

Closed-Form Phase Prediction and Spectral-Leakage Modeling for the Explicit Formula of $L(s, \chi_3)$, with Blind FFT Detection

Stefka Georgieva

VPR Research Ltd · October 2026 · Preprint

DOI: 10.5281/zenodo.23180120 · Code on GitHub

ABSTRACT

The nontrivial zeros of an L-function appear as spectral peaks in the logarithmic Fourier transform of the prime-power sum $\psi(x, \chi)/\sqrt{x}$. We show that, for the quadratic character χ_3 , the *complex* FFT spectrum of this signal can be predicted from the zero list alone, with no free parameter. Two closed-form ingredients make this possible: an explicit expression for the phase of the FFT coefficient at each bin (the analytic phase of the explicit formula, a sampling-origin term, and the linear phase of the Kaiser analysis window), and a multi-zero leakage model that sums the windowed contribution of every zero using the exact numerically evaluated DTFT of the window. Both are elementary consequences of the explicit formula and of Fourier-window geometry; what we contribute is their explicit form for this setting and a systematic quantitative validation.

At $N = 5 \times 10^8$, blind peak detection recovers peaks at all 17 reference zeros ($\gamma \leq 47.5$), each within 0.62 FFT bins. For isolated zeros the phase formula matches the measured phase to a median of 4×10^{-4} **rad**, and, once the window response at the actual bin offset is included, the single-zero amplitude formula holds to $\pm 0.1\%$ (if that response is ignored the amplitude error reaches 6%). The multi-zero model with $C = 1$ reproduces the complex spectrum at all 17 zeros, including pairs only about three bins apart, with a mean relative residual of **0.0009**, against 0.039 for a single-zero model without window response. The residual (0.0008–0.0014) and the fitted normalisation constant ($|C| = 1.0000 \pm 3.4 \times 10^{-4}$) are stable across ten $(N, x_{\min})$ configurations; the factor by which the multi-zero model improves on the single-zero baseline (18–65 $\times$) is not, because it depends on how the FFT bins fall relative to the zeros.

The same code applied to $\zeta(s)$ (60 zeros, $\gamma \leq 163$) shows the same behaviour. The relative residual grows approximately in proportion to γ and, at equal height ($\gamma < 50$), is the same for the two functions (0.0008 and 0.0009). The model is a calibrated forward model of the log-space spectrum: it can serve as an end-to-end consistency check of prime-sum computations, as a basis for refining peak positions and separating close pairs, and for choosing analysis parameters in advance. It does not locate new zeros and does not test the Riemann hypothesis. Code, per-zero tables, and the robustness sweeps are provided.

KEY RESULTS AT A GLANCE

- **What is modelled.** The complex FFT spectrum of $\psi(x, \chi_3)/\sqrt{x}$ in log-space, predicted from the zero list with no fitted parameter (Eqs. 5, 8, 11).
- **Detection.** 17/17 reference zeros matched within 0.62 bins at $N = 5 \times 10^8$; for $\gamma \leq 55$ the 20 detected peaks and the 20 known zeros coincide one-to-one (Fig. 1).
- **Phase.** Median error 4×10^{-4} rad for isolated zeros; Eq. (8) contains no fitted constant.
- **Amplitude and leakage.** Complex-spectrum residual 0.0009 (multi-zero, $C = 1$) against 0.039 (single zero, no window); the improvement decomposes into 7.4 $\times$ from the window response of the same zero and 5.8 $\times$ from neighbouring zeros (Section 3.4).
- **Second L-function.** $\zeta(s)$, 60 zeros: same behaviour; the residual scales approximately with γ and equals that of χ_3 at equal height (Appendix C).
- **Not claimed.** No new zeros, no test of GRH, real characters only; the origin of the 10^{-3} residual floor is not analysed.

1. Introduction

For a Dirichlet character χ modulo q , the associated L-function is defined by the Dirichlet series [2]

$$L(s, \chi) = \sum_{n=1}^{\infty} \frac{\chi(n)}{n^s},$$

which converges absolutely for $\Re(s) > 1$ and extends meromorphically to the entire complex plane. For a nontrivial character, $L(s, \chi)$ is entire and satisfies a functional equation analogous to that of the Riemann zeta function. We focus here on the **nontrivial quadratic character modulo 3**, $\chi_3(n) = (n/3)$, the Legendre symbol:

$$\chi_3(n) = \begin{cases} +1 & \text{if } n \equiv 1 \pmod{3}, \\ -1 & \text{if } n \equiv 2 \pmod{3}, \\ 0 & \text{if } 3 \mid n. \end{cases}$$

This character governs the "race" between primes $\equiv 2 \pmod{3}$ and $\equiv 1 \pmod{3}$, which is equivalent to the classification $6n - 1$ versus $6n + 1$, since every prime > 3 is of the form $6n \pm 1$ with $6n - 1 \equiv 2 \pmod{3}$ and $6n + 1 \equiv 1 \pmod{3}$.

The Chebyshev-type function associated with χ_3 is

$$\psi(x, \chi_3) = \sum_{p^k \leq x} \chi_3(p^k) \log p,$$

where χ_3 is completely multiplicative, so $\chi_3(p^k) = \chi_3(p)^k$. Unlike the Riemann zeta case, there is **no main linear term** to subtract: $\psi(x, \chi_3)$ is directly the fluctuation. Under the Generalized Riemann Hypothesis it is $O(\sqrt{x} \log^2 x)$ [1].

The explicit formula for $L(s, \chi)$ (see [3]) relates $\psi(x, \chi)$ to the nontrivial zeros $\rho = 1/2 + i\gamma$ of $L(s, \chi)$:

$$\psi(x, \chi) = - \sum_{\rho} \frac{x^{\rho}}{\rho} - \dots,$$

where the sum runs over the nontrivial zeros. Each zero contributes an oscillatory term

$$\frac{x^{\rho}}{\rho} = \frac{\sqrt{x} e^{i\gamma \ln x}}{\rho},$$

which, after dividing by $\sqrt{x}$, becomes a pure sinusoid in $u = \ln x$ with angular frequency γ . A Fourier transform of $\psi(x, \chi_3)/\sqrt{x}$ therefore exposes the zeros γ directly as spectral peaks.

1.1 Scope and contributions

The paper consists of four layers, which we keep separate throughout:

1. **Analytic derivation.** From the pair-sum representation of the explicit formula, each zero $\rho_j = 1/2 + i\gamma_j$ contributes a sinusoid with a complex coefficient $A_0(\gamma_j) = -1/(1/2 + i\gamma_j)$, whose modulus and argument are explicitly determined (Eqs. 4–5).
2. **Discrete Fourier geometry.** When the signal is sampled in log-space and analyzed with a Kaiser window, the phase of the resulting FFT coefficient at bin k acquires three additional contributions: the sampling origin ($\gamma u_{\min}$), the window's linear-phase (centering) term $(\omega_0 - \omega_k)(N_{\log} - 1)/2$, and the analytic phase $\pi - \theta$ from the coefficient itself (Eqs. 6–8). The amplitude at a bin is likewise weighted by the window response at the bin offset (Eq. 5).
3. **Multi-component spectral model.** Because the window acts linearly, the FFT coefficient at bin k is the sum of the windowed contributions of every zero. Using the exact numerically evaluated DTFT of the Kaiser window, we obtain a quantitative prediction for the complex spectrum at each bin (Eqs. 9–11).
4. **Numerical validation.** The predictions are checked against the 17 reference zeros of $L(s, \chi_3)$ with $\gamma \leq 47.5$ in ten $(N, x_{\min})$ configurations, and the pipeline is replicated on $\zeta(s)$ with 60 zeros up to $\gamma \leq 163$.

Each ingredient is individually standard: the analytic phase follows directly from the explicit formula, the window linear-phase term is classical in spectral analysis [7], and the structural form of the multi-zero sum follows from linearity. The contributions claimed here are:

- i. the explicit closed form of the FFT-bin phase for this log-space setting (Eq. 8), validated to 4×10^{-4} rad;
- ii. an amplitude formula that includes the window response at the bin offset (Eq. 5), which removes a scalloping error of up to 6% in the naive single-zero amplitude;

- iii. a decomposition of the spectral model into the window response of the same zero and the leakage from neighbouring zeros, with parameter-free versions of both (Section 3.4);
- iv. the observation that the residual is stable in N and $x_{\min}$, grows approximately in proportion to γ , and is the same for χ_3 and ζ at equal height (Section 3.6 and Appendix C).

1.2 What is and is not blind in this work

Our numerical procedure has two stages with different epistemic status, and we are explicit about the distinction here to avoid overstating the results.

Blind detection stage. The FFT is computed without any knowledge of the zeros of $L(s, \chi_3)$. Peaks are detected purely from the spectrum of $\psi(x, \chi_3)/\sqrt{x}$, using a fixed amplitude threshold. The reference zero list is not consulted, and the detected peak positions are determined entirely from the data. This is a genuine blind test of the detection step.

Reference-based validation stage. Once the peaks have been detected, we test whether the explicit formula correctly predicts their *amplitudes*, *phases*, and *complex spectral residuals*. Because the formula involves the zeros explicitly, these tests necessarily use the reference zero list as input to the model. They are *controlled numerical tests* of a theoretical prediction, not blind discoveries. We report both stages separately and label them accordingly throughout.

1.3 Replication on $\zeta(s)$

The same pipeline — log-space FFT with a Kaiser window, the closed-form phase (Eq. 8), and the multi-zero model (Eq. 11) — is applied to the Riemann zeta function, using the Chebyshev function $\psi(x) - x$ in place of $\psi(x, \chi_3)$. The code is shared; the L -function-specific inputs are the signal and the reference zero list. The $\zeta(s)$ runs use 60 zeros ($\gamma \leq 163$), against 17 for χ_3 ($\gamma \leq 47.5$), a matching tolerance of 3% for both, and the same threshold, window, and grid. All $\zeta(s)$ numbers in this paper were recomputed with the released code and with a zero list generated independently [8]. Two L -functions, both with real coefficients and symmetric zero sets, do not establish generality; see Section 3.6 and Appendix C.

1.4 What the model is for

The zeros of $L(s, \chi_3)$ and $\zeta(s)$ are known to high precision, and they have been verified directly to far greater heights than are accessible here. This work therefore does not locate new zeros and does not test the Riemann hypothesis. Its purpose is to turn the log-space Fourier transform of ψ into a calibrated measurement: given the zeros, the complex spectrum is predicted to a relative accuracy of about 10^{-3} , with no fitted parameter. Intended uses are:

- **End-to-end consistency check.** Agreement at the 10^{-3} level between the measured spectrum and the prediction validates the chain prime sieve $\rightarrow \psi(x, \chi) \rightarrow$ resampling $\rightarrow$ windowed FFT $\rightarrow$ zero list in a single test.
- **Peak refinement and close pairs.** Because the shape of each peak, including the leakage of its neighbours, is known, fitting Eq. (11) to the measured spectrum is a natural route to positions finer than the bin spacing and to separating pairs that the bare spectrum merges. We demonstrate the first ingredient (for χ_3 pairs about three bins apart are reproduced to 10^{-3} ; for ζ the two closest pairs, about two bins apart, to 4×10^{-3} , Appendix C); the fitting itself is future work.
- **A priori parameter choice.** The window response predicts the scalloping loss, the leakage between zeros, and hence which values of N and $x_{\min}$ give a given accuracy at a given bin alignment (Section 3.5).
- **A baseline for departures from the model.** Any component not of the form (4) would appear as structure above the residual floor. The floor is a numerical-setting statement, not a competitive bound on the location of zeros.

2. Method

2.1 The Chebyshev function for χ_3

For a given upper bound N , we compute

$$\psi(x, \chi_3) = \sum_{\substack{p^k \leq x \\ p \text{ prime}, k \geq 1}} \chi_3(p^k) \log p \quad (0 \leq x \leq N),$$

using a standard sieve of Eratosthenes to enumerate primes up to N and accumulating $\chi_3(p^k) \log p$ at each prime power $p^k \leq N$. Note that $\chi_3(3) = 0$, so $p = 3$ contributes nothing. Because χ_3 is completely multiplicative, $\chi_3(p^k) = \chi_3(p)^k$, so the contribution of p^k alternates in sign with k . A cumulative sum yields $\psi(x, \chi_3)$ at every integer x .

2.2 Log-space resampling and the resolution criterion

Because the oscillations in the explicit formula are pure sinusoids in $u = \ln x$, we resample $\psi(x, \chi_3)/\sqrt{x}$ at $N_{\log} = 2^{17}$ points on the half-open log-grid $u_n = u_{\min} + n \Delta u$ with $\Delta u = (u_{\max} - u_{\min})/N_{\log}$ and $n = 0, \dots, N_{\log} - 1$, where $u_{\min} = \ln x_{\min}$ and $u_{\max} = \ln N$. This is the sampling convention assumed by the discrete Fourier transform: the sampled sequence is treated as periodic with period $N_{\log} \Delta u = u_{\max} - u_{\min}$. We then subtract the mean and apply a Kaiser window [5] with shape parameter $\beta = 14$. The spacing of the FFT bins in γ is

$$\Delta\gamma = \frac{2\pi}{\ln N - \ln x_{\min}} \quad (1)$$

(for $N = 5 \times 10^8$ and $x_{\min} = 150$, $\Delta\gamma = 0.418$). *Implementation.* Because ψ is constant between integers, the resampled value at $x = e^{u_n}$ is exactly $\psi(\lfloor e^{u_n} \rfloor)$, and the normalisation uses $\sqrt{\lfloor x \rfloor}$. The window is the symmetric Kaiser window of length $N_{\log}$ (`scipy.signal.windows.kaiser`, $\beta = 14$), which the phase formula of Section 2.5 assumes.

Resolution criterion. The bin spacing (1) is not the resolution: the effective resolution is set by the main lobe of the window, whose full width at half amplitude is $\kappa \Delta\gamma$ with $\kappa \approx 2.87$ for $\beta = 14$ (equivalent noise bandwidth 2.16 bins). The zero-counting function of a primitive character of conductor q is $N(T, \chi) \sim (T/2\pi) \log(qT/2\pi e)$ (see [3]), so the *mean* spacing of consecutive zeros at height T is

$$\delta(T, \chi) \approx \frac{2\pi}{\ln(qT/2\pi)}. \quad (2)$$

Requiring $\kappa \Delta\gamma < \delta(T, \chi)$ gives

$$\boxed{x_{\min} < N \left(\frac{2\pi}{qT} \right)^{\kappa}}. \quad (3)$$

For $\kappa = 1$ this reduces to the weaker bound $2\pi N/(qT)$ obtained from the bin spacing alone. Equation (3) is a heuristic about the *mean* spacing, not a sufficient condition: consecutive zeros with an anomalously small gap can remain unresolved even when (3) holds (Section 3.1). For $q = 3$, $T = 50$ it gives $x_{\min} \lesssim 5.5 \times 10^4$ at $N = 5 \times 10^8$ and 5.5×10^3 at $N = 5 \times 10^7$. All values used here ($x_{\min} \leq 300$) are at least 18 times below this bound, so the sweep in Appendix B shows insensitivity to $x_{\min}$ inside the admissible region but does not probe the bound itself.

2.3 Blind peak detection and matching

Applying a real FFT to the windowed signal yields a spectrum $S[k]$. We restrict to positive frequencies f_k (cycles per unit u) and convert them to $\gamma = 2\pi f_k$. Local maxima above a fraction $\theta = 5\%$ of the maximum amplitude are retained as candidate zeros (82 maxima for $N = 5 \times 10^8$, $x_{\min} = 150$). **No information about the true zeros is used at this stage.**

Matching to the reference zeros. To compare with the 17 reference zeros we use a one-to-one assignment (Hungarian algorithm, minimising the total $|\Delta\gamma|$) in which every local maximum is a candidate and a zero is left unmatched if no candidate lies within 3% of its γ . Amplitudes play no role in the assignment. Detected peaks lie on the FFT bin grid, so we report deviations both in percent and in bins; a tolerance in percent alone is a poor criterion at low γ , where 1% of γ is smaller than half a bin.

2.4 Amplitude prediction

Once the peaks have been detected, we compare their amplitudes against a closed-form prediction. As with the phase prediction in the next section, **this stage uses the reference zero list** as input to the model; it is a controlled test, not a blind discovery.

Derivation. From the pair-sum representation (6) derived in the next section, each zero contributes a sinusoid in $u = \ln x$ with complex coefficient

$$A_0(\gamma) = -\frac{1}{|\rho|} e^{-i\theta}, \quad |\rho| = \sqrt{1/4 + \gamma^2}, \quad \theta = \arctan(2\gamma). \quad (4)$$

A single sinusoid of this form, windowed and transformed, gives at bin k the coefficient $C A_0(\gamma) e^{i\gamma u_{\min}} W(\omega_k - \omega_0)$ with W the window DTFT (Eq. 10), $\omega_0 = \gamma \Delta u$ and $\omega_k = 2\pi k / N_{\log}$. Its modulus predicts the amplitude of the FFT peak at an isolated zero:

$$|S[k]| \approx \frac{|C| |W(\omega_k - \omega_0)|}{\sqrt{1/4 + \gamma^2}}. \quad (5)$$

The peak bin is generally offset from the zero by up to half a bin, so $|W(\omega_k - \omega_0)| < |W(0)|$: the *scalping loss* is about 1.3% at 0.2 bins and 6.7% at 0.46 bins for the Kaiser window with $\beta = 14$. A naive amplitude formula replaces the window factor by its peak value $|W(0)|$, i.e. $\text{amp}(\gamma) \approx C_{\text{amp}} / \sqrt{1/4 + \gamma^2}$ with a single real constant C_{amp} ; we call this the *naive* form and report both in Section 3.2. The constants are calibrated from **one isolated zero** (the first matched zero with isolation gap > 2.0 , $\gamma = 8.04$), which is excluded from the reported statistics. The two calibrations are related by $C_{\text{amp}} = |C| |W(\omega_{\text{cal}} - \omega_{0,\text{cal}})|$; this is an identity of the same calibration point, not an independent test.

2.5 Phase prediction: a closed-form formula for $\arg(S[k])$

We now derive the closed-form prediction for the complex phase of each FFT peak. **This stage also uses the reference zero list** as input to the model: each zero's known γ appears explicitly in the phase formula. It is therefore not blind, but a controlled test: given the theoretical positions, does the FFT reproduce the theoretical phases?

Derivation. For a real primitive character, the zeros on the critical line come in conjugate pairs $\rho = 1/2 \pm i\gamma$, and the pair sums to a real cosine term:

$$\frac{\psi(x, \chi_3)}{\sqrt{x}} \sim -2 \sum_{\gamma > 0} \frac{\cos(\gamma u - \theta)}{|\rho|}, \quad u = \ln x, \quad \theta = \arctan(2\gamma), \quad |\rho| = \sqrt{1/4 + \gamma^2}. \quad (6)$$

In the frequency domain, each such term is a sinusoid $e^{i\gamma u}$ with complex coefficient $A_0(\gamma)$ given by (4), so

$$\arg A_0(\gamma) = \pi - \theta. \quad (7)$$

Three additional phase contributions must be tracked to match the FFT convention exactly:

1. **Sampling origin.** The signal is sampled starting at $u = u_{\min}$, not at $u = 0$. A sinusoid $e^{i\gamma u}$ sampled this way acquires a phase $+\gamma u_{\min}$.
2. **Kaiser-window centering.** A symmetric real window w_n of length $N_{\log}$ has DTFT $W(\omega) = e^{-i\omega(N_{\log}-1)/2} R(\omega)$ with R real and even. The windowed FFT at bin k therefore acquires a linear-phase term $(\omega_0 - \omega_k)(N_{\log} - 1)/2$, where $\omega_0 = \gamma \cdot \Delta u$ is the true digital frequency and $\omega_k = 2\pi k / N_{\log}$ is the actual bin frequency. This linear-phase contribution is classical in spectral analysis [7].
3. **Main-lobe response.** For a well-resolved, isolated zero, the main-lobe response $R(\omega_k - \omega_0)$ is positive near the peak, so it contributes no additional phase.

Combining (7) with the three corrections, we obtain our closed-form phase prediction:

$$\arg(S[k]) \approx (\pi - \theta) + \gamma u_{\min} + (\omega_0 - \omega_k) \frac{N_{\log} - 1}{2} \pmod{2\pi}, \quad (8)$$

with $\theta = \arctan(2\gamma)$, $\omega_0 = \gamma \cdot \Delta u$, $\omega_k = 2\pi k / N_{\log}$, and $\Delta u = (\ln N - \ln x_{\min}) / N_{\log}$.

Equation (8) is our phase-prediction formula. Each of its three components is individually standard: the analytic phase follows directly from (7), the sampling-origin term is immediate, and the window linear-phase term is classical [7]. We are not aware of an earlier explicit statement of this FFT-bin phase formula for the log-space transform of the explicit-formula signal. Eq. (8) contains no fitted parameter, so every matched zero is a legitimate test, including the zero used to calibrate the amplitude. For a well-resolved isolated zero it gives the predicted phase at an FFT bin k in the resolved main-lobe neighborhood; the assumption of a positive main-lobe response is precisely what fails for close-neighbor zeros, which is why those cases show larger phase deviations in Section 3.3. For isolated zeros the observed median absolute error is 4×10^{-4} rad (0.02°).

2.6 Multi-zero spectral-leakage model

For a zero with a close neighbor, the single-zero predictions (5) and (8) are no longer sufficient. We now derive a quantitative model for the neighbor's contribution. **This stage again uses the reference zero list**, as a controlled test of whether the explicit formula explains the observed spectrum given the known zeros.

Derivation. The windowed FFT at bin k is a sum over all frequencies present in the windowed signal. Because the signal is a superposition of contributions from all zeros, and because the window acts linearly, we may write

$$S[k] = \sum_i A_0(\gamma_i) e^{i\gamma_i u_{\min}} W(\omega_k - \omega_i) + \varepsilon_k, \quad (9)$$

where $W(\omega)$ is the *exact* DTFT of the Kaiser window used in the FFT, evaluated numerically, and $\omega_i = \gamma_i \Delta u$ is the digital frequency of the i -th zero.

$$W(\omega) = \sum_{n=0}^{N_{\log}-1} w_n e^{-i\omega n}. \quad (10)$$

Introducing a single complex constant C to absorb any residual normalization mismatch (FFT scaling, the $\sqrt{x}$ factor, the exact window definition), we obtain the multi-zero model:

$$S[k] \approx C \sum_i A_0(\gamma_i) e^{i\gamma_i u_{\min}} W(\omega_k - \omega_i). \quad (11)$$

Equation (11) is our multi-zero leakage model. The structural form of the sum is standard Fourier-window geometry. The contribution claimed here is the specific predictive model obtained by using the analytic complex coefficients $A_0(\gamma_i)$ together with the exact numerically evaluated DTFT (10), summed over *every* known zero rather than only the nearest one, and its validation without a fitted parameter ($C = 1$).

Models compared. Throughout Section 3 we compare five predictions of the complex coefficient $S[k]$ at the detected bin of zero j , scored by the relative residual $|S_{\text{meas}} - S_{\text{pred}}|/|S_{\text{meas}}|$:

- **A** — single zero, window response ignored: amplitude $C_{\text{amp}}/|\rho_j|$ (naive Eq. 5) and phase from Eq. (8);
- **B** — single zero with window response: $CA_0(\gamma_j)e^{i\gamma_j u_{\min}}W(\omega_k - \omega_j)$;
- **C** — multi-zero, Eq. (11), C fitted;
- **D, E** — models B and C with $C = 1$ (no fitted parameter).

The step $A \rightarrow B$ isolates the window response at the bin offset (scalloping); the step $B \rightarrow C$ isolates the contribution of neighboring zeros.

Relation to Eq. (9). Equation (11) models the dominant positive-frequency contribution of the zeros in the reference list. The omitted terms of Eq. (9) — negative-frequency conjugate leakage, explicit-formula remainder terms, zeros above the reference list, finite-window truncation, and numerical sampling effects — are represented by the residual between the measured spectrum and the prediction. Extending the sum to all 85 zeros below $\gamma = 160$ (tested at $N = 5 \times 10^7$) leaves the residual of the first 17 zeros at the same 10^{-3} level. The origin of this residual floor is not analysed here; it is independent of N , $x_{\min}$, and of the length of the zero list, and it grows approximately in proportion to γ (Section 3.6).

Equation (11) uses the known low-lying zeros on the critical line, $\rho_j = 1/2 + i\gamma_j$, rather than assuming GRH for the full zero set; if an off-critical-line zero $\rho = \beta + i\gamma$ with $\beta \neq 1/2$ existed, its contribution to $\psi(x, \chi_3)/\sqrt{x}$ would not be a constant-amplitude sinusoid, and the model would need modification. The pair-sum form (6) also relies on the zero set of a real character being symmetric about the real axis; complex characters would need both signs of γ .

The constant C is fit from one isolated zero. If the analytic normalization is consistent with the measured spectrum, $C = 1$: the fits give $|C|$ between 0.99985 and 1.00034 over ten configurations (Section 3.5), and the parameter-free models D and E perform identically to B and C. The calibration is therefore a consistency check rather than a necessary step.

3. Results

3.1 Detection of the reference zeros

At $N = 5 \times 10^8$ and $x_{\min} = 150$ the FFT produces 82 local maxima above the 5% threshold. Matching them one-to-one to the 17 reference zeros of $L(s, \chi_3)$ with $\gamma \leq 47.51$ (Section 2.3) gives:

Metric	Value
Local maxima above threshold	82
Reference zeros	17
Matched (one-to-one, tolerance 3%)	17
Median absolute deviation $ \Delta\gamma /\gamma$	0.44%
Best match	0.036% ($\gamma \approx 33.90$)
Worst match	1.22% ($\gamma \approx 15.70$)
Largest offset, in FFT bins ($\Delta\gamma = 0.418$)	0.62 ($\gamma \approx 46.27$)

Every matched peak therefore lies within one bin of its zero; the deviations in percent are limited by the bin grid. Full per-zero results are in Appendix A.

Peaks beyond the reference list. The reference set is not the whole spectrum. For $\gamma \leq 55$ (Figure 1) there are 20 detected peaks and 20 known zeros — the 17 reference zeros and $\gamma = 50.375, 52.497, 54.194$, taken from an independent computation [8] — and each peak lies at a known zero. A wider blind check at $N = 5 \times 10^7$, $x_{\min} = 150$ against all 85 zeros with $\gamma < 160$ gives 69 of 70 detected peaks ($\gamma \leq 159.7$) within one bin of a known zero (the remaining peak is at 1.02 bins), while 13 of the 83 known zeros up to $\gamma = 156.7$ do not produce a separate peak. These are close pairs that merge; for instance, the 25th zero ($\gamma = 63.18$) lies 0.97 (about two bins) from its neighbour at 62.21. At these parameters, pairs closer than about two bins are therefore not separated by the peak finder, although the multi-zero model (11) can still describe them once the zeros are given (Section 3.4).




Figure 1. FFT amplitude spectrum of $\psi(x, \chi_3)/\sqrt{x}$ in log-space at $N = 5 \times 10^8$ and $x_{\min} = 150$. Black dots: local maxima above 5% of the maximum amplitude (82 in total, 20 with $\gamma \leq 55$). Orange dashed lines: the 17 reference zeros; grey dotted lines: further known zeros ($\gamma = 50.375, 52.497, 54.194$). Every detected peak with $\gamma \leq 55$ lies at a known zero, and every known zero in this range has a detected peak. Zero positions are used only for comparison after detection.

3.2 Amplitude

For the 16 matched zeros other than the calibration zero we compare the measured peak amplitude with the naive prediction (the window factor replaced by its peak value) and with Eq. (5):

Ratio measured / predicted ($n = 16$)	Naive (offset ignored)	Eq. (5), with window factor
Mean	0.9714	1.0012
Standard deviation	0.0359	0.0080
Minimum	0.8914	0.9878
Maximum	1.0190	1.0279

The deviations of the naive ratio are not random: they correlate with the offset between the zero and the centre of the detected bin ($r = 0.85$ between $|1 - \text{ratio}|$ and $|\text{offset}|$) and are accounted for by the scalloping loss. For example, the zero at $\gamma = 15.70$ lies 0.46 bins from its bin; the window response alone predicts a ratio of 0.948, and 0.9485 is measured. For the six isolated zeros, Eq. (5) reproduces the amplitude to between 0.9998 and 1.0008. The remaining deviations (at most 2.8%) occur at close pairs, where the neighbor contributes: 44.12 (1.0279), 47.51 (0.9878), 42.62 (1.0087), 46.27 (0.9937). The largest naive deviation, at $\gamma = 46.27$ (ratio 0.891), arises mainly because the detected peak lies 0.62 bins from the zero, and not primarily from the neighbor.

3.3 Phase verification

For isolated zeros (isolation gap > 2.0), the measured phase matches the closed-form prediction (8) to a median of 4×10^{-4} **rad** (0.02°):

Statistic	Isolated (gap > 2.0)	All matched
Mean absolute phase error	0.0003 rad	0.0018 rad
Median absolute phase error	0.0004 rad	0.0004 rad
Sample size	$n = 7$	$n = 17$

Equation (8) contains no fitted parameter, so all 17 zeros are valid tests (the isolated set of 7 includes the zero used for amplitude calibration). For close pairs the phase error is larger, up to 0.010 rad (26.58: +0.0093; 28.22: −0.0104), as expected when the neighbor's leakage perturbs the main-lobe phase. The fitted phase of the calibration constant, $\arg C = +0.0005$ rad, equals the phase error at the calibration zero by construction. We are not aware of an earlier direct comparison of FFT-bin phases with a closed-form prediction derived from the explicit formula in this form. The agreement confirms that the FFT reproduces not only the positions and weights but also the **phases** of the oscillatory terms.

3.4 Complex spectrum: window response and neighboring zeros

The calibration of C from one isolated zero gives $|C| = 0.99988$ and $\arg C = +0.0005$ rad. Using the five models of Section 2.6, the mean relative residual on the complex spectrum of the 16 non-calibration zeros is:

Subset	n	A: single, no window	B: single + window	C: multi-zero (C fitted)	E: multi-zero ($C = 1$)
All matched	16	0.0394	0.0053	0.0009	0.0009
Isolated (gap > 2.0)	6	0.0385	0.0006	0.0006	0.0005
Close neighbors (gap < 2.0)	10	0.0399	0.0082	0.0011	0.0011

Model D (B with $C = 1$) gives 0.0053, 0.0005, and 0.0081 for the three rows. Ratios quoted in the text are computed from unrounded means (the rounded means above would give $43.8\times$ for all zeros).

The overall reduction of the residual by a factor 42.6 decomposes into a factor 7.4 from the window response of the same zero ($A \rightarrow B$) and a factor 5.8 from the neighboring zeros ($B \rightarrow C$). For isolated zeros the window response alone brings the residual from 0.0385 to 0.0006 and the neighbors add nothing ($B = C$). For close pairs the window response gives $4.9\times$ and the multi-zero sum a further $7.3\times$. The parameter-free model E equals the fitted model C, so a calibration constant is not needed.

The 17 reference zeros contain pairs with gaps down to 1.24 (≈ 3 bins) but no cluster of three or more zeros within a few bins. Pairs at higher γ , as in $\zeta(s)$, are denser; see Appendix C.




Figure 2. Relative residual of the complex spectrum at the detected bin of each reference zero (calibration zero excluded; $N = 5 \times 10^8$, $x_{\min} = 150$) for model A (single zero, no window response), model B (single zero with window response), and model C (multi-zero). (a) versus γ ; (b) versus the isolation gap; (c) versus the distance, in FFT bins, from the detected bin to the nearest neighboring zero; the dashed curve is the window response $|W(d)|/|W(0)|$. For isolated zeros B and C coincide (only C is visible). The residual of C lies between 5×10^{-5} and 3.3×10^{-3} (mean 9×10^{-4}).

3.5 Robustness to N and $x_{\min}$

We repeat the full analysis for two values of the upper bound ($N = 5 \times 10^7$ and 5×10^8) and five values of the log-window lower bound ($x_{\min} = 50$ to 300): ten configurations, listed in Appendix B (Table B.1), whose highlighted row is the configuration of Sections 3.1–3.4. The position match is 17/17 in every configuration. *Stable quantities*: the residual of the multi-zero model, 0.0008–0.0014 (identical, 0.0008–0.0014, with $C = 1$); the fitted $|C|$, 0.99985–1.00034 with $\arg C$ between -3.4×10^{-4} and $+4.7 \times 10^{-4}$ rad; the median phase error for isolated zeros, 0.2 – 0.6×10^{-3} rad; and the naive amplitude ratio, within 3.2% of unity on average.

Quantities that are not stable: the factor by which the multi-zero model improves on the single-zero model ($18.2\times$ to $65.1\times$), and its split into window response (1.2 – $8.3\times$) and neighbors (2.7 – $35.6\times$). The reason is that both effects depend on how the FFT bins fall relative to the zeros. For the pair $\gamma = 46.27/47.51$, the detected bin of the first zero lies 3.58 bins from its neighbor at $N = 5 \times 10^8$ ($|W|/|W(0)| = 0.006$, model-B residual 0.0069), but 2.17 bins from it at $N = 5 \times 10^7$ ($|W|/|W(0)| = 0.196$, residual 0.193). Panel (c) of Figure 2 shows that the model-B residual follows the window response. The improvement factor is therefore not a property of the method; the stable quantity is the absolute residual of the multi-zero model.

3.6 A second L -function, and dependence on height

Applied to $\zeta(s)$ with the same parameters ($N = 5 \times 10^8$, $x_{\min} = 150$; 60 zeros, $14.13 \leq \gamma \leq 163.03$), all 60 zeros are matched. The multi-zero model with $C = 1$ reaches a mean residual of 0.0036 (fitted: 0.0037), against 0.0375 for model A; $|C| = 0.9995$. For isolated zeros ($n = 18$) model B reproduces the amplitude to 0.9994 ± 0.0022 , and the median phase error is 1.1×10^{-3} rad (calibration zero excluded; including it, 0.9×10^{-3} rad, $n = 19$). The residual is larger than for χ_3 , but it grows approximately in proportion to γ (about $3.7 \times 10^{-5}\gamma$ for $\gamma > 50$), as expected if the absolute error is roughly constant while the amplitude falls as $1/\gamma$. A comparison between L -functions therefore has to be made at equal height: for $\gamma < 50$ the residuals are 0.0008 (ζ , $n = 9$) and 0.0009 (χ_3 , $n = 16$). Over ten ($N, x_{\min}$) configurations the multi-zero residual stays within 0.0029–0.0039, $|C|$ within 5.2×10^{-4} of 1, and the amplitude ratio of Eq. (5) for isolated zeros between 0.9992 and 1.0011. Details are in Appendix C.

4. Discussion

The results are a quantitative numerical check that the explicit formula for $L(s, \chi_3)$, and for $\zeta(s)$ in the replication, describes the complex log-space spectrum of the prime-power sum down to a relative accuracy of about 10^{-3} at heights $\gamma \lesssim 50$. We structure the discussion to make the scope of the claims explicit.

What is new, and what is not. The three components of Eq. (8) — analytic phase, sampling origin, window linear phase — are each individually standard, and Eq. (11) is a direct consequence of the linearity of the windowed Fourier transform. What we add is (i) their explicit combination for this log-space setting, with a validation that needs no fitted parameter; (ii) the separation of the single-zero amplitude error into a window response (scalloping) and a neighbor term, which shows that the naive single-zero amplitude formula is accurate to 0.1% for isolated zeros once the window response is included; (iii) the observation that the achievable residual is a property

of the transform ($\approx 10^{-3}$, growing with γ) while the improvement factor over the naive model is a property of the bin alignment; and (iv) the cross-check on a second L -function at equal height.

Uses and their limits. The uses listed in Section 1.4 follow from the model being a calibrated forward model. Their limits should be stated plainly. The zeros of $L(s, \chi_3)$ and $\zeta(s)$ are known; the model neither finds nor improves them, and the 10^{-3} residual is far from competitive with direct verification of the zeros. The statement that the model detects departures from the explicit formula is a statement about the numerical setting: it concerns the amplitude-and-phase structure of the transform, not the zeros themselves. Peak refinement and the separation of merged pairs require fitting Eq. (11), which we do not carry out here; we show only that the model reproduces pairs two to three bins apart. Extension to other real primitive characters needs only the zero list; complex characters, whose zeros are not symmetric about the real axis, and L -functions of higher degree require a modified coefficient A_0 and pair structure (Eq. 6), and are not tested here. Fourier-analytic uses of the zeros in a different setting are discussed in [6].

On the word "blind". Only the peak-detection stage is blind in the strict sense — the FFT and the peak-finding algorithm use no information about the zeros. The matching to the reference zeros (Section 2.3) and the amplitude, phase, and leakage tests necessarily use the reference list, because they test whether the explicit formula, which involves the zeros, predicts the observed spectrum. "Blind verification" means *blind peak detection followed by reference-based quantitative validation*.

Scope of the model. Equation (11) models the dominant positive-frequency contribution of the zeros in the reference list. The residual absorbs the omitted negative-frequency conjugate leakage, explicit-formula remainder terms, zeros beyond the list, finite-window truncation, and sampling effects. The model uses known zeros on the critical line rather than assuming GRH; an off-line zero would require modification.

Robustness. The position match is 17/17 for two values of N and five values of $x_{\min}$. The multi-zero residual stays within 0.0008–0.0014, the fitted $|C|$ within 3.4×10^{-4} of 1, and the isolated-zero phase error below 0.6×10^{-3} rad (median). The improvement factor varies between $18\times$ and $65\times$ because of bin alignment (Section 3.5). Only two values of N were used.

Limitations.

- The reference set is the first 17 zeros of $L(s, \chi_3)$ ($\gamma \leq 47.51$). The next zero, $\gamma = 50.375$, is known and appears as a detected peak, but is not part of the validation tables. A set of 25 zeros was also tried at $N = 5 \times 10^7$: 24 were matched and the multi-zero residual stayed at the same level (0.0015 against 0.0013 for 17 zeros), the 25th (63.18) merging with its neighbor.
- Close pairs closer than about two bins are not separated by the peak finder (Section 3.1).
- The residual floor of $\approx 10^{-3}$ is not explained here; it is independent of N and $x_{\min}$ and grows approximately in proportion to γ .
- Only real characters (χ_3) and $\zeta(s)$ are treated; the two share real coefficients and symmetric zero sets.
- The resolution criterion (3) is a mean-spacing heuristic, and the sweep does not probe it (Section 2.2).

5. Conclusion

We have presented a parameter-free forward model of the complex log-space FFT spectrum of $\psi(x, \chi_3)/\sqrt{x}$, built from a closed-form FFT-bin phase (Eq. 8), a window-corrected single-zero amplitude (Eq. 5), and a multi-zero leakage model with the exact Kaiser-window DTFT (Eq. 11). At $N = 5 \times 10^8$ all 17 reference zeros ($\gamma \leq 47.5$) are detected within 0.62 bins; the phase error for isolated zeros has a median of 4×10^{-4} rad; and the multi-zero model with $C = 1$ reproduces the complex spectrum with a mean residual of 0.0009, against 0.039 for the naive single-zero model. The improvement splits into a window response of the same zero ($7.4\times$) and neighboring zeros ($5.8\times$); the first alone accounts for isolated zeros, the second is needed for pairs. The absolute residual is stable across ten configurations, the improvement factor is not. The same pipeline on $\zeta(s)$ (60 zeros, $\gamma \leq 163$) gives the same behaviour, with a residual that grows approximately in proportion to γ and coincides with that of χ_3 at equal height. The model can be used as an end-to-end consistency check, as a basis for peak refinement and separation of close pairs, and for parameter choice; it does not locate new zeros or test GRH. Source code, per-zero tables, and the robustness sweeps are provided.

Code Availability

The complete Python source is available at:

- $L(s, \chi_3)$ and $\zeta(s)$ analysis: github.com/model-vpr/dirichlet-l-fft-verification

The repository contains the shared module `fft_explicit_formula_lib.py` (FFT, window DTFT, phase formula, matching, multi-zero model); `chi3_fft_verification.py` and `zeta_fft_verification.py` (per-zero tables);

`chi3_model_comparison.py` and `zeta_model_comparison.py` (models A–E, isolation and height splits); `generate_chi3_figures.py` (Figures 1–2 and Appendix A); `generate_zeta_figures.py` (Figures C.1–C.2); and the zero lists (`chi3_zeros_below_160.txt`, `zeta_reference_zeros.py`), computed independently with `mpmath`. The scripts depend on `numpy`, `scipy`, and `matplotlib` (plots only; `mpmath` only to regenerate the zero lists). A run at $N = 5 \times 10^8$ needs about 5–6 GB of memory (the prime sieve and the ψ array dominate); $N = 5 \times 10^7$ reproduces the qualitative results much faster.

Acknowledgments

The author thanks the developers of NumPy, SciPy, and mpmath. The 17 reference zeros of $L(s, \chi_3)$ were checked against an independent evaluation of $L(s, \chi_3)$ with mpmath [8] (agreement better than 10^{-8}); the first zero also agrees with the LMFDB entry [4].

References

1. J. B. Conrey, *The Riemann Hypothesis*, Notices of the AMS **50** (2003), no. 3, 341–353.
 2. 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*, Abhandlungen der Königlich Preußischen Akademie der Wissenschaften zu Berlin (1837), 45–81.
 3. H. Davenport, *Multiplicative Number Theory*, 3rd ed., Springer, 2000.
 4. The LMFDB Collaboration, *The L-functions and Modular Forms Database*, L-function 1-3-3.2-r1-0-0 (Dirichlet character $\chi_3(2, \cdot)$), <https://www.lmfdb.org/L/rational/1/3>, accessed October 2026.
 5. J. F. Kaiser, *Nonrecursive digital filter design using the I_0 -sinh window function*, Proc. IEEE Int. Symp. Circuits Syst. (1974), 20–23.
 6. A. Bondarenko, D. Radchenko, K. Seip, *Fourier Interpolation with Zeros of Zeta and L-Functions*, arXiv:2209.05918.
 7. J. O. Smith III, *Mathematics of the Discrete Fourier Transform (DFT): with Audio Applications*, 2nd ed., W3K Publishing, 2007.
 8. F. Johansson et al., *mpmath: a Python library for arbitrary-precision floating-point arithmetic*, version 1.3, 2023, <https://mpmath.org>.
-

Appendix A. Per-Zero Results

The table lists the 17 reference zeros of $L(s, \chi_3)$ ($\gamma \leq 47.51$) with the detected peak position, the bin offset, the amplitude ratios, the phase error, and the complex spectral residual under models A, B, and C (Section 2.6). The first row is the zero used to calibrate C and C_{amp} : its residuals under models B and C vanish by construction, while model A has the residual 0.0005, equal to the phase error of Eq. (8) at this zero.

Column definitions: *Offset* is $(\gamma_{\text{det}} - \gamma_{\text{known}})/\Delta\gamma$ in FFT bins ($\Delta\gamma = 0.418$). *Amp (naive)* is the measured peak amplitude divided by the naive prediction $C_{\text{amp}}/\sqrt{1/4 + \gamma^2}$; *Amp (Eq. 5)* divides by the prediction including the window factor. *Phase diff* is the signed difference between the measured phase and Eq. (8). *Res A, B, C* are the relative residuals $|S_{\text{meas}} - S_{\text{pred}}|/|S_{\text{meas}}|$ of the models of Section 2.6.

γ (known)	γ (detected)	Gap	Offset (bins)	Amp (naive)	Amp (Eq. 5)	Phase diff (rad)	Res A	Res B	Res C	Note
8.0397	7.9484	3.21	-0.218	1.0000	1.0000	0.0005	0.0005	0.0000	0.0000	calibration
11.2492	11.2951	3.21	+0.110	1.0116	0.9998	-0.0001	0.0115	0.0006	0.0006	
15.7046	15.8968	2.56	+0.459	0.9485	1.0007	-0.0000	0.0543	0.0009	0.0009	
18.2620	18.4068	2.19	+0.346	0.9770	1.0003	0.0004	0.0236	0.0003	0.0003	
20.4558	20.4984	2.19	+0.102	1.0122	0.9999	0.0000	0.0121	0.0004	0.0004	
24.0594	24.2635	2.52	+0.488	0.9402	1.0008	0.0009	0.0636	0.0009	0.0009	
26.5779	26.7735	1.64	+0.468	0.9460	1.0006	0.0093	0.0579	0.0089	0.0017	close neighbor
28.2182	28.0285	1.64	-0.453	0.9479	0.9983	-0.0104	0.0560	0.0110	0.0012	close neighbor
30.7450	30.5385	2.53	-0.494	0.9380	1.0004	0.0004	0.0661	0.0004	0.0004	
33.8974	33.8852	1.71	-0.029	1.0153	0.9999	-0.0026	0.0153	0.0031	0.0033	close neighbor
35.6084	35.5585	1.71	-0.119	1.0101	0.9991	-0.0002	0.0100	0.0011	0.0001	close neighbor
37.5518	37.6502	1.93	+0.235	0.9984	1.0009	0.0000	0.0016	0.0010	0.0009	close neighbor
39.4852	39.3236	1.93	-0.386	0.9683	1.0012	0.0007	0.0327	0.0012	0.0010	close neighbor
42.6164	42.6702	1.50	+0.129	1.0190	1.0087	-0.0001	0.0186	0.0086	0.0005	close neighbor
44.1206	43.9252	1.50	-0.467	0.9720	1.0279	-0.0000	0.0288	0.0272	0.0010	close neighbor
46.2741	46.0169	1.24	-0.615	0.8914	0.9937	0.0031	0.1219	0.0069	0.0014	close neighbor
47.5141	47.6903	1.24	+0.421	0.9466	0.9878	-0.0017	0.0564	0.0126	0.0001	close neighbor

Table A.1. Per-zero results at $N = 5 \times 10^8$ and $x_{\min} = 150$. The calibration point (first row) is excluded from the summary statistics in Section 3.

Appendix B. Robustness Table

Results of the pipeline at ten combinations of $(N, x_{\min})$. The position match is 17/17 in every case. The multi-zero residual (C, E) and $|C|$ are stable; the improvement factors are not (Section 3.5). Residuals are means over the 16 non-calibration zeros; improvement factors are ratios of unrounded means.

N	$x_{\min}$	Position	Amp, naive (mean)	Phase (median, isolated)	A	B	C	E ($C = 1$)	A→C	A→B	B→C	$ C $
5×10^7	50	17 / 17	1.0140	0.5×10^{-3}	0.0389	0.0228	0.0012	0.0012	32.2×	1.7×	18.9×	1.00024
5×10^7	100	17 / 17	0.9932	0.3×10^{-3}	0.0406	0.0294	0.0008	0.0008	48.2×	1.4×	34.9×	0.99992
5×10^7	150	17 / 17	1.0119	0.6×10^{-3}	0.0418	0.0322	0.0013	0.0013	32.7×	1.3×	25.1×	1.00023
5×10^7	200	17 / 17	0.9809	0.3×10^{-3}	0.0483	0.0387	0.0011	0.0011	42.3×	1.2×	34.0×	1.00000
5×10^7	300	17 / 17	1.0314	0.3×10^{-3}	0.0616	0.0337	0.0009	0.0009	65.1×	1.8×	35.6×	1.00028
5×10^8	50	17 / 17	1.0202	0.3×10^{-3}	0.0268	0.0056	0.0014	0.0013	19.6×	4.8×	4.1×	1.00029
5×10^8	100	17 / 17	0.9873	0.3×10^{-3}	0.0247	0.0037	0.0014	0.0014	18.2×	6.7×	2.7×	1.00009
5×10^8	150	17 / 17	0.9714	0.4×10^{-3}	0.0394	0.0053	0.0009	0.0009	42.6×	7.4×	5.8×	0.99988
5×10^8	200	17 / 17	0.9686	0.2×10^{-3}	0.0380	0.0046	0.0013	0.0011	30.4×	8.3×	3.7×	1.00034
5×10^8	300	17 / 17	0.9962	0.5×10^{-3}	0.0285	0.0043	0.0012	0.0012	24.4×	6.7×	3.7×	0.99985

Table B.1. Robustness sweep for $L(s, \chi_3)$. The highlighted row is the configuration used in Sections 3.1–3.4. A, B, C, E: mean relative residual of models A, B, C, E (Section 2.6). The phase median is over the 7 isolated zeros (including the calibration zero).

Appendix C. Replication on $\zeta(s)$

The same pipeline — the same FFT routine, Kaiser window, amplitude and phase formulas (Eqs. 5, 8), and multi-zero model (Eq. 11) — was applied to $\zeta(s)$ with the signal $\psi(x) - x$. Reference zeros are the first 60 nontrivial zeros ($14.13 \leq \gamma \leq 163.03$), generated with mpmath [8]; matching tolerance 3%; threshold 5%; $N_{\log} = 2^{17}$; the calibration zero is the first isolated matched zero ($\gamma = 14.13$ for ζ , 8.04 for χ_3).

L -function	n	γ range	A	B	C	E ($C = 1$)	$ C $	Amp ratio, Eq. (5), isolated	Phase (median, isolated)
χ_3	16	11.2–47.5	0.0394	0.0053	0.0009	0.0009	0.99988	0.9998–1.0008 ($n = 6$)	4×10^{-4} ($n = 7$)
ζ	59	21.0–163.0	0.0375	0.0310	0.0037	0.0036	0.99952	0.9994 ± 0.0022 ($n = 18$)	1.1×10^{-3} ($n = 18$)

Table C.1. Both at $N = 5 \times 10^8$, $x_{\min} = 150$; calibration zero excluded. For ζ : isolated (gap > 2) $n = 18$, A 0.0127, B = C = 0.0031; close $n = 41$, A 0.0484, B 0.0433, C 0.0039. Improvement A→C: $10.2 \times$ (A→B $1.2 \times$, B→C $8.5 \times$). The phase median for ζ excludes the calibration zero ($n = 18$); for χ_3 the 7 isolated zeros include it, as in Table B.1. With the calibration zero the ζ value is 0.9×10^{-3} ($n = 19$).




Figure C.1. Top: FFT amplitude spectrum of $\psi(x) - x$ in log-space at $N = 5 \times 10^8$, $x_{\min} = 150$. Black dots: local maxima above 5% of the maximum amplitude; orange dashed lines: the 60 reference zeros; grey dotted lines: further known zeros up to $\gamma = 170$. Bottom: the two shaded windows enlarged, around the closest pairs 111.03/111.87 and 150.05/150.93 (Table C.4). Blue: measured $|S[k]|$ at every FFT bin; orange squares: multi-zero model, Eq. (11), with the fitted C ($|C| = 0.9995$). The two curves coincide to within the line width, including the bins between the peaks.

Dependence on height. The relative residual is the ratio of an approximately constant absolute error to an amplitude that falls as $1/\gamma$, and indeed it grows with γ . For ζ in four height ranges:

L -function, γ range	n	mean γ	A	B	C	E ($C = 1$)	C / mean γ
χ_3 , 11–48	16	31.4	0.0394	0.0053	0.0009	0.0009	2.9×10^{-5}
ζ , 0–50	9	36.6	0.0115	0.0008	0.0008	0.0006	2.2×10^{-5}
ζ , 50–100	19	76.9	0.0238	0.0138	0.0028	0.0028	3.6×10^{-5}
ζ , 100–150	23	124.8	0.0438	0.0411	0.0048	0.0048	3.9×10^{-5}
ζ , 150–163	8	156.3	0.0810	0.0770	0.0057	0.0056	3.7×10^{-5}

Table C.2. Mean relative residual of models A, B, C, E by height ($N = 5 \times 10^8$, $x_{\min} = 150$). For $\gamma > 50$ the ratio of C to the mean height is approximately constant.




Figure C.2. Relative residual of the complex spectrum at the detected bin of each of the 59 reference zeros other than the calibration zero ($\zeta(s)$, $N = 5 \times 10^8$, $x_{\min} = 150$) for models A (single zero, no window response), B (single zero with window response), and C (multi-zero). (a) versus γ ; the dashed line is the guide $3.7 \times 10^{-5}\gamma$. (b) versus the isolation gap.

Robustness for ζ . Ten configurations with 60 zeros, with the same matching tolerance (3%), threshold, window, and grid as Table C.1; the highlighted row is Table C.1:

N	$x_{\min}$	Position	A	B	C	E ($C = 1$)	A→C	$ C $	Amp ratio, Eq. (5), isolated	Phase (median, isolated)
5×10^7	50	59 / 60	0.0715	0.0446	0.0035	0.0035	20.7×	1.00003	1.0000 ± 0.0019	1.3×10^{-3}
5×10^7	100	60 / 60	0.0716	0.0399	0.0033	0.0033	21.8×	0.99975	0.9999 ± 0.0013	1.0×10^{-3}
5×10^7	150	60 / 60	0.0644	0.0430	0.0033	0.0033	19.4×	0.99996	1.0000 ± 0.0017	1.4×10^{-3}
5×10^7	200	60 / 60	0.0736	0.0432	0.0029	0.0029	25.2×	0.99992	1.0000 ± 0.0015	0.6×10^{-3}
5×10^7	300	60 / 60	0.0801	0.0607	0.0037	0.0036	21.8×	1.00024	0.9999 ± 0.0019	0.8×10^{-3}
5×10^8	50	60 / 60	0.0387	0.0224	0.0038	0.0037	10.2×	0.99948	1.0011 ± 0.0022	0.8×10^{-3}
5×10^8	100	60 / 60	0.0309	0.0183	0.0039	0.0039	7.8×	1.00035	0.9992 ± 0.0029	1.3×10^{-3}
5×10^8	150	60 / 60	0.0375	0.0310	0.0037	0.0036	10.2×	0.99952	0.9994 ± 0.0022	1.1×10^{-3}
5×10^8	200	59 / 60	0.0433	0.0262	0.0032	0.0032	13.6×	1.00017	1.0002 ± 0.0013	1.2×10^{-3}
5×10^8	300	59 / 60	0.0488	0.0358	0.0034	0.0034	14.4×	1.00005	1.0007 ± 0.0023	0.9×10^{-3}

Table C.3. Robustness for $\zeta(s)$ (60 zeros; residuals are means over the matched zeros other than the calibration zero, $\gamma = 14.13$, which is also excluded from the phase medians; improvement factors from unrounded means; isolated: $n = 18$). The multi-zero residual stays within 0.0029–0.0039 (identical with $C = 1$), $|C|$ within 5.2×10^{-4} of 1, and the amplitude ratio of Eq. (5) for isolated zeros between 0.9992 and 1.0011. The improvement factor A→C, $7.8 \times$ – $25.2 \times$ (7.8 – $14.4 \times$ at $N = 5 \times 10^8$, 19 – $25 \times$ at 5×10^7), varies with N and $x_{\min}$ as for χ_3 . Unmatched zeros: 95.87 (gap 1.22) at 5×10^7 , $x_{\min} = 50$; 150.93 (gap 0.87) at 5×10^8 , $x_{\min} = 200, 300$.

The closest pairs. At $N = 5 \times 10^8$, $x_{\min} = 150$ the bin spacing is 0.418, and $\zeta(s)$ has two pairs only about two bins apart among the 60 zeros. They are resolved as separate peaks here, and the multi-zero model reproduces them to a few 10^{-3} , whereas the single-zero model fails (residual up to 0.52, phase error up to 0.47 rad):

γ	Neighbor gap (γ units / bins)	Phase diff (rad)	Res A	Res C
111.0295	0.85 / 2.0	−0.129	0.1616	0.0031
111.8747	0.85 / 2.0	+0.472	0.5164	0.0044
150.0535	0.87 / 2.1	−0.348	0.3437	0.0025
150.9253	0.87 / 2.1	+0.155	0.1557	0.0035

Table C.4. The two closest pairs of the 60 zeros of $\zeta(s)$ ($N = 5 \times 10^8$, $x_{\min} = 150$). Res A: single zero, no window response; Res C: multi-zero model. At other values of $x_{\min}$ the pair near 150.9 is not always resolved (Table C.3).

Position at low height. A percent tolerance is a poor criterion for the lowest zeros: at $N = 5 \times 10^7$, $x_{\min} = 150$, a 1% tolerance leaves $\gamma = 14.13$ and 21.02 unmatched, because 1% of γ is smaller than half a bin (0.25), while a 3% tolerance matches all 60. The zeros that remain unmatched at 3% (Table C.3) are members of close pairs with gaps of 0.87–1.22, i.e. two to three bins, which the peak finder merges.