

Zeros of a whole family: an accessible guide to the $7/8$ argument for $\zeta(s)$

Claude Opus 5.5 and OpenAI Astra
at the behest of Michael O. Rubinstein

A note from Michael Rubinstein. These are working notes, generated by Claude and Codex to help me read the papers on the quasi-Riemann hypothesis that OpenAI released.

Abstract

We describe, in classical terms and with the averages written out, the argument of [1] showing that $\zeta(s)$ has no zeros in $\Re s > 7/8$. Unlike the $11/12$ argument of the companion guide [3], it is a contradiction argument about the rightmost zero of a whole family of L -functions over $\mathbb{Q}(\sqrt{-3})$. A single average (the “probe”) is evaluated in two exact ways. The $7/8$ step makes the dual sum longer; short sums over primes, arranged so that a scalar contribution cancels, make the longer sum affordable, and the rare bad rows are counted with two different moment bounds.

The proof in six steps.

1. Let β_* be the supremum of the real parts of the zeros (in $1/2 \leq \Re s \leq 1$) of all L -functions $L(s, \eta)$ in a fixed family over the Eisenstein integers; the family contains $\zeta(s)L(s, \chi_{-3})$. Assume $\beta_* > 7/8$ and aim for a contradiction.
2. For every member η , one builds an explicit average $I(Z)$, the *probe*, of cubic Gauss sums against sextic residue symbols and Gauss sums, and evaluates it in two exact ways.
3. The first way (Cauchy–Schwarz; a planar large sieve for one factor and, for the other, the automorphy of the cubic theta function followed by a quadratic large sieve) shows that the probe is small: at most $Z^{3/16+\varepsilon}$. Small averaging scales help here, but they make the dual sum in the second way long.
4. The second way (Poisson summation) writes the probe as a sum over “rows” u ; the row $u = 1$ contains $1/L(s, \eta)$ and would be of size $Z^{\beta_*-11/16}$ if $L(s, \eta)$ had a zero near β_* , while every other row involves twisted L -functions, and a row can be large only if one of *them* has a zero of bounded height well to the right of the critical line.
5. A zero forces two Dirichlet polynomials attached to that row (one with Möbius coefficients, one without) to be unusually large, and two moment estimates show that such rows are rare. Short sums over primes built into the probe, arranged so that a row-independent term cancels, are what make the long dual sum affordable: they have full size on the row $u = 1$, but on the other rows they are only square-root-size prime polynomials, which also sharpen the count.
6. Comparing the two evaluations shows that $1/L(s, \eta)$ continues a fixed distance to the left of β_* for every η , contradicting the definition of β_* ; hence $\beta_* \leq 7/8$, and in particular $\zeta(s) \neq 0$ for $\Re s > 7/8$.

Scope. This is a guide to [1], not an independent proof. We quote the averages and exponents of [1] as stated there (equation and line references are to its T_EX source), simplify notation, and explain why each step is needed. We did not re-verify every line of [1], which runs to about 16,700 lines of T_EX. Independently, the final statement for ζ and for all Dirichlet L -functions has a Lean formalization in the same repository; we compiled it and replayed it through Lean’s kernel (it uses only the standard axioms). The companion guide [3] explains the simpler 11/12 argument of [2], which is organized differently.

1 The family and the number β_*

Let $\omega = e^{2\pi i/3}$ and let $\mathbb{Z}[\omega]$ be the Eisenstein integers, with norm $N(a + b\omega) = a^2 - ab + b^2$. An element is *primary* if it is $\equiv 1 \pmod{3}$; every ideal prime to 3 has exactly one primary generator. For a prime $\pi \nmid 6$, the sextic residue symbol $(u/\pi)_6$ is the sixth root of unity congruent to $u^{(N(\pi)-1)/6}$ modulo π (0 if $\pi \mid u$); for primary n prime to 6, $(u/n)_6 = \prod_{\pi \in n} (u/\pi)_6$. We write $\chi_n(u) = (u/n)_6$, so that χ_n^3 and χ_n^2 are the quadratic and cubic symbols.

The family. A *finite-order Hecke character* η of $\mathbb{Q}(\sqrt{-3})$ is, concretely, a completely multiplicative function on primary n that depends only on n modulo some fixed modulus (and is 0 when n is not coprime to it); through primary generators it is a function on ideals. Its L -function is

$$L(s, \eta) = \sum_{\mathfrak{a} \neq 0} \frac{\eta(\mathfrak{a})}{N(\mathfrak{a})^s} \quad (\Re s > 1),$$

summed over the nonzero ideals of $\mathbb{Z}[\omega]$ (for ideals prime to 3 this is the sum over primary n ; the prime above 3 enters in the usual way). Examples:

- $\eta = 1$: $L(s, 1) = \zeta(s)L(s, \chi_{-3})$, the product studied in [3];
- $\eta(n) = \chi(N(n))$ for a Dirichlet character χ : $L(s, \eta) = L(s, \chi)L(s, \chi\chi_{-3})$, up to finitely many Euler factors;
- $\eta(n) = \nu(n) \overline{(u/n)_6}$ for fixed u and a residue-class character ν : a *sextic twist*. These are the L -functions that appear as “rows” below. In general they are not products of Dirichlet L -functions.

All of these L -functions continue meromorphically (entire unless η is principal, when there is a simple pole at $s = 1$) and satisfy functional equations (Hecke). Define

$$\beta_* = \sup(\{1/2\} \cup \{\Re \rho : L(\rho, \eta) = 0, 1/2 \leq \Re \rho \leq 1, \eta \text{ primitive of finite order}\}) \quad (1)$$

[1, L378–385]. Every zero of ζ is a zero of $L(s, 1)$, so $\beta_* \leq 7/8$ implies the theorem for ζ .

Why the whole family is unavoidable. The second evaluation of the probe produces sextic twists of the target, and the only a priori information about their zeros is $\Re \rho \leq \beta_*$. Moreover, the contradiction at the end is produced by *whichever* member has a zero close to β_* . So even for ζ every estimate must hold for every target η , with margins that do not depend on η .

2 The continuation criterion

Everything is reduced to the following statement [1, Prop. lem:continuation-criterion].

Proposition 1. *Fix $\sigma_0 \in (1/2, 1)$, suppose $\beta_* > \sigma_0$, and let $C(s) = s + c$ with c fixed. Suppose there are $0 < \omega < \beta_* - \sigma_0$ and $\sigma > 0$, independent of η , such that for every primitive η there are a finite set $\mathcal{S}$ of primes, a holomorphic H_η with $|H_\eta(s) - 1| \leq 1/2$ on $\Re s > \sigma_0$, and a function $J_\eta(Z)$ with*

$$|J_\eta(Z)| \ll_\eta Z^{C(\sigma_0)+\omega} \quad (\text{low estimate}), \quad (2)$$

$$|J_\eta(Z) - f_\eta(Z)| \ll_\eta Z^{C(\beta_*)-\sigma} \quad (\text{high estimate}), \quad (3)$$

where the signal is

$$f_\eta(Z) = \frac{1}{2\pi i} \int_{\Re s=2} Z^{C(s)} e^{(s-5/6)^2} \frac{H_\eta(s)}{L^\mathcal{S}(s, \eta)} ds. \quad (4)$$

Then we have a contradiction.

Why. Together, (2) and (3) give $|f_\eta(Z)| \ll_\eta Z^{C(\beta_*)-\epsilon_*}$ for large Z , with $\epsilon_* = \min(\beta_* - \sigma_0 - \omega, \sigma) > 0$ independent of η and $\beta_* - \epsilon_* > \sigma_0$. Moving the contour in (4) to the right (the Gaussian decays on vertical lines, and H_η and $1/L^\mathcal{S}$ are bounded on $\Re s \geq 2$) gives $f_\eta(Z) \ll_{\eta, B} Z^B$ as $Z \rightarrow 0$, for every B . Hence the Mellin transform $\int_0^\infty f_\eta(Z) Z^{-C(s)} dZ/Z$ is holomorphic in $\Re s > \beta_* - \epsilon_*$, and by Mellin inversion it equals $e^{(s-5/6)^2} H_\eta(s)/L^\mathcal{S}(s, \eta)$ for $\Re s > 1$. Since $|H_\eta| \geq 1/2$ on $\Re s > \sigma_0$ and the deleted Euler factors do not vanish in $\Re s > 0$, $1/L(s, \eta)$ continues to $\Re s > \beta_* - \epsilon_*$ for every primitive η . But by (1) some primitive η has a zero with real part in $(\beta_* - \epsilon_*, \beta_*]$, a contradiction.

Note the asymmetry: the low estimate is measured against $C(\sigma_0)$ and the high estimate against $C(\beta_*)$. Since C has slope one, $C(\beta_*) - C(\sigma_0) = \beta_* - \sigma_0$. The argument in Part I of [1] uses $\sigma_0 = 11/12$ and $C(s) = s - 2/3$; the 7/8 argument uses

$$\sigma_0 = \frac{7}{8}, \quad C(s) = s - \frac{11}{16}, \quad C(7/8) = \frac{3}{16}.$$

3 The probe

3.1 The formula

Fix the target η (for ζ , $\eta = 1$). Besides η , the probe contains a few fixed factors that depend only on residue classes at a finite set S of bad primes (the primes above 2 and 3, those dividing the modulus of η , and all primes of norm below a fixed bound): a character ξ modulo a fixed b_* supported on S , its Gauss sum τ , a reciprocity sign $\mathcal{R}(\cdot, \cdot)$ and a phase G . The character ξ removes, after Poisson summation, the zero frequency and the frequencies not prime to S ; τ , $\mathcal{R}$ and G cancel unit and reciprocity phases (between cn^3 and s , and inside cn^3), so that the coefficients after Poisson summation factor prime by prime. We collect them into c_s and $\kappa_s(cn^3)$ below (both of absolute value one). With

$$\gamma_2(c) = N(c)^{-1/2} \sum_{x \bmod c} (x/c)_3 e(x/c), \quad \alpha(c) = c/|c|, \quad e(z) = \exp(4\pi i \Im z / \sqrt{3}),$$

the *base probe* [1, eq. general-probe, L3397] is, after evaluating its Mellin integral,

$$I_\eta(X, Y, Z) = \frac{1}{Y} \sum_s \frac{c_s W_1(N(s)/Y)}{\sqrt{N(s)X}} \sum_{m \in \mathbb{Z}[\omega]} \xi(m) \frac{g_s(-m)}{\sqrt{N(s)}} W_0\left(\frac{N(m)}{N(s)X}\right) \times \sum_{\substack{c \text{ squarefree} \\ n}} \frac{\gamma_2(c) \overline{\alpha(cn^3)} \eta(cn^3) \kappa_s(cn^3) \chi_{cn^3}(m)}{\sqrt{N(c)N(n)}} V\left(\frac{N(c)N(n)^3}{Z}\right). \quad (5)$$

Here s, c, n run over primary elements prime to S (with s not necessarily squarefree, and c, n allowed to share primes), m runs over all of $\mathbb{Z}[\omega]$,

$$g_s(-m) = \sum_{d \bmod s} (d/s)_6 e(-md/s)$$

is a sextic Gauss sum, $W_0, W_1 \geq 0$ are fixed smooth bumps, and $V(y) = (2\sqrt{\pi})^{-1} e^{-(\log y)^2/4}$ is a log-Gaussian, so that $N(cn^3) \approx Z$. (We have absorbed the norm of b_* into X ; in (7) and in the principal row below, X is the paper's unabsorbed scale, which changes only bounded factors.)

What the three sums are.

- The innermost sum over the *completed index* cn^3 is a twisted sum of Fourier coefficients of Kubota's cubic theta function, whose coefficient at cn^3 is (up to normalization) $\gamma_2(c)N(n)^{1/2}$. “Completed” means the cube factor n^3 is kept; the theta function needs it.
- The sum over s , with the Gauss sum $g_s(-m)$, is an additive polynomial in the averaging variable m .
- The averaging variable m has $N(m) \approx XY$; it ties the two together through $\chi_{cn^3}(m)$ and $g_s(-m)$.

In Part I of [1], $X = Y = Z^{1/2}$.

3.2 First evaluation: the probe is small

Expanding W_0 by Mellin inversion separates m from s , and splitting s into finitely many residue classes removes the remaining dependence of κ_s on s ; this gives an exact bilinear form in the rows m [1, eq. low-separated, L3470]:

$$I_\eta(X, Y, Z) \approx Q^{-1/2} \sum_{N(m) \asymp Q} A_m B_m, \quad Q = XY, \quad (6)$$

(an average over finitely many residue classes and one Mellin variable is suppressed), where

$$A_m = \frac{1}{Y} \sum_{N(s) \asymp Y} c(s) \frac{g_s(-m)}{\sqrt{N(s)}}, \quad B_m = \sum_{c, n} \frac{\gamma_2(c) \overline{\alpha(cn^3)} \eta(cn^3) \kappa(cn^3) \chi_{cn^3}(m)}{\sqrt{N(c)N(n)}} V\left(\frac{N(cn^3)}{Z}\right),$$

with bounded coefficients $c(s)$. Cauchy–Schwarz gives $|I| \leq Q^{-1/2} (\sum |A_m|^2)^{1/2} (\sum |B_m|^2)^{1/2}$, and two large sieves bound the two mean squares.

- *Additive large sieve.* A_m is a sum of $e(-md/s)$ over the reduced fractions d/s , which are $\gg 1/Y$ apart in $\mathbb{C}/\mathbb{Z}[\omega]$. The planar large sieve [1, Lemma lem:planar-additive-sieve], with separation $\asymp 1/Y$ and coefficient mass $O(1/Y)$, gives [1, Lemma lem:balanced-additive-norm] $\sum_{N(m) \leq Q} |A_m|^2 \ll (Q + Y^2)/Y$, which is Q/Y when $Y^2 \leq Q$, i.e. $Y \leq X$.
- *Theta function plus quadratic large sieve.* The automorphy of the cubic theta function (with Patterson's coefficients as recorded in [5]) turns B_m into a dual sum twisted by a *quadratic* symbol attached to m , of length about $N(m)^2/Z$ (no longer than Z when $N(m) \leq Z$), as in [3, Section 7]; no descent is needed, since B_m is already completed. The quadratic large sieve [6], with the powerful part of m counted separately, gives [1, Lemma lem:unmarked-completed-moment] $\sum_{N(m) \leq Q} |B_m|^2 \ll Z^\varepsilon (Q + Q^2/Z)$, i.e. the diagonal bound $\ll Q Z^\varepsilon$ when $Q \leq Z$.

Hence, when $Y \leq X$ and $XY \leq Z$,

$$|I_\eta(X, Y, Z)| \ll Q^{-1/2} \left(\frac{Q}{Y}\right)^{1/2} (Q Z^\varepsilon)^{1/2} = X^{1/2} Z^{\varepsilon/2},$$

and with $X = Y = Z^{1/2}$ this is [1, Prop. prop:balanced-low] $|I_\eta(Z^{1/2}, Z^{1/2}, Z)| \ll Z^{1/4+\varepsilon}$.

3.3 Second evaluation: rows and $1/L(s, \eta)$

Poisson summation in m replaces m by a dual frequency. The character ξ kills the frequency 0 and everything not prime to S , and every remaining frequency is uniquely $u a^6$ with u a sixth-power-free element, unit factor included (a *row*), and a primary. After summing over everything else, the probe becomes [1, L3831 and eq. probe-high]

$$\boxed{I_\eta(X, Y, Z) = \frac{1}{(2\pi i)^3} \iiint X^{1/2-z} Z^{x+z-1} Y^{w-1} e^{(x+z-1)^2} M(z) \widehat{W}_1(w) \times \sum_u^{(6)} \frac{\overline{\xi(u)}}{N(u)^z} \frac{\zeta_K^S(6z) L^S(w, \chi_\bullet(u))}{L^S(x, \eta \chi_\bullet(u))} \mathcal{H}_{\eta, u}(x, w, z) dz dw dx,} \quad (7)$$

where $\chi_\bullet(u)$ is the character $n \mapsto \chi_n(u) = (u/n)_6$, $\zeta_K = \zeta \cdot L(\cdot, \chi_{-3})$, the superscript S removes finitely many Euler factors, $\mathcal{H}_{\eta, u}$ is an absolutely convergent Euler product in explicit regions, and $M(z) = \int_0^\infty \widehat{W}_0(r) r^{z-1} dr$ and $\widehat{W}_1(w) = \int_0^\infty W_1(r) r^{w-1} dr$ are Mellin transforms of the weights ($\widehat{W}_0$ is the planar Fourier transform of $W_0(|y|^2)$). Each row is therefore a ratio of L -functions: the sum over a gives $\zeta_K(6z)$, the sum over s gives the numerator $L(w, \chi_\bullet(u))$, and the squarefree variable c , whose local factor carries a sign -1 coming from the Gauss sums (as in the Möbius–Gauss identity of [3, Section 6]), gives the reciprocal $1/L(x, \eta \chi_\bullet(u))$.

The principal row. Only $u = 1$ has a principal numerator (for a unit $u \neq 1$ the numerator character is a nontrivial Kummer character). Its two simple poles, at $w = 1$ (from $L^S(w, \chi_\bullet(1)) = \zeta_K^S(w)$) and at $z = 1/6$ (from $\zeta_K^S(6z)$), give the double residue

$$c_S X^{1/3} Z^{x-5/6} e^{(x-5/6)^2} \frac{H_\eta(x)}{L^S(x, \eta)}, \quad H_\eta(x) = \mathcal{H}_{\eta, 1}(x, 1, 1/6),$$

with a constant $c_S > 0$ [1, eq. shared-principal-data]. After division by c_S and with $X = Z^{1/2}$, this is $Z^{x-2/3} e^{(x-5/6)^2} H_\eta(x)/L^S(x, \eta)$: exactly the signal (4) with $C(s) = C_1(s) := s - 2/3$. This is where C_1 and the Gaussian $e^{(s-5/6)^2}$ come from.

The other rows: zero detection. A row $u \neq 1$ is large only if its twisted L -functions have zeros near the contour. The paper assigns to each row a bin $a \leq \beta_*$: the largest real part of a zero of height at most a small power Z^τ among the finitely many twists $\nu\chi_\bullet(u)^{\pm 1}$ attached to the row, rounded down to a fine grid with floor $51/100$. It writes $\delta = 2a - 1$. A zero ρ of $L(s, \psi)$, ψ the row character, produces [1, Prop. prop:detector-witness] two large Dirichlet polynomials with the same character and height: for $U = N(u)$,

$$M_r = U^{-r/2} \sum_n \mu(n) \psi(n) W\left(\frac{N(n)}{U^r}\right), \quad S_m = U^{-m/2} \sum_l \psi(l) W'\left(\frac{N(l)}{U^m}\right), \quad (8)$$

with $t - \frac{1}{2} \leq r \leq t$, $0 \leq m \leq \frac{1}{2}$ and $r + m \geq t$ for any chosen $t \in [1, 3/2]$, and

$$|M_r|^2 \gg U^{\delta r - \varepsilon}, \quad |S_m|^2 \gg U^{\delta m - \varepsilon}.$$

(Here M_r is a truncation of $1/L$ and S_m of L ; at a zero both are as large as the zero-free region allows.) Part I uses only M_r with $t = 1$. A large sieve for sextic characters (proved in [1, Lemma lem:sextic-large-sieve] by the method of [4]) bounds, for weights fixed independently of the row, the mean square

$$\sum_{N(u) \asymp U} |M_r(u)|^2 \ll (UD)^\varepsilon \{U + D + (UD)^{2/3}\}, \quad D = U^r,$$

for squarefree rows (rows with a square-full part are handled by fixing that part and counting it separately, with the same final result; the dependence of the witness weights on the row costs only a fixed power of Z^τ). Dividing by the spike $U^{\delta r}$ and taking the worst length $r \in [1/2, 1]$ (the endpoints $r = 1/2$ and $r = 1$) shows that at most $U^{R(\delta)}$ rows can be bad, with $R(\delta) = \min\{1, \max(1 - \delta/2, 4/3 - \delta)\}$ [1, Prop. prop:sextic-row-count].

3.4 Why Part I stops at 11/12

At the balanced scales $X = Y = Z^{1/2}$ the low estimate is $Z^{1/4}$ and the signal exponent is $C_I(\beta_*) = \beta_* - 2/3$. The criterion needs $1/4 < \beta_* - 2/3$, i.e. $\beta_* > 11/12$; so

$$\frac{11}{12} = \frac{2}{3} + \frac{1}{4}.$$

At these scales the non-principal rows have room to spare: by our computation from the paper's formulas, the row bins alone would allow any $\beta_* > 5/6$. (The paper's high-side lemmas are, however, stated only for $\sigma_0 \geq 7/8$; the Euler region it uses for the principal row and for H_η is $\Re s \geq 7/8$.) But the balanced scales are a choice. By Section 3.2 the low estimate is $X^{1/2}$, while the signal carries only $X^{1/3}$, so shrinking X improves the comparison on the low side. The price is a longer dual sum: rows u up to norm about Z/X (Section 4). Part I needs only a bootstrap and does not pursue this trade-off; the 7/8 argument is about how far it can be pushed.

4 The 7/8 probe

4.1 The exponent formula

With general scales $X = Z^{l_x}$, $Y = Z^{l_y} = Z^{l_x + b}$ and short prime sums of total log-length ℓ (below), the principal row gives [1, Def. def:stage-high-data]

$$C(s) = s - \frac{5}{6} + \frac{l_x}{3} + \frac{\ell}{6}, \quad h := 1 - l_x + \ell,$$

where h is the effective log-length of the dual rows ($N(u) \lesssim Z^h$), and, for the paper's geometry (in particular under the relation $l_x + l_y + \ell = 1$ below), the low estimate becomes $Z^{l_x/2+b/12}$. Equating the two,

$$\sigma_0 = 1 - \frac{h}{6} + \frac{b}{12}.$$

So σ_0 is governed by the length h of the dual rows: Part I has $h = 1/2$, and $7/8$ needs $h - b/2 = 3/4$. The paper's design relation $l_x + l_y + \ell = 1$, under which its low estimate is proved, makes this $\sigma_0 = 11/12 - \ell/4$, independent of b . Its choices are

$$b = \frac{1}{8}, \quad \ell = \frac{1}{6}, \quad h = \frac{13}{16}, \quad l_x = \frac{17}{48}, \quad l_y = \frac{23}{48}, \quad \sigma_0 = \frac{11}{12} - \frac{1}{24} = \frac{7}{8}, \quad C(s) = s - \frac{11}{16}. \quad (9)$$

Where the gain comes from. The low side does not need the prime sums. By Section 3.2, the unmodified probe with $X = Y = Z^{1/4}$ already has low estimate $Z^{1/8+\varepsilon}$, and the formula above gives $C(s) = s - 3/4$; so $\sigma_0 = 7/8$, with dual rows of length $h = 3/4$. What fails is the high side. There are about Z^h rows, and even the rows whose L -functions have no zeros to the right of $1/2$ must add up to less than the signal. By our computation from the paper's per-bin exponent [1, eq. stage-high-exponent], at real part $1/2$ with every row counted and no prime factors, these *floor* rows contribute $Z^{C(\sigma_0)+E}$, where

$$E = 8 - 9\sigma_0 + \frac{b}{6} - \ell \quad \left(\text{when } \sigma_0 = 1 - \frac{h}{6} + \frac{b}{12} \right).$$

So without slots the paper's row-by-row bounds cannot reach below $\sigma_0 = 8/9 + b/54$, for any choice of scales; at $X = Y = Z^{1/4}$ they give $E = +1/8$.

The prime slots pay for the longer rows. At fixed h and b they let X and Y both be Z^ℓ larger ($l_x = 1 + \ell - h$). This leaves σ_0 unchanged, since the low estimate and $C(s)$ both rise by $\ell/2$. But it lowers the exponent of every non-principal bin by at least ℓ (our computation from the same formula), because on such a row a slot contributes only a square-root-normalized prime polynomial (Section 4.4), while at the principal row it has its full size $Z^{\ell/6}$. The low side limits the slots: its estimate is proved under $l_x + l_y + \ell = 1$, which together with $\sigma_0 = 7/8$ forces $\ell = 1/6$, whatever b is. With the paper's geometry the floor rows have $E = -1/48$, and the binding bins are those with zeros near real part 0.69 (Section 6). Everything else in Part II exists to make the high estimate survive rows of length $h = 13/16$.

4.2 Prime slots and the compensated probe

Fix K disjoint "slots" of primes: slot i consists of primes $p \notin S$ in a fixed residue class (the identity class of a fixed finite group) with $N(p) \asymp P_i = Z^{\ell_i}$, weighted by a smooth bump Ω_i , where $\sum \ell_i = \ell = 1/6$. Write $I_{\eta;D}$ for the probe (5) with the condition $D \mid cn^3$ inserted in the completed index. The *compensated probe* is [1, eq. compensated-probe-definition, L6894]

$$I_{\text{mod}}(Z) = \sum_{p_1, \dots, p_K} \prod_{i=1}^K \Omega_i \left(\frac{N(p_i)}{P_i} \right) \sum_{J \subseteq \{1, \dots, K\}} (-1)^{|J|} N(p_J)^{-3/2} \overline{\eta(p_{J^c})} \times I_{\eta; p_{J^c}} \left(\frac{X}{N(p_J)}, \frac{Y}{N(p_J)}, Z N(p_{J^c}) \right), \quad (10)$$

with $p_J = \prod_{i \in J} p_i$. With one slot it reads

$$I_{\text{mod}}(Z) = \sum_p \Omega_1 \left(\frac{N(p)}{P_1} \right) \left[\underbrace{\overline{\eta(p)} I_{\eta; p}(X, Y, Z N(p))}_{\text{marked}} - \underbrace{N(p)^{-3/2} I_{\eta}(X/N(p), Y/N(p), Z)}_{\text{rescaled}} \right].$$

The *marked* term restricts the completed index to multiples of p and lengthens its range to $ZN(p)$. The *rescaled* term shrinks both averaging scales by $N(p)$.

4.3 Low side: still small

On the low side the marked primes are summed *inside* the completed row B_m , whose mean square stays diagonal, $\ll Q Z^\varepsilon$ [1, Lemma lem:probe-row-norm]; the proof now also uses Heath-Brown's cubic large sieve [7] for the marked primes [1, Lemma lem:hybrid-energy], and the relation $l_x + l_y + \ell = 1$. The additive mean square acquires a factor from the unequal scales, and each rescaled slot of length d gains a factor Z^{-d} (from the count Z^d , the coefficient $Z^{-3d/2}$ and $X^{1/2} \mapsto X^{1/2} Z^{-d/2}$). The result [1, Prop. prop:probe-low] is

$$|I_{\text{mod}}(Z)| \ll Z^{l_x/2+b/12+\varepsilon} = Z^{17/96+1/96+\varepsilon} = Z^{3/16+\varepsilon}.$$

The term $b/12$ is the price of the unequal scales. With $P_a = Y^2/Q = Y/X \approx Z^b = Z^{1/8}$, the paper's Gram bound gives $(Q/Y)(1 + P_a^{1/6} + P_a^2/Y)$ for the additive mean square [1, Prop. prop:probe-gram] (the planar sieve alone would give $(Q/Y)(1 + P_a)$); the middle term counts the rare frequencies that are sixth powers up to factors supported on the common divisor and on S . Since $P_a^2/Y \ll P_a^{1/6}$, the middle term dominates, and after Cauchy-Schwarz it becomes $P_a^{1/12} = Z^{b/12}$.

4.4 High side: what the compensation cancels

In the Poisson evaluation (7), a slot prime $p \nmid u$ multiplies row u by (roughly)

$$\text{marked: } -\overline{\chi_p(u)} + N(p)^{-w}, \quad \text{rescaled: } N(p)^{-w},$$

(times a common factor $N(p)^{z-1}$). The difference is $-\overline{\chi_p(u)}$ plus small terms [1, eq. old-eq:5.13b, Prop. prop:probe-errors]. (For the few slot primes $p \mid u$ the main term vanishes and $-N(p)^{-w}$ survives; the paper controls them through the drop in the row's conductor [1, eqs. old-eq:5.13d, old-eq:5.13e].) So each slot contributes, in row u , the factor $-P_i^{z-1/2} Q_i(u)$, where

$$Q_i(u) = P_i^{-1/2} \sum_{p \in \text{slot } i} \Omega_i\left(\frac{N(p)}{P_i}\right) \left(\frac{N(p)}{P_i}\right)^{z-1} \overline{\left(\frac{u}{p}\right)_6} \quad (11)$$

is a centrally normalized *prime polynomial*, typically of size $O(1)$. Without the subtraction, the row-independent term $N(p)^{-w}$ would act on *every* row like a prime polynomial of maximal size, and could not be controlled by any average over rows.

At the principal row $u = 1$ each slot gives -1 per prime. Hence the signal survives, multiplied by the nonzero normalizer

$$A_T(Z) = (-1)^K Z^{-\ell/6} \prod_{i=1}^K \sum_{p \in \text{slot } i} \Omega_i\left(\frac{N(p)}{P_i}\right) N(p)^{-5/6},$$

with $|A_T(Z)| \asymp (\log Z)^{-K}$, so that dividing by it costs only Z^ε ($A_T(Z) > 0$ when K is even), and the normalized probe $J_{\text{II},\eta} = I_{\text{mod}}/(c_S A_T(Z))$ has signal (4) with $C(s) = s - 11/16$; the factor $Z^{\ell/6}$ in C is the slots evaluated at $z = 1/6$.

5 Counting the bad rows

5.1 Three averages

A row in bin δ has prime polynomials of *mean amplitude* q : the length-weighted mean of the exponents g_i with $|Q_i| \approx P_i^{g_i}$. The bin ceiling $|Q_i| \ll P_i^{\delta/2+o(1)}$ gives $0 \leq q \leq \delta/2$, and suitably selected slots of total base- U length λ satisfy $|\prod_{\text{sel}} Q_i|^2 \geq U^{2q\lambda-o(1)}$. Such a row is detected by large values of M_r , S_m and the Q_i . Three averages over rows bound how many such rows exist.

1. Inverse moment with prime factors [1, Lemma lem:marked]. With row range H , length D , slot sizes P_i , $P = \prod_i P_i$, and fixed residue-class characters ν and coefficients $a_i(p)$,

$$\sum_{0 < N(u) \ll H} \left| D^{-1/2} \sum_n \mu(n) \nu(n) \left(\frac{u}{n} \right)_6 W \left(\frac{N(n)}{D} \right) \prod_i P_i^{-1/2} \sum_{p \in P_i} a_i(p) \nu(p) \left(\frac{u}{p} \right)_6 \Omega_i \left(\frac{N(p)}{P_i} \right) \right|^2 \ll H^{1+\varepsilon}, \quad (12)$$

provided $DP^2 \leq HZ^{-c_1}$ and $D^2P^8 \leq H^3Z^{-c_2}$ (fixed $c_1, c_2 > 0$). With no prime factors this is exactly the mean-square estimate behind the alternate 11/12 proof of [2] (explained in [3]; Part I of [1] does not use it), $\sum_{N(u) \leq D^{1+\vartheta}} |\sum_n \mu(n) (u/n)_6 W(N(n)/D)|^2 \ll D^{2+\vartheta+\varepsilon}$, and its proof follows the same route: Poisson in u , the identity converting Möbius coefficients into cubic Gauss sums, the cubic theta function with the quadratic large sieve, and the two-Poisson descent [3, Sections 6–8]. With prime factors the paper adds a reduction for slot primes dividing n and, in the terminal step, Heath-Brown’s cubic large sieve for the slot primes [1, Lemma lem:hybrid-energy].

2. Plain fourth moment with prime factors [1, Lemma lem:plain]. With lengths Z^{n_1}, Z^{n_2} and rows $N(k) \ll Z^{m'}$ (excluding rows whose character is principal, or lies in a fixed finite group, when prime factors are present),

$$\sum_k \left| \prod_{j=1}^2 Z^{-n_j/2} \sum_l \tau(l) \left(\frac{k}{l} \right)_6 W_j \left(\frac{N(l)}{Z^{n_j}} \right) \cdot \prod_i P_i^{-1/2} \sum_p \tau(p) \nu_i(p) \left(\frac{k}{p} \right)_6 \Omega_i \left(\frac{N(p)}{P_i} \right) \right|^2 \ll Z^{m'+\varepsilon}, \quad (13)$$

for all lengths when there are no prime factors, and otherwise, for $3/4 \leq \kappa \leq 1$ and slots of sufficiently small individual length, when $n_1 + n_2 + 6\kappa\lambda \leq m'$, where $\kappa = 2\beta_* - 1$ (we state the case used in the application, where the paper’s moving-support parameter is 0). This is new. It is an induction on the effective width: the functional equation of the sextic twists $L(s, \psi)$ (Hecke reflection) shortens the sums, and two finite Poisson transforms per step return products of the same kind at smaller width. Transformed rows whose character lies in the fixed finite group can carry volume-sized main terms, which an equal-scale centered difference cancels. The zero-free hypothesis $\beta_* \leq (1 + \kappa)/2$ (which holds, with equality, for $\kappa = 2\beta_* - 1$) bounds the prime factors pointwise, $|\prod_i Q_i|^2 \ll Z^{\kappa\lambda+\varepsilon}$.

3. Sixth-power amplification [1, Lemma lem:inverse-amplification]. Replacing a sixth-power-free row u ($N(u) \asymp U$) by ua^6 for the $(H/U)^{1/6}$ ideals a with $N(a) \leq (H/U)^{1/6}$, where $H = \max(2U, D^{1+c})$, and applying (1) without prime factors at row range H (the same trick as in the deduction of 11/12 in [2, 3]), gives $\sum_{N(u) \asymp U} |M_r(u)|^2 \ll U^{e(r)+\varepsilon}$ with $D = U^r$ and $e(r) = \max\{1, (1 + 5r)/6\}$.

5.2 Combining the counts

Apply the moments at base U : rows have $N(u) \asymp U$, so the row exponent is 1, and r, m, λ are base- U lengths. Dividing each moment by the guaranteed size of a bad row ($|M_r|^2 \gg U^{\delta r}$, $|S_m|^4 \gg U^{2\delta m}$,

$|\prod_{\text{sel}} Q_i|^2 \gg U^{2q\lambda}$) gives, for $\#\mathcal{B}$ rows of norm $\asymp U$ in one bin,

$$\#\mathcal{B} \ll U^{1-\delta r-2q\lambda+\varepsilon} \quad (\text{inverse}), \quad \#\mathcal{B} \ll U^{1-2\delta m-2q\lambda+\varepsilon} \quad (\text{plain, two copies of } S_m),$$

with the prime-factor “capacities” $\lambda \leq (1-r)/2$ (from $DP^2 < H$; the second condition then needs $r > 1/2$, and the inverse count is used only for $r \geq 23/37$) and $\lambda \leq (1-2m)/(6\kappa)$ (from $2m+6\kappa\lambda \leq 1$). The inverse bound is good when the Möbius witness is long, the plain bound when the plain witness is long; since $r+m \geq t$, at least one witness is long, and the better of the two bounds is used. Writing $q = x\delta$ with $0 \leq x \leq 1/2$, taking the better of the two, and using the amplified count $L(t) = 1 - \delta + (\alpha - \delta)(t - 1)$, $\alpha = 5/6$, for long inverse witnesses, gives [1, Prop. prop:detector-counts] (assuming the available prime length exceeds $7/37$ by a fixed margin; at the longest rows it is $\ell/h = 8/39$)

$$\#\mathcal{B} \ll U^{\max\{R_{\text{short}}(t), L(t)\} + \Delta/4 + \varepsilon}, \quad R_{\text{short}}(t) = 1 - \delta + \frac{\delta P_x}{D_x} \left(\frac{3}{2} - t \right),$$

where $D_x = 3 - 17x/9$, $P_x = (2 - 8x/9)(1 - x)$, and $\Delta = \beta_* - 7/8$.

6 The final inequality

A bin of rows of norm $U = Z^d$ contributes at most $Z^{C(7/8)+E(d)+o(1)}$ to the probe, where [1, eq. common-high-exponent]

$$E(d) = -\frac{1}{48} + \frac{2}{3}\delta + \frac{q}{6} - h(1 - R) + (d - h) \left(R + \frac{\delta}{2} - \frac{17}{50} \right).$$

Its terms record the Mellin weight on the central contour, the size $U^{\delta/2}$ of the numerator $L(w, \chi_\bullet(u))$ on the reflected line, the number U^R of rows, and the prime polynomials. The slope in d is positive, so the worst rows are the longest, $d = h = 13/16$. There one takes $t = 1 + \delta P_x/(2\mathcal{J})$ with $\mathcal{J} = (\alpha - \delta)D_x + \delta P_x$, so that $R = R_* = 1 - \delta + (\alpha - \delta)\delta P_x/(2\mathcal{J})$, and the paper proves [1, Lemma lem:balanced-endpoint]

$$-E(h) \geq \frac{49}{440640} \quad \text{for all } 0 \leq \delta \leq \frac{5}{6}, \quad 0 \leq x \leq \frac{1}{2}.$$

Including the small loss from $\kappa = 2\beta_* - 1$, the comparison actually needed is

$$E_{\text{actual}}(h) - \Delta \leq -\frac{49}{440640} - \frac{51}{64}\Delta + \epsilon_{\text{total}} < 0,$$

where ϵ_{total} collects the adjustable losses (chosen small after the certificate); i.e. every non-principal bin is smaller than the signal size $Z^{C(\beta_*)} = Z^{C(7/8)+\Delta}$ by a fixed power. Rows slightly longer than Z^h ($h < d \leq h + \zeta$) cost at most 2ζ , absorbed by taking ζ small. The top bin $\delta = 5/6$ (no selected primes; $E(h) \leq -1/48 - \delta/16$) and the floor bin ($E \leq -7/1200$) are separate easy cases, and shorter rows ($d \leq 1/2$, no prime factors needed), very short rows and very long rows ($d > h + \zeta$) are handled with room to spare. All margins are chosen before the target η , so they are uniform in η , and Proposition 1 with $\sigma_0 = 7/8$ gives the contradiction. Hence $\beta_* \leq 7/8$, and $\zeta(s) \neq 0$ for $\Re s > 7/8$.

The worst case. At $x = 1/2$ (prime polynomials as large as the bin allows) and $\delta \approx 0.387$ (zeros with real part about 0.69) one finds $R_* \approx 0.668$ and $E(h) \approx -2.3 \times 10^{-4}$: the inequality holds, but only just.

7 What each ingredient does

Ingredient	Role
Family supremum β_*	Gives a priori control of every twisted L -function that appears as a row; the contradiction may come from any member.
Bootstrap $\beta_* \leq 11/12$	From Part I. Gives $\Delta \leq 1/24$ and $\kappa = 2\beta_* - 1 \leq 5/6$; since $\delta \leq 2\beta_* - 1$ by the definition of β_* , $\delta \leq 5/6$ for every bin, as used in the counts and the final certificate.
Prime slots, $\ell = 1/6$	Make the long dual rows ($h = 13/16$) affordable: at fixed h they let X, Y be Z^ℓ larger, which leaves $\sigma_0 = 1 - h/6 + b/12$ unchanged and lowers every non-principal bin; full size at $u = 1$, square-root-size prime polynomials elsewhere. The low side caps them at $\ell = 1/6$.
Compensation (marked – rescaled)	Removes the row-independent term $N(p)^{-w}$, leaving the row polynomial $-\chi_p(u)$; free on the low side; keeps the principal row nonzero.
Unequal scales, $b = 1/8$	Cost $Z^{b/12}$ on the low side, gain on the high side; b cancels from σ_0 under $l_x + l_y + \ell = 1$.
Two witnesses M_r, S_m	Complementary detectors from one zero ($r + m \geq t$): a long Möbius witness is counted by the inverse moment (or by sixth-power amplification when $r \geq 1$), a long plain witness by the plain fourth moment.
Inverse moment	The machinery of the alternate 11/12 proof [2] (cubic theta, quadratic large sieve, two-Poisson descent), extended to prime factors with Heath-Brown’s cubic large sieve.
Plain fourth moment	New: an induction combining Hecke reflection with two finite Poisson transforms (and a centered difference for rows in the fixed finite group); prime factors bounded pointwise using $\beta_* \leq (1 + \kappa)/2$, which holds with equality for $\kappa = 2\beta_* - 1$.

How tight is it? The certificate guarantees a margin of $49/440640 \approx 1.1 \times 10^{-4}$; the true worst case is about 2.3×10^{-4} . We re-evaluated the paper’s general exponent [1, eq. stage-high-exponent] at $d = h$ with the same row counts (which do not depend on ℓ or b while $\ell/h > 7/37$), together with $\sigma_0 = 1 - h/6 + b/12$ and $l_x + l_y + \ell = 1$ (our computation, not a claim of [1]). The endpoint inequality then survives only for $\ell \lesssim 0.1668$ when $b = 1/8$, and only for $0.094 \lesssim b \lesssim 0.152$ when $\ell = 1/6$; at $\ell = 0.17$ or $b = 0$ the worst bin exceeds the signal by about $Z^{0.004}$. So $7/8$ appears to be close to the limit of this particular design.

References

- [1] OpenAI, *The Quasi-Riemann Hypothesis: A Zero-Free Half-Plane $\Re s > 7/8$* , manuscript dated September 30, 2026, github.com/openai/math. Line references are to `build/paper.tex`.
- [2] OpenAI, *The Quasi-Riemann Hypothesis*, manuscript dated October 5, 2026 (an alternate proof of the 11/12 half-plane), same repository.

- [3] Claude Opus 5.5 and OpenAI Astra, *Sixth powers in a mean square: an accessible guide to the 11/12 argument for $\zeta(s)$* , companion note, https://www.math.uwaterloo.ca/~mrubinst/quasi_rh_notes/zeta_11_12_guide.pdf.
- [4] V. Blomer, L. Goldmakher and B. Louvel, *L-functions with n -th-order twists*, Int. Math. Res. Not. IMRN (2014), no. 7, 1925–1955.
- [5] A. Dunn and M. Radziwiłł, *Bias in cubic Gauss sums: Patterson’s conjecture*, Ann. of Math. (2) **200** (2024), 967–1057.
- [6] L. Goldmakher and B. Louvel, *A quadratic large sieve inequality over number fields*, Math. Proc. Cambridge Philos. Soc. **154** (2013), 193–212.
- [7] D. R. Heath-Brown, *Kummer’s conjecture for cubic Gauss sums*, Israel J. Math. **120** (2000), 97–124.