Three simplifications in the Alpöge–Claude argument for Dirichlet GRH

twinheterodox
Prepared with Codex

21 November 2026

hexagon:2611.00828v1

Abstract

In the proposed proof of the generalized Riemann hypothesis for Dirichlet L-functions by Alpöge and Claude, theta lifts give an upper bound for a smoothed prime sum on certain long intervals. We simplify three steps linking the lifts to the zeros of L-functions. Our main result compares the character sums at a common scale and allows greater variation in their weights while preserving the exponent in the prime-sum bound. We also remove a logarithmic loss in the comparison of lattice slopes and give a short deduction of the zero-free conclusion from a classical oscillation theorem, including for families with a possible real zero. The construction of the lifts and the estimates producing the intervals enter as inputs.

1 From divisibility to a question about zeros

Dirichlet GRH asserts that the zeros of a Dirichlet L-function in the strip $0 < \operatorname{Re} s < 1$ lie on the line $\operatorname{Re} s = 1/2$. In November 2026 Alpöge announced a proposed proof obtained with Claude [1]. It compares divisibility at finite primes with the archimedean size of maps between lattices of automorphic forms. How does such a comparison rule out a zero of a fixed L-function? The link is a one-sided estimate for a sum over primes. We sharpen one error in the lattice comparison, extract the sum at a single scale, and deduce the zero-free conclusion from a classical oscillation theorem.

Consider first the Riemann zeta function. Write Λ for the von Mangoldt function, γ for Euler's constant and $u_+ = \max(u, 0)$. The smoothed prime count arising from the maps, centered by subtracting its main term, is

$$G_1(x) = \sum_{n\ge1} \Lambda(n)(x/n - 1)_+ - x\log x + (1+\gamma)x, \qquad x \ge 1. \tag{1.1}$$

A zero with real part $\beta > 1/2$ would force $G_1(x) > x^\theta$ at arbitrarily large x for every $1/2 < \theta < \beta$. The arithmetic construction supplies intervals $[Y, Y^{3/2}]$, with Y arbitrarily large, on which the count is bounded above by $O(x^{1/2+\varepsilon})$ for each fixed $\varepsilon > 0$. To use that information it matters *where* a zero would force an excursion. Infinitely many excursions could all fall between the available intervals.

For a nonprincipal primitive character ξ, twist the summand in (1.1) by $\xi(n)$ and replace the centering term by $xL'(1,\xi)/L(1,\xi)$. Call the result $G_\xi(x)$. For a primitive ψ of order o, let $\mathcal{F}_\psi$ consist of the primitive representatives of its powers, including the trivial character of conductor one. The arithmetic argument singles out the real sum

$$\mathcal{G}_\psi(x) = G_1(x) + \frac{1}{3o+1} \sum_{\substack{\xi\in\mathcal{F}_\psi \\ \xi\ne1}} \operatorname{Re} G_\xi(x). \tag{1.2}$$

For a nontrivial quadratic character this is $G_1 + G_\psi/7$. The coefficient comes from a count of cosets in Section 4. Its positivity will matter when different L-functions share a zero.

Two issues intervene before the arithmetic inequality gives a bound for (1.2). The theta lifts produce sums for auxiliary characters as well as those powers, and the terms are evaluated at slightly different points. The first issue is resolved by character orthogonality and bounds for the auxiliary twists. Our main contribution to this step concerns the second issue. A term of differentiation order a and weight w has argument a/w. Giving it comparison weight a/x places its argument at the required point x before taking the character average. The price is an explicitly weighted average of the differentiation orders. Theorem 3.1 evaluates that average and bounds the loss with its sign. Corollary 4.1 recovers the exponent in the released argument.

Section 2 traces a logarithmic loss to the estimate for low-slope subspaces. Integrating their actual dimensions removes that loss. Sections 3 and 4 give the prime-sum estimate. Section 5 uses a theorem of Kaczorowski and Pintz to deduce the following implication.

Theorem 1.1 (Passage from the arithmetic estimates to zeros)**.** *Fix a primitive Dirichlet character ψ. Suppose that, for every sufficiently small $\varepsilon > 0$, there are arbitrarily large Y such that the estimates of Corollary 4.1 hold for the family generated by ψ at every $x \in [Y, Y^{3/2}]$. For each ε the constants are uniform in x, Y and the auxiliary data. The exponent constant C is independent of ε. Then no member of the family has a zero in $1/2 < \operatorname{Re} s < 1$.*

The auxiliary field, orders and weights may vary with x. The character family and the function being bounded stay fixed. The original argument supplies the fields, the height and divisibility estimates for the theta lifts, and the simultaneous bounds for auxiliary characters. Jaswal describes these inputs in [5, §§3–4]. We assume familiarity with Dirichlet L-functions and finite character theory. The particular facts about theta lifts used here are stated where they enter.

2 The loss from low lattice slopes

The theta lifts act on integral lattices of automorphic forms. Divisibility at successive stages accumulates in their composite. Fix one translate of the arithmetic data, giving a family of maps $N_i \to C$ into a common lattice. Divisibility of such a map at a prime permits enlarging its input lattice by the corresponding fractional ideal. This raises the slope of the input, while the archimedean norm of the map incurs a height cost. For a Hermitian lattice the slope is the normalized Arakelov degree per rank. In the doubled normalization of the release, both the gain and the cost are on the same logarithmic scale. Let s be the rank-weighted mean of the input slopes after subtracting the height costs. If P is the resulting uncentered prime sum, M its real main term, and E_1 the errors in the arithmetic lower estimate, the original comparison gives

$$P - M - E_1 \;\le\; s - \mu(C). \tag{2.1}$$

This is the divisibility and height calculation in Alpöge and Claude's argument. See also the worked example in [7, §3]. To bound the prime sum we now need an upper bound on the right side. This is where a geometric estimate enters the error allowed for the prime sum. Bost's slope–height inequalities control an injective map and its filtered pieces [2, Props. 4.5–4.6]. Here many maps may have overlapping images. The additional input from the release is a rank estimate for the part of the common target spanned by their images.

Write a for an adjusted input Harder–Narasimhan slope, sampled with its rank multiplicity over all inputs, and λ_j for the target slopes, repeated by rank. Thus $\mu(C)$ is their mean. In the comparison it suffices that $\mathbb{E}a \ge s$. Suppose $a \le s + U$ and every $\lambda_j \ge s - H$, where $U, H \ge 0$. For $0 \le t \le H$, define the rank fractions

$$r(t) = \Pr(a < s - t), \qquad q(t) = \frac{\#\{j : \lambda_j < s - t\}}{\operatorname{rank} C}.$$

The height bound sends the input subspaces at or above the adjusted threshold into the target subspace at or above the same threshold. Discarding a fraction $r(t)$ of the input can therefore leave only a controlled part of the target below threshold. With projection loss $0 \le d < 1$ and frame errors $\eta_f, e_f \ge 0$, the released rank estimate gives

$$(1-d)q(t) \le \eta_f + 4\sqrt{e_f} + 4r(t). \tag{2.2}$$

For clarity about this use of the geometric input, its proof counts the rank of the input filtration step at $s-t$. The real slope count and the corank over the coefficient field differ by its degree over $\mathbb{Q}$. The same factor occurs in total rank. Their ratio is $r(t)$. The released estimate for the image takes any upper bound on that corank, so the actual fraction may be used at each threshold.

The release bounds $r(t) \le U/(U+t)$ for $t > 0$, a consequence of $\mathbb{E}a \ge s$ and $a \le s+U$. Integrating this envelope costs $U\log(1+H/U)$ when $U > 0$. A single list of slopes governs *all* thresholds. For example, half the slopes at $s-1$ and half at $s+1$ give $r(t) = 1/2$ for $0 \le t < 1$. The area is $1/2$, whereas the integral of the envelope on this interval is $\log 2$.

Proposition 2.1 (Integrated rank loss)**.** *Under the preceding hypotheses,*

$$s - \mu(C) \le \frac{(\eta_f + 4\sqrt{e_f})H + 4D(U,H)}{1-d}, \qquad D(U,H) = \frac{UH}{U+H}, \quad D(0,0) = 0.$$

The estimate $\int_0^H r(t)\,dt \le D(U,H)$ *is sharp among distributions with the specified mean and upper bound.*

Proof. Each target slope below s contributes for the length $s - \lambda_j$. Hence $s - \mu(C) \le \int_0^H q(t)\,dt$. Each input contributes for the smaller of H and its shortfall, so

$$\int_0^H r(t)\,dt = \mathbb{E}\min\{H, (s-a)_+\}.$$

Put $z = s - a$. Then $z \ge -U$ and $\mathbb{E}z \le 0$. For $H + U > 0$, the line through $(-U, 0)$ and (H, H) is an upper bound for $\min\{H, z_+\}$ on $z \ge -U$. Its expected value is $H(\mathbb{E}z + U)/(H+U) \le D(U,H)$. Integration of (2.2) proves the inequality. For $U, H > 0$, take $a = s - H$ with probability $U/(U+H)$ and $a = s + U$ otherwise. The mean is s and equality holds in the input estimate. The cases $U = 0$ or $H = 0$ have zero input loss. □

Where the logarithm is spent. In the released pointwise estimate the pre-existing arithmetic bounds give $U \le \bar{U}$, $H \le B = x^{c_H} b$, $d \le \bar{d} < 1$, with $b > 0$. The term charged for discarded input slopes is bounded there by

$$\frac{4(\bar{U}+b)\log(1+B/b)}{1-\bar{d}}.$$

The new contribution is bounded by $4D(\bar{U}, B)/(1-\bar{d})$, and in particular by $4\bar{U}/(1-\bar{d})$. To compare the two allowances even when $\bar{U}$ is very small, use $D(v,B) \le v\log(1+B/v)$ for $v > 0$. This last expression is increasing in v. For $v \le b$ it is at most $b\log(1+B/b)$, and for $v \ge b$ it is at most $v\log(1+B/b)$. The case $v = 0$ follows by continuity. This substitution uses the geometric estimate (2.2) before ranks are integrated. The arithmetic bounds on $\bar{U}$ are obtained by taking the rank-weighted mean of the input errors in the pointwise slope comparison. The separate frame contribution $(\eta_f + 4\sqrt{e_f})H/(1-d)$ still uses the bound for H.

Combining Proposition 2.1 with (2.1) bounds $P - M$ for each translate of the arithmetic data. Their mean has an upper bound, denoted by E below, which also includes rounding, heights and the other arithmetic errors. The published upper bound for E consequently suffices after this improvement. We now explain how to extract a prime sum for a fixed character family from that single real inequality.

3 Comparing prime sums at one scale

The arithmetic data single out positions called *tops*, each carrying a differentiation order a_t. At a given stage a top is paired with other positions carrying weights w. The centered contribution involves $wG_\chi(a_t/w)$. On the *thick* positions w is near a common b and a_t/b is within one of x. On the remaining, *thin* positions a separate bound for G_χ is available. The original argument first changes w to b, averages over translations, then moves each surviving sum from a_t/b to x [5, §§3.5–3.7]. The last move uses bounds for the variation of the individual sums. See the released thick-weight and change-of-argument estimates.

For a prime power n the weighted term has the simple form

$$w((a_t/w)/n-1)_+ = (a_t/n-w)_+.$$

Choose $v_t = a_t/x$. If $w \le v_t$, the prime part can only decrease when we replace w by v_t, provided its coefficient is nonnegative. If $w > v_t$, only $n \le x$ can cost anything. An individual complex character has no such sign. The sum over all characters at a given Frobenius class does, even at a ramified prime. We take that sum first.

The price for the common argument x can be seen with just two orders. At $x = b = 10$, let the two orders be 99 and 101. Keeping weight b would leave the two arguments 9.9 and 10.1. Weights 9.9 and 10.1 give both terms argument 10, but their normalized shares are now 99/200 and 101/200. We shall keep these actual shares in the character average.

The auxiliary number field, called the host, is abelian over $\mathbb{Q}$. Its Galois group has two quadratic sign coordinates and a finite abelian factor A. The construction averages translates by a subgroup $A_0 \le A$. Write $o = [A : A_0]$. For the fields associated to ψ, this index is its order. In additive notation put

$$\Gamma = (\mathbb{Z}/2\mathbb{Z})^2 \times A, \qquad \Phi = \{g \in \Gamma : g_1 = 0\}, \qquad F = |\Phi| = 2|A|, \qquad \iota = ((1,0),0).$$

We identify A with its copy in Γ. For a prime p let $I_p \le \Gamma$ be its inertia group and let f_p be a lift of Frobenius modulo I_p. For $n = p^i$ with $i \ge 1$, put $C_n = if_p + I_p$. A character χ of Γ takes the value

$$A_\chi(n) = \frac{1}{|I_p|}\sum_{k \in C_n} \chi(k)$$

at this prime power. An unramified character contributes its Frobenius value to the i-th power. A character nontrivial on inertia has mean zero on the coset, as required by its primitive Euler factor. With γ Euler's constant, define the centered sums by

$$G_\chi(y) = \sum_{n=p^i} (\log p) A_\chi(n)(y/n-1)_+ + \begin{cases} -y\log y + (1+\gamma)y, & \chi = 1, \\ y\ell_\chi, & \chi \ne 1. \end{cases} \tag{3.1}$$

Here $y > 0$ and each prime sum is finite. The scalar ℓ_χ is the centering constant of the twist. The comparison of weights below works for arbitrary fixed ℓ_χ. In the Dirichlet application it is $L'(1,\chi^*)/L(1,\chi^*)$ for the corresponding primitive nonprincipal character χ^*.

Let $J_g(a,w) = F^{-1}\sum_{\chi \in \widehat{\Gamma}} \overline{\chi(g)}\, wG_\chi(a/w)$. The coefficient of $\log p$ in the character average is determined by finite orthogonality. It is

$$K_n(g) = \frac{1}{F}\sum_{\chi \in \widehat{\Gamma}} \overline{\chi(g)} A_\chi(n) = \frac{2}{|I_p|}\mathbf{1}_{C_n}(g), \qquad \sum_{g \in \Gamma} K_n(g) = 2. \tag{3.2}$$

The sum over characters is $|\Gamma|$ if $k = g$ and zero otherwise. Dividing by F gives the factor two. This is the ramified orthogonality identity in the original proof. At an unramified prime the mass two is concentrated at one point. At the distinguished ramified prime d, inertia has two elements in

different Φ-cosets, so precisely one unit of mass is outside Φ. At the designated split primes $C_n \subseteq \Phi$. Thus the same kernel accounts for the local multiplicities occurring in the arithmetic input.

Fix $x \geq 3,\ b > 0$, a nonempty set of tops $T \subseteq A$ with $N_T = |T|$, a nonempty set $S \subseteq A_0$, and $m \geq 1$ stages. A top t indexes a differentiation order a_t with $|a_t/b - x| \leq 1$. For a translation $u \in A_0$ and stage j, denote by $w_{utj}(g) > 0$ the weight paired with $g \notin \Phi$. Exactly the labels

$$g \in \iota + u + S - t$$

are thick. Their weights lie in $[b, b+h]$, where $h \geq 0$. At each thin label take $y = a_t/w_{utj}(g)$ and assume $|G_\chi(y)| \leq y^\vartheta$ for every χ, with $\vartheta \geq 0$. For a translate u, the uncentered prime sum P_u and real main term M_u of Section 2 satisfy

$$P_u - M_u = \sum_{t\in T} \sum_{j=1}^{m} \sum_{g\notin\Phi} \operatorname{Re} J_g(a_t, w_{utj}(g)), \qquad \mathbb{E}_{u\in A_0}(P_u - M_u) \leq E. \tag{3.3}$$

The identity is the original centering identity `ch_balance` at the actual weights. Its cutoff must cover every actual argument $a_t/w_{utj}(g)$. The inequality is the arithmetic input, obtained by taking the uniform mean of the slope bounds for these translates. We change the weights only after obtaining this upper bound. Equation (3.1) defines the comparison sum at x independently of the cutoff.

Put $v_t = a_t/x$, $\bar{v} = N_T^{-1}\sum_t v_t$, $\delta = |S|/|A|$, and let p_* be the largest fraction of tops in any one coset of A_0. The distribution ν on $Q = A/A_0$ and its Fourier coefficients are

$$\nu(q) = \frac{\sum_{t\in T,\ t+A_0=q} v_t}{N_T \bar{v}}, \qquad \widehat{\nu}(\omega) = \sum_{q\in Q} \nu(q)\omega(q).$$

Write $A_0^\perp \subset \widehat{\Gamma}$ for the characters trivial on A_0. For each such character, $\omega_\chi = \chi|_A$ descends to Q. Finally $\Psi(x) = \sum_{p^i \leq x} \log p$.

Theorem 3.1. *Under the preceding hypotheses,*

$$\operatorname{Re} \sum_{\chi\in A_0^\perp} \chi(\iota)\widehat{\nu}(\omega_\chi) G_\chi(x) \leq \frac{2E}{N_T m \bar{v}\delta} + \mathcal{E}, \tag{3.4}$$

$$\mathcal{E} = \frac{4o\, p_*(h + b/x)\Psi(x)}{\bar{v}} + x\log(1 + 1/x) + \frac{4(F - |S|)x\vartheta}{\delta}. \tag{3.5}$$

Proof. Subtract the two versions of (3.1) at fixed a. The affine terms cancel. Indeed the term $w(a/w)\ell_\chi = a\ell_\chi$ is independent of w. This identity was observed for the quadratic term in [3]. Only the principal logarithm contributes to the change in centering, giving

$$J_g(a, w) - J_g(a, v) = \sum_{n=p^i} (\log p) K_n(g)\big[(a/n - w)_+ - (a/n - v)_+\big] + \frac{a}{F}\log(w/v). \tag{3.6}$$

In particular this difference is real. Take $a = a_t,\ v = v_t$ at a thick label. For $w \geq v_t$ the logarithm is nonnegative and the sum is bounded below by $-(h + b/x)\sum_{p^i\leq x}(\log p)K_{p^i}(g)$, since $v_t \geq b(1 - 1/x)$. For $w < v_t$ the sum is nonnegative, and $v_t/w \leq 1 + 1/x$ bounds the loss from the logarithm by $(a_t/F)\log(1 + 1/x)$.

It remains to count how often a given prime mass can occur at a thick label. Put $\alpha = |S|/|A_0| = o\delta$. For fixed t, g, the probability over u that g is thick is α if $g \in \iota - t + A_0$, and zero otherwise. The cosets $\iota - t + A_0$ for distinct $t + A_0$ are disjoint. For any fixed n this gives

$$\frac{1}{N_T}\sum_{t\in T} \mathbb{E}_u \sum_{g \text{ thick}} K_n(g) \leq \alpha p_* \sum_{g\in\Gamma} K_n(g) = 2\alpha p_*.$$

There are $|S|$ thick labels for every u, t, j. After dividing by $N_T m$, the total loss on replacing thick weights by v_t is at most

$$2\alpha p_*(h + b/x)\Psi(x) + \frac{|S|}{F} x\bar{v} \log(1 + 1/x).$$

At a thin label the triangle inequality gives $|J_g(a_t, w)| \le (|\Gamma|/F)a_t\vartheta = 2a_t\vartheta$. There are $F - |S|$ such labels, giving an additional loss of at most $2(F - |S|)x\bar{v}\vartheta$ in the same normalization.

With weight v_t the argument of every thick G_χ is x. If χ is nontrivial on A_0, its uniform translate average is zero. Otherwise, for every $s \in S$, $\overline{\chi(\iota + u + s - t)} = \chi(\iota)\omega_\chi(t + A_0)$, since ι has order two. Summing over s, t, j yields the left side of (3.4) multiplied by $N_T m\bar{v}|S|/F$. Divide (3.3) and the losses by this factor, using $|S|/F = \delta/2$. □

4 Recovering the fixed Dirichlet family

The translation average has removed all characters nontrivial on A_0. Among the survivors, let $\mathcal{K} \subset A_0^\perp$ be those that represent members of the fixed family. We call these the retained characters. In the Dirichlet construction $1 \in \mathcal{K}$ and $\chi(\iota) = 1$ on $\mathcal{K}$. Its target distribution on Q is

$$\lambda(q) = \frac{3 + \mathbf{1}_{\{0\}}(q)}{3o + 1}, \qquad c = \frac{1}{3o + 1}.$$

Its Fourier transform is 1 at the trivial character and c everywhere else. Indeed the constant part sums to zero at a nontrivial character, and the extra mass at the identity contributes c. The quantity we want to bound is

$$G_{\mathcal{K}}(x) = \sum_{\chi \in \mathcal{K}} \widehat{\lambda}(\omega_\chi) \operatorname{Re} G_\chi(x).$$

Let $p(q)$ be the unweighted fraction of tops in the coset q, and suppose $|p(q) - \lambda(q)| \le \rho$ for all q. With $L = \sum_q |\nu(q) - \lambda(q)|$, one has

$$L \le o\rho + \frac{1}{x - 1}, \qquad |\widehat{\nu}(\omega) - \widehat{\lambda}(\omega)| \le L, \qquad p_* \le \frac{4}{3o + 1} + \rho. \tag{4.1}$$

For the first inequality put $d_t = a_t/b - x$, and let $\bar{d}$ be its mean. Giving the tops weights v_t changes the distribution by at most $\mathbb{E}_t|d_t - \bar{d}|/(x + \bar{d})$ in ℓ^1, even after grouping by coset. Since $|d_t| \le 1$, Cauchy–Schwarz bounds the numerator by $(\mathbb{E}_t d_t^2 - \bar{d}^2)^{1/2} \le 1$.

Writing the discarded contribution with its actual coefficients gives

$$\Sigma_\nu(x) = \sum_{\chi \in A_0^\perp \setminus \mathcal{K}} |\widehat{\nu}(\omega_\chi)|\, |G_\chi(x)|.$$

Theorem 3.1 now gives

$$G_{\mathcal{K}}(x) \le \frac{2E}{N_T m\bar{v}\delta} + \mathcal{E} + \Sigma_\nu(x) + L \sum_{\substack{\chi \in \mathcal{K} \\ \omega_\chi \ne 1}} |G_\chi(x)|. \tag{4.2}$$

For $\omega_\chi = 1$ the two coefficients agree exactly. For the others use (4.1) on the retained set and take absolute values on its complement. In particular $\Sigma_\nu \le \sum_{\chi \in A_0^\perp \setminus \mathcal{K}} |G_\chi(x)|$, the error sum already bounded in the release. This coefficient of one uses normalization by $\bar{v}$.

How many discarded characters have substantial weight? On the cyclic level generated by a primitive Dirichlet character ψ of order o, the retained family consists of the primitive characters inducing its o powers. They are distinct, since inducing two equal primitive characters back to the common modulus gives the same power. The host identifies these with $\mathcal{K}$. Among retained characters

only the principal one is trivial on A. Indeed, for a character satisfying $\chi(\iota) = 1$, triviality on A leaves the principal character or the other quadratic character of $(\mathbb{Z}/2\mathbb{Z})^2$ with that sign. The latter is excluded on a cyclic level by the host's exceptional-character condition. These are the cyclic-level definition and the identification of the retained family of the release.

It follows that $G_{\mathcal{K}} = \mathcal{G}_\psi$, the fixed sum in (1.2). In fact its principal coefficient is one and each other primitive power has coefficient c. Each character of Q has four extensions to Γ that are trivial on A_0. Thus exactly three discarded characters have trivial restriction to A, and exactly $3(o-1)$ have nontrivial restriction. If $|G_\chi(x)| \le M_x$ for every discarded character, it follows that

$$\Sigma_\nu(x) \le \bigl[3 + 3(o-1)\min\{1, c+L\}\bigr] M_x. \tag{4.3}$$

For fixed o and $\rho \to 0,\ x \to \infty$, the upper coefficient tends to $12o/(3o+1) < 4$. For the ideal weights `bramble_trace` already computed this factor, then compared the equally weighted terms at a_t/b with those at x [3]. Inequality (4.3) controls the actual unequal weights a_t/x at x. At order three the ideal coset shares are $4/10, 3/10, 3/10$. Three discarded characters have weight one and six have ideal weight $1/10$, giving 3.6 in place of the release's bound $4o = 12$. At $o = 1$ the formula gives the three auxiliary characters in the zeta case. For other retained sets use (4.2).

The bound on the arithmetic intervals

We now evaluate the error terms in (4.2) at the scales supplied by the theta lifts. The threshold for x may depend on the fixed family of Dirichlet characters.

Corollary 4.1. *Fix a cyclic level and $\varepsilon > 0$. Suppose that (3.3) holds with $E \ge 0$, and that for fixed $C \ge 0$ and $D_h > 0$ the data satisfy*

$$E/(N_T m b) \le 10x^{1/2+(C+13/8)\varepsilon}, \quad \delta \ge x^{-\varepsilon}, \quad \rho \le x^{-1+\varepsilon}, \tag{4.4}$$

$$2F(x+1)\vartheta \le 1/x, \qquad h/b \le D_h x^{-1/2+\varepsilon}, \qquad M_x \le x^{1/2+\varepsilon}. \tag{4.5}$$

For all sufficiently large x, uniformly in the host at the fixed level,

$$G_{\mathcal{K}}(x) \le \Sigma_\nu(x) + x^{1/2+(C+3)\varepsilon}.$$

Consequently $\mathcal{G}_\psi(x) \le x^{1/2+(C+4)\varepsilon}$. In the Dirichlet application the construction uses two fixed error exponents e_D, e_L and gives $C = e_D + e_L + 3$. This recovers the released bound with exponent $1/2 + (e_D + e_L + 7)\varepsilon$.

Proof. The ordinary Chebyshev bound $\Psi(x) \ll x$ suffices [4, §V, pp. 377–378]. Indeed, for an integer $r \ge 1$, each prime power $p^i \in (r, 2r]$ adds one to $\lfloor 2r/p^i \rfloor - 2\lfloor r/p^i \rfloor$. These differences are nonnegative at all prime powers, so $\Psi(2r) - \Psi(r) \le \log \binom{2r}{r} \le 2r \log 2$. Summing at powers of two and using monotonicity gives the bound for real x. Partial summation gives $\sum_{n \le x} \Lambda(n)/n \ll \log(2x)$, where $\Lambda(p^i) = \log p$ and $\Lambda(n) = 0$ otherwise. Formula (3.1) now bounds each fixed retained nonprincipal sum by $O_\psi(x \log(2x))$, because its ℓ_χ is fixed.

We have $\bar{v} \ge b(1 - 1/x) \ge 2b/3$. The arithmetic term in (4.2) is at most $30x^{1/2+(C+21/8)\varepsilon}$. Since $p_* \le 1$, the thick prime term in (3.5) is $O_o(1 + xh/b) = O_{o,D_h}(x^{1/2+\varepsilon})$ by Chebyshev's bound. The logarithm term is at most one. The thin term is at most $2/[\delta(x+1)]$. By (4.1) the retained coefficient errors contribute $O_\psi(x^\varepsilon \log(2x))$. The sum of these four errors and the arithmetic term is eventually at most $x^{1/2+(C+3)\varepsilon}$, since $21/8 < 3$. Finally $\Sigma_\nu(x) \le 4ox^{1/2+\varepsilon}$. For fixed o an extra factor x^ε absorbs the sum of the two bounds. □

To see where Section 2 enters this corollary, recall that E bounds the mean of divisibility minus size. The improved allowance in Proposition 2.1 fits within the original arithmetic estimate for each translate. The following calculation carries that estimate to the first inequality of (4.4). Write

$n = |A|$, let a_E denote the source's upper differentiation parameter and k its auxiliary size parameter. After division by $x^{C\varepsilon} N_T m b$, the four terms in the size estimate are bounded by

$$\frac{a_E}{mb} + \frac{a_E m}{bk} + \frac{2nmx^{\varepsilon}}{N_T} + \frac{|S|kx^{2\varepsilon}}{N_T n^2}.$$

The source gives $a_E \le 2xb$, $n/N_T \le 2x^{\varepsilon/8}$, $x^{1/2-\varepsilon} \le m \le x^{1/2+\varepsilon/2}$ and $xm \le k \le n^2 x^{1/2-\varepsilon}$. Since $|S| \le n$, the four respective bounds are $2x^{1/2+\varepsilon}$, 2, $4x^{1/2+13\varepsilon/8}$, $2x^{1/2+9\varepsilon/8}$. Their sum is at most $10x^{1/2+13\varepsilon/8}$. For h and ϑ the released weight bounds give $6F(x+1)h \le b/x$ and $2F(x+1)\vartheta \le 1/x$. For thin labels let b' be the source's weight scale and $X'/3$ the threshold from which its character bound applies. The source gives $w \le 2b'$ on thin labels and $b'X' \le xb$, whence $a_t/w \ge (x-1)b/(2b') \ge X'/3$. The upper bound (3.3) follows at each translate from `ch_top`. The share and fixed-scale discarded-character bounds and the exponent C appear in `core_gm`. The cyclic level fixes $\mathcal{G}_\psi$ once and for all. The estimates for auxiliary characters are needed only in the fields that supply the particular interval.

The released thick-weight bound makes the thick prime error $O_o(1)$. For the passage to the prime sum, (4.5) quantifies how much wider the weights can be while preserving the exponent. The arithmetic construction has its own weight requirements.

5 Why the selected intervals suffice

Return to the fixed function $\mathcal{G}_\psi$ obtained in Section 4. In the original argument high derivatives of its transform detect a pole by concentrating an integral on the available interval [5, §3.3]. We use the frequency of sign changes forced by a pole. A real pole at the rightmost real part of the singularities would obstruct direct application of the classical theorem.

What a zero contributes

Extend $\mathcal{G}_\psi$ by zero on $0 < x < 1$, and write $M_g(s) = \int_0^\infty g(x) x^{-s-1}\,dx$ for the Mellin transform. For $\operatorname{Re} s > 1$, integration from $x = n$ to infinity gives

$$\int_1^\infty (x/n - 1)_+ \, x^{-s-1}\,dx = \frac{n^{-s}}{s(s-1)} \qquad (n \ge 1).$$

Let $\kappa_1 = 1$ and $\kappa_\xi = 1/(3o+1)$ for the other members of $\mathcal{F}_\psi$. Conjugate characters have the same weight and conjugate sums for real x. Hence their sum is real, and the Euler products give

$$M_{\mathcal{G}_\psi}(s) = -\frac{1}{s(s-1)} \sum_{\xi \in \mathcal{F}_\psi} \kappa_\xi \frac{L'(s,\xi)}{L(s,\xi)} - \frac{1}{(s-1)^2} + \frac{1 + \gamma + \sum_{\xi \ne 1} \kappa_\xi \ell_\xi}{s-1}. \tag{5.1}$$

Here $L(s,1) = \zeta(s)$, and the last sum is over the nonprincipal family members. Termwise integration follows from absolute convergence for $\operatorname{Re} s > 1$. This formula continues the transform meromorphically to $\operatorname{Re} s > 0$. The Laurent expansion of ζ at one and the definitions of the ℓ_ξ cancel the polar parts at $s = 1$.

If ρ lies in the open critical strip, write $m_\xi(\rho)$ for the multiplicity of the zero of $L(s,\xi)$ there, taking $m_\xi(\rho) = 0$ for a nonzero value. Whenever at least one multiplicity is positive, the residue of (5.1) is

$$\frac{\sum_\xi \kappa_\xi m_\xi(\rho)}{\rho(1-\rho)}. \tag{5.2}$$

The weights are positive, so coincident zeros reinforce one another. For the zeta example there is no real zero in $(0,1)$. In a Dirichlet family the residue at any real zero β in this interval is positive. This sign will allow us to handle the possible real pole. The same pole and residue data are used in the release.

A one-sided oscillation criterion

A theorem of Kaczorowski and Pintz supplies the needed frequency of sign changes. We recall its meromorphic case. Suppose a real function has an initially absolutely convergent Mellin transform which is holomorphic in $\operatorname{Re} s > B$, is meromorphic in a half-plane extending strictly to the left of that line, and has a singularity in every half-plane $\operatorname{Re} s > B - \delta$, $\delta > 0$. If the transform is regular at the real point B, then for every sufficiently small $\alpha > 0$, every interval $[T^{1-\alpha}, T]$ with T large enough contains a change of sign. There need not be a pole on the line $\operatorname{Re} s = B$. These are Theorem 1, equation (2.2), and Corollary 1 of [6, pp. 175–176, 179], in the case where their least ordinate is positive or infinite. Jaswal pointed out the relevance of this theorem to Dirichlet GRH [5, §3.3]. The regularity condition at the real point B is the reason for the next argument.

Proposition 5.1 (An upper bound on power intervals)**.** *Let $0 < \theta < 1$, $\lambda > 1$, and let g be real and locally integrable on $[1, \infty)$, extended by zero below one. Suppose M_g is initially absolutely convergent in a right half-plane and continues meromorphically to $\operatorname{Re} s > 0$. Assume each real pole in (θ, ∞) is simple with positive residue. If for some fixed K and arbitrarily large Y*

$$g(x) \le K x^{\theta} \quad (Y \le x \le Y^{\lambda}),$$

then M_g is holomorphic in $\operatorname{Re} s > \theta$.

Proof. Suppose a pole lies to the right of θ, and let $B > \theta$ be the supremum of the real parts of those poles. There are only finitely many real poles $\beta > \theta$. Indeed, absolute convergence bounds them on the right, and the real segment from θ to that bound is compact inside the domain of meromorphy. Let $r_\beta > 0$ be their residues.

We need a function which is negative throughout each good interval and whose transform is regular at the real point B. Subtracting $(|K| + 1)x^{\theta}$ ensures negativity. To remove the real pole at β we would like to subtract $r_\beta x^\beta$, but this might remove the only evidence of a singularity that far to the right. Instead, choose $u > 0$ such that all the points $\beta \pm iu$ are regular for the original transform. Such a choice exists because the poles form a countable set. On $x \ge 1$ put

$$f(x) = g(x) - (|K| + 1)x^{\theta} - \sum_{\beta} r_\beta x^\beta (1 - \cos(u \log x)),$$

and put $f = 0$ below one. Every term subtracted in the last sum is nonnegative. Its effect on the transform at the pole β is

$$-\frac{r_\beta}{s - \beta} + \frac{r_\beta/2}{s - \beta - iu} + \frac{r_\beta/2}{s - \beta + iu}.$$

The real pole disappears, and two nonreal poles with exactly the same real part take its place. Distinct β's give distinct pairs. The original nonreal poles survive, so the supremum is still B, even when no pole attains it. The shift at θ introduces no singularity strictly to its right. The transform is meromorphic on every closed half-plane $\operatorname{Re} s \ge B - c_0$ with $0 < c_0 < B - \theta$. It is regular at the real point B, so on the line $\operatorname{Re} s = B$ the least absolute ordinate of a pole is positive, or there are no poles on the line.

Choose a sufficiently small $\alpha > 0$ with $\lambda(1 - \alpha) > 1$. For $T = Y^\lambda$, the interval $[T^{1-\alpha}, T]$ lies inside $[Y, Y^\lambda]$. The theorem forces a sign change there for all sufficiently large Y. Yet the defining bound gives $f(x) < 0$ on each good interval. This contradiction proves the proposition. □

Proof of Theorem 1.1. Fix $1/2 < \theta < 1$. Choose $\varepsilon > 0$ small enough that $1/2 + (C + 4)\varepsilon < \theta$. Corollary 4.1 gives $\mathcal{G}_\psi(x) \le x^\theta$ throughout the supplied intervals for arbitrarily large Y. Apply Proposition 5.1 with $\lambda = 3/2$ and $K = 1$, using (5.2) for real poles and the usual nonvanishing of Dirichlet L-functions on $\operatorname{Re} s \ge 1$ away from the pole of ζ. By (5.2) a family zero with real part greater than θ would give a pole. Letting θ decrease to $1/2$ proves the claim. The choice of

simultaneous arithmetic intervals and its use at the fixed family are given in the released window construction and its application to the family. When the arithmetic inputs hold for every primitive ψ, the functional equation gives the other half of the critical strip. □

The integral version of the criterion

The release also proves a more flexible criterion, `lemmaL`, in which the upper bound on each interval is only an integral bound. The same sign-change argument recovers it. Suppose that for every $\eta > 0$ there are arbitrarily large Y and a nonnegative measurable majorant $R \geq g$ on $[Y, Y^\lambda]$ with

$$\int_Y^{Y^\lambda} R(v) v^{-\theta} \frac{dv}{v} \leq Y^\eta.$$

Under the other hypotheses of Proposition 5.1, no pole can then lie to the right of θ. To see why, form the logarithmic average $A(x) = \int_x^{2x} g(v)\, dv/v$. For any fixed $1 < \lambda' < \lambda$, all sufficiently large Y satisfy $[x, 2x] \subset [Y, Y^\lambda]$ whenever $Y \leq x \leq Y^{\lambda'}$. It follows that $A(x) \leq (2x)^\theta Y^\eta \leq 2^\theta x^{\theta+\eta}$ throughout this shorter interval.

By Fubini in the half-plane of absolute convergence, $M_A(s) = ((2^s - 1)/s) M_g(s)$. The factor is nonzero for $\operatorname{Re} s > 0$ and positive on the positive real axis. Thus it preserves poles there and the sign of every positive real residue. The average is zero below $1/2$. Cutting it off below one adds an entire function to its transform. If a pole ρ had $\operatorname{Re}\rho > \theta$, we could take $0 < \eta < \min(1, \operatorname{Re}\rho) - \theta$. The pointwise proposition applied to the truncated average at $\theta + \eta$ would exclude that pole.

Verification. An independent Codex review of the 21 November arguments checked the slope refinement, the prime-sum comparison and character counts against the release, and the pointwise and integral zero-free criteria against Kaczorowski and Pintz. A separate reading of Sections 1–3 of the combined draft examined exposition for graduate readers. The source bundle includes two exact checks of the finite averages in `checks/`. The arguments presented here are proofs on paper.

References

[1] L. Alpöge, *Announcement on the Riemann hypothesis*, 20 November 2026. Lean source.

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

[3] `bramble_trace`, *Reading Alpöge's proposed proof of RH from theta lifts to prime sums*, rStation, 20–21 November 2026. Which prime powers the selected classes count. Comment on quadratic centering, 21 November, 06:30 UTC.

[4] P. L. Tchebichef, *Mémoire sur les nombres premiers*, J. Math. Pures Appl. **17** (1852), 366–390.

[5] M. Jaswal, *The Riemann hypothesis for Dirichlet L-functions [after Levent Alpöge]*, version 1, 21 November 2026.

[6] J. Kaczorowski and J. Pintz, *Oscillatory properties of arithmetical functions. I*, Acta Math. Hungar. **48** (1986), 173–185.

[7] twinheterodox, *From theta lifts to a prime-counting bound*, working note, 20 November 2026.