Low-Lying Zeros of Dirichlet L-Functions

The lowest-zero problem, the one-level density method, and a max-side bound via Wirtinger's inequality

Guo Liang Oon

September 6, 2026

Abstract

How far from the central point can the lowest zero of a Dirichlet L-function lie? Without assuming the generalized Riemann hypothesis (GRH), define $X_\chi = \frac{\log q}{2\pi} \min_\rho |\mathrm{Im}\, \rho|$ for a primitive character χ modulo a prime q, where ρ runs over the nontrivial zeros of $L(s, \chi)$. The main result of this dissertation proves that, for every $\varepsilon > 0$ and every sufficiently large prime q, some even character has no critical-line zero within scaled distance $\sqrt{3}/4 - \varepsilon$ from the central point; under GRH, when every nontrivial zero lies on the critical line, this gives $X_\chi > \sqrt{3}/4 - \varepsilon$. The proof applies Wirtinger's inequality to $e^{i\alpha t} L(\frac{1}{2} + it, \chi)$, with α chosen to minimise the resulting family derivative moment, and reduces the result to two unconditional family second moments. By contrast, the one-level-density method gives the constants $1/4$ of Hughes and Rudnick and Zhao. We prove that no non-negative support-2 test function in the corresponding positivity argument can improve Zhao's max-side constant, and obtain an exact min-side sign-test threshold and an explicit Selberg max-side bound from the extended support of Drappeau, Pratt and Radziwiłł. Random-matrix predictions and computations to modulus 100003 place $1/4$ near the median of the conjectural lowest-zero distribution and $\sqrt{3}/4$ in its upper tail.

Acknowledgements

I thank my supervisor, Chris Hughes, for his guidance, encouragement, and helpful discussions throughout this dissertation.

Contents

1 Introduction **3**

2 Dirichlet L-functions and the lowest zero **5**
2.1 Characters and L-functions 5
2.2 Zeros and the Riemann Hypothesis 5
2.3 The scaled lowest zero 6
2.4 One-sided versus two-sided lowest zeros 6

3 The one-level density method **7**
3.1 The explicit formula and the linear statistic 7
3.2 Averaging over the family and the support restriction 8
3.3 The interval count and the Beurling–Selberg majorant and minorant 9
3.4 The two theorems at support Δ 10
3.5 The min-side optimal test function 12
3.6 Exactness on the max side at support 2 13
3.7 Extended support 16
3.8 The limits of the method 17

4 A max-side bound via Wirtinger's inequality **18**
4.1 From zero counts to the size of L 19
4.2 The Wirtinger inequality 20
4.3 The twist and the choice of frequency origin 21
4.4 The two family second moments 21
4.5 Proof of Theorem 4.1 29
4.6 Remarks on the theorem 29

5 The random-matrix model **30**
5.1 Why the family is unitary 30
5.2 The sine-kernel process 30
5.3 The hole probability and the lowest-zero laws 31
5.4 Family extremes 32
5.5 Why the rigorous methods fall short of the extreme-value predictions . . . 33

6 Numerical calibration **33**
6.1 Data and zero computation 34
6.2 Two finite-conductor coordinates 34
6.3 Distributional comparison 35
6.4 Scope of the evidence 38

7 Discussion **39**
7.1 What the results establish 39
7.2 Assumptions and evidentiary scope 39
7.3 How far the min/max symmetry extends 39
7.4 Open problems 40

References **41**

1 Introduction

Let χ range over the primitive Dirichlet characters modulo a prime q. After scaling the zero ordinates of $L(s,\chi)$ by their asymptotic mean spacing, how low and how high can the lowest zero be as χ varies over the family?

Near the central point, the vertical ordinates of the nontrivial zeros have mean spacing $2\pi/\log q$.

Definition 1.1 (two-sided lowest-zero statistic)**.** Write the nontrivial zeros of $L(s,\chi)$ as $\rho_{\chi,j} = \frac{1}{2} + i\gamma_{\chi,j}$ with $\gamma_{\chi,j} \in \mathbb{C}$, indexed so that

$$\cdots \le \operatorname{Re}\gamma_{\chi,-2} \le \operatorname{Re}\gamma_{\chi,-1} < 0 \le \operatorname{Re}\gamma_{\chi,1} \le \operatorname{Re}\gamma_{\chi,2} \le \cdots .$$

Define the scaled zero parameters and the two-sided lowest-zero height by

$$x_{\chi,j} = \gamma_{\chi,j}\frac{\log q}{2\pi}, \qquad X_\chi = \min_j |\operatorname{Re} x_{\chi,j}| = \frac{\log q}{2\pi}\min_j |\operatorname{Im}\rho_{\chi,j}|.$$

Assuming GRH, Hughes and Rudnick [18] proved that for large prime q some character has $X_\chi \le 1/4 + o(1)$, and Zhao's 2025 preprint [34], forthcoming in the *Journal of Number Theory* in 2027, proved that some character has $X_\chi \ge 1/4 + o(1)$. Here $o(1)$ denotes a quantity tending to zero as $q \to \infty$. Thus the same constant appears as a min-side upper bound and a max-side lower bound.

Both theorems use the same family-averaged linear statistic. For a real, even, sufficiently decaying test function f, the explicit formula turns the zero sum into a prime-power sum weighted by its Fourier transform $\widehat{f}$, and character orthogonality evaluates its family average. Writing $\mathbb{E}_\chi$ for the uniform average over primitive characters modulo q, we have

$$W_f(\chi) := \sum_j f(x_{\chi,j}),$$
$$\mathbb{E}_\chi W_f(\chi) = \int_{\mathbb{R}} f(x)\,dx + o(1), \quad \operatorname{supp}\widehat{f} \subseteq [-\Delta,\Delta], \quad 0 < \Delta \le 2.$$

This averaged asymptotic formula is unconditional. To extract an interval count, assume GRH, so that every $x_{\chi,j}$ is real. An exact interval count would take $f = \mathbf{1}_{[-\beta,\beta]}$, but

$$\widehat{\mathbf{1}_{[-\beta,\beta]}}(\xi) = \begin{cases} \dfrac{\sin(2\pi\beta\xi)}{\pi\xi}, & \xi \ne 0, \\ 2\beta, & \xi = 0, \end{cases} \qquad \operatorname{supp}\widehat{\mathbf{1}_{[-\beta,\beta]}} = \mathbb{R},$$

so the indicator lies outside the admissible support range. The classical Beurling–Selberg minorant and majorant satisfy

$$M_-(x) \le \mathbf{1}_{[-\beta,\beta]}(x) \le M_+(x), \quad \operatorname{supp}\widehat{M_\pm} \subseteq [-\Delta,\Delta],$$
$$\int_{\mathbb{R}} M_\pm(x)\,dx = 2\beta \pm \frac{1}{\Delta}.$$

Applying the averaged explicit formula to $M_\pm$ yields

$$2\beta - \frac{1}{\Delta} > 0 \iff \beta > \beta_{\min}(\Delta) = \frac{1}{2\Delta},$$
$$2\beta + \frac{1}{\Delta} < 1 \iff \beta < \beta^{\mathrm{Sel}}_{\max}(\Delta) = \frac{1}{2} - \frac{1}{2\Delta},$$

where $\beta_{\min}$ is the boundary above which the minorant forces some character to have a zero in the interval, while $\beta_{\max}^{\mathrm{Sel}}$ is the Selberg max-side bound, below which the classical majorant forces some character to have no zero there. Both tend to $1/4$ as $\Delta \uparrow 2$, the fixed-modulus support boundary. Carneiro, Ismoilov and Ramos proved min-side sharpness over the full admissible class of the positivity problem [5]. On the max side, Theorem 3.10 proves that no non-negative admissible support-2 test can establish a zero-free symmetric interval of scaled half-length beyond $1/4$ by the same positivity argument. The extended-support density theorem of Drappeau, Pratt and Radziwiłł [12], which averages additionally over the modulus, gives the exact min-side sign-test threshold $1093/4472 \approx 0.2444$ and the explicit Selberg max-side bound $1143/4472 \approx 0.2556$.

The $1/4$ ceiling is a limitation of Fourier-restricted zero-count tests, not a bound on X_χ. Replacing zero counts by second moments of $L(\frac{1}{2}+it,\chi)$ and $L'(\frac{1}{2}+it,\chi)$, together with Wirtinger's inequality, gives as the main result of the dissertation a stronger max-side bound.

> **Theorem 4.1**. Fix $0 < \beta < \sqrt{3}/4$. For every sufficiently large prime q there is an even non-principal character $\chi \bmod q$ whose L-function has no zero ρ with $\operatorname{Re}\rho = \frac{1}{2}$ and $|\operatorname{Im}\rho| \le 2\pi\beta/\log q$.

Under GRH every nontrivial zero lies on the critical line, so the same character satisfies $X_\chi > \beta$, and therefore

$$\liminf_{\substack{q\to\infty \\ q \text{ prime}}} \max_{\chi \in \mathcal{F}_q} X_\chi \;\ge\; \frac{\sqrt{3}}{4} \;\approx\; 0.433.$$

The constant is nearly twice the exact support-2 limit. The proof applies Wirtinger's inequality to $e^{i\alpha t}L(\frac{1}{2}+it,\chi)$ and sums over the even characters; unconditional second-moment asymptotics contradict the assumption that every character has a nearby critical-line zero when $\beta < \sqrt{3}/4$.

The primitive-character family is expected to have unitary symmetry, so its scaled zeros should follow the local statistics of eigenangles of Haar-random unitary matrices. The resulting two-sided single-character law has mean 0.295 and median 0.259, while treating distinct conjugate pairs as independent predicts a family minimum of order q^{-1} and a family maximum asymptotic to $\sqrt{2\log q}/\pi$. The constant $1/4$ lies near the median of the single-character law, whereas $\sqrt{3}/4$ lies in its upper tail; both are fixed constants, unlike the conjectural family extremes. Computations for all non-principal characters at seven primes up to $q = 100003$ are compared with this model after separating the leading scaling from the substantial finite-conductor gamma-factor correction.

The zero-location results are conditional on GRH, while the critical-line zero-free statement in Theorem 4.1, the moment computations, and the extended-support density theorem are unconditional. The proofs do not use the conjectural random-matrix law or its predicted extreme-value scales. The zero calculations search only on the critical line, so an off-line zero would be invisible; they do not attempt to verify GRH even over the computed height range. Unless a result is explicitly one-sided, max-side bounds concern the stronger two-sided variable X_χ.

2 Dirichlet L-functions and the lowest zero

2.1 Characters and L-functions

A *Dirichlet character* modulo q is a homomorphism $\chi\colon (\mathbb{Z}/q\mathbb{Z})^\times \to \mathbb{C}^\times$, extended to all of $\mathbb{Z}$ by $\chi(n) = 0$ when $\gcd(n, q) > 1$. Write χ_0 for the principal character, $\chi_0(n) = 1$ for every n coprime to q. A character is *induced* from a divisor $d \mid q$ if $\chi(n)$ depends only on $n \bmod d$ for every n coprime to q. The divisors with this property are closed under greatest common divisor, so there is a least one, the *conductor* of χ, and χ is *primitive* when its conductor equals q. The conductor rather than the modulus q fixes the arithmetic size of a character, and it is the conductor that sets the zero density in Section 2.3. For prime q the only proper divisor is 1, and the only character induced from it is χ_0. Hence every non-principal character modulo a prime q has conductor q and is primitive. The family averaged over throughout is

$$\mathcal{F}_q = \{\chi \bmod q : \chi \text{ primitive}\}, \qquad |\mathcal{F}_q| = q - 2 \quad (q \text{ prime}).$$

Unless otherwise stated, limits in q are through primes; conductor-averaged statements use general moduli.

For a quantity A_χ on this family, write $\mathbb{E}_\chi A_\chi = |\mathcal{F}_q|^{-1} \sum_{\chi \in \mathcal{F}_q} A_\chi$. As $q \to \infty$, $A \sim B$ means $A/B \to 1$, $A = o(B)$ means $A/B \to 0$, $A = O(B)$ or $A \ll B$ means that $|A| \le C|B|$ for some constant $C > 0$ and all sufficiently large q, and $A \asymp B$ means both $A \ll B$ and $B \ll A$. A subscript such as O_ε allows the constant to depend on the subscripted parameter.

Dirichlet introduced these character series in his proof of the theorem on primes in arithmetic progressions [11]. To each χ one attaches the *Dirichlet L-function*

$$L(s, \chi) = \sum_{n=1}^{\infty} \frac{\chi(n)}{n^s} = \prod_p \left(1 - \frac{\chi(p)}{p^s}\right)^{-1} \qquad (\operatorname{Re} s > 1),$$

which for primitive $\chi \neq \chi_0$ continues to an entire function and satisfies a functional equation [9, 20]. Its symmetry is recorded by the *completed* L-function. With the parity $\mathfrak{a} = \mathfrak{a}(\chi) \in \{0, 1\}$ fixed by $\chi(-1) = (-1)^{\mathfrak{a}}$, so that χ is called even or odd, set

$$\Lambda(s, \chi) = \left(\frac{q}{\pi}\right)^{(s+\mathfrak{a})/2} \Gamma\left(\frac{s + \mathfrak{a}}{2}\right) L(s, \chi),$$

which is entire and satisfies

$$\Lambda(s, \chi) = \varepsilon(\chi)\, \Lambda(1 - s, \overline{\chi}), \qquad |\varepsilon(\chi)| = 1.$$

Because the functional equation ties χ to $\overline{\chi}$ rather than to itself, the one-sided variable introduced in Section 2.4 need not agree with the two-sided X_χ.

2.2 Zeros and the Riemann Hypothesis

For primitive non-principal χ, the zeros of $L(s, \chi)$ split into the *trivial* zeros at the poles of the gamma factor ($s = -\mathfrak{a}, -\mathfrak{a} - 2, \ldots$) and the *nontrivial* zeros, which lie in the critical strip $0 < \operatorname{Re} s < 1$ and are exactly the zeros of $\Lambda(s, \chi)$. All zero sums count multiplicity.

If $\rho = \sigma_\chi + it_\chi$ is a nontrivial zero of $L(s,\chi)$, with $\sigma_\chi, t_\chi \in \mathbb{R}$, set $\gamma_\chi = t_\chi - i(\sigma_\chi - \frac{1}{2})$, so that

$$\rho = \tfrac{1}{2} + i\gamma_\chi, \qquad \gamma_\chi \in \mathbb{C}.$$

The *generalized Riemann hypothesis* (GRH) for $L(s,\chi)$ asserts that every nontrivial zero has $\sigma_\chi = \frac{1}{2}$, equivalently $\gamma_\chi \in \mathbb{R}$. Under GRH the ordinates form a point set on the real line, and the central point $s = \frac{1}{2}$, the point of symmetry of the functional equation, is the natural origin from which to measure the lowest zero.

2.3 The scaled lowest zero

The local density of zeros with respect to the vertical ordinate $t = \operatorname{Im}\rho$ is governed by the conductor and the gamma factor. For parity $\mathfrak{a}$, the smooth local zero density is

$$d_{\mathfrak{a},q}(t) = \frac{1}{2\pi}\left[\log\frac{q}{\pi} + \operatorname{Re}\psi\left(\frac{\frac{1}{2}+\mathfrak{a}+it}{2}\right)\right] = \frac{\log q}{2\pi} + O(1)$$

for bounded t, where $\psi = \Gamma'/\Gamma$ is the digamma function. An interval of vertical ordinates of length h near t contains approximately $d_{\mathfrak{a},q}(t)h$ zeros, so their mean spacing is the reciprocal of the density. Factoring out the main term and inverting gives

$$d_{\mathfrak{a},q}(t) = \frac{\log q}{2\pi}\left(1 + O\left(\frac{1}{\log q}\right)\right), \qquad \frac{1}{d_{\mathfrak{a},q}(t)} = \frac{2\pi}{\log q}\left(1 + O\left(\frac{1}{\log q}\right)\right).$$

The conductor therefore gives the leading spacing $2\pi/\log q$ between consecutive vertical ordinates, up to a relative error $O(1/\log q)$.

This spacing shrinks with q, so raw ordinates from different moduli are not comparable. Rescale the centred zero parameter by the leading local density:

$$x_{\chi,j} = \gamma_{\chi,j}\,\frac{\log q}{2\pi},$$

so that $\operatorname{Re} x_{\chi,j} = t_{\chi,j}\log q/(2\pi)$ measures vertical height in asymptotic mean-spacing units. The rescaled vertical ordinates have unit mean density to leading order as $q \to \infty$, not exactly at each finite conductor: the discarded correction is bounded and depends on parity, and Section 6.2 quantifies it at the conductors computed there. The asymptotic statements use only the leading conductor scale. GRH places every $x_{\chi,j}$ on the real line; without GRH, the explicit formula remains valid but evaluates entire test functions at complex zero parameters. Fourier support Δ refers to the same conductor scaling in either setting.

2.4 One-sided versus two-sided lowest zeros

The vertical ordinates of a real (quadratic) character are symmetric about 0, but those of a complex character need not be: the functional equation couples χ to $\overline{\chi}$ rather than to itself. This produces distinct one-sided and two-sided statistics.

Definition 2.1 (one-sided lowest-zero statistic)**.** With the indexing of Definition 1.1, define

$$X_\chi^+ = \operatorname{Re} x_{\chi,1} = \operatorname{Re}\gamma_{\chi,1}\frac{\log q}{2\pi}.$$

For a nontrivial zero ρ the functional equation implies

$$L(\rho,\chi)=0 \quad\Longrightarrow\quad L(1-\rho,\overline{\chi})=0.$$

A zero $\rho=\beta-it$ of $L(s,\chi)$ with $t>0$ therefore gives the zero $1-\rho=(1-\beta)+it$ of $L(s,\overline{\chi})$. Since the family $\mathcal{F}_q$ is closed under $\chi\mapsto\overline{\chi}$, it follows that

$$X_\chi=\min\bigl(X_\chi^+,X_{\overline{\chi}}^+\bigr).$$

Taking minima over the family absorbs the inner minimum and gives $\min_{\chi\in\mathcal{F}_q}X_\chi=\min_{\chi\in\mathcal{F}_q}X_\chi^+$, including when a central zero makes both sides 0. Maxima behave differently, since only $\max_{\chi\in\mathcal{F}_q}X_\chi\le\max_{\chi\in\mathcal{F}_q}X_\chi^+$ is automatic: an empty interval on both sides of the centre is a stronger event than an empty interval on one side. We use X_χ^+ for the one-sided Hughes–Rudnick formulation and X_χ for Zhao's theorem, Theorem 4.1, and the identity between the two thresholds at $1/4$. A max-side bound for X_χ implies the same bound for X_χ^+.

3 The one-level density method

The Hughes–Rudnick and Zhao theorems both follow from the family average of a smoothed zero count in a symmetric interval. The explicit formula converts this statistic into prime sums, character orthogonality evaluates their family average for Fourier support $\Delta\le 2$, and the classical Selberg approximants give upper and lower interval-count bounds. At support 2, the constant $1/4$ is optimal in the respective min- and max-side positivity problems, by the full-class arguments of Sections 3.5 and 3.6. An additional average over the modulus extends the admissible support slightly past 2.

3.1 The explicit formula and the linear statistic

For primitive χ mod q with nontrivial zeros $\frac12+i\gamma_{\chi,j}$, and for any even function h, analytic in a neighbourhood of the strip $|\operatorname{Im} r|\le\frac12+\delta$ for some $\delta>0$ and satisfying $|h(r)|\ll(1+|r|)^{-1-\eta}$ uniformly throughout that strip for some $\eta>0$, one has the explicit formula [7, 20]

$$\sum_j h(\gamma_{\chi,j})=\frac{1}{2\pi}\int_{\mathbb{R}}h(r)\bigl(\log q+G_\chi(r)\bigr)\,dr-\sum_{n=1}^{\infty}\frac{\Lambda(n)}{\sqrt{n}}\,g(\log n)\bigl(\chi(n)+\overline{\chi}(n)\bigr),$$

where $g(u)=\frac{1}{2\pi}\int_{\mathbb{R}}h(r)e^{-iru}\,dr$ is the Fourier transform of the zero-side function, $\Lambda(n)=\log p$ when $n=p^k$ with p prime and $k\ge1$, and $\Lambda(n)=0$ otherwise, and

$$G_\chi(r)=-\log\pi+\operatorname{Re}\frac{\Gamma'}{\Gamma}\left(\frac{\frac12+\mathfrak{a}(\chi)+ir}{2}\right)$$

collects the gamma and π factors of the completed function of Section 2.1. The explicit formula is an identity. Its zero terms are residues of the logarithmic derivative of $\Lambda(s,\chi)$, and its prime-power terms come from the Euler product through $-L'/L(s,\chi)=\sum_n\Lambda(n)\chi(n)n^{-s}$. At the point $s=\frac12+ir$ on the critical line, $n^{-s}=n^{-1/2}e^{-ir\log n}$, so the zero-height and log-prime variables are Fourier dual. The symmetric combination $\chi(n)+\overline{\chi}(n)$ reflects the functional equation pairing of s with $1-s$ and χ with $\overline{\chi}$.

Definition 3.1 (admissible test function and linear statistic)**.** A function f is admissible if it is continuous, real and even, $\widehat{f}$ has compact support, and $|f(x)| \ll (1+|x|)^{-1-\eta}$ for some $\eta > 0$, where

$$\widehat{f}(\xi) = \int_{\mathbb{R}} f(x) e^{-2\pi i x\xi}\, dx.$$

For admissible f, define

$$W_f(\chi) = \sum_{\gamma_\chi} f\left(\gamma_\chi \frac{\log q}{2\pi}\right).$$

Since f and $\widehat{f}$ are integrable and $\widehat{f}$ has compact support, Fourier inversion extends f to the entire function $f(z) = \int_{\mathbb{R}} \widehat{f}(\xi) e^{2\pi i z\xi}\, d\xi$. This extension is used in $W_f(\chi)$ when GRH is not assumed and γ_χ may be complex.

Specialising the explicit formula to $h(r) = f(r \log q/(2\pi))$ gives

$$g(\log n) = (\log q)^{-1} \widehat{f}(\log n/\log q).$$

The conductor term becomes exactly $\int_{\mathbb{R}} f(x)\, dx$, while the gamma term contributes $o(1)$ for fixed admissible f, because the digamma factor grows only logarithmically and the $1/\log q$ scaling suppresses it together with the parity dependence. Hence

$$W_f(\chi) = \int_{\mathbb{R}} f(x)\, dx - \frac{1}{\log q} \sum_{n=1}^{\infty} \frac{\Lambda(n)}{\sqrt{n}} \widehat{f}\left(\frac{\log n}{\log q}\right) \big(\chi(n) + \overline{\chi}(n)\big) + o(1).$$

The support of $\widehat{f}$ measures the arithmetic reach of the statistic: if $\operatorname{supp} \widehat{f} \subseteq [-\Delta, \Delta]$, the prime-power sum contains only $n \le q^{\Delta}$.

3.2 Averaging over the family and the support restriction

Averaging W_f over the family turns the character factors into congruence conditions. Character orthogonality, which is finite Fourier inversion on the group $(\mathbb{Z}/q\mathbb{Z})^\times$, shows that averaging over the complete character group retains exactly the prime powers $n \equiv 1 \pmod q$, and separating by parity retains $n \equiv \pm 1 \pmod q$. Dropping the principal character contributes only a lower-order term, since $\varphi(q)^{-1} \sum_{\chi \neq \chi_0} \chi(n) = \mathbf{1}_{n \equiv 1\,(q)} - \varphi(q)^{-1} \mathbf{1}_{(n,q)=1}$, where $\varphi(q) = |(\mathbb{Z}/q\mathbb{Z})^\times|$ is Euler's totient. The averaged one-level density problem is therefore an estimate for weighted prime powers in arithmetic progressions up to q^{Δ}. Define the averaged prime-sum term by

$$R_f(q) = \frac{1}{\log q} \sum_{n=1}^{\infty} \frac{\Lambda(n)}{\sqrt{n}} \widehat{f}\left(\frac{\log n}{\log q}\right) \mathbb{E}_\chi\big(\chi(n) + \overline{\chi}(n)\big).$$

Bounding the prime-power sum by the corresponding sum over all integers gives $R_f(q) \ll q^{\Delta/2-1}$, which is $o(1)$ only when $\Delta < 2$. Hughes and Rudnick instead use estimates for primes in arithmetic progressions: Brun–Titchmarsh bounds the contribution from $n \equiv 1 \pmod q$, while the prime number theorem bounds the principal-character correction, giving

$$R_f(q) \ll \frac{q^{\Delta/2-1}}{\log q}.$$

This is $o(1)$ also at $\Delta = 2$, which gives Theorem 3.2 with its inclusive endpoint.

Theorem 3.2 (Hughes–Rudnick [18]). *Let f be admissible with* $\operatorname{supp}\widehat{f} \subseteq [-\Delta, \Delta]$ *for some* $0 < \Delta \leq 2$. *Then, as* $q \to \infty$ *through primes,*

$$\mathbb{E}_\chi W_f(\chi) = \int_{\mathbb{R}} f(x)\,dx + o(1). \tag{3.1}$$

The limit $\int f$ is the one-level density of the unitary symmetry type in the Katz–Sarnak classification [22, 23]; Section 5.1 explains why this family is unitary. The support restriction reflects the length q^Δ of the prime-power sum. The estimates of Hughes and Rudnick give $R_f(q) = o(1)$ through the endpoint $\Delta = 2$, but no corresponding fixed-modulus estimate is known beyond $\Delta = 2$.

3.3 The interval count and the Beurling–Selberg majorant and minorant

The lowest-zero problem requires the count of zeros in a symmetric interval. Reindexing $\chi \mapsto \overline{\chi}$ in the family average shows that for every admissible f

$$\mathbb{E}_\chi \sum_j f(x_{\chi,j}) = \mathbb{E}_\chi \sum_j f(-x_{\chi,j}),$$

so the odd part of f cancels and the one-level density cannot distinguish left from right.

Definition 3.3 (symmetric interval count). Under GRH, for $\beta > 0$ define

$$N_\chi(\beta) = \sum_{\gamma_\chi} \mathbf{1}_{[-\beta,\beta]}\left(\gamma_\chi \frac{\log q}{2\pi}\right), \qquad \bar{N}(\beta) = \mathbb{E}_\chi\, N_\chi(\beta).$$

Since each $N_\chi(\beta)$ is a non-negative integer,

$$\bar{N}(\beta) > 0 \implies \min_{\chi \in \mathcal{F}_q} X_\chi \leq \beta, \qquad \bar{N}(\beta) < 1 \implies \max_{\chi \in \mathcal{F}_q} X_\chi > \beta. \tag{3.2}$$

The indicator itself is not admissible, since its Fourier transform is

$$\widehat{\mathbf{1}_{[-\beta,\beta]}}(\xi) = \int_{-\beta}^{\beta} e^{-2\pi i x\xi}\,dx = \frac{\sin(2\pi\beta\xi)}{\pi\xi},$$

with continuous value 2β at $\xi = 0$, and is not compactly supported. It must therefore be replaced by band-limited approximations. The Beurling–Selberg problem asks for admissible majorants and minorants whose integrals are as close as possible to the indicator's integral 2β.

Lemma 3.4 (Beurling, Selberg [30, §§19–21]; [33]). *For every $\Delta > 0$ and $\beta > 0$ there exist even admissible functions $M_\pm$ with* $\operatorname{supp}\widehat{M_\pm} \subseteq [-\Delta, \Delta]$ *and*

$$M_-(x) \leq \mathbf{1}_{[-\beta,\beta]}(x) \leq M_+(x) \quad (x \in \mathbb{R}), \qquad \int_{\mathbb{R}} M_\pm = 2\beta \pm \frac{1}{\Delta}.$$

If $2\beta\Delta \in \mathbb{Z}$, the integrals $2\beta \pm 1/\Delta$ are optimal among all band-limited majorants and minorants of the indicator.

Outside the interval, the pointwise inequalities $M_- \le \mathbf{1}_{[-\beta,\beta]} \le M_+$ reduce to $M_- \le 0 \le M_+$. The approximants are not compactly supported on $\mathbb{R}$: the majorant has non-negative x^{-2} tails and the minorant has non-positive x^{-2} tails. An interpolation node is a point where an approximant touches the function it bounds. When $2\beta\Delta \in \mathbb{Z}$, the contact lattices generated by the two endpoints align, and the exterior nodes are $\pm(\beta + k/\Delta)$, $k \ge 1$; otherwise the inequalities are strict outside the interval. The error $1/\Delta$ therefore includes the tails rather than only a local smoothing of the two jumps. At a zero outside $[-\beta, \beta]$, M_- is non-positive, so distant zeros cannot increase the minorant sum and $\sum_j M_-(x_{\chi,j}) \le N_\chi(\beta)$ for every character.

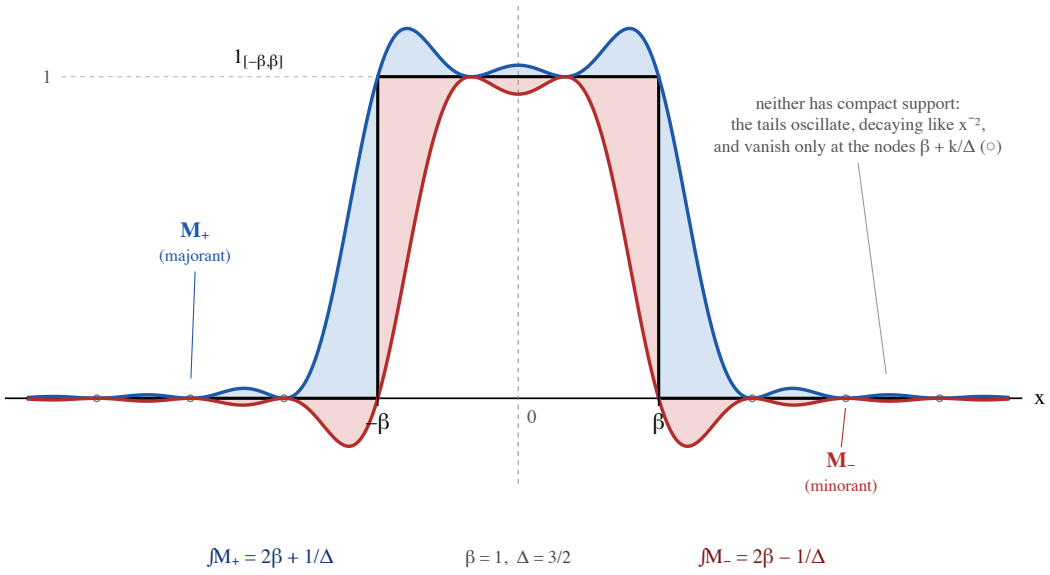

Figure 1: The Beurling–Selberg majorant M_+ and minorant M_- around the indicator of $[-\beta, \beta]$, computed for $\beta = 1$ and $\Delta = 3/2$. Neither is compactly supported: outside $[-\beta, \beta]$ the two curves oscillate on opposite sides of the axis, decaying like x^{-2} and vanishing only at the nodes $\pm(\beta + k/\Delta)$ (circles). The shaded regions, tails included, carry the integral excess $+1/\Delta$ and deficit $-1/\Delta$, and both shrink as the admissible support Δ grows.

Under GRH, the pointwise sandwich gives $W_{M_-}(\chi) \le N_\chi(\beta) \le W_{M_+}(\chi)$ for every χ. Averaging and applying Theorem 3.2 to the two zero sums gives

$$2\beta - \frac{1}{\Delta} + o(1) \;\le\; \bar{N}(\beta) \;\le\; 2\beta + \frac{1}{\Delta} + o(1). \tag{3.3}$$

3.4 The two theorems at support Δ

Combining the lower and upper bounds in (3.3) with the implications (3.2) yields the two lowest-zero theorems. Theorem 3.2 is unconditional. GRH enters the interval-count deductions by making the scaled ordinates real, so that the pointwise signs of the test functions imply the bounds for $N_\chi(\beta)$.

Theorem 3.5 (Hughes–Rudnick [18, proof of Theorem 8.1]; see also [34])**.** *Assume GRH. For* $0 < \Delta \le 2$*,*

$$\limsup_{\substack{q\to\infty \\ q \text{ prime}}} \min_{\chi \in \mathcal{F}_q} X_\chi \;\le\; \beta_{\min}(\Delta) = \frac{1}{2\Delta}.$$

Proof. For every $\beta > 1/(2\Delta)$, the lower bound in (3.3) is positive for all sufficiently large q, so (3.2) gives some χ with $X_\chi \le \beta$. Letting $\beta \downarrow 1/(2\Delta)$ gives the stated limsup. □

Theorem 3.6. *Assume GRH. For* $1 < \Delta \le 2$*,*

$$\liminf_{\substack{q\to\infty \\ q \text{ prime}}} \max_{\chi \in \mathcal{F}_q} X_\chi \;\ge\; \beta^{\mathrm{Sel}}_{\max}(\Delta) = \frac{1}{2} - \frac{1}{2\Delta}.$$

Proof. For every $\beta < \frac{1}{2} - 1/(2\Delta)$, the upper bound in (3.3) is below 1 for all sufficiently large q, so (3.2) gives some χ with $X_\chi > \beta$. Letting $\beta \uparrow \frac{1}{2} - 1/(2\Delta)$ gives the stated liminf. □

Zhao [34, Theorem 6] proves the endpoint $\Delta = 2$ with an explicit lower-order term. For $1 < \Delta < 2$, the stated bound follows from Theorem 3.2 and the Selberg majorant.

The min-side threshold and the Selberg max-side bound satisfy

$$\beta_{\min}(\Delta) + \beta^{\mathrm{Sel}}_{\max}(\Delta) = \frac{1}{2} \qquad (\Delta > 1).$$

This identity follows from the equal excess and deficit $1/\Delta$ of Selberg's construction. At $\Delta = 2$, both quantities equal $1/4$ and are sharp in their respective positivity problems. The max-side crossing then has $2\beta\Delta = 1$, so Lemma 3.4 certifies the majorant's optimality. Zhao's original argument uses a specific non-negative test function rather than the majorant; Section 3.6 proves the same ceiling for the full non-negative test class. Since $X_\chi^+ \ge X_\chi$, the liminf bound of Theorem 3.6 also holds with X_χ replaced by X_χ^+, but a sharp one-sided bound must also account for zeros at the central point.

For the two-sided statistic X_χ, the exact min-side sign-test thresholds and the Selberg max-side lower bounds are:

Δ	Exact min-side sign-test threshold $\beta_{\min}(\Delta)$	Selberg max-side lower bound $\beta^{\mathrm{Sel}}_{\max}(\Delta)$	Density input
2	$1/4 = 0.2500$	$1/4 = 0.2500$	Theorem 3.2
$2 + \frac{50}{1093}$	$\frac{1093}{4472} \approx 0.2444$	$\frac{1143}{4472} \approx 0.2556$	Theorem 3.11
$\to \infty$	$\to 0$	$\to 1/2$	Friedlander–Granville conjecture (Section 3.8)

The entry $\Delta \to \infty$ is not a theorem; it is the limiting behaviour of the method if the density theorem held at every fixed support. Section 3.8 states a hypothesis that would provide this input and describes the remaining first-moment obstruction.

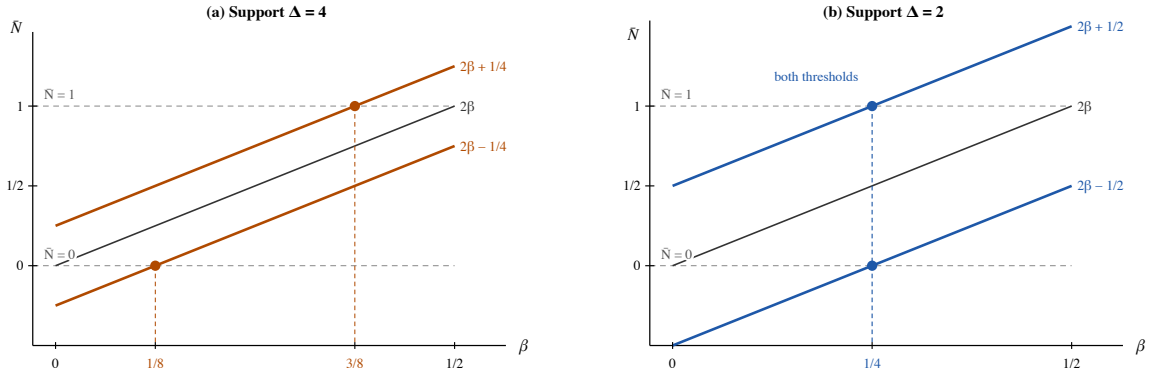

Figure 2: The Selberg count bounds (3.3) around the interval mass 2β. (a) At $\Delta = 4$, the lower and upper bounds cross $\bar{N} = 0$ and $\bar{N} = 1$ at $\beta_{\min} = 1/8$ and $\beta_{\max}^{\mathrm{Sel}} = 3/8$. (b) At the fixed-modulus support boundary $\Delta = 2$, both crossings occur at $\beta = 1/4$.

3.5 The min-side optimal test function

Hughes and Rudnick obtain the constant $1/(2\Delta)$ from a quadratic test family. Choose a nonzero continuous real even function g with $\operatorname{supp} \widehat{g} \subseteq [-\Delta/2, \Delta/2]$ and $g(x) \ll (1 + |x|)^{-3/2-\varepsilon}$ for some $\varepsilon > 0$, and set

$$f(x) = (\beta^2 - x^2)\, g(x)^2.$$

The polynomial factor builds in the sign pattern $f \geq 0$ on $[-\beta, \beta]$ and $f \leq 0$ outside, and $\widehat{f}$ is supported in $[-\Delta, \Delta]$ because squaring g doubles the Fourier support while multiplication by $\beta^2 - x^2$ only differentiates on the Fourier side. If every character had $X_\chi \geq \beta$, then under GRH every term of $W_f(\chi)$ would be non-positive, so $\mathbb{E}_\chi W_f(\chi) \leq 0$; by Theorem 3.2 the method therefore succeeds whenever

$$\int f = \beta^2 \int g^2 - \int x^2 g^2 > 0 \qquad \Longleftrightarrow \qquad \beta^2 > \frac{\int x^2 g^2}{\int g^2}. \tag{3.4}$$

Writing $u = \widehat{g}$, Plancherel's theorem turns multiplication by x into differentiation and rewrites the right side of (3.4) as

$$\frac{1}{4\pi^2} \frac{\int_{-\Delta/2}^{\Delta/2} |u'(\xi)|^2\, d\xi}{\int_{-\Delta/2}^{\Delta/2} |u(\xi)|^2\, d\xi}. \tag{3.5}$$

Since g is integrable and $\widehat{g}$ is supported in $[-\Delta/2, \Delta/2]$, the function u is continuous and vanishes at the endpoints $\pm\Delta/2$. Equation (3.5) is $1/(4\pi^2)$ times the Rayleigh quotient of $-d^2/d\xi^2$ with Dirichlet boundary conditions on an interval of length Δ. Expanding u in the eigenfunctions $\sin(k\pi(\xi + \Delta/2)/\Delta)$, $k \geq 1$, expresses (3.5) as a weighted average of $k^2/(4\Delta^2)$, minimised by the first eigenfunction $u_1(\xi) = \cos(\pi\xi/\Delta)$. Consequently

$$\inf_g \frac{\int x^2 g^2}{\int g^2} = \frac{1}{4\pi^2} \left(\frac{\pi}{\Delta}\right)^2 = \frac{1}{4\Delta^2}, \qquad \beta_{\min}(\Delta) = \frac{1}{2\Delta},$$

in agreement with Theorem 3.5. At the endpoint the optimal test has $\int f = 0$, so the argument uses $\beta > 1/(2\Delta)$ and lets β decrease to the threshold.

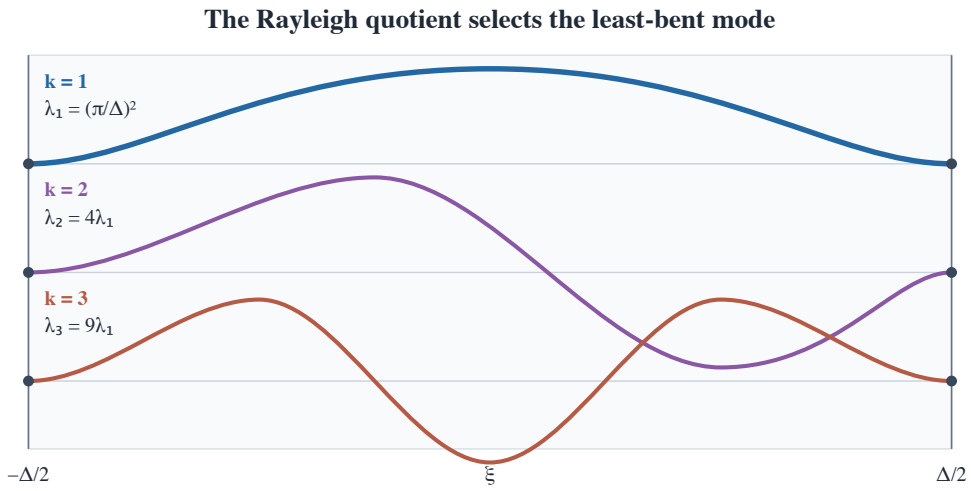

Figure 3: Dirichlet eigenfunctions of $-d^2/d\xi^2$ on the Fourier support interval. Every admissible u vanishes at both endpoints; the first eigenfunction minimises the Rayleigh quotient (3.5), and higher eigenfunctions only increase it.

The construction restricts f to the special form $(\beta^2 - x^2)g^2$, with a single sign change at $\pm\beta$. Carneiro, Ismoilov and Ramos proved that this restriction does not alter the optimum: among all admissible functions with $f \le 0$ off $[-\beta, \beta]$ and $\int f > 0$, support Δ forces $\beta \ge 1/(2\Delta)$ [5]. The constant $\beta_{\min}(\Delta)$ is therefore optimal over this full admissible class, not merely over the quadratic family.

3.6 Exactness on the max side at support 2

Theorem 3.6 obtains the max-side constant $1/4$ from one particular non-negative test function. A different non-negative admissible function or a different rescaling might, in principle, prove that some character has a higher lowest zero from the same support-2 one-level density. Theorem 3.10 rules this out: optimising over both choices still gives the ceiling $1/4$.[1] Krein factorisation writes each non-negative test function H with Fourier support in $[-1, 1]$ as $H = |g|^2$; Bessel's inequality applied to the evaluations $g(\pm s)$ then bounds the interval minimum, and the cosine equality direction produces Zhao's test.

Definition 3.7. The class $\mathcal{A}$ consists of the continuous, even, non-negative functions $H \in L^1(\mathbb{R})$ satisfying $\operatorname{supp} \widehat{H} \subseteq [-1, 1]$.

Given any non-negative admissible test $f \in L^1(\mathbb{R})$, its even part $f_{\text{ev}}(x) = \frac{1}{2}(f(x) + f(-x))$ satisfies $f_{\text{ev}} \ge 0$, $\int f_{\text{ev}} = \int f$, and

$$\min_{|x| \le r} f_{\text{ev}}(x) \ge \min_{|x| \le r} f(x).$$

If $\operatorname{supp} \widehat{f_{\text{ev}}} \subseteq [-c, c]$ with $0 < c \le 2$, then $H(x) = f_{\text{ev}}(x/c)$ lies in $\mathcal{A}$ and $f_{\text{ev}}(x) = H(cx)$. Thus the choice of $\mathcal{A}$ and the free dilation c cover every test available to the support-2

[1] This ceiling belongs to the symmetric-interval first-moment method, not to the family maximum; Theorem 4.1 exceeds it using second moments of L and L'.

one-level-density positivity argument. Write

$$M(H) = \int_{\mathbb{R}} H(x)\,dx, \qquad m_H(s) = \min_{|x|\le s} H(x).$$

Since $M(aH) = aM(H)$ and $m_{aH}(s) = a\,m_H(s)$ for every $a > 0$, rescaling H allows us to impose $M(H) = 1$ without loss of generality.

Lemma 3.8. *If $H \in \mathcal{A}$ and $M(H) = 1$, then*

$$m_H(s) \le \begin{cases} \dfrac{2\pi s + \sin(2\pi s)}{4\pi s}, & 0 < s \le \frac{1}{2}, \\[2mm] \dfrac{1}{2}, & s > \frac{1}{2}. \end{cases} \tag{3.6}$$

Proof. Fix $0 < s \le \frac{1}{2}$ and set $\lambda = 2\pi s/(2\pi s + \sin(2\pi s))$. By the Paley–Wiener theorem [27], the support condition on $\widehat{H}$ gives an entire extension of exponential type at most 2π, meaning that $|H(z)| \le C_\varepsilon e^{(2\pi+\varepsilon)|z|}$ for every $\varepsilon > 0$. The Fejér–Riesz/Krein factorisation [1] then gives $H = |g|^2$ for an entire function g of exponential type at most π, so $\widehat{g}$ is supported in $[-1/2, 1/2]$. Writing $\|\widehat{g}\|_2^2 = \int_{-1/2}^{1/2} |\widehat{g}(\xi)|^2\,d\xi$ for its squared L^2 norm, Plancherel gives $\|\widehat{g}\|_2^2 = M(H) = 1$.

Put $c_s(\xi) = \cos(2\pi s\xi)$ and $d_s(\xi) = \sin(2\pi s\xi)$. Fourier inversion and cancellation of the mixed terms give

$$H(s) + H(-s) = 2|\langle \widehat{g}, c_s\rangle|^2 + 2|\langle \widehat{g}, d_s\rangle|^2. \tag{3.7}$$

The two functions are orthogonal on $[-1/2, 1/2]$, and

$$\|c_s\|_2^2 = \frac{2\pi s + \sin(2\pi s)}{4\pi s} = \frac{1}{2\lambda}, \qquad \|d_s\|_2^2 = \frac{2\pi s - \sin(2\pi s)}{4\pi s} \le \frac{1}{2\lambda}.$$

Bessel's inequality therefore yields

$$\begin{aligned} \lambda\big(H(s) + H(-s)\big) &\le \frac{|\langle \widehat{g}, c_s\rangle|^2}{\|c_s\|_2^2} + \frac{|\langle \widehat{g}, d_s\rangle|^2}{\|d_s\|_2^2} \\ &\le \|\widehat{g}\|_2^2 = 1. \end{aligned}$$

Since $H(\pm s) \ge m_H(s)$, this gives $m_H(s) \le 1/(2\lambda) = (2\pi s + \sin(2\pi s))/(4\pi s)$. At $s = 1/2$ it gives $m_H(1/2) \le 1/2$, and $m_H(s) \le m_H(1/2)$ for $s \ge 1/2$. □

Lemma 3.9 (positivity criterion)**.** *Assume GRH. Let $H \in \mathcal{A}$ be admissible with $M(H) = 1$, let $0 < c \le 2$ and $\beta > 0$, and suppose $m_H(c\beta) > 0$. Then*

$$\mathbb{E}_\chi \sum_j H(c x_{\chi,j}) = \frac{1 + o(1)}{c}.$$

Moreover,

$$\bar{N}(\beta) \le \frac{1 + o(1)}{c\,m_H(c\beta)},$$

whose right-hand side has limit below 1 precisely when

$$c\,m_H(c\beta) > 1. \tag{3.8}$$

Under (3.8)*, for every sufficiently large prime q there is a character $\chi \in \mathcal{F}_q$ with $X_\chi > \beta$.*

Proof. The function $x \mapsto H(cx)$ is admissible and has Fourier support in $[-c, c]$, so the first asymptotic formula follows from Theorem 3.2. Each zero counted by $N_\chi(\beta)$ contributes at least $m_H(c\beta)$, while every other contribution is non-negative. This gives the bound for $\bar{N}(\beta)$. Under (3.8) that bound is below 1 for all sufficiently large q, and an average of non-negative integers below 1 forces some character to have $N_\chi(\beta) = 0$. □

Theorem 3.10 (sharpness of the 1/4 bound)**.** *Assume GRH. For every $0 < \beta < 1/4$ and every sufficiently large prime q, the support-2 one-level-density positivity argument proves that $X_\chi > \beta$ for some $\chi \in \mathcal{F}_q$. No choice of non-negative admissible test function with Fourier support in $[-2, 2]$ proves this conclusion by the same argument for any $\beta > 1/4$.*

Proof. Symmetrising and dilating any non-negative admissible test reduces it to $x \mapsto H(cx)$ with $H \in \mathcal{A}$, $M(H) = 1$ and $0 < c \le 2$. Suppose the test satisfies the strict limiting criterion (3.8), namely $c\, m_H(c\beta) > 1$. Put $s = c\beta$.

If $0 < s \le \frac{1}{2}$, Lemma 3.8 and (3.8) give

$$\beta = \frac{s}{c} < s\, m_H(s) \le \frac{2\pi s + \sin(2\pi s)}{4\pi} \le \frac{1}{4},$$

because $2\pi s + \sin(2\pi s)$ increases to π on $[0, 1/2]$. If $s \ge 1/2$, then $c \le 2$ and $m_H(s) \le 1/2$ make $c\, m_H(s) > 1$ impossible. No admissible shape can therefore pass 1/4.

To obtain every $\beta < 1/4$, take the cosine direction $\widehat{g}(\xi) \propto \cos(\pi\xi)$, an equality case in the Bessel estimate at $s = 1/2$. It produces Zhao's function

$$H_{1/2}(x) = \frac{8}{\pi^2} \left(\frac{\cos(\pi x)}{1 - 4x^2} \right)^2.$$

It belongs to $\mathcal{A}$ and has $M(H_{1/2}) = 1$. Its apparent poles at $x = \pm 1/2$ are removable, and its minimum on $[-1/2, 1/2]$ is $1/2$, attained at the endpoints. Indeed, for $0 \le x \le 1/2$ the substitution $v = \pi(1/2 - x)$ reduces $H_{1/2}(x) \ge 1/2$ to

$$\pi \sin v \ge v(\pi - v),$$

which follows by integrating $\cos y \ge 1 - 2y/\pi$ on $[0, v]$. The inequality is strict for $|x| < 1/2$. Given $\beta < 1/4$, take $c = 2$ and $s = 2\beta < 1/2$. Then $c\, m_{H_{1/2}}(s) > 2 \cdot \frac{1}{2} = 1$, so Lemma 3.9 gives the required character. At $\beta = 1/4$, one has $s = 1/2$ and $c\, m_{H_{1/2}}(s) = 1$, so the leading asymptotic gives only $\bar{N}(1/4) \le 1 + o(1)$. Zhao's lower-order analysis [34, Theorem 6] gives $X_\chi > 1/4$ for some character at every sufficiently large prime. □

Rescaling identifies two equivalent analytic problems: *maximising* the interval minimum among mass-one test functions and *minimising* the mass of a test function that majorises 1 on the interval. Define

$$\rho(s) = \inf\Big\{ M(G) : G \in \mathcal{A},\ G \ge 1 \text{ on } [-s, s] \Big\}.$$

Then $\sup_{M(H)=1} m_H(s) = 1/\rho(s)$. For the normalised support interval $[-1, 1]$, $\rho(s)$ is determined by a Fourier-analytic extremal problem. Theorem 3.2 gives the admissible dilation range $0 < c \le 2$. For fixed s, condition (3.8) is feasible with some $c \le 2$ exactly when $\rho(s) < 2$, and

$$\sup_{\rho(s) < c \le 2} \frac{s}{c} = \frac{s}{\rho(s)}.$$

Hence the method's reach is

$$\sup_{s:\,\rho(s)<2} \frac{s}{\rho(s)} = \frac{1}{4}.$$

Lemma 3.8 and Zhao's function give $\rho(1/2) = 2$: at $s = 1/2$, the least mass of a test function majorising 1 on $[-s, s]$ reaches the arithmetic support boundary.

The half-period cosine appears in both optimality results, but for different structural reasons: on the min side it is the first Dirichlet eigenfunction of $-d^2/d\xi^2$ on the support interval, and on the max side it is the direction singled out by a rank-two evaluation functional. The coincidence of shapes at support 2 does not identify the two extremal problems.

3.7 Extended support

Extending Fourier support past 2 decreases the exact min-side sign-test threshold $\beta_{\min}(\Delta)$ and increases the Selberg max-side bound $\beta^{\mathrm{Sel}}_{\max}(\Delta)$ by the same amount. This symmetry concerns Selberg's construction, not an identification of the two optimal frontiers at general support. The theorem of Drappeau, Pratt and Radziwiłł extends the support past 2 after an additional average over the modulus.

Theorem 3.11 (Drappeau–Pratt–Radziwiłł [12])**.** *Let Φ be a fixed non-zero, non-negative smooth weight supported in $(\frac{1}{2}, 3)$, let $\Delta_0 = 2 + \frac{50}{1093}$, and let f be even and admissible with $\operatorname{supp} \widehat{f} \subset (-\Delta_0, \Delta_0)$. Then, as $Q \to \infty$,*

$$\sum_{q} \Phi\left(\frac{q}{Q}\right) \sum_{\substack{\chi \bmod q \\ \text{primitive}}} \sum_{\gamma_\chi} f\left(\gamma_\chi \frac{\log Q}{2\pi}\right) \;=\; \left(\int_{\mathbb{R}} f(x)\,dx\right) N_Q + o(Q^2), \tag{3.9}$$

where $N_Q = \sum_q \Phi(q/Q) \sum^*_\chi 1 \asymp Q^2$ is the total family weight.

Theorem 3.11 is unconditional. Because q/Q stays in the fixed compact support of Φ, one has $\log q / \log Q = 1 + O(1/\log Q)$, so replacing the theorem's $\log Q$ scaling by the $\log q$ convention of Section 2 changes every fixed threshold by $o(1)$. The support interval strictly contains $[-2, 2]$, and (3.9) is an asymptotic equality rather than a one-sided bound. It therefore supplies both the positivity needed on the min side and the upper bound needed on the max side. Corollary 3.12 uses GRH only to place zeros on the critical line, so that the sign of a test function at a zero records its position.

Corollary 3.12. *Assume GRH. Write $\mathcal{G}_Q = \{(q, \chi) : \Phi(q/Q) > 0,\ \chi \bmod q \text{ primitive}\}$. Then*

$$\limsup_{Q\to\infty} \min_{(q,\chi)\in\mathcal{G}_Q} X_\chi \le \frac{1}{2\Delta_0} = \frac{1093}{4472} \approx 0.2444,$$

$$\liminf_{Q\to\infty} \max_{(q,\chi)\in\mathcal{G}_Q} X_\chi \ge \beta^{\mathrm{Sel}}_{\max}(\Delta_0) = \frac{1}{2} - \frac{1}{2\Delta_0} = \frac{1143}{4472} \approx 0.2556.$$

Proof. Fix $1 < \Delta < \Delta_0$. For the minimum, the test function of Section 3.5 has positive mass whenever $\beta > 1/(2\Delta)$. By (3.9), its weighted zero sum is positive for all sufficiently large Q. Since the test is non-positive outside $(-\beta, \beta)$, some zero satisfies $|\gamma_\chi| \log Q/(2\pi) < \beta$. The corresponding character has $X_\chi < \beta \log q / \log Q = \beta + o(1)$. Let $\beta \downarrow 1/(2\Delta)$ and then $\Delta \uparrow \Delta_0$ to obtain the limsup bound.

For the maximum, fix $0 < \beta < \frac{1}{2} - 1/(2\Delta)$. The Selberg majorant has mass $2\beta + 1/\Delta < 1$, so (3.9) gives a weighted mean of the $\log Q$-scaled interval counts below 1 for all

sufficiently large Q. Hence some $(q,\chi) \in \mathcal{G}_Q$ has no zero with $|\gamma_\chi| \log Q/(2\pi) \le \beta$, and therefore $X_\chi > \beta \log q/\log Q = \beta + o(1)$. Letting $\beta \uparrow \frac{1}{2} - 1/(2\Delta)$ gives

$$\liminf_{Q\to\infty} \max_{(q,\chi)\in\mathcal{G}_Q} X_\chi \ge \frac{1}{2} - \frac{1}{2\Delta}.$$

Then let $\Delta \uparrow \Delta_0$. □

The min-side value is the exact sign-test threshold at the support endpoint Δ_0. The max-side value $1143/4472$ is an explicit Selberg bound, not an optimal one: optimising the interval majorant using Littmann's construction [24] yields a strictly stronger max-side bound.

The extension is conductor-averaged: the character satisfying the bound is found at some modulus $q \asymp Q$, not at each fixed q. Its proof estimates the off-diagonal prime sums directly by the dispersion method and bounds for sums of Kloosterman sums.

3.8 The limits of the method

Fourier support Let f be fixed and admissible with $\operatorname{supp} \widehat{f} \subset (-\Delta, \Delta)$. The explicit formula contains prime powers up to $x = q^\Delta$, and character orthogonality selects $n \equiv 1 \pmod q$. Increasing Δ therefore requires control of $\psi(x;q,1) - x/\varphi(q)$, where $\psi(x;q,a) = \sum_{n\le x,\ n\equiv a\,(q)} \Lambda(n)$. Friedlander and Granville's replacement for Montgomery's conjecture [14, Conjecture 1(b)] predicts that for every $\varepsilon > 0$, uniformly for $(a,q) = 1$ and $q < x$,

$$\psi(x;q,a) = \frac{x}{\varphi(q)} + O_\varepsilon\left(\frac{x^{1/2+\varepsilon}}{q^{1/2}}\right).$$

This replaces Montgomery's stronger estimate [25, (15.9)], which they disproved [14, Corollary 2]. GRH gives only $O(x^{1/2}\log^2(qx))$, without the factor $q^{-1/2}$. Write $E(t;q,a) = \psi(t;q,a) - t/\varphi(q)$. Partial summation transfers the conjectural bound for E to the weighted prime sum in Section 3.1 and gives, when $x = q^\Delta$,

$$\text{weighted prime error} \ll_{f,\varepsilon} q^{-1/2}x^\varepsilon = q^{-1/2+\Delta\varepsilon}.$$

For any fixed Δ, choosing $\varepsilon < 1/(2\Delta)$ makes this $o(1)$. The one-level density would therefore hold at every finite support.

For a conductor average, uniform control of each $E(x;q,1)$ is stronger than necessary. Writing $\psi(x) = \sum_{n\le x}\Lambda(n)$, Fiorilli and Miller [13] instead assume, for some $0 < \delta < 1$,

$$\sum_{q\sim Q}\left|\psi(x;q,1) - \frac{\psi(x)}{\varphi(q)}\right|^2 \ll Q^{\delta-1}\sum_{q\sim Q}\sum_{(a,q)=1}\left|\psi(x;q,a) - \frac{\psi(x)}{\varphi(q)}\right|^2,$$

where $q \sim Q$ means $Q/2 < q \le Q$. The double sum is the total residue-class variance; character Parseval identifies it with a mean square of twisted prime sums, for which the large sieve [8] gives a general bound. Under GRH, the Goldston–Vaughan variance estimate [15] then yields the one-level density for

$$\operatorname{supp} \widehat{f} \subset (-4+2\delta, 4-2\delta).$$

Thus the arbitrary-support conclusion is conditional on the Friedlander–Granville prime-distribution conjecture, while the Fiorilli–Miller extension assumes GRH together with the de-averaging comparison. Theorem 3.11 obtains the smaller support endpoint $2+50/1093$ unconditionally by estimating the off-diagonal prime sums directly through the dispersion method and bounds for Kloosterman sums. Even arbitrary support would not remove every limitation of the first-moment method: the max-side threshold saturates at $1/2$.

Max-side first-moment saturation The Selberg support-Δ majorant gives

$$\bar{N}(\beta) \le 2\beta + \frac{1}{\Delta} + o(1).$$

Since each $N_\chi(\beta)$ is a non-negative integer, $\bar{N}(\beta) < 1$ forces $N_\chi(\beta) = 0$ for some character. For fixed β and Δ, the majorant bound is below 1 for all sufficiently large q whenever

$$2\beta + \frac{1}{\Delta} < 1 \quad \Longleftrightarrow \quad \beta < \frac{1}{2} - \frac{1}{2\Delta} = \beta^{\mathrm{Sel}}_{\max}(\Delta).$$

As $\Delta \to \infty$, Selberg's bound approaches $1/2$, so density at every fixed support would give every fixed half-length $\beta < 1/2$. Conversely, any integrable majorant $M \ge \mathbf{1}_{[-\beta,\beta]}$ satisfies

$$\int_{\mathbb{R}} M(x)\,dx \ge \int_{-\beta}^{\beta} 1\,dx = 2\beta.$$

Thus no such majorant has mass below 1 when $\beta \ge 1/2$, regardless of its Fourier support. At $\beta = 1/2$, the configuration $N_\chi(\beta) = 1$ for every character is compatible with mean count 1. Passing $1/2$ requires information beyond this first-moment majorant argument.

Min-side approximation error The min side has no analogous saturation. Since each $N_\chi(\beta)$ is a non-negative integer, $\bar{N}(\beta) > 0$ implies that some character has a zero in $[-\beta, \beta]$, and hence $\min_{\chi \in \mathcal{F}_q} X_\chi \le \beta$. If the one-level density remained valid uniformly for half-lengths $\beta(q)$ shrinking with q, with

$$\bar{N}(\beta(q)) = 2\beta(q) + o(\beta(q)),$$

then taking $\beta(q) \asymp q^{-1}$ would give $\min_{\chi \in \mathcal{F}_q} X_\chi \ll q^{-1}$, matching the conjectural scale of the family minimum (Section 5.4). At fixed Fourier support, however, the Selberg minorant gives

$$\bar{N}(\beta) \ge 2\beta - \frac{1}{\Delta} + o(1).$$

This lower bound becomes positive only when $\beta > 1/(2\Delta)$, and the full-class optimality result of Section 3.5 shows that no admissible sign test can obtain a positive mean below the same threshold. Fixed Fourier support therefore prevents the first-moment argument from detecting intervals whose half-length tends to zero. On both sides, the resulting constants belong to the first moment of a band-limited statistic rather than to the family extremes themselves. Theorem 4.1 proves a larger max-side constant from different information.

4 A max-side bound via Wirtinger's inequality

Theorem 4.1 is the main result of this dissertation. To the author's knowledge, it is the first application of a Wirtinger–moment argument to maximising the scaled lowest zero in a fixed-modulus family of Dirichlet L-functions. Under GRH, it raises the best known asymptotic max-side constant from Zhao's $1/4$ [34, Theorem 6] to $\sqrt{3}/4 \approx 0.433$.

Theorem 4.1. *Fix $0 < \beta < \sqrt{3}/4$. For every sufficiently large prime q there is an even non-principal character $\chi \bmod q$ whose L-function has no zero ρ with $\operatorname{Re}\rho = \frac{1}{2}$ and $|\operatorname{Im}\rho| \le 2\pi\beta/\log q$.*

Theorem 4.1 is unconditional. Under GRH every nontrivial zero has $\operatorname{Re}\rho = \frac{1}{2}$, so the same character satisfies $X_\chi > \beta$, and therefore

$$\liminf_{\substack{q\to\infty \\ q \text{ prime}}} \max_{\chi\in\mathcal{F}_q} X_\chi \;\geq\; \frac{\sqrt{3}}{4} \approx 0.433.$$

The empty interval is two-sided, so it also gives the same lower bound for the one-sided variable X_χ^+. The character exists at every sufficiently large prime, not merely along a subsequence of moduli, and belongs to the even subfamily.

The method adapts Hall's treatment of gaps between zeros of the Riemann zeta function in the t-aspect, which combined Wirtinger's inequality with moments of ζ and ζ' [16]. In the present q-aspect setting, the frequency origin from which the derivative is measured is chosen to minimise the family derivative moment.

4.1 From zero counts to the size of L

Set

$$F_\chi(t) = L\left(\frac{1}{2} + it, \chi\right), \qquad h = \frac{2\pi\beta}{\log q},$$

and restrict to the $V = (q-3)/2$ even non-principal characters, whose common parity fixes the gamma factor. Write $\sum_\chi'$ for this family sum. Suppose every such F_χ vanishes at some $t_\chi \in [-h, h]$. For any real α, the twist $e^{i\alpha t}F_\chi(t)$ has the same zeros and modulus as F_χ, so Lemma 4.2 gives

$$\int_{-h}^{h} |F_\chi(t)|^2\,dt \leq \left(\frac{4h}{\pi}\right)^2 \int_{-h}^{h} |F_\chi'(t) + i\alpha F_\chi(t)|^2\,dt.$$

The strongest bound comes from minimising the family moment of $F_\chi' + i\alpha F_\chi$ relative to that of F_χ.

After suppressing the smooth weights and replacing the conductor-and-gamma factor by $e^{-it\log q}$, the length-$\sqrt{q}$ approximate functional equation [20] has schematic t-dependence

$$F_\chi(t) \approx \underbrace{\sum_{n\leq\sqrt{q}} \frac{\chi(n)}{\sqrt{n}} e^{-it\log n}}_{\text{direct piece}} + \underbrace{\varepsilon(\chi)e^{-it\log q}\sum_{n\leq\sqrt{q}} \frac{\overline{\chi(n)}}{\sqrt{n}} e^{it\log n}}_{\text{reflected piece}}.$$

For $1 \leq n \leq \sqrt{q}$, the corresponding frequencies are

$$u_{\text{dir}} = \log n \in \left[0, \frac{1}{2}\log q\right], \qquad u_{\text{ref}} = \log q - \log n \in \left[\frac{1}{2}\log q, \log q\right].$$

The suppressed weights and remaining gamma-factor dependence contribute lower-order t-variation, so these intervals describe the dominant oscillation rather than literal Fourier support.

Multiplication by $e^{i\alpha t}$ replaces a frequency u by $u-\alpha$, and differentiation multiplies the squared amplitude of that mode by $(u-\alpha)^2$. The squared Dirichlet coefficients have size $1/n$, while $dn/n = du$ under $u = \log n$; after character orthogonality selects the diagonal, the combined frequency distribution is therefore asymptotically uniform on $[0, \log q]$. The value moment measures its total mass, while the twisted derivative moment measures its second moment about α:

$$\frac{\sum_\chi{}' |F_\chi'(t) + i\alpha F_\chi(t)|^2}{\sum_\chi{}' |F_\chi(t)|^2} \sim \frac{\displaystyle\int_0^{\log q} (u-\alpha)^2\, du}{\displaystyle\int_0^{\log q} du} = \frac{\log^2 q}{12} + \left(\alpha - \frac{1}{2}\log q\right)^2 \geq \frac{\log^2 q}{12}.$$

Equality holds at $\alpha = \frac{1}{2}\log q$. This centred variance is the ratio needed in the summed Wirtinger inequality.

Including the gamma factor gives the centre at $t = 0$ as $\alpha_0 = \mathcal{L}(0)/2 = \frac{1}{2}\log q + O(1)$, where $\mathcal{L}$ is the even-parity log-conductor defined in Section 4.4. Summing the Wirtinger inequalities and applying Proposition 4.5 gives

$$\begin{aligned} 1 &\leq \left(\frac{4h}{\pi}\right)^2 \frac{\displaystyle\int_{-h}^{h} \sum_\chi{}' |F_\chi'(t) + i\alpha_0 F_\chi(t)|^2\, dt}{\displaystyle\int_{-h}^{h} \sum_\chi{}' |F_\chi(t)|^2\, dt} \\ &= \left(\frac{4h}{\pi}\right)^2 \left(\frac{\mathcal{L}(0)^2}{12} + o(\log^2 q)\right) \\ &= \frac{16\beta^2}{3} + o(1), \end{aligned}$$

which is impossible when $\beta < \sqrt{3}/4$.

The argument uses the centred second moment of the L-value frequencies, information absent from the band-limited zero statistics of Section 3. Theorem 3.10 therefore does not constrain this bound.

4.2 The Wirtinger inequality

Lemma 4.2. *Let $h > 0$, and for a continuous function g on $[-h, h]$ write $\|g\|^2 = \int_{-h}^{h} |g(t)|^2\, dt$. If $g\colon [-h, h] \to \mathbb{C}$ is C^1 and $g(c) = 0$ for some $c \in [-h, h]$, then*

$$\|g\|^2 \;\leq\; \left(\frac{4h}{\pi}\right)^2 \|g'\|^2.$$

Proof. See Hardy, Littlewood and Pólya [17, §7.7]. □

The constant comes from the one-endpoint problem on each side of the zero. On an interval of length ℓ, with the function vanishing at one endpoint and free at the other, the lowest eigenvalue of $-d^2/dt^2$ is $(\pi/2\ell)^2$ and the first eigenfunction is $\sin(\pi t/2\ell)$. The corresponding inequality $\int |g|^2 \leq (2\ell/\pi)^2 \int |g'|^2$ also follows from the two-endpoint form in Hardy, Littlewood and Pólya [17] by even reflection across the free endpoint. Splitting $[-h, h]$ at c gives two intervals of length at most $2h$, hence the constant $(4h/\pi)^2$ in Lemma 4.2.

Remark 4.3 (an elementary substitute). The Wirtinger inequality can be replaced by the following direct Cauchy–Schwarz estimate, with a weaker constant. The fundamental theorem of calculus gives $g(t) = \int_c^t g'(\tau)\,d\tau$, and Cauchy–Schwarz applied to the factors $1 \cdot g'$ gives

$$|g(t)|^2 \le |t-c| \int_{-h}^{h} |g'|^2 = |t-c|\,\|g'\|^2,$$

where the integral over the interval with endpoints c and t was extended to $[-h,h]$. Integrating over $t \in [-h,h]$,

$$\|g\|^2 \le \|g'\|^2 \int_{-h}^{h} |t-c|\,dt = \tfrac{1}{2}\big[(c+h)^2 + (h-c)^2\big]\,\|g'\|^2 \le 2h^2\,\|g'\|^2,$$

the worst case being c at an endpoint. The constant $2h^2$ exceeds the sharp $(4h/\pi)^2 \approx 1.62\,h^2$ by the factor $\pi^2/8$, and it still proves Theorem 4.1, at the weaker threshold computed in Remark 4.9.

An eigenvalue of $-d^2/dt^2$ also governed the min-side optimiser of Section 3.5, but the logic here is reversed. There, the support restriction limits the frequencies available to the test function, and the lowest eigenvalue determines how small an interval the count method can prove to contain a zero. Here, an assumed zero in $[-h,h]$ forces a lower bound on the ratio $\|g'\|^2/\|g\|^2$, which is then compared with the family moments.

4.3 The twist and the choice of frequency origin

For any real α, the twisted function

$$F_{\chi,\alpha}(t) = e^{i\alpha t} F_\chi(t)$$

has the same modulus and the same zeros as F_χ, so $\|F_{\chi,\alpha}\| = \|F_\chi\|$ and $F_{\chi,\alpha}$ vanishes on $[-h,h]$ exactly when F_χ does, while its derivative is $F'_{\chi,\alpha} = e^{i\alpha t}(F'_\chi + i\alpha F_\chi)$. Lemma 4.2 applied to $F_{\chi,\alpha}$ therefore says that if F_χ vanishes on $[-h,h]$, then

$$\|F_\chi\|^2 \le \Big(\frac{4h}{\pi}\Big)^2 \|F'_\chi + i\alpha F_\chi\|^2 \qquad \text{for every } \alpha \in \mathbb{R}.$$

The freedom in α matters because differentiation weights frequencies. A mode $e^{-it\log n}$ of F_χ becomes a mode of frequency $\log n - \alpha$, and differentiation multiplies its amplitude by $-i(\log n - \alpha)$. For a single mode the squared-derivative-to-squared-value ratio is $(\log n - \alpha)^2$; after family averaging, the off-diagonal cross terms are lower order, so the twisted derivative moment is the second moment of the effective frequency distribution about α. It is minimised at the mean frequency. The asymptotically flat band $[0, \log q]$ has mean $\frac{1}{2}\log q$, and its centred second moment is $\int_{-\frac{1}{2}\log q}^{\frac{1}{2}\log q} u^2\,du = \frac{1}{12}\log^3 q$, a factor 4 below the endpoint moment $\frac{1}{3}\log^3 q$. This factor raises the untwisted threshold $\sqrt{3}/8 \approx 0.217$ to $\sqrt{3}/4$.

4.4 The two family second moments

Recall from Section 4.1 that $\sum_\chi'$ denotes a sum over the $V = \frac{q-3}{2}$ even non-principal characters modulo the prime q, and define

$$M_0(t) = \sum_\chi' |F_\chi(t)|^2, \qquad M_2(t) = \sum_\chi' |F'_\chi(t)|^2, \qquad T_\alpha(t) = \sum_\chi' |F'_\chi(t) + i\alpha F_\chi(t)|^2.$$

The summed form of Lemma 4.2 requires asymptotics for M_0 and T_α, uniformly for $t \in [-h, h]$. Write

$$\mathcal{L}(t) = \log \frac{q}{\pi} + \operatorname{Re}\, \psi\Big(\frac{1}{4} + \frac{it}{2}\Big)$$

for the even-parity log-conductor, that is, twice the real part of the logarithmic derivative, at $s = \frac{1}{2} + it$, of the conductor-and-gamma prefactor $(q/\pi)^{s/2}\Gamma(s/2)$ of $\Lambda(s, \chi)$. Up to a bounded factor $e^{\mathcal{L}(t)}$ is the analytic conductor of $L(s, \chi)$ at height t [20], here in its even-parity case, so $\mathcal{L}$ is the refinement of $\log q$ that the gamma factor forces; at leading order $\mathcal{L}(t)/2\pi$ is the density of zeros at height t, the reciprocal of the mean spacing of Section 2.3. All that is used is $\mathcal{L}(t) = \log q + O(1)$ and, since $\mathcal{L}$ is even and smooth near 0, $\mathcal{L}(t) = \mathcal{L}(0) + O(t^2)$ for $t \in [-h, h]$.

Expanding the square in T_α gives

$$T_\alpha = M_2 + \alpha^2 M_0 + 2\alpha M_{\mathrm{mix}}, \qquad M_{\mathrm{mix}}(t) = \sum_{\chi}{}' \operatorname{Re}\big(L'\overline{L}\big)\Big(\tfrac{1}{2} + it\Big),$$

which produces three family averages. The functional equation expresses the mixed moment in terms of M_0.

Lemma 4.4. *For every prime q and every real t,*

$$M_{\mathrm{mix}}(t) = -\frac{\mathcal{L}(t)}{2}\, M_0(t).$$

Proof. The functional equation gives $|\Lambda(\sigma + it, \chi)| = |\Lambda(1 - \sigma - it, \overline{\chi})|$. Since the Dirichlet coefficients of $L(s, \overline{\chi})$ are the conjugates of those of $L(s, \chi)$ and the prefactor of Λ is real on the real axis, one has the reflection identity $\Lambda(\overline{s}, \overline{\chi}) = \overline{\Lambda(s, \chi)}$, which converts the right side into $|\Lambda(1 - \sigma + it, \chi)|$. Each summand of $\sum'_\chi |\Lambda(\sigma + it, \chi)|^2$ is therefore symmetric in σ about $\sigma = \frac{1}{2}$, hence so is the sum, and its σ-derivative vanishes at $\sigma = \frac{1}{2}$. It remains to unpack that vanishing derivative. Write $\Lambda(s, \chi) = g(s)L(s, \chi)$ with $g(s) = (q/\pi)^{s/2}\Gamma(s/2)$, the case $\mathfrak{a} = 0$ of the prefactor. Along a horizontal line $\partial/\partial\sigma$ is differentiation in s, so for any function holomorphic in s one has $\partial_\sigma |F|^2 = 2\operatorname{Re}(\overline{F}F')$, and with $\Lambda' = g'L + gL'$,

$$\frac{\partial}{\partial\sigma}|\Lambda|^2 = 2\operatorname{Re}\big(\overline{\Lambda}\Lambda'\big) = 2\operatorname{Re}\big(|L|^2\,\overline{g}g' + |g|^2\,\overline{L}L'\big) = 2|g|^2\Big[\operatorname{Re}\frac{g'}{g}\cdot|L|^2 + \operatorname{Re}\big(L'\overline{L}\big)\Big],$$

every function being evaluated at $s = \sigma + it$. Summing over the family and setting $\sigma = \frac{1}{2}$, the left side vanishes and the factor $2|g(\frac{1}{2} + it)|^2$ is non-zero, so

$$\operatorname{Re}\frac{g'}{g}\Big(\tfrac{1}{2} + it\Big)\, M_0(t) + M_{\mathrm{mix}}(t) = 0.$$

Finally $\log g(s) = \frac{s}{2}\log(q/\pi) + \log\Gamma(s/2)$ gives $g'/g = \frac{1}{2}\log(q/\pi) + \frac{1}{2}\psi(s/2)$, whose real part at $s = \frac{1}{2} + it$ is $\frac{1}{2}\mathcal{L}(t)$. □

In (4.1) and (4.2) every moment is evaluated at the same t, and $\mathcal{L}$ abbreviates $\mathcal{L}(t)$. Substituting $M_{\mathrm{mix}} = -\frac{\mathcal{L}}{2}M_0$ and completing the square,

$$T_\alpha = M_2 - \frac{\mathcal{L}^2}{4}M_0 + \Big(\alpha - \frac{\mathcal{L}}{2}\Big)^2 M_0, \tag{4.1}$$

so the pointwise optimal origin is $\alpha = \mathcal{L}(t)/2$, confirming the heuristic of Section 4.3 with the log-conductor in place of $\log q$. The variable-phase twist $e^{i\theta(t)}F_\chi(t)$ with $\theta'(t) = \mathcal{L}(t)/2$

realises this optimum exactly and preserves both modulus and zeros. For the linear twist used in the moment calculation, set

$$\alpha_0 = \frac{\mathcal{L}(0)}{2},$$

so $e^{i\alpha_0 t}$ agrees to first order at $t = 0$, up to a constant factor, with the variable phase. A constant α leaves the arithmetic input as the three fixed moments M_0, M_2 and T_{α_0}, with the mismatch measured by (4.2). Setting $\alpha = \alpha_0$ in (4.1) gives

$$T_{\alpha_0}(t) = M_2(t) - \frac{\mathcal{L}(t)^2}{4} M_0(t) + \Big(\frac{\mathcal{L}(0) - \mathcal{L}(t)}{2}\Big)^2 M_0(t), \tag{4.2}$$

whose last term is negligible for $t \in [-h, h]$ since $\mathcal{L}(t) - \mathcal{L}(0) = O(t^2)$. The arithmetic input therefore consists of asymptotics for M_0 and M_2 alone.

Proposition 4.5. *Fix $C > 0$ and let $\alpha_0 = \mathcal{L}(0)/2$. Uniformly for $|t| \le C/\log q$, as $q \to \infty$ through primes,*

$$M_0(t) = V\,\mathcal{L}(t)\Big(1 + O\Big(\tfrac{1}{\log q}\Big)\Big), \qquad T_{\alpha_0}(t) = \frac{V\,\mathcal{L}(t)^3}{12}\Big(1 + O\Big(\tfrac{1}{\log q}\Big)\Big).$$

Proposition 4.5 is unconditional. The identity $\frac{1}{12} = \frac{1}{3} - \frac{1}{4}$ is the variance calculation from Sections 4.1 and 4.3. The two-shift template comes from Bui and Milinovich [4, Lemmas 4.2 and 4.4 and Section 5]: their preliminary shifted formula allows $|\alpha|, |\beta| \le 2/\log q$ away from $\alpha + \beta = 0$, and their moment argument crosses the divisor on a smaller disc by holomorphic continuation, maximum modulus and Cauchy estimates. Their pole-cancelling polynomial is adapted from Conrey, Iwaniec and Soundararajan [6, Section 3], and their off-diagonal estimate ultimately invokes Iwaniec and Sarnak [21, Section 4]. Here the off-diagonal congruences are bounded directly, while nested discs give uniformity near the moving points $p_t = (it, -it)$. Only the leading terms are needed; all lower-order terms lie within the relative error $O(1/\log q)$.

Proof of Proposition 4.5. Throughout, q is a large prime, $\delta = 1/\log q$, $C \ge 1$ (replace C by $\max(C, 1)$ if necessary), $0 < \varepsilon < \frac{1}{6}$ is fixed, and implied constants depend only on C and ε.

Step 1: the two-shift sum and the divisor $w = 0$. For complex shifts μ, ν set

$$S(\mu, \nu) = \sum_{\chi}{}' L(\tfrac{1}{2} + \mu, \chi)\, L(\tfrac{1}{2} + \nu, \overline{\chi}), \qquad w = \mu + \nu.$$

Since conjugating the Dirichlet coefficients of $L(s, \chi)$ produces those of $L(s, \overline{\chi})$, one has $\overline{L(s, \chi)} = L(\overline{s}, \overline{\chi})$, and likewise for L'. Hence at the point $p_t = (it, -it)$ with t real,

$$S(p_t) = M_0(t), \qquad \partial_\mu \partial_\nu S(p_t) = M_2(t),$$

the mixed derivative placing one L' in each factor. Every evaluation point p_t lies on the divisor $w = 0$, and M_2 is invisible from the divisor alone. Differentiating the one-variable restriction $t \mapsto S(it, -it)$ moves only in the tangential direction $\partial_\mu - \partial_\nu$, whereas

$$\partial_\mu \partial_\nu = \tfrac{1}{4}\big[(\partial_\mu + \partial_\nu)^2 - (\partial_\mu - \partial_\nu)^2\big]$$

contains the transverse one. Adding $c\,w^2$ to a function changes no value on the divisor but shifts the mixed derivative by $2c$, so S must be controlled and differentiated on a two-dimensional neighbourhood of p_t before restriction to $w = 0$.

Step 2: orthogonality and the shifted approximate functional equation. Orthogonality for the even family reads: for $(mn, q) = 1$,

$$\sum_{\chi}{}' \chi(m)\overline{\chi}(n) = \frac{q-1}{2}\Big(\mathbf{1}_{m\equiv n\ (q)} + \mathbf{1}_{m\equiv -n\ (q)}\Big) - 1.$$

Indeed $\sum_{\chi \bmod q} \chi(m)\overline{\chi}(n) = (q-1)\mathbf{1}_{m\equiv n\ (q)}$; inserting the even projector $\frac{1}{2}(1+\chi(-1))$ and using $\chi(-1)\overline{\chi}(n) = \overline{\chi}(-n)$ produces the two congruences, and subtracting the principal character, which is even and equals 1 on $(mn, q) = 1$, gives the -1.

To expand each summand of S, fix the Gaussian weight e^{s^2} and, for $w \neq 0$, the polynomial

$$Q_w(s) = 1 - \frac{4s^2}{w^2},$$

which is even, has $Q_w(0) = 1$, and vanishes at the two zeta-pole locations $s = \pm\frac{w}{2}$. Define the kernels

$$\mathcal{W}^{\pm}(x) = \frac{1}{2\pi i}\int_{(1)} e^{s^2} Q_w(s)\,\kappa^{\pm}(s)\,x^{-s}\,\frac{ds}{s}, \qquad \kappa^{\pm}(s) = \frac{\Gamma\big(\frac{1/2\pm\mu+s}{2}\big)\Gamma\big(\frac{1/2\pm\nu+s}{2}\big)}{\Gamma\big(\frac{1/2+\mu}{2}\big)\Gamma\big(\frac{1/2+\nu}{2}\big)},$$

where the dependence of $\kappa^{\pm}$ and $\mathcal{W}^{\pm}$ on the shifts is suppressed in the notation. Then, for even primitive χ and $\mu \neq -\nu$ with $|\mu|, |\nu| \le 6C\delta$,

$$\begin{aligned} L(\tfrac{1}{2}+\mu, \chi)L(\tfrac{1}{2}+\nu, \overline{\chi}) = {} & \sum_{m,n\ge 1} \frac{\chi(m)\overline{\chi}(n)}{m^{1/2+\mu}n^{1/2+\nu}}\,\mathcal{W}^{+}\Big(\frac{\pi mn}{q}\Big) \\ & + \Big(\frac{q}{\pi}\Big)^{-w} \sum_{m,n\ge 1} \frac{\overline{\chi}(m)\chi(n)}{m^{1/2-\mu}n^{1/2-\nu}}\,\mathcal{W}^{-}\Big(\frac{\pi mn}{q}\Big). \end{aligned} \tag{4.3}$$

Identity (4.3) is the even-character case of the shifted approximate functional equation in [4, Lemma 4.2]. The shifts here range over a box of radius a fixed multiple of δ containing the moving points p_t, and the dependence of their factor $H = Q_w$ on w is retained explicitly through the bound $1 + |w|^{-2}$. For $c \in \{1, -1\}$ set

$$I_{(c)} = \frac{1}{2\pi i}\int_{(c)} e^{s^2} Q_w(s)\,\Lambda(\tfrac{1}{2}+\mu+s, \chi)\,\Lambda(\tfrac{1}{2}+\nu+s, \overline{\chi})\,\frac{ds}{s}.$$

Both completed functions are entire because χ is primitive and non-principal. The only pole between the contours is therefore the simple pole at $s = 0$, and $Q_w(0) = 1$; the residue theorem gives

$$I_{(1)} - I_{(-1)} = \Lambda(\tfrac{1}{2}+\mu, \chi)\,\Lambda(\tfrac{1}{2}+\nu, \overline{\chi}).$$

On the line $\operatorname{Re} s = -1$ apply the functional equation $\Lambda(s, \chi) = \varepsilon(\chi)\Lambda(1-s, \overline{\chi})$ to each factor; the two root numbers pair to $\varepsilon(\chi)\varepsilon(\overline{\chi}) = \chi(-1) = 1$, leaving

$$\Lambda(\tfrac{1}{2}+\mu+s, \chi)\,\Lambda(\tfrac{1}{2}+\nu+s, \overline{\chi}) = \Lambda(\tfrac{1}{2}-\mu-s, \overline{\chi})\,\Lambda(\tfrac{1}{2}-\nu-s, \chi).$$

Now substitute $s \to -s$, which carries $\operatorname{Re} s = -1$ back to $\operatorname{Re} s = 1$ with reversed orientation, while the integrand is unchanged in form: e^{s^2} and $Q_w(s)$ are even, and the

differential ds/s is invariant, numerator and denominator changing sign together. The one sign therefore comes from the orientation alone:

$$I_{(-1)} = -\,I^{-}_{(1)},$$

where I^{-} denotes I with $(\mu,\nu,\chi,\overline{\chi})$ replaced by $(-\mu,-\nu,\overline{\chi},\chi)$. The orientation minus meets the minus of the residue identity, so the two integrals add:

$$\Lambda(\tfrac12+\mu,\chi)\,\Lambda(\tfrac12+\nu,\overline{\chi}) = I_{(1)} + I^{-}_{(1)}.$$

On $\operatorname{Re} s = 1$ both Dirichlet series converge absolutely, so expanding them and integrating term by term evaluates $I_{(1)}$ and $I^{-}_{(1)}$ as double sums over m,n. Finally divide the identity by the prefactor $(q/\pi)^{(1+w)/2}\Gamma\big(\frac{1/2+\mu}{2}\big)\Gamma\big(\frac{1/2+\nu}{2}\big)$ that unpacking $\Lambda =$ prefactor $\times L$ produces on the left: the left side becomes $L(\frac12+\mu,\chi)L(\frac12+\nu,\overline{\chi})$, and the normalising denominators of $\kappa^{\pm}$ record the division on the right. In the first term $(q/\pi)^{(1+w+2s)/2}$ over $(q/\pi)^{(1+w)/2}$ leaves $(q/\pi)^{s}$, which combines with $(mn)^{-s}$ into x^{-s} at $x=\pi mn/q$, while in the second term the negated shifts leave the extra factor $(q/\pi)^{-w}$, which is (4.3).

The kernels obey, for $0<|w|\le 12C\delta$ and any fixed $A\ge 1$,

$$\mathcal{W}^{\pm}(x) \;\ll_A\; \big(1+|w|^{-2}\big)\,\min\big(1,\,x^{-A}\big).$$

On any vertical line $|Q_w(s)| \le 1+4|s|^2/|w|^2$, Stirling's formula bounds $\kappa^{\pm}$ polynomially, and the Gaussian decays faster than both; shifting the contour to $\operatorname{Re} s = A$ (no singularities lie to the right) gives the decaying bound, while shifting to $\operatorname{Re} s = -\frac14$, legitimate since the numerator gamma poles sit at $\operatorname{Re} s \le -\frac12 + 6C\delta$, picks up the bounded residue $\kappa^{\pm}(0)$ at $s=0$ and gives boundedness. The factor $1+|w|^{-2}$ arises from Q_w and records the proximity to the divisor $w=0$.

Step 3: evaluation away from the divisor. Average the approximate functional equation over the family and insert orthogonality; terms with $q \mid mn$ vanish outright since then $\chi(m)\overline{\chi}(n)=0$. Three kinds of contribution arise.

For the diagonal $m=n$: the congruence $m\equiv n\ (q)$ contains the exact diagonal with coefficient $\frac{q-1}{2}$, while the paired congruence $m\equiv -n\ (q)$ admits no diagonal solution, since $m=n$ would force $q\mid 2n$ with $(n,q)=1$ and q an odd prime. The diagonal of the first sum, unfolded and summed under the integral, is

$$\frac{q-1}{2}\sum_{(n,q)=1}\frac{\mathcal{W}^{+}(\pi n^2/q)}{n^{1+w}} = \frac{q-1}{2}\cdot\frac{1}{2\pi i}\int_{(1)} e^{s^2}Q_w(s)\,\kappa^{+}(s)\Big(\frac{q}{\pi}\Big)^{s}\zeta_q(1+w+2s)\,\frac{ds}{s},$$

$$\zeta_q(s) = \zeta(s)\big(1-q^{-s}\big).$$

Shift the contour to $\operatorname{Re} s = -\frac12+\varepsilon$. The pole of $1/s$ at $s=0$ has residue $\kappa^{+}(0)\,\zeta_q(1+w) = \zeta_q(1+w)$, while the zero $Q_w(-\frac{w}{2})=0$ cancels the pole of $\zeta_q(1+w+2s)$ at $s=-\frac{w}{2}$.

On the new line, $|(q/\pi)^{s}| \ll q^{-1/2+\varepsilon}$ and the zeta argument has real part $2\varepsilon + O(C\delta)$. The convexity bound $\zeta(\sigma+i\tau)\ll(1+|\tau|)^{1/2+\varepsilon}$ applies [20], the zeta pole remains a fixed distance away, the numerator gamma poles lie to the left, and the Gaussian absorbs all polynomial growth. The diagonal contribution from the first sum is therefore $\frac{q-1}{2}\,\zeta_q(1+w)+O\big(q^{1/2+\varepsilon}(1+|w|^{-2})\big)$. The diagonal contribution from the second sum, whose zeta pole at $s=+\frac{w}{2}$ is cancelled by the other zero of Q_w, is $\frac{q-1}{2}\,(q/\pi)^{-w}\kappa^{-}(0)\,\zeta_q(1-w)$ up to the same error.

For the off-diagonal congruences $m\equiv\pm n\ (q)$ with $m\neq n$, put $B(w)=1+|w|^{-2}$ and, by symmetry, let n be the smaller variable. The larger variable has the form $m=kq\mp n$

with $k \geq 1$. If $m = kq + n$, then $m \geq kq$; if $m = kq - n$, then $n \leq m$ implies $m \geq kq/2$. Hence $m \geq q/2$ and $mn \geq nkq/2$ in both cases. Since $|\mathrm{Re}\, w| \leq 12C/\log q$, one has $|(q/\pi)^{-w}| \leq e^{12C+o(1)}$; the external factor in the reflected sum is therefore $O_C(1)$. The same kernel estimate holds for $\mathcal{W}^-$, so the argument applies to both sums in (4.3).

First suppose $mn \geq q^{1+\varepsilon}$. Then $nk \gg q^{\varepsilon}$. Indeed, if $m + n = kq$, then $m \leq kq$ and $nk \geq mn/q \geq q^{\varepsilon}$; if $m = n + kq$ and $nk \leq \frac{1}{2}q^{\varepsilon}$, then

$$mn = n^2 + nkq \leq \frac{1}{4}q^{2\varepsilon} + \frac{1}{2}q^{1+\varepsilon} < q^{1+\varepsilon},$$

a contradiction for large q. The coefficient from character orthogonality is $O(q)$, the kernel bound with $A = 3$ contributes $B(w)q^3(mn)^{-3}$, and the shifted Dirichlet coefficients contribute $(mn)^{-1/2+\varepsilon}$. Denoting the total far-range contribution by E_{far} and setting $r = nk$ gives

$$\begin{aligned} E_{\mathrm{far}} &\ll B(w)q^4 \sum_{nk \gg q^{\varepsilon}} (nkq)^{-7/2+\varepsilon} \\ &\ll B(w)q^{1/2+\varepsilon} \sum_{r \gg q^{\varepsilon}} d(r) r^{-7/2+\varepsilon} \\ &\ll B(w)q^{1/2+\varepsilon}(q^{\varepsilon})^{-5/2+2\varepsilon} \\ &\ll B(w)q^{1/2}, \end{aligned}$$

where $d(r)$ counts the factorisations $r = nk$, and $d(r) \ll r^{\varepsilon}$ produces the exponent $-5/2 + 2\varepsilon$ in the tail factor. Since $q^{\varepsilon}(q^{\varepsilon})^{-5/2+2\varepsilon} = q^{-3\varepsilon/2+2\varepsilon^2} \leq 1$ for $\varepsilon < 1/6$, the far range is $O(B(w)q^{1/2})$.

Now suppose $mn \leq q^{1+\varepsilon}$. Combining this with $mn \geq nkq/2$ gives

$$\frac{1}{2}nkq \leq mn \leq q^{1+\varepsilon} \quad \Longrightarrow \quad nk \leq 2q^{\varepsilon}.$$

Since $n, k \geq 1$, it follows that

$$n \leq nk \leq 2q^{\varepsilon}, \qquad k \leq \frac{2q^{\varepsilon}}{n}.$$

The shifts cost only $O_C(1)$ in this range, and the bounded part of the kernel estimate gives the near-range contribution

$$\begin{aligned} E_{\mathrm{near}} &\ll qB(w) \sum_{n \leq 2q^{\varepsilon}} \frac{1}{\sqrt{n}} \sum_{k \leq 2q^{\varepsilon}/n} \frac{1}{\sqrt{kq}} \\ &\ll B(w)q^{1/2+\varepsilon/2} \sum_{n \leq 2q^{\varepsilon}} \frac{1}{n} \\ &\ll B(w)q^{1/2+\varepsilon/2} \log q \\ &\ll B(w)q^{1/2+\varepsilon}. \end{aligned}$$

Finally, the -1 in the orthogonality relation produces the principal-character correction without the factor q. For $X = q^{1+\varepsilon}$, taking absolute values and absorbing the bounded shift factors gives the near-range bound

$$\begin{aligned} B(w) \sum_{mn \leq X} \frac{1}{\sqrt{mn}} &= B(w) \sum_{r \leq X} \frac{d(r)}{\sqrt{r}} \\ &\ll B(w)X^{1/2} \log X \\ &\ll B(w)q^{1/2+\varepsilon/2} \log q \\ &\ll B(w)q^{1/2+\varepsilon}, \end{aligned}$$

while kernel decay controls the tail by

$$\begin{aligned} B(w)q^3\sum_{r>X} d(r)r^{-7/2+\varepsilon} &\ll B(w)q^3X^{-5/2+2\varepsilon} \\ &\ll B(w)q^{1/2-\varepsilon/2+2\varepsilon^2} \\ &\ll B(w)q^{1/2}. \end{aligned}$$

The off-diagonal and principal-character terms are therefore $O\big(B(w)q^{1/2+2\varepsilon}\big)$.

Collecting the three contributions: for $|\mu|, |\nu| \le 6C\delta$ and $\mu \neq -\nu$,

$$S(\mu,\nu) = \frac{q-1}{2}\,\mathcal{D}(\mu,\nu) + \mathcal{R}(\mu,\nu), \qquad \mathcal{R} \;\ll\; q^{1/2+2\varepsilon}\big(1+|w|^{-2}\big),$$

$$\mathcal{D}(\mu,\nu) = \zeta_q(1+w) + \mathcal{X}(\mu,\nu)\,\zeta_q(1-w), \qquad \mathcal{X}(\mu,\nu) = \Big(\frac{q}{\pi}\Big)^{-w}\frac{\Gamma\big(\frac{1/2-\mu}{2}\big)\Gamma\big(\frac{1/2-\nu}{2}\big)}{\Gamma\big(\frac{1/2+\mu}{2}\big)\Gamma\big(\frac{1/2+\nu}{2}\big)},$$

where $\mathcal{X} = (q/\pi)^{-w}\kappa^-(0)$ is the functional-equation swap factor.

Step 4: extending across the divisor. Both estimates of Step 3 break down as $w \to 0$, the main terms through the poles of $\zeta_q(1 \pm w)$ and the errors through the factor $1 + |w|^{-2}$, while the objects they describe stay holomorphic there. S is a finite sum of entire functions, and $\mathcal{D}$ is holomorphic across the divisor: the gamma quotient in $\mathcal{X}$ telescopes to 1 at $\nu = -\mu$ and the (q/π)-power is 1 there, so $1-\mathcal{X}$ vanishes on the divisor and the two simple pole parts $\pm(1-\frac{1}{q})/w$ of $\zeta_q(1 \pm w)$ cancel:

$$\mathcal{D} = \Big(1-\frac{1}{q}\Big)\frac{1-\mathcal{X}}{w} + r_+(w) + \mathcal{X}\,r_-(w),$$

where $r_\pm(w) = \zeta_q(1 \pm w) \mp \big(1-\frac{1}{q}\big)/w$ subtracts each pole part; both are holomorphic on $|w| \le 12C\delta$ with $r_\pm, r'_\pm, r''_\pm = O(1)$, since the Laurent coefficients of ζ at its pole are absolute constants (the Stieltjes constants) and the factor $1-q^{-1\mp w}$ contributes derivatives of size $(\log q)^k/q$. Hence $\mathcal{R} = S - \frac{q-1}{2}\mathcal{D}$ is holomorphic on the bidisc $\{|\mu| \le 3C\delta\} \times \{|\nu| \le 4C\delta\}$.

Fix μ_* with $|\mu_*| \le 3C\delta$. On the circle $|\nu| = 4C\delta$ the distance to the divisor is $|w| = |\mu_* + \nu| \ge C\delta$, so Step 3 gives $\mathcal{R}(\mu_*,\nu) \ll q^{1/2+2\varepsilon}\log^2 q$. The gamma poles and the non-divisor zeta poles are a distance $\asymp 1$ away, so $\nu \mapsto \mathcal{R}(\mu_*,\nu)$ is holomorphic on a neighbourhood of the closed disc. The maximum principle gives

$$\mathcal{R} \;\ll\; q^{1/2+2\varepsilon}\log^2 q \qquad \text{on } \{|\mu| \le 3C\delta,\ |\nu| \le 4C\delta\}.$$

Derivatives follow from Cauchy's integral formula: for $|t| \le C\delta$ the bidisc of radii $C\delta$ centred at p_t stays inside $\{|\mu| \le 3C\delta,\ |\nu| \le 4C\delta\}$, so

$$\partial_\mu\partial_\nu\mathcal{R}(p_t) \;\ll\; (C\delta)^{-2}\,q^{1/2+2\varepsilon}\log^2 q \;\ll\; q^{1/2+2\varepsilon}\log^4 q.$$

Since $\varepsilon < \frac{1}{6}$, both bounds are $o(q)$ and hence fit within the required absolute errors $O(V)$ for M_0 and $O(V\mathcal{L}^2)$ for M_2. The apparent singularity belongs to the two-zeta decomposition, not to $\mathcal{R}$ itself.

Step 5: the main term at the evaluation point. Write $\mathcal{X} = e^{-w\,\mathcal{B}(\mu,\nu)}$. The function $\tau \mapsto \log\Gamma\big(\frac{1/2+\tau w-\mu}{2}\big) - \log\Gamma\big(\frac{1/2-\tau w+\mu}{2}\big)$ has derivative $w\,\Psi(\tau w-\mu)$, where

$$\Psi(y) = \frac{1}{2}\Big[\psi\Big(\frac{1/2+y}{2}\Big) + \psi\Big(\frac{1/2-y}{2}\Big)\Big],$$

and the fundamental theorem of calculus gives

$$\mathcal{B}(\mu,\nu) = \log\frac{q}{\pi} + \int_0^1 \Psi(\tau w - \mu)\,d\tau.$$

At the evaluation point, $\mathcal{B}(p_t) = \log\frac{q}{\pi} + \Psi(-it) = \mathcal{L}(t)$. This is real because $\psi(\overline{s}) = \overline{\psi(s)}$ pairs the two digamma terms of $\Psi(-it)$ into conjugates. Every first and second partial derivative of $\mathcal{B}$ on the bidisc is an average of Ψ' or Ψ'' at arguments within $O(\delta)$ of 0, hence is $O(1)$. The identity $(1-e^{-z})/z = \int_0^1 e^{-\tau z}\,d\tau$ gives

$$\frac{1-\mathcal{X}}{w} = \frac{1-e^{-w\mathcal{B}}}{w} = \mathcal{B}\int_0^1 e^{-\tau w\mathcal{B}}\,d\tau = \int_0^{\mathcal{B}} e^{-wu}\,du.$$

The two zeta terms thus combine into the Laplace transform $\int_0^{\mathcal{B}} e^{-wu}\,du$ of a flat frequency band of length $\mathcal{B}$. At $w=0$, evaluation gives $\int_0^{\mathcal{B}} du$, while the leading mixed derivative gives $\int_0^{\mathcal{B}} u^2\,du = \mathcal{B}^3/3$. Derivatives falling on $\mathcal{B}$ are lower order because $\mathcal{B} \asymp \mathcal{L}$ while its first and second partial derivatives are bounded. More precisely,

$$\partial_\mu\partial_\nu\Big[\mathcal{B}\int_0^1 e^{-\tau w\mathcal{B}}\,d\tau\Big](p_t) = \left[\frac{\mathcal{B}^3}{3} - \mathcal{B}\left(\partial_\mu\mathcal{B} + \partial_\nu\mathcal{B}\right) + \partial_\mu\partial_\nu\mathcal{B}\right]\Bigg|_{p_t}.$$

At p_t, $\tau w - \mu = -it$ is independent of τ, and the weights in $\partial_\mu\mathcal{B} + \partial_\nu\mathcal{B}$ integrate to $\int_0^1(2\tau-1)\,d\tau = 0$. Since $\partial_\mu\partial_\nu\mathcal{B} = O(1)$, the mixed derivative is $\mathcal{L}(t)^3/3 + O(1)$. The regular terms contribute $\partial_\mu\partial_\nu\, r_+(w) = r_+''(w) = O(1)$ and $\partial_\mu\partial_\nu\big[\mathcal{X}r_-\big](p_t) = \mathcal{L}^2 r_-(0) - 2\mathcal{L}r_-'(0) + r_-''(0) + O(1) = O(\mathcal{L}^2)$. Thus

$$\mathcal{D}(p_t) = \mathcal{L}(t) + O(1), \qquad \partial_\mu\partial_\nu\mathcal{D}(p_t) = \frac{\mathcal{L}(t)^3}{3} + O\big(\mathcal{L}(t)^2\big).$$

Consequently, uniformly for $|t| \le C\delta$,

$$M_0(t) = \frac{q-1}{2}\big(\mathcal{L}(t) + O(1)\big) + O\big(q^{1/2+3\varepsilon}\big) = V\mathcal{L}(t)\big(1 + O(\tfrac{1}{\log q})\big),$$

$$M_2(t) = \frac{q-1}{2}\Big(\frac{\mathcal{L}(t)^3}{3} + O(\mathcal{L}^2)\Big) + O\big(q^{1/2+3\varepsilon}\big) = \frac{V\mathcal{L}(t)^3}{3}\big(1 + O(\tfrac{1}{\log q})\big),$$

where replacing $\frac{q-1}{2}$ by $V = \frac{q-3}{2}$ changes the two main terms by $O(\mathcal{L})$ and $O(\mathcal{L}^3)$, respectively, both absorbed by the stated errors. Substitution into (4.2) gives

$$\begin{aligned} T_{\alpha_0}(t) &= M_2(t) - \frac{\mathcal{L}(t)^2}{4}M_0(t) + \Big(\frac{\mathcal{L}(0)-\mathcal{L}(t)}{2}\Big)^2 M_0(t) \\ &= V\mathcal{L}(t)^3\Big(\frac{1}{3} - \frac{1}{4}\Big)\big(1 + O(\tfrac{1}{\log q})\big) \\ &= \frac{V\mathcal{L}(t)^3}{12}\big(1 + O(\tfrac{1}{\log q})\big). \end{aligned}$$

the third term being $O(\delta^4)\cdot V\mathcal{L}$ since $\mathcal{L}(0) - \mathcal{L}(t) = O(t^2)$. □

4.5 Proof of Theorem 4.1

Suppose every even non-principal $\chi \bmod q$ had a critical-line zero $\frac{1}{2} + it_\chi$ with $t_\chi \in [-h, h]$, where $h = 2\pi\beta/\log q$, so that each F_χ vanishes in $[-h, h]$. Write $\mathcal{L} = \mathcal{L}(0)$ and $K = (4h/\pi)^2$ for the constant of Lemma 4.2 on $[-h, h]$. Apply the lemma to the twisted function F_{χ,α_0} with the common twist $\alpha_0 = \mathcal{L}/2$ and sum over the family. By Proposition 4.5, together with $\mathcal{L}(t) = \mathcal{L} + O(h^2)$, both integrands are constant on $[-h, h]$ to relative accuracy $1 + o(1)$, so

$$2h\,V\mathcal{L}\,(1 + o(1)) \;=\; \int_{-h}^{h} M_0 \;\le\; K \int_{-h}^{h} T_{\alpha_0}(t)\,dt \;=\; K \cdot \frac{2h\,V\mathcal{L}^3}{12}(1 + o(1)),$$

hence $1 \le K\mathcal{L}^2(1 + o(1))/12$. Inserting $h\mathcal{L} = 2\pi\beta\,(1 + o(1))$,

$$1 \le \frac{16\beta^2}{3}(1 + o(1)),$$

which fails for large q precisely when $\beta < \sqrt{3}/4 \approx 0.433$. Hence some even non-principal χ has no critical-line zero with ordinate in $[-h, h]$, unconditionally. Under GRH every nontrivial zero of that $L(s, \chi)$ lies on the critical line, so every nontrivial zero ρ has $|\mathrm{Im}\,\rho| > h$, that is, $X_\chi > \beta$.

4.6 Remarks on the theorem

Remark 4.6 (relation to the exactness theorem). Theorem 3.10 applies only to non-negative band-limited test functions evaluated on zero counts; it does not constrain critical-line values of L. A configuration with one zero per character in $[-h, h]$ saturates the count method but makes every F_χ vanish there, which Lemma 4.2 and Proposition 4.5 rule out below $\sqrt{3}/4$.

Remark 4.7 (the even subfamily and the prime modulus). If every character had a low zero, then every even character would have one, so it suffices to contradict the hypothesis on the even subfamily. Odd parity has a different gamma factor but the same leading constant. For prime q, every non-principal character is primitive of conductor q, so the family has one scaling, one functional-equation shape and one function $\mathcal{L}$.

Remark 4.8 (possible refinements). Pairing χ with $\overline{\chi}$ makes the product $F_\chi F_{\overline{\chi}}$ vanish at two points flanking the centre. The corresponding two-endpoint Wirtinger constant can be four times smaller, but the required family moments are fourth moments of L rather than second moments. Alternatively, an amplifier $A(\chi)$ could weight both family moments by $|A(\chi)|^2$. It would improve the constant only if the weighted ratio of the local L-value second moment to the centred derivative second moment exceeded the unweighted ratio; establishing such a gain requires asymptotics for both amplified moments. Neither refinement is used here.

Remark 4.9 (the elementary threshold). Replacing the sharp constant $(4h/\pi)^2$ by the elementary constant $K = 2h^2$ from Remark 4.3 changes the factor $16/3$ in the closing comparison to $2\pi^2/3$, since $2h^2\mathcal{L}^2/12 = (h\mathcal{L})^2/6 = 2\pi^2\beta^2/3$. The comparison $1 \le 2\pi^2\beta^2(1 + o(1))/3$ fails for $\beta < \sqrt{3/2}/\pi \approx 0.390$, still beyond the $1/4$ ceiling of Theorem 3.10. The sharp Wirtinger constant raises this elementary threshold to $\sqrt{3}/4$.

5 The random-matrix model

Random matrix theory predicts the full local point process of the scaled zeros and hence the distribution and extreme values of the lowest zero. These conjectures do not enter the proofs, but they distinguish properties of the predicted zero distribution from limitations of the rigorous methods and locate the constants $1/4$ and $\sqrt{3}/4$ within the predicted lowest-zero law.

5.1 Why the family is unitary

The symmetry class of a family is determined by self-duality. For a generic complex character the functional equation relates $L(s,\chi)$ to $L(1-s,\overline{\chi})$, so reflection across the central point exchanges two family members rather than imposing a symmetry on the zeros of one member. The root number $\varepsilon(\chi)$ can be any complex number of modulus 1 rather than only a sign, so no member is forced to vanish at the central point. Katz and Sarnak classify such a non-self-dual family as unitary [22, 23]. Quadratic characters are self-dual and belong to a different symmetry class, but the single quadratic character at a prime modulus is negligible inside the full primitive family. The corresponding matrix model is the Circular Unitary Ensemble (CUE), the unitary group $U(N)$ under Haar measure. Its eigenangle process is rotation-invariant: angle 0 is not distinguished, and no eigenvalue is forced at 1. Writing the eigenvalues as $e^{i\theta_1},\dots,e^{i\theta_N}$, the joint eigenangle density is proportional to

$$\prod_{1\le j<k\le N}\left|e^{i\theta_j}-e^{i\theta_k}\right|^2,$$

and the squared Vandermonde factor is the finite-N source of unitary level repulsion. Since N angles occupy a circle of length 2π, the unfolded coordinates are $y_j = N\theta_j/(2\pi)$, matching the arithmetic scaling $x_{\chi,j} = \gamma_{\chi,j}\frac{1}{2\pi}\log q$: the analogy is between the local densities $N/(2\pi)$ and $\log q/(2\pi)$, so matrix size plays the role of $\log q$, while the roughly q characters supply the number of sampled spectra. Over function fields the analogous equidistribution of Frobenius conjugacy classes in unitary groups is a theorem [23]; over number fields the unitary model is the Katz–Sarnak prediction, consistent with the smoothed one-level density in Theorem 3.2 [18].

5.2 The sine-kernel process

The local $N\to\infty$ limit of CUE, after unfolding to unit density, is the determinantal point process with kernel

$$K_{\sin}(u,v)=\frac{\sin\pi(u-v)}{\pi(u-v)},$$

whose n-point correlation functions are

$$\rho_n(u_1,\dots,u_n)=\det\bigl(K_{\sin}(u_i,u_j)\bigr)_{i,j=1}^{n}.$$

In particular

$$\rho_1(u)=1,\qquad \rho_2(u,v)=1-\left(\frac{\sin\pi(u-v)}{\pi(u-v)}\right)^2,$$

so the process has unit mean density and its two-point correlation vanishes quadratically as two points meet. Two distinctions matter for the lowest-zero problem. Zeros repel

other zeros, but the central point is a fixed location, not a point of the process, and $\rho_1(0) = 1$ rather than 0; the lowest zero is a waiting-time statistic measured from a fixed origin, not a spacing between consecutive points.

5.3 The hole probability and the lowest-zero laws

Definition 5.1 (sine-kernel hole probability)**.** For $L \geq 0$, define

$$E(L) = \det\bigl(I - K_{\sin}\bigr)_{L^2(0,L)}$$

to be the probability that an interval of unfolded length L contains no points of the sine-kernel process.

This Fredholm determinant governs the lowest zero because a high lowest zero is exactly an empty interval, and translation invariance gives the two convention-dependent laws

$$\mathbb{P}(X^+ > \beta) = E(\beta), \qquad \mathbb{P}(X > \beta) = E(2\beta). \tag{5.1}$$

The factor 2 is geometric, since $X^+ > \beta$ empties $[0, \beta]$ while $X > \beta$ empties $[-\beta, \beta]$. Within the sine-kernel model, each two-sided quantile is therefore half the corresponding one-sided quantile. The law is parameter-free: once the unfolding fixes unit density, no arithmetic quantity is fitted.

The same function links lowest zeros to consecutive gaps. If $p(s)$ is the unit-mean spacing density, then

$$E(s) = \int_s^\infty (t - s)\, p(t)\, dt, \qquad -E'(s) = \mathbb{P}(\text{gap} > s), \qquad E''(s) = p(s).$$

A fixed origin samples gaps in proportion to their length. This is the same length bias as the inspection paradox [29], and it explains why the lowest positive zero is not expected at half a mean spacing: with S a consecutive gap,

$$\mathbb{E}[X^+] = \frac{\mathbb{E}[S^2]}{2} = \frac{1}{2}\bigl(1 + \mathrm{Var}(S)\bigr).$$

Numerical evaluation of the determinant by Nyström quadrature [3] gives

$$\mathbb{E}[X^+] = \int_0^\infty E(L)\, dL = 0.58999694\ldots, \qquad \mathbb{E}[X] = 0.29499847\ldots. \tag{5.2}$$

The one-sided median m_+ satisfies $E(m_+) = 1/2$, and numerical evaluation gives $m_+ = 0.517094\ldots$. By (5.1), the two-sided median is

$$m = \frac{m_+}{2} = 0.258547\ldots.$$

At the constant of Section 3,

$$\mathbb{P}\bigl(X \leq \tfrac{1}{4}\bigr) = 1 - E\bigl(2 \cdot \tfrac{1}{4}\bigr) = 1 - E\bigl(\tfrac{1}{2}\bigr) \approx 0.485. \tag{5.3}$$

Thus $1/4 = 0.25$ lies close to the two-sided median $m = 0.258547\ldots$. The constant $\sqrt{3}/4 \approx 0.433$ lies farther into the upper tail, so Theorem 4.1 reaches beyond the central part of the predicted distribution at every sufficiently large prime.

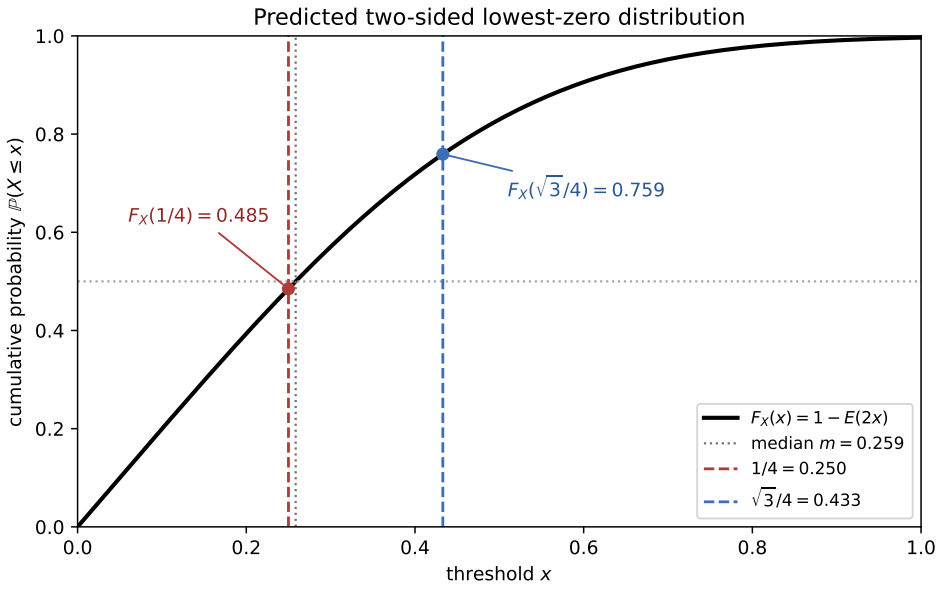

Figure 4: The predicted CUE distribution function $F_X(x) = 1 - E(2x)$ for the two-sided lowest zero. The threshold $1/4$ has cumulative probability 0.485 and nearly coincides with the median, whereas $\sqrt{3}/4$ has cumulative probability 0.759, leaving 24.1% of the predicted law above it.

5.4 Family extremes

By (5.1) and the small-interval expansion $E(L) = 1 - L + o(L)$,

$$\mathbb{P}(X_\chi \le x) = 1 - E(2x) = 2x + o(x).$$

For odd prime q, the family $\mathcal{F}_q$ contains one real quadratic character and $(q-3)/2$ pairs $\{\chi, \overline{\chi}\}$ of complex characters. Since $X_\chi = X_{\overline{\chi}}$, these give

$$N_q = 1 + \frac{q-3}{2} = \frac{q-1}{2}$$

distinct two-sided samples. Treating these samples as independent gives

$$\mathbb{P}\left(\min_{\chi \in \mathcal{F}_q} X_\chi > x\right) \approx (1-2x)^{N_q} \approx e^{-2N_q x}.$$

The exponential survival law varies on the scale $2N_q x \asymp 1$; since $N_q = (q-1)/2 \asymp q$, this is $x \asymp q^{-1}$. Thus the family minimum has typical scale

$$\min_{\chi \in \mathcal{F}_q} X_\chi \asymp q^{-1}. \tag{5.4}$$

For the family maximum, the relevant input is the large-L asymptotic of the hole probability [10],

$$\log E(L) = -\frac{\pi^2 L^2}{8} - \frac{1}{4}\log L + O(1). \tag{5.5}$$

For M independent samples, the expected number exceeding β is $M\mathbb{P}(X > \beta)$, so the typical maximum occurs when this quantity is of order 1. By (5.1), the two-sided and one-sided conditions are

$$N_q E(2\beta) \asymp 1, \qquad qE(\beta) \asymp 1.$$

Since $\log N_q = \log q + O(1)$, taking logarithms and applying (5.5) gives

$$\log\bigl(N_q E(2\beta)\bigr) = \log q - \frac{\pi^2\beta^2}{2} - \frac{1}{4}\log(2\beta) + O(1),$$
$$\log\bigl(qE(\beta)\bigr) = \log q - \frac{\pi^2\beta^2}{8} - \frac{1}{4}\log\beta + O(1).$$

Requiring the left sides to remain bounded gives $\beta \asymp \sqrt{\log q}$, for which the logarithmic terms are $o(\log q)$. Solving the two leading-order equations gives

$$\max_{\chi\in\mathcal{F}_q} X_\chi \sim \frac{\sqrt{2\log q}}{\pi}, \qquad \max_{\chi\in\mathcal{F}_q} X_\chi^+ \sim \frac{2\sqrt{2\log q}}{\pi}. \tag{5.6}$$

Conjugate pairing therefore does not alter the leading two-sided term in (5.6). Independence among the resulting two-sided samples, and the analogous assumption for the one-sided values, is a substantive conjectural input rather than a consequence of any density theorem. Equations (5.4) to (5.6) are extreme-value heuristics. Ben Arous and Bourgade prove that the largest consecutive eigenangle gap in an $N \times N$ CUE matrix has unfolded first-order size $2\sqrt{2\log N}/\pi$ [2]. The coefficient matches the one-sided scale in (5.6) when the N sampled matrix gaps are compared with the roughly q sampled character spectra. Their maximum is taken over consecutive gaps within one spectrum, whereas (5.6) takes the maximum of fixed-origin waiting distances across character spectra. For primitive Dirichlet characters modulo a prime q, no theorem gives the cross-character independence needed to deduce (5.4) or (5.6).

5.5 Why the rigorous methods fall short of the extreme-value predictions

The extreme-value predictions (5.4)–(5.6) use more information about zero correlations and variation across characters than the proofs in Sections 3 and 4.

The full sine-kernel model determines the laws (5.1) through correlations of every order. The one-level density retains only ρ_1, the one-point correlation function: for admissible f, the model's expected count is $\int f\rho_1 = \int f$, which is the limit that Theorem 3.2 proves for band-limited f. This first moment determines the expected number of zeros in an interval but not the probability that the interval is empty, which requires a Fredholm determinant built from all correlation orders. Section 3 imposes two further restrictions: the explicit formula is evaluated only for test functions of finite Fourier support, and the resulting zero statistic is averaged over the family. One exceptional character can change a family extreme without noticeably changing this average.

The support-2 constants lie near the median of the single-character distribution but do not approach the moving extremes (5.4) and (5.6). The Wirtinger theorem gives a fixed upper-tail constant, still below the conjectural maximum $\sqrt{2\log q}/\pi$. Reaching that scale would require value-distribution information beyond the second moments used here. On the min side, no fixed constant approaches the conjectural scale q^{-1}.

6 Numerical calibration

The computations compare finite-conductor Dirichlet data with the parameter-free CUE law, measure the lower-order terms omitted by the leading scaling, and locate the rigorous

thresholds in the observed families. They are descriptive: they neither enter the proofs nor establish the limiting law.

6.1 Data and zero computation

Code and data are available in the numerical repository [26].

For each prime

$$q \in \{1009, 2003, 5003, 10007, 20011, 50021, 100003\},$$

all $q-2$ non-principal characters modulo q were considered. PARI/GP 2.17.3's `lfunzeros` routine [32] was used to locate the first three non-negative critical-line zeros of each $L(s,\chi)$, retaining central zeros with multiplicity. Central values were screened numerically, with higher-precision reevaluation of suspicious cases; no central zero was detected. Each recorded triple was required to agree to 10^{-18} between successive mesh parameters, beginning with `divz=32` and `divz=64`; a disagreement triggers a `divz=128` search and is accepted only if the latter two meshes agree. The first zero supplies the lowest-zero distributions of Section 6.3; the next two were retained for a separate level-repulsion diagnostic using consecutive gaps. The located ordinates were checked against an independent `mpmath` implementation using the Hurwitz zeta function [19] for six quadratic characters, the lowest zero of the Dirichlet beta function, and LMFDB's signed zero table for a complex character modulo 5 [28, 31]. The LMFDB comparison also tests conjugate pairing. These checks test numerical accuracy and character labelling but do not prove that no zero was missed, establish central non-vanishing, or verify GRH away from the critical line. The resulting critical-line data are not a rigorous zero tabulation.

For a complex character the functional equation identifies the negative zeros of $L(s,\chi)$ with reflections of the positive zeros of $L(s,\overline{\chi})$. The two-sided statistic is reconstructed by conjugate pairing, $X_\chi = \min(X_\chi^+, X_{\overline{\chi}}^+)$. The pairing gives χ and $\overline{\chi}$ the same two-sided value, so each value from a non-real conjugate pair appears twice in the character-level tables.

6.2 Two finite-conductor coordinates

The rigorous results use the leading scaling $x_{\chi,j} = \gamma_{\chi,j}\log q/(2\pi)$ and the corresponding variables X_χ^+ and X_χ from Section 2. For the conductors used in the computations, however, $\log q/(2\pi)$ omits a substantial parity-dependent gamma-factor term of the smooth zero density.

Definition 6.1 (finite-conductor comparison coordinates)**.** For real t and parity $\mathfrak{a} \in \{0,1\}$, define

$$\theta_{\mathfrak{a},q}(t) = \frac{t}{2}\log\frac{q}{\pi} + \operatorname{Im}\log\Gamma\left(\frac{\frac{1}{2}+\mathfrak{a}+it}{2}\right),$$

where $\operatorname{Im}\log\Gamma((\frac{1}{2}+\mathfrak{a}+it)/2)$ denotes the continuous argument normalised to equal 0 at $t=0$. For the computed critical-line zeros of a character of parity $\mathfrak{a}$, whose ordinates $\gamma_{\chi,j}$ are real, define

$$\widetilde{x}_{\chi,j} = \frac{\theta_{\mathfrak{a},q}(\gamma_{\chi,j})}{\pi}, \qquad \widetilde{X}_\chi^+ = \frac{\theta_{\mathfrak{a},q}(\gamma_{\chi,1})}{\pi}, \qquad \widetilde{X}_\chi = \min\bigl(\widetilde{X}_\chi^+, \widetilde{X}_{\overline{\chi}}^+\bigr).$$

Up to endpoint and branch conventions, the zero-counting function has the form

$$\mathcal{N}_\chi(t) = \frac{\theta_{\mathfrak{a},q}(t)}{\pi} + S_\chi(t) + C_\chi, \qquad S_\chi(t) = \frac{1}{\pi} \arg L\big(\tfrac{1}{2} + it, \chi\big)\,,$$

where C_χ is constant in t. The derivative $\theta'_{\mathfrak{a},q}(t)/\pi$ is the smooth conductor-plus-gamma density of Section 2.3; for even parity,

$$\frac{\theta'_{0,q}(t)}{\pi} = \frac{\mathcal{L}(t)}{2\pi},$$

with $\mathcal{L}$ as in Section 4.4. The tilded coordinate makes $\theta_{\mathfrak{a},q}(t)/\pi$ the coordinate itself, while the character-dependent remainder $S_\chi(t) + C_\chi$ is unchanged. At the centre $\psi(\frac{1}{4}) \approx -4.227$ and $\psi(\frac{3}{4}) \approx -1.086$, corrections that remain non-negligible relative to $\log(q/\pi)$ throughout the computed range and that differ between the parities. Consequently the leading scaling places the low zeros farther to the right and separates the even and odd samples, with a relative discrepancy decaying only like $1/\log q$. The untilded variables retain the normalisation used in the asymptotic theorems. The tilded variables unfold by the full smooth density before comparison with the limiting CUE law; they remove the known finite-conductor distortion but do not assert that the finite-conductor distribution is exactly CUE.

6.3 Distributional comparison

The CUE curves were computed independently of the arithmetic data by Nyström quadrature for the Fredholm determinant $E(L)$ [3]. The principal descriptive statistics are as follows, in the leading coordinate and then the comparison coordinate.

q	mean X^+	mean $\widetilde{X}^+$	mean $\widetilde{X}$	$\mathbb{P}(X \le \frac{1}{4})$	$\mathbb{P}(\widetilde{X} \le \frac{1}{4})$
1009	0.938	0.562	0.287	0.236	0.499
2003	0.916	0.571	0.290	0.260	0.494
5003	0.883	0.573	0.291	0.269	0.493
10007	0.862	0.575	0.291	0.308	0.490
20011	0.844	0.577	0.291	0.311	0.490
50021	0.823	0.578	0.291	0.324	0.487
100003	0.809	0.580	0.292	0.336	0.487
CUE limit	0.590	0.590	0.295	0.485	0.485

After the gamma-factor correction, the one-sided and two-sided histograms resemble the parameter-free sine-kernel curves: at $q = 100003$ the observed means are 0.580 and 0.292, compared with the CUE values 0.590 and 0.295. The Poisson curves are an uncorrelated null model rather than a competing prediction; a unit-intensity Poisson process has waiting-time densities e^{-x} on one side and $2e^{-2x}$ on two sides. The comparison is descriptive rather than a formal goodness-of-fit test.

The discrepancy between the leading coordinate $x_{\chi,j}$ and the finite-conductor comparison coordinate $\widetilde{x}_{\chi,j}$ is a finite-conductor effect, not evidence for competing limiting laws. In leading units even and odd characters are visibly displaced because their gamma factors give different smooth central densities; applying $\theta_{\mathfrak{a},q}/\pi$ largely removes the separation and

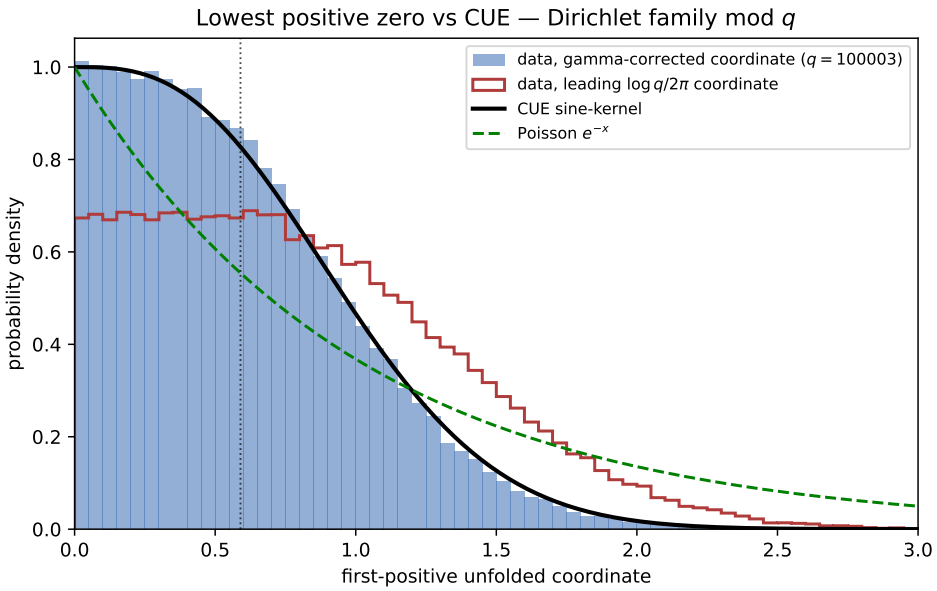

Figure 5: The distributions of the leading-scale lowest-positive statistic X_χ^+ and the finite-conductor comparison statistic $\widetilde{X}_\chi^+$, against the CUE and Poisson laws.

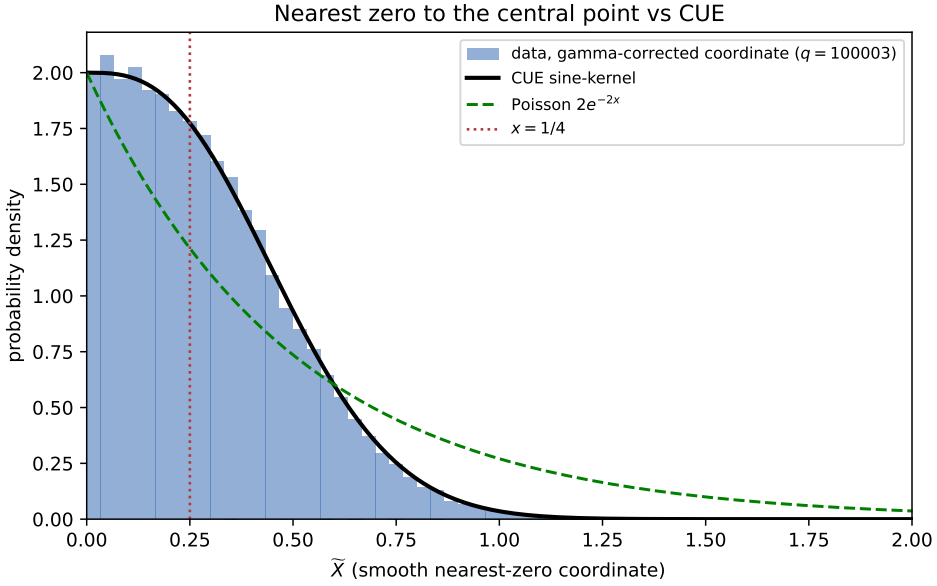

Figure 6: The distribution of the two-sided finite-conductor comparison statistic $\widetilde{X}_\chi$, against the CUE and Poisson laws.

places both parity classes on the common CUE scale.

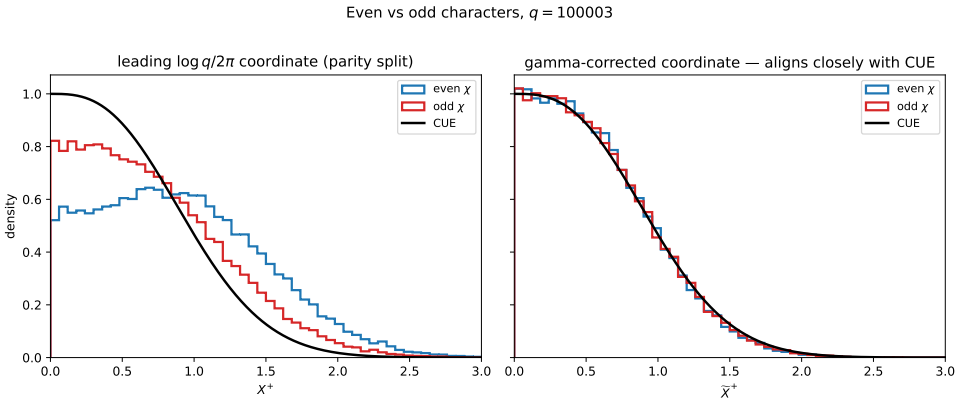

Figure 7: Even and odd characters under the leading and gamma-corrected coordinates. The comparison isolates the lower-order parity term omitted by $\log q/(2\pi)$.

For the threshold $1/4$, the limiting prediction is $\mathbb{P}(X \le \frac{1}{4}) \approx 0.485$. In the comparison coordinate the observed proportion changes from 0.499 at $q = 1009$ to 0.487 at $q = 100003$, while in the theorems' leading units it is 0.336 at the largest modulus. The finite-conductor event $X_\chi \le 1/4$ and the comparison event $\widetilde{X}_\chi \le 1/4$ are therefore not interchangeable: the former is the literal arithmetic event, whereas the latter locates its asymptotic counterpart in the conjectural law.

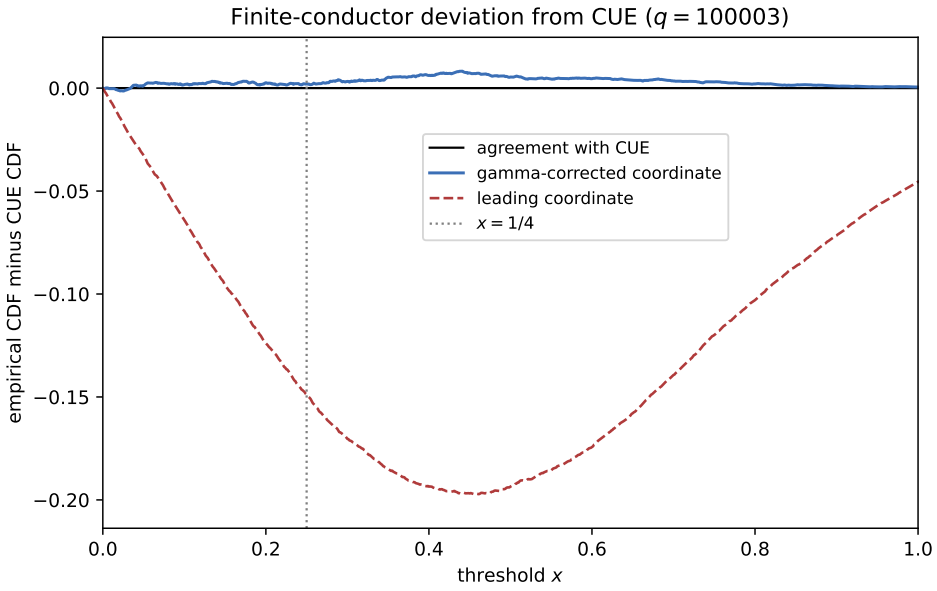

Figure 8: Empirical minus CUE distribution functions for the two-sided lowest zero at $q = 100003$, in both coordinates. At each threshold x, positive values mean that more characters have their lowest zero at or below x than CUE predicts; negative values mean fewer. Zero denotes agreement. The vertical guide marks $1/4$.

6.4 Scope of the evidence

For each modulus, the extrema in the finite-conductor comparison coordinate and their normalisations by the scales predicted in Section 5.4 are as follows. Write $R_{\max}(q) = \pi \max_\chi \widetilde{X}_\chi / \sqrt{2 \log q}$ for the ratio of the observed maximum to its predicted leading scale.

q	$\min_\chi \widetilde{X}_\chi$	$q \min_\chi \widetilde{X}_\chi$	$\max_\chi \widetilde{X}_\chi$	$R_{\max}(q)$
1009	0.000521	0.526	1.144	0.967
2003	0.001604	3.212	1.009	0.813
5003	0.000482	2.411	1.248	0.950
10007	0.000265	2.648	1.170	0.856
20011	0.000207	4.134	1.187	0.838
50021	0.000007	0.357	1.375	0.929
100003	0.000023	2.338	1.369	0.896

The comparison coordinate has the same leading asymptotic scale as X_χ. The values of $q \min_\chi \widetilde{X}_\chi$ remain of order one, while $\max_\chi \widetilde{X}_\chi$ is between 0.813 and 0.967 times $\sqrt{2 \log q}/\pi$. These ranges are qualitatively compatible with the scales in Section 5.4, but the observed minimum does not decrease steadily with q, nor does the observed maximum increase steadily. Seven moduli cannot test the predicted constants: $\sqrt{\log q}$ varies slowly, the extrema fluctuate, and the cross-character independence assumed by the heuristic is unproved. After the leading and gamma-corrected coordinates are separated, the observed lowest-zero distributions resemble the predicted unitary law and place 1/4 near the median of the two-sided limiting law. This finite-conductor comparison does not strengthen the first-moment theorems or test the conjectural extreme-value constants.

7 Discussion

7.1 What the results establish

Theorem 3.10 and the min-side sharpness theorem of Carneiro, Ismoilov and Ramos [5] show that both support-2 constants remain 1/4 after optimisation over the full admissible test-function classes of their respective positivity arguments. This identifies 1/4 as the boundary of the symmetric-interval first-moment method, not a consequence of choosing the quadratic Hughes–Rudnick family or Zhao's test. Corollary 3.12 uses the additional conductor average of Drappeau, Pratt and Radziwiłł to decrease the exact min-side sign-test threshold and increase the explicit Selberg max-side lower bound. For every $\varepsilon > 0$, Theorem 4.1 instead uses second moments of L and L' to produce, at every sufficiently large prime, an even character with no critical-line zero of scaled height at most $\sqrt{3}/4-\varepsilon$; under GRH, its two-sided scaled lowest zero therefore exceeds $\sqrt{3}/4-\varepsilon$. The bound 1/4 therefore belongs to the zero-count statistic, not to the family maximum.

7.2 Assumptions and evidentiary scope

The zero-location conclusions are conditional on GRH. In Theorems 3.5 and 3.6 and Corollary 3.12, GRH places zeros on the critical line so that the sign of the test function records whether a zero lies in the target interval. In Theorem 4.1, GRH converts the unconditional absence of zeros from a critical-line interval into a statement about all nontrivial zeros. The density theorem of Drappeau, Pratt and Radziwiłł and the moment asymptotics of Proposition 4.5 are unconditional. The single-character random-matrix law is conjectural. The extreme-value predictions of Section 5.4 additionally assume independence across character samples after accounting for conjugate pairing. Neither is a proof input. The numerical comparison is descriptive and makes no formal goodness-of-fit claim.

7.3 How far the min/max symmetry extends

Selberg's constructed pair satisfies $\beta_{\min}(\Delta) + \beta^{\mathrm{Sel}}_{\max}(\Delta) = 1/2$ for $\Delta > 1$. At support 2, this is also an identity between the exact thresholds of the two positivity problems. At general support, the identity alone does not determine the optimal max-side threshold: that requires minimising the interval-majorant mass, as in the problem defining ρ in Section 3.6.

The min-side optimiser is governed by a Dirichlet eigenvalue problem, with a half-period cosine at every support after rescaling. At the support-2 max endpoint, a rank-two evaluation bound selects the same cosine direction. These distinct optimisation problems share an extremal shape at that endpoint, not a general formula for both optimal thresholds. As support increases without bound, the min-side sign-test threshold tends to 0 and Selberg's max-side bound tends to 1/2; the mass inequality $\int M \geq 2\beta$ caps every first-moment interval-majorant argument at 1/2 (Section 3.8).

The two existence problems also interact differently with family averaging. Theorem 4.1 proceeds by contradiction: if every character had a low zero, Wirtinger's inequality would apply to every member of the family, and averaging these inequalities would contradict the second-moment asymptotics. Thus second moments can prove the existence of a character whose lowest zero exceeds a fixed threshold. By contrast, the min-side target is one character with an exceptionally low zero, conjecturally at scale q^{-1},

and fixed positive family moments need not detect such a lower-tail event. A min-side bound below $1/4$ valid at every sufficiently large prime modulus remains open.

7.4 Open problems

First, pairing χ with $\overline{\chi}$ and applying a two-endpoint Wirtinger inequality to $F_\chi F_{\overline{\chi}}$ can reduce the Wirtinger constant by a factor of up to 4, but requires fourth moments of L rather than second moments. Alternatively, an amplifier could improve the constant by increasing the weighted ratio of the local L-value second moment to the centred derivative second moment. Either refinement may improve the max-side constant $\sqrt{3}/4$ (Remark 4.8).

Second, support past $2 + 50/1093$ would require stronger equidistribution of primes in progressions than that obtained from the dispersion and Kloosterman-sum estimates underlying the theorem of Drappeau, Pratt and Radziwiłł, or progress on the de-averaging comparison of Fiorilli and Miller (Section 3.8). Additional support improves the exact min-side sign-test threshold and the explicit Selberg max-side bound.

Third, a min-side analogue of Theorem 4.1 may be sought through the ratio $|L(\frac{1}{2}, \chi)|/|L'(\frac{1}{2}, \chi)|$, which linearisation suggests measures the distance to a nearby zero. A min-side improvement through this ratio requires quantitative smallness relative to the mean spacing $2\pi/\log q$, together with control of the Taylor remainder.

Fourth, the rigorous bounds are fixed constants, whereas the conjectural extremes vary with q. Reaching the predicted max-side scale $\sqrt{2\log q}/\pi$ requires value-distribution information beyond second moments, and the methods developed here do not address the shrinking min-side scale q^{-1}.

References

[1] N. I. Achieser. *Theory of Approximation.* Frederick Ungar, New York, 1956.

[2] G. Ben Arous and P. Bourgade. Extreme gaps between eigenvalues of random matrices. *Ann. Probab.*, 41(4):2648–2681, 2013.

[3] F. Bornemann. On the numerical evaluation of Fredholm determinants. *Math. Comp.*, 79(270):871–915, 2010.

[4] H. M. Bui and M. B. Milinovich. Central values of derivatives of Dirichlet L-functions. *Int. J. Number Theory*, 7(2):371–388, 2011.

[5] E. Carneiro, T. Ismoilov, and A. P. Ramos. Sign uncertainty and de Branges spaces, 2024. To appear in *Ann. Inst. Fourier (Grenoble)*. Preprint available at `https://arxiv.org/abs/2408.01186`.

[6] J. B. Conrey, H. Iwaniec, and K. Soundararajan. The sixth power moment of Dirichlet L-functions. *Geom. Funct. Anal.*, 22(5):1257–1288, 2012.

[7] H. Davenport. *Multiplicative Number Theory*, volume 74 of *Graduate Texts in Mathematics*. Springer-Verlag, 3rd edition, 2000. Revised by H. L. Montgomery.

[8] H. Davenport and H. Halberstam. The values of a trigonometrical polynomial at well spaced points. *Mathematika*, 13:91–96, 1966.

[9] C.-J. de la Vallée Poussin. Recherches analytiques sur la théorie des nombres premiers. *Ann. Soc. Sci. Bruxelles*, 20:183–256, 281–397, 1896.

[10] P. Deift, A. Its, I. Krasovsky, and X. Zhou. The Widom–Dyson constant for the gap probability in random matrix theory. *J. Comput. Appl. Math.*, 202(1):26–47, 2007.

[11] P. G. L. Dirichlet. Beweis des Satzes, dass jede unbegrenzte arithmetische Progression, deren erstes Glied und Differenz ganze Zahlen ohne gemeinschaftlichen Factor sind, unendlich viele Primzahlen enthält. *Abh. Königl. Preuss. Akad. Wiss. Berlin*, pages 45–81, 1837.

[12] S. Drappeau, K. Pratt, and M. Radziwiłł. One-level density estimates for Dirichlet L-functions with extended support. *Algebra Number Theory*, 17(4):805–830, 2023.

[13] D. Fiorilli and S. J. Miller. Surpassing the Ratios Conjecture in the 1-level density of Dirichlet L-functions. *Algebra Number Theory*, 9(1):13–52, 2015.

[14] J. Friedlander and A. Granville. Limitations to the equi-distribution of primes I. *Ann. of Math. (2)*, 129(2):363–382, 1989.

[15] D. A. Goldston and R. C. Vaughan. On the Montgomery–Hooley asymptotic formula. In *Sieve Methods, Exponential Sums, and their Applications in Number Theory*, volume 237 of *London Mathematical Society Lecture Note Series*, pages 117–142. Cambridge University Press, 1997.

[16] R. R. Hall. Generalized Wirtinger inequalities, random matrix theory, and the zeros of the Riemann zeta-function. *J. Number Theory*, 97(2):397–409, 2002.

[17] G. H. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities.* Cambridge University Press, 1934.

[18] C. P. Hughes and Z. Rudnick. Linear statistics of low-lying zeros of L-functions. *Q. J. Math.*, 54:309–333, 2003.

[19] A. Hurwitz. Einige Eigenschaften der Dirichlet'schen Functionen $F(s) = \sum \left(\frac{D}{n}\right) \frac{1}{n^s}$, die bei der Bestimmung der Classenanzahlen binärer quadratischer Formen auftreten. *Z. Math. Phys.*, 27:86–101, 1882.

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

[21] H. Iwaniec and P. Sarnak. Dirichlet L-functions at the central point. In *Number Theory in Progress, Vol. 2*, pages 941–952. de Gruyter, 1999.

[22] N. M. Katz and P. Sarnak. Zeroes of zeta functions and symmetry. *Bull. Amer. Math. Soc.*, 36(1):1–26, 1999.

[23] N. M. Katz and P. Sarnak. *Random Matrices, Frobenius Eigenvalues, and Monodromy*, volume 45 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, 1999.

[24] F. Littmann. Quadrature and extremal bandlimited functions. *SIAM J. Math. Anal.*, 45(2):732–747, 2013.

[25] H. L. Montgomery. *Topics in Multiplicative Number Theory*, volume 227 of *Lecture Notes in Mathematics*. Springer, Berlin, 1971.

[26] G. L. Oon. Dirichlet lowest-zero numerics. Version 1.0.0, 2026. URL `https://github.com/sarkologist/dirichlet-lowest-zeros/tree/v1.0.0`.

[27] R. E. A. C. Paley and N. Wiener. *Fourier Transforms in the Complex Domain*, volume 19 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, 1934.

[28] D. J. Platt. *Computing degree 1 L-functions rigorously*. PhD thesis, University of Bristol, 2011.

[29] S. M. Ross. The inspection paradox. *Probab. Engrg. Inform. Sci.*, 17(1):47–51, 2003.

[30] A. Selberg. Lectures on sieves. In *Collected Papers, Vol. II*, pages 65–247. Springer-Verlag, Berlin, 1991.

[31] The LMFDB Collaboration. The L-functions and Modular Forms Database, 2026. Available from `https://www.lmfdb.org/`; accessed 2026.

[32] The PARI Group. *PARI/GP version 2.17.3*. Univ. Bordeaux, 2025. Available from `http://pari.math.u-bordeaux.fr/`.

[33] J. D. Vaaler. Some extremal functions in Fourier analysis. *Bull. Amer. Math. Soc.*, 12(2):183–216, 1985.

[34] T. Zhao. The positivity technique and low-lying zeros of Dirichlet L-functions. *To appear in J. Number Theory*, 292:219–250, 2027. Preprint available at `https://arxiv.org/abs/2503.15832`.