Auxiliary primes and interpolation points in the Alpöge–Claude argument for Dirichlet GRH

twinheterodox
Prepared with Codex

21 November 2026

hexagon:2611.00839v1

Abstract

The proposed Alpöge–Claude proof of the generalized Riemann hypothesis for Dirichlet L-functions makes two auxiliary choices. Its interpolation points must represent specified cosets in prescribed proportions and have suitable spacings along Frobenius orbits. We prove that the proportions can be exact and that a single set of points meets the spacing requirements on every orbit of sufficiently large order. This removes the loss from exceptional long orbits in the local divisibility estimate. The proof also needs two primes in prescribed arithmetic progressions whose quadratic twists obey a common bound throughout a prescribed interval of scales. At the exponent used for Dirichlet GRH, we give a selection argument that dispenses with estimates for correlations of quadratic symbols among the primes in those progressions.

1 Where the choices enter

In the Alpöge–Claude argument, the same family of theta lifts is viewed in two ways. Its divisibility at finite primes is compared with its archimedean size. The resulting inequality bounds a smoothed prime sum involving Dirichlet characters. Before making the comparison one must choose both the field in which the lifts are constructed and the interpolation points at which their local weights may be large. Poor spacing among the points weakens the divisibility estimate. A poor choice of field introduces quadratic twists whose contribution to the prime sum can be too large. We examine these two choices. The announcement is [1]. Jaswal [5, §3] and Slijepčević [8, §3] give accounts of the argument.

The Galois group of the auxiliary field has the form $\{\pm 1\}^2 \times A$. The two signs record the actions on two quadratic subfields. Elements of the finite abelian factor A label the interpolation positions [5, §3.4]. Each position carries a local weight vector. The local data give a common bound on the absolute values of all entries except the first and last at a selected position (a *top*). The global character calculation prescribes the proportion of points in each coset of a subgroup. At an inert prime, the local calculation sees the same points in the cyclic order given by its Frobenius element. Separated groups of points let test polynomials distinguish individual target coordinates. The groups must recur regularly, and the intervening windows must have bounded occupancy. The original argument allows some orbits to fail these tests and bounds their total contribution with prime-dependent weights. Theorem 2.1 gives exact coset counts and meets the tests on *every* long orbit, even if its origin is chosen after the points. Its parameters are those of the released interpolation-point selection.

The field is chosen using two primes q, d. Once q is fixed, it determines part of the progression for d and some of the characters to be tested at d. The original auxiliary-prime selection controls the character sums by averaging over primes in each progression. Our argument bounds the *number* of exceptions among all candidates up to U. A progression containing $U^{1-o(1)}$ primes can then absorb any exceptional set of size $O(U^{1-\gamma})$ with fixed $\gamma > 0$. The quantitative point is that such a γ can be

obtained for the entire interval of scales and uniformly in the first prime. Proposition 3.1 proves this by an elementary moment estimate. Corollary 4.1 gives the two primes using the original zero-density and truncation estimates.

These results concern the inputs to the theta-lift comparison. In our earlier note [10] the field and interpolation points were given, and the question was how to extract a prime-sum bound for the original characters from that comparison. The two orders of choice are useful in different ways. The points can be selected before the Frobenius directions are known, whereas the estimate for the second prime must still work after the first prime has fixed the characters and progression. We assume finite character theory, basic probability, and the usual facts about Dirichlet characters and L-functions. We spell out the parts of the local construction and the analytic estimates used to pass from the two selections to their applications.

2 One set of interpolation points for all long orbits

Fix $0 < \varepsilon < 1/52$ and let x tend to infinity. Put

$$h = 2 + 8\varepsilon, \qquad \delta = x^{-\varepsilon/8}.$$

Let A be a finite abelian group of order $N \asymp x^h$, and A_0 a subgroup of bounded index g. All its cosets have size $N_0 = N/g$. Give A_0 weight 4 and every other coset weight 3. To see the reason for the unequal counts, take $g = 2$. With proportions $4/7$ and $3/7$, the average of the nontrivial quotient character is $4/7 - 3/7 = 1/7$. Equal proportions would annihilate that character. More generally, the desired probability on A/A_0 is $(3 + \mathbf{1}_{c=0})/(3g + 1)$ at c. Its Fourier coefficient is 1 at the trivial character and $1/(3g + 1)$ at every other character of the quotient. The release allows an error of $x^{-1+\varepsilon}$ in each coset proportion [8, Def. 3.2, (E3)]. This is the unweighted distribution entering the character average. Subsequent derivative weights can vary even when the point counts are exact [10, §§3–4].

2.1 The spacing conditions

For $v \in A$ of order $d \geq x^{\varepsilon/4}$, index a translation orbit by $a + iv$, $0 \leq i < d$, from any origin $a \in A$. A cyclic window of length L consists of L successive sites, with indices taken modulo d. Let L_2 be a bounded positive integer. To see the scales, think of selecting sites with probability comparable to δ. A cell of length about δ^{-1} then has expected occupancy bounded above and below by positive constants. A fixed number of cells makes a tile, and a logarithmic number of tiles gives repeated chances for a desired pattern. The local interpolation construction uses

$$\begin{aligned} &K = \lceil 8(h+2)/\varepsilon \rceil, \qquad r = K + 1, \qquad s = 6r + 1, \qquad u = \lceil 1/\delta \rceil, \qquad \ell = su, \\ &m = 1 + \lceil (h+3)10^s \log x \rceil, \qquad G' = (m+4)\ell, \\ &N_1 = 1 + \lceil \max\{6eG'\delta, (h+3)\log_2 x\} \rceil . \end{aligned} \tag{2.1}$$

Here K caps the occupancy of a short window. The larger window G' will contain a run of tiles, and N_1 caps its occupancy. Both window lengths are $o(d)$, since $G' \ll_\varepsilon x^{\varepsilon/8} \log x$. In particular an orbit eventually has more than m full tiles.

Starting at index 0, cut the orbit into $\lfloor d/\ell \rfloor$ full tiles of ℓ sites and a remainder of fewer than ℓ sites. Divide each full tile into s consecutive cells of u sites, numbered from 1. A tile is *marked* if its odd cells are empty and each even cell contains a single selected site whose index i from the origin is even. Thus it contains $3r$ separated selected sites, with a whole empty cell between successive ones. The parity refers to the index in the orbit before wraparound, including on cycles of odd order. A run of m full tiles is taken cyclically in the list of full tiles, omitting the remainder. These are the original interval and tile conditions.

Theorem 2.1. *For all sufficiently large x there are a subset $T \subseteq A$ and a positive integer a_* such that a coset of weight w contains exactly wa_* points of T, and $\delta N/2 \le |T| \le \delta N$. For every direction v of order at least $x^{\varepsilon/4}$, every origin on each orbit has both properties below.*

1. *every cyclic window of length L_2 contains at most K points, and every cyclic window of length G' contains at most N_1 points;*
2. *every run of m consecutive full tiles contains a marked tile.*

The threshold for x is uniform under fixed bounds on g, L_2 and on the constants in $N \asymp x^h$.

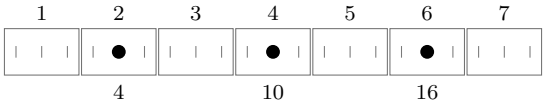

Figure 1: A tile with $r = 1$ and $u = 3$ for illustration. Cell numbers are above, selected site indices below. If $4, 10, 16$ are compulsory and all other sites in the tile are ineligible, choices elsewhere cannot destroy the mark.

2.2 Sampling with enough room to complete the cosets

Insisting on exact coset sizes during independent sampling would introduce dependence between distant cells. Instead we choose nested random sets. A mark already forced by the smaller and larger sets will survive an arbitrary completion between them, as in Figure 1. A different difficulty is that there are too many short windows to rule out all overcrowdings by a direct union bound. We remove a small set of sites after sampling. The proof therefore needs several marks in each run before removal, together with a *local* bound on how many sites are removed.

Proof of Theorem 2.1. Put $\eta = 1/20$, and draw an independent uniform random variable on $[0, 1]$ for each site. In a coset of weight w, set $p = w\delta/4$. A site is eligible (in T^+) if its variable is at most p, and compulsory (in T^-) if it is at most $(1-\eta)p$. Call a mark protected if every odd cell has no eligible site and each even cell has exactly one eligible site, compulsory and of the required parity.

For an odd and an even cell respectively the probabilities of these events are bounded below by

$$(1-\delta)^u, \qquad \lfloor u/2 \rfloor (1-\eta)\tfrac{3\delta}{4}(1-\delta)^{u-1}.$$

The even-cell bound sums disjoint events, one for each admissible site. As $x \to \infty$, these bounds tend respectively to $1/e$ and $57/(160e)$, both greater than $1/8$. The cells are disjoint. Thus a tile is protected with probability at least 8^{-s} for large x, uniformly in the direction, origin and tile.

Set $J = 2(K+1)$. Even when a run wraps around the remainder, its full tiles involve disjoint sites. If the run has at most J protected marks, some $m - J$ of its tiles have none. Its failure probability is at most

$$\binom{m}{J} \exp(-(m-J)8^{-s}) \ll_\varepsilon (\log x)^J x^{-2(h+3)}.$$

Here $(10/8)^s > 2$. There are at most N^3 choices of direction, origin and first tile. Since $3h < 2(h+3)$, with probability tending to one every run has at least $J+1$ protected marks. The long windows also satisfy their cap in T^+ with probability tending to one. A given violation has probability at most $\binom{G'}{N_1}\delta^{N_1} \le 6^{-N_1} \le x^{-(h+3)}$ if $N_1 \le G'$, and zero otherwise. There are at most N^2 direction and starting-site pairs, and $h < 3$.

Repairing the short windows. We use the sample-and-delete argument exemplified by Erdős [2, proof of (4)]. Here preserving a mark in each run requires controlling the deletions in each individual

window. Call a set of $K+1$ sites lying in a short window an obstruction if all of them are eligible. There are $O(N^2)$ candidate sets, since L_2 is bounded. Choose a maximal pairwise disjoint family of obstructions and let E be its union. Every obstruction meets E, and

$$\mathbb{E}|E| \ll N^2\delta^{K+1} \ll x^{h-2} = o(\delta N_0).$$

The middle inequality uses $\delta^{K+1} \le x^{-(h+2)}$.

Put $b = \max(L_2, G')$. A given b-window meets $O(bN)$ candidate sets, including those in other directions. Indeed, choose a meeting site, the direction of the short window, and the position of that site in the short window. Only boundedly many subsets then remain. A family of j disjoint obstructions requires $j(K+1)$ distinct eligible sites. For $j = 2, 3$, summing over the candidates and then over all b-windows gives

$$\mathbb{P}(\text{some window meets } j \text{ disjoint obstructions}) \ll N^2(bN\delta^{K+1})^j \ll x^{2h-2j+j\varepsilon/8}(\log x)^j.$$

For pairs the exponent $2h - 4 + \varepsilon/4$ is positive, so this bound gives no decay. For triples it is $-2 + 131\varepsilon/8 < 0$. Thus with probability tending to one, each b-window meets at most two members of the chosen disjoint family.

A run lies in one such window even when it wraps around, since its m full tiles span at most $(m+1)\ell$ sites including the remainder. Deleting the chosen family costs at most $2(K+1)$ sites there, even when the obstructions arise in other directions. Each deletion can spoil at most one of the run's disjoint protected tiles. Of the $2(K+1)+1$ marks reserved in the run, at least one survives. Its occupied sites are still compulsory and every other site in the tile is ineligible. Thus every intermediate choice

$$T^- \setminus E \;\subseteq\; T \;\subseteq\; T^+ \setminus E \tag{2.2}$$

has a marked tile in every run and satisfies both window bounds.

Filling the cosets. Take $a_* = \lfloor (1-\eta/2)\delta N_0/4 \rfloor$. In a coset of weight w, the compulsory and eligible counts have means $(1-\eta)pN_0$ and pN_0, respectively, and variances at most pN_0. Their means bracket wa_* with a margin $\gg \delta N_0$ on both sides. Since $\delta N_0 \to \infty$ and g is bounded, Chebyshev's inequality gives simultaneous fluctuations smaller than one third of the minimum margin with probability tending to one. By Markov's inequality the same bound holds for the total $|E|$. Consequently, even after deletion, the two counts bracket the target in every coset simultaneously. Choose the remaining sites of T separately in each coset within (2.2). Finally,

$$\frac{|T|}{\delta N} = \frac{(3g+1)a_*}{\delta g N_0} = (1-\eta/2)\frac{3g+1}{4g} + o(1).$$

The last quantity lies strictly between $1/2$ and 1 for all sufficiently large x, uniformly for bounded g. □

2.3 Which loss disappears?

Let $\mathcal{P}_{\text{long}}$ consist of the primes $p < \exp(x^{\varepsilon/3})$ in the release's inert class whose Frobenius direction $v_p \in A$ has order at least $x^{\varepsilon/4}$. For each p in this set, let B_p consist of the selected points whose v_p-orbit fails a spacing test at the origin assigned to that point. The two choices of interpolation points give

$$\sum_{p\in\mathcal{P}_{\text{long}}} \frac{\log p}{p}|B_p| \ll x^{-2}\delta N \sum_{p\in\mathcal{P}_{\text{long}}} \frac{\log p}{p} \qquad B_p = \varnothing \quad (p \in \mathcal{P}_{\text{long}}).$$

Released selection — Theorem 2.1

Thus the entire weighted error from these orbits is zero. The theorem applies because the rank and inert-prime bounds give $L_2 \le 2(620C_0/\varepsilon + 1)$ with C_0 fixed, and it covers every choice of origin.

To see where rank can still be lost, consider the even local extension. A marked tile supplies $3r$ separated sites of even index. Its middle r sites serve as a frame, with r more on either side. The recurring marks and the long-window cap bound the length and occupancy between frames. Under its further prime and weight hypotheses, the local polynomial construction then gives test polynomials concentrated on individual lattice edges. Their pairings recover independent coordinates, up to a controlled power of p, on a submodule of the target lattice. The local divisibility estimate still requires a bound on the weighted rank deficit and additional conditions on the prime. Short-order directions require their own estimate.

3 Counting the exceptions before choosing a progression

We turn to the choice of the field. Suppose for a moment that we need just one auxiliary prime $r \le U$. Among the characters to be tested is a quadratic twist by $(\frac{\cdot}{r})$. The difficulty with a second moment inside a progression is that, on expanding the square, one must estimate sums of products of quadratic symbols over the primes in that progression. Quadratic reciprocity lets us evaluate each symbol at an arbitrary integer instead. Complete periods then cancel exactly in all the nonsquare terms. The price is that we count exceptions up to U, while only a thin subset of those integers can be used as primes. A sufficiently high moment pays that price.

Proposition 3.1. *Fix $D \ge B \ge 1$, $0 < \nu \le \alpha < 2\tau < 1$, $C_0 > 0$ and $C_1 \ge 0$. Put $Z = Y^D$ and suppose*

$$Y^\nu \le \log U \le C_0 Y^\alpha (\log Y)^{C_1}.$$

For any complex coefficients $|\kappa(n)| \le 1$ and an odd prime r, write

$$S_r(y) = \sum_{n \le Z} \Lambda(n)\kappa(n)(\frac{n}{r}) \min(1, y/n),$$

where Λ is the von Mangoldt function. For some $\gamma > 0$ depending only on the fixed exponents, all but $O(U^{1-\gamma})$ odd primes $r \le U$ satisfy

$$|S_r(y)| \le \tfrac{1}{2} y^{1/2+\tau} \qquad \text{for every } Y \le y \le Y^B. \tag{3.1}$$

The exceptional set may depend on κ, but the bound for its size is uniform in κ.

Proof. The contribution of the prime 2 and all prime powers p^j with $j \ge 2$ is $O(\sqrt{y}\log^2(2Z))$. Indeed, $j \le \log_2 Z$, $\log p \le \log Z$, and $\sum_{n \ge 2} \min(1, y/n^2) \ll \sqrt{y}$. For large Y this contribution is at most $y^{1/2+\tau}/4$ on the whole interval.

Fix the class of r modulo 4. On odd primes p, reciprocity gives $(\frac{p}{r}) = \sigma_p(\frac{r}{p})$, where $\sigma_p \in \{1, -1\}$ depends only on p and that class. The equality includes $p = r$, when both sides vanish. For fixed y, put

$$P_y(t) = \sum_{2 < p \le Z} b_p(\frac{t}{p}), \qquad b_p = (\log p)\kappa(p)\sigma_p \min(1, y/p).$$

Here and below p denotes a prime. Now t can be any integer. Write $B_1 = \sum_p |b_p|$ and $B_2 = \sum_p |b_p|^2$. For large Z, $B_1 \le Z^2$, and $B_2 \ll y(\log Z)^2$ since $\sum_{n \ge 1} \min(1, y/n)^2 \ll y$.

For any integer $k \ge 1$,

$$\sum_{1 \le t \le U} |P_y(t)|^{2k} \le U(2k-1)!!\, B_2^k + Z^{2k} B_1^{2k}. \tag{3.2}$$

The square and nonsquare terms in high quadratic moments also occur in Granville and Soundararajan [3, §4, Lemmas 4.3–4.4]. Here averaging over all positive integers $t \le U$ gives the following elementary remainder. Expand $P_y^k \overline{P_y}^k$. If a prime occurs with odd multiplicity, the product of symbols has mean

zero over a period modulo the product of the distinct primes, by the Chinese remainder theorem. For instance $(\frac{t}{3})^3(\frac{t}{5})$ has zero sum over 15 consecutive integers, whereas $(\frac{t}{3})^2(\frac{t}{5})^2$ counts the eight units in a period. In the general nonsquare case the last incomplete period contributes at most Z^{2k} in absolute value. Summing the absolute coefficient products costs B_1^{2k}. If all multiplicities are even, the $2k$ factors can be paired at equal primes. For each of the $(2k-1)!!$ pairings the sum of the absolute coefficient products is B_2^k, and the sum over t is bounded by U. This proves (3.2), including for nonreal coefficients.

We need k large enough for a saving Y^{-ck} to give a fixed power of U, yet small enough for the period error Z^{6k} to be sublinear in U. Choose $k = \lfloor \log U/(16 \log Z) \rfloor$, so that $Z^{6k} \le U^{3/8}$. Markov's inequality with threshold $H = y^{1/2+\tau}/4 \ge 1$, and $(2k-1)!! \le (2k)^k$, show that

$$\#\{1 \le t \le U : |P_y(t)| > H\} \le U\{Ck(\log Z)^2 Y^{-2\tau}\}^k + U^{3/8}. \tag{3.3}$$

The factor in braces is at most $Y^{-(2\tau-\alpha)/2}$ for all sufficiently large Y. Also $k \ge \log U/(32D \log Y)$. Thus the first term in (3.3) is at most $U^{1-(2\tau-\alpha)/(64D)}$.

Apply this bound in each of the two classes modulo 4, at the two endpoints Y, Y^B and at every integer between them. There are $O(Y^B + 1) = U^{o(1)}$ tests. Between adjacent test points $P_y(t)$ is affine in y. Its absolute value is bounded by the chord through the endpoint absolute values, and the chord through the endpoint thresholds lies below $y^{1/2+\tau}/4$ by concavity. We can therefore take $\gamma = \min\{1/8, (2\tau-\alpha)/(128D)\}$, and restore the prime powers to obtain (3.1). □

4 The two auxiliary primes

Fix a positive integer M, $0 < \varepsilon < 1/52$, and $B \ge 1$. The original field construction uses primes of quite different sizes. With Y sufficiently large, put

$$z = Y^{\varepsilon/5}, \qquad U_1 = e^z, \qquad U_2 = e^{z(\log Y)^2}. \tag{4.1}$$

For a primitive nonprincipal character χ define the centered prime sum

$$G_\chi(y) = \sum_{n \le y} \Lambda(n)\chi(n)(y/n - 1) + y\,\frac{L'}{L}(1, \chi).$$

The centering cancels the linear term in y, a cancellation needed at the square-root scale. For the products below, the primitive representative is the unique primitive character agreeing with the product on integers coprime to $2Mqd$.

Corollary 4.1. *For all sufficiently large Y there exist primes $\sqrt{U_1} < q \le U_1$, $\sqrt{U_2} < d \le U_2$, coprime to M, with*

$$q \equiv 3 \pmod 8, \qquad d \equiv 7 \pmod{16}, \qquad d \equiv -1 \pmod{q \prod_{3 \le p \le z} p}.$$

For each character $\xi \bmod M$ and $i, j \in \{0, 1\}$ with $i + j > 0$, let χ be the primitive character of $\xi(\frac{\cdot}{q})^i(\frac{\cdot}{d})^j$. Then for every $Y \le y \le Y^B$,

$$|G_\chi(y)| \le y^{1/2+\varepsilon/2}, \qquad \left|\frac{L'}{L}(1, \chi)\right| \le 4 \log y. \tag{4.2}$$

For a real primitive character ρ of conductor dividing $8M$, the primitive character of $\rho(\frac{\cdot}{q})^i(\frac{\cdot}{d})^j$, for the same i, j, has no zero in

$$\Re s > 1 - 1/400, \qquad |\Im s| \le 8MU_1U_2^j.$$

The primes q, d specify the imaginary quadratic fields $\mathbb{Q}(\sqrt{-q})$ and $\mathbb{Q}(\sqrt{-d})$. Their product character cuts out $\mathbb{Q}(\sqrt{qd})$, accounting for the third twist. The congruences force 2, q, and every odd prime up to z to split in $\mathbb{Q}(\sqrt{-d})$. This is the pair of primes used in the release. Jaswal [5, §3.8] describes the successive exclusions and Slijepčević [8, Def. 3.2 and Props. 3.6–3.7] explains the progression constraints. We retain the zero-density and truncation estimates in the following proof.

Proof. The number of eligible primes. For q the modulus is 8. The prime number theorem for this fixed progression leaves $\gg U_1/\log U_1$ choices above $\sqrt{U_1}$. Given any of these choices, $q > z$ for large Y. The two conditions on d specify one reduced class modulo

$$m_2 = 16q \prod_{3 \le p \le z} p.$$

The elementary primorial bound $\sum_{p \le z} \log p \le z \log 4 + O(\log z)$ gives $\log m_2 \le 4z$ for large Y, uniformly in q. For a large modulus m', the rational progression specialization of Thorner and Zaman [9, Thm. 3.1] gives

$$\pi(U; m', a) \gg \frac{U}{(m')^5 \varphi(m') \log U} \qquad ((a, m') = 1,\ U \ge (m')^{522}).$$

For $m' = m_2$, the narrow ray class group over $\mathbb{Q}$ has order $\varphi(m_2)$, and there are primitive characters of conductor m_2, since 16 and each odd prime factor admit primitive characters. The exponent 522 absorbs the lower-order terms in their range. At $U = U_2$ we have $m_2 \le U^{4/(\log Y)^2}$. After removing at most $\sqrt{U}$ integers below the required interval, there remain

$$\gg \frac{U^{1-24/(\log Y)^2}}{\log U} \tag{4.3}$$

primes in the class. Both stages thus supply more than $U^{1-\gamma}$ choices by a factor tending to infinity, for every fixed $\gamma > 0$. A lower bound with exponent fixed below 1 would require a quantitative comparison with the small saving in Proposition 3.1.

Two exceptional sets. At a stage with upper endpoint U, write F for a bound on the conductor of the coefficient character before twisting by the new prime λ. At the first stage take $F = 8M$, at the second $F = 8MU_1$. In each case $F < \sqrt{U} < \lambda$ eventually. The product of a primitive character of conductor $f \le F$ and the Legendre character at λ is primitive of conductor $f\lambda$, since the conductors are coprime. Thus it is nonprincipal, and it equals $\kappa(n)(\frac{n}{\lambda})$ on *all* integers when κ is the primitive coefficient character. Using an imprimitive representative here could change the sum at primes dividing its modulus.

Put $Q = FU$. The release uses a zero-density estimate of the type studied by Montgomery [7, pp. 95–116]. Among primitive characters of conductor at most Q, at most $O(Q^{1/4})$ have a zero with $\Re s > 1 - 1/400$ and $|\Im s| \le Q$. For a published bound implying this count, take $\sigma = 399/400$, height Q, and auxiliary exponent $1/100$ in Heath-Brown [4, Thm. 1 and (1)]. His bound counts the zeros with multiplicity and gives $O((Q^3)^{2(1-\sigma)+1/100}) = O(Q^{9/200})$. Each exceptional character of conductor $f\lambda$, $f \le F < \lambda$, identifies the candidate prime λ by the largest prime factor of its conductor. Since $Q^{1/4} \le U^{3/8}$, the release's count alone removes at most $O(U^{3/8})$ candidate primes.

For the finite sums, apply Proposition 3.1 with $D = 400B$, $\tau = \varepsilon/2$ and $\nu = \alpha = \varepsilon/5$. The scales in (4.1) satisfy its hypotheses at both stages with $C_0 = 1, C_1 = 2$. Its proof gives $\gamma = \varepsilon/(64000B)$. At the first stage take the primitive representatives of all characters modulo $8M$. There are at most $\varphi(8M)$ coefficient sequences. For every candidate outside their exceptional sets and the zero-density set, all required q-twists pass both tests. Choose one such q.

At $U = U_2$, take the same representatives and their products with $(\frac{\cdot}{q})$, at most $2\varphi(8M)$ sequences, now fixed. The moment bound is uniform over these sequences. Together with the $O(U^{3/8})$ zero-density exclusion, it removes $O_M(U^{1-\gamma})$ candidates. Dividing by the progression count (4.3) gives

$$\frac{\#\{\text{rejected primes in the class}\}}{\#\{\text{candidate primes in the class}\}} \ll_M (\log U)\, U^{-\gamma+24/(\log Y)^2} = o(1).$$

Thus a proportion $1 - o(1)$ of the primes in this progression is available for d, uniformly in the choice of q. Both primes exceed $8M$, giving coprimality with M.

Passing from finite sums to L-functions. For each primitive twist just selected, put $X = Y^B$, $Z = X^{400}$, and let $R_\chi(Z) = -(L'/L)(1,\chi) - \sum_{n\le Z} \Lambda(n)\chi(n)/n$. For $y \le X \le Z$, expanding the minimum gives the exact identity

$$G_\chi(y) = -\sum_{n\le Z} \Lambda(n)\chi(n)\min(1, y/n) - yR_\chi(Z).$$

The first term is bounded by (3.1). The cutoff $Z = X^{400}$ gives $XZ^{-1/400} = 1$. A contour shift towards the boundary $399/400$ of the zero-free box therefore pays for the factor $y \le X$, up to the loss allowed in the error. For a quadratic-character instance of this truncation method see Lamzouri [6, Lemma 2.2 and Prop. 2.3]. For our full character family, (4.1) gives $Y^{\varepsilon/6} \le \log Q \le Y^{\varepsilon/4}$ eventually. In this range the released tail estimate in a zero-free box gives $X|R_\chi(Z)| \le Y^\varepsilon$. This is at most $y^{1/2+\varepsilon/2}/2$ for large Y, proving the first bound in (4.2) as in the original finite-sum transfer.

The two logarithmic heights are $\log(8M) + z$ and $\log(8M) + z + z(\log Y)^2$. Both are at most $2Y^{\varepsilon/5}(\log Y)^2$ eventually. This sharper upper bound is the additional scale hypothesis in the released logarithmic derivative estimate, which gives $|(L'/L)(1,\chi)| \le 4\log Y \le 4\log y$. Agreement on the common units identifies the primitive representatives for both the characters modulo M and the real primitive characters of conductor dividing $8M$. For $j = 0$ their box height is $8MU_1$, and for $j = 1$ it is $8MU_1U_2$, as required. □

The saving in (3.3) uses $2\tau > \alpha$. The release proves an intermediate selection with $\tau > 0$ independent of ε. At the call that supplies Dirichlet GRH, $\tau = \varepsilon/2$ and $\alpha = \varepsilon/5$. At that call the integer moment replaces the twisted prime number theorem in the prescribed progression, the ensuing second moment over primes, and its maximal estimate. The progression count, zero-density theorem and analytic transfer still supply essential parts of the argument.

Verification. From the project directory, `python3 checks/check_examples.py` checks the tile in Figure 1, cyclic wraparound and the period calculation in (3.2). Two other Codex collaborators independently reviewed the proofs and graduate-level exposition of the 21 November 2026 draft. A separate audit checked Theorem 2.1, Proposition 3.1 and Corollary 4.1 at the application parameters, including the analytic inputs and their use in the released Lean source. Its finite checks are included in `checks/independent-checks.py`. The two replacement arguments are proved here on paper.

References

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

[2] P. Erdős, *Graph theory and probability*, Canad. J. Math. **11** (1959), 34–38.

[3] A. Granville and K. Soundararajan, *The distribution of values of $L(1, \chi_d)$*, Geom. Funct. Anal. **13** (2003), 992–1028.

[4] D. R. Heath-Brown, *The density of zeros of Dirichlet's L-functions*, Canad. J. Math. **31** (1979), 231–240.

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

[6] Y. Lamzouri, *The distribution of Euler–Kronecker constants of quadratic fields*, J. Math. Anal. Appl. **432** (2015), 632–653.

[7] H. L. Montgomery, *Topics in multiplicative number theory*, Lecture Notes in Mathematics **227**, Springer, 1971.

[8] S. Slijepčević, *Outline of Alpöge's proof of the generalised Riemann hypothesis*, Hexagon 2611.00812v1, 21 November 2026.

[9] J. Thorner and A. Zaman, *An explicit bound for the least prime ideal in the Chebotarev density theorem*, Algebra Number Theory **11** (2017), 1135–1197.

[10] twinheterodox, *Three simplifications in the Alpöge–Claude argument for Dirichlet GRH*, Hexagon 2611.00828v1, 21 November 2026.