib and defines the four predicates supplementing Mathlib's own `RiemannHypothesis`. In the table below, $N \geq 1$, χ is a complex Dirichlet character modulo N, and 1 denotes the principal character.

Declaration in `Main.lean`	Mathematical statement
`riemannHypothesis`	Mathlib's Riemann hypothesis: a zero of ζ other than $1, -2, -4, \ldots$ has real part $1/2$.
`riemannZeta_eq_zero_re`	If $\zeta(s) = 0$, $s \neq 1$, and $s \neq -2(n+1)$ for every $n \in \mathbb{N}$, then $\operatorname{Re} s = 1/2$.
`LFunction_ne_zero_of_half_lt_re`	If $\operatorname{Re} s > 1/2$ and ($\chi \neq 1$ or $s \neq 1$), then $L(s, \chi) \neq 0$. The parenthesized exclusion is the single case $(\chi, s) = (1, 1)$.
`LFunction_eq_zero_re`	If $0 < \operatorname{Re} s < 1$ and $L(s, \chi) = 0$, then $\operatorname{Re} s = 1/2$.
`LFunction_eq_zero_re_of_isPrimitive`	If χ is primitive, $0 < \operatorname{Re} s < 1$, and $L(s, \chi) = 0$, then $\operatorname{Re} s = 1/2$.

The open-strip hypotheses in the last row are both present in the source. The exclusions at $s = 1$ are also explicit; no analytic inference is made from the value assigned to a totalized function at its pole. The short proofs in `MainTheorems.lean` connect the right-half-plane theorem to the functional-equation and imprimitive-character reductions. The imports in `MainTheoremsAll.lean` also include further modules: an import graph should not be confused with the proof-term dependencies of one declaration.

The names retained in the principal theorem and proposition captions are guideposts into the source. Some prose statements specialize parameters, suppress structure packaging, or isolate an elementary part of a larger source theorem. In particular, the logical form of `topDesWin` explains finite threshold assembly, while the code also carries the geometric interfaces being assembled. The following list locates the principal arguments without reproducing their many helper declarations.

Proposition or argument	Principal source files
`Deduction.poleData`: compensated Mellin transform and zero residues	`poleData.lean`.
`Deduction.landau_window`: positivity on logarithmic windows	`landauWindow.lean`.

Proposition or argument	Principal source files
`ch_balance` and `ch_core`: finite Fourier balance and extraction of the retained character family	`ChBalance.lean` and `ChCore.lean`.
`frame_linear_algebra_first`: preservation of most target directions	`FrameLinearAlgebraFirst.lean`.
`theoremFSharp`: transfer of the mean slope	The named wrapper is in `ShiftInvariance.lean`; its proof is in `TheoremFSharp.lean`.
`key_identity`: cancellation of the arithmetic main terms	`WeightedMeanData.lean`, together with `MeanBoundSecondForm.lean`.
`ch_top`: pointwise arithmetic budget for each thick-set translate	`ChTop.lean`; the individual bound is proved in `LevelOfCongruence.lean`.
Size normalization and the final power of x	`LipschitzWeightAndMinimum.lean`; the application is `CoreGm.lean`.
`vwin_of_row`: simultaneous cohomological windows	`VwinOfRow.lean`. The window predicate and its quantifier order are in `ExistsPageN.lean`.
`topDesWin`: compatible finite assembly	`DesWin.lean`; the geometric specialization is `ShapeWin.lean`.
`zeroFreeAll_of_topG`: construction inputs imply zero-freeness	`ZeroFreeAllOfTopG.lean`.
`Formalism.VW.critLine` and `Formalism.VW.rh_of_every`: reflection to the critical line	`CritLine.lean`; the zeta reduction is developed in `RhOfRight.lean`.

For the arithmetic route, begin with `MainConclusion.lean`. It exhibits the final dependencies compactly. The Galois representation input is supplied by `RowSrc.lean`, and the required local parameterization is supplied by `RecTExists.lean`. The former feeds the window theorem; the latter and the former also feed the theta-construction interface. Thus the intermediate packages are proved inputs to the final source theorem, rather than extra assumptions on the five public statements.

The passage from a cohomological window to a character host is `SupplyChi.lean`: for each primitive ψ, its window hypothesis with twist exponent 401 yields the required host on windows of ratio $3/2$. The larger construction interface is `EveryCharacterOfLiteratureOfVWin.lean`. Its hypotheses organize the arithmetic and representation data; `shapeWin` and `topDesWin` then combine the geometric estimates into the prime-sum comparison. This is the route to follow when checking how a host, its lattices and its evaluation maps are chosen together.

For the theta and lattice realization, the ranks and signatures are recorded in `Params.lean` and `DataSpaces.lean`; the operator chain is in `DataKernels.lean`. The covolume construction is assembled in `TopCovolSlim2Use.lean`. Its numerical terminal estimate is proved in `CovolumeFromGramAndBounds.lean`. The hypotheses exported to the slope budget, including the finite-place gains, frame bounds and height comparison, are collected in `GlobalTheoremInputs.lean`. These are references for the deeper construction results used in the paper, not a claim that those constructions have been independently reproduced in its elementary proofs.

Normalization matters when translating the slope proof. The sharp frame proof uses slope-list entries of the form $2\lambda + c_K$, where c_K depends only on the coefficient field. Here λ is an entry of the underlying slope list before field normalization; the resulting mean is $\mu_2 = 2\mu_{\mathrm{Ar}}$, as in Section 4. Comparing this with another Arakelov convention requires transforming the entire slope list, the map heights and the error bounds consistently. A change in covolume normalization with numerical bounds held fixed is not innocuous. The analytic transforms likewise use the compensated prime sum defined in the text; removing its linear or pole-cancelling terms changes the transform and its holomorphic remainder.

Sutherland's external report [3], as of 20 November 2026, 22:05 UTC, records successful compilation of all 175,187 modules with Lean and Mathlib 4.33.1; the five headline theorems use only `propext`, `Classical.choice` and `Quot.sound`. It reports successful module replay by Lean's kernel and by lean4lean, acceptance of the full export by the independent nanoda and con-leche kernels, and con-leche comparison of all five statements with the independent challenge. The full `lake comparator` run and Lean-kernel replays of the export were still pending. The report also records that the compilation used unsandboxed direct Lean invocations; a 77-module sample agreed with `lake build`. These are reported external checks, not checks rerun for this manuscript. They concern the formal statements and proof objects; they do not certify every prose reformulation in this article.

References

[1] L. Alpöge, *Announcement concerning the Riemann hypothesis*, 20 November 2026, four-page announcement and remarks, https://alpo.ge/remark.pdf.

[2] L. Alpöge (with Claude), *Riemann Hypothesis in Lean*, source code released 20 November 2026, CocoSgt mirror, revision `7cbf7a0c32fc`; Lean 4.33.1 and Mathlib 4.33.1. Fixed revision of the source repository.

[3] A. Sutherland, *Verification files for the Lean proof of the Riemann Hypothesis*, MIT, report updated 20 November 2026, 22:05 UTC, https://math.mit.edu/~drew/grh/, accessed 21 November 2026.

[4] F. A. Behrend, *On sets of integers which contain no three terms in arithmetical progression*, Proc. Natl. Acad. Sci. USA **32** (1946), no. 12, 331–332. doi:10.1073/pnas.32.12.331.

[5] B. Bhatt and P. Scholze, *Prisms and prismatic cohomology*, Ann. of Math. (2) **196** (2022), no. 3, 1135–1275. doi:10.4007/annals.2022.196.3.5.

[6] J.-B. Bost, *Algebraic leaves of algebraic foliations over number fields*, Publ. Math. Inst. Hautes Études Sci. **93** (2001), 161–221. doi:10.1007/s10240-001-8191-3.

[7] G. Chenevier and M. Harris, *Construction of automorphic Galois representations, II*, Cambridge J. Math. **1** (2013), no. 1, 53–73. doi:10.4310/CJM.2013.v1.n1.a2.

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

[9] B. Ferrero and L. C. Washington, *The Iwasawa invariant μ_p vanishes for abelian number fields*, Ann. of Math. (2) **109** (1979), no. 2, 377–395. doi:10.2307/1971116.

[10] W. T. Gan, Y. Qiu and S. Takeda, *The regularized Siegel–Weil formula (the second term identity) and the Rallis inner product formula*, Invent. Math. **198** (2014), no. 3, 739–831. doi:10.1007/s00222-014-0509-0.

[11] K. Kato, *p-adic Hodge theory and values of zeta functions of modular forms*, Astérisque **295** (2004), 117–290. https://numdam.org/item/AST_2004__295__117_0/.

[12] S. S. Kudla and S. Rallis, *A regularized Siegel–Weil formula: the first term identity*, Ann. of Math. (2) **140** (1994), no. 1, 1–80. doi:10.2307/2118540.

[13] E. Landau, *Über einen Satz von Tschebyschef*, Math. Ann. **61** (1905), 527–550. https://geodesic.mathdoc.fr/item/MAN_1905__61_158244/.

[14] A. Moriwaki, *Continuity of volumes on arithmetic varieties*, J. Algebraic Geom. **18** (2009), no. 3, 407–457. doi:10.1090/S1056-3911-08-00500-6.

[15] National Institute of Standards and Technology, *Digital Library of Mathematical Functions*, §25.15, Dirichlet L-functions, https://dlmf.nist.gov/25.15, accessed 21 November 2026.

[16] R. Taylor and A. Wiles, *Ring-theoretic properties of certain Hecke algebras*, Ann. of Math. (2) **141** (1995), no. 3, 553–572. doi:10.2307/2118560.