

Arithmetic Structure in the Fourier Spectrum of Riemann-Zero Pair Correlations

Brian Scannell

September 2026

Abstract

Montgomery’s pair-correlation conjecture predicts that suitably normalised zeros of the Riemann zeta function have the same local two-point statistics as eigenvalues of large random unitary matrices. The limiting correlation is smooth, but finite-height zero data contain smaller arithmetic corrections. We revisit a numerical observation first reported in 2020: when a long-range empirical pair-correlation curve is Fourier analysed, a family of narrow peaks appears near frequency one. Using the finite-height off-diagonal theory of Bogomolny–Keating and its arithmetic form developed by Keating–Smith, we interpret these peaks as a direct Fourier resolution of the known rationally indexed arithmetic structure. Their predicted frequencies are

$$f_{n,t}(T) = 1 + \frac{\log(n/l)}{\log(T/2\pi)}.$$

In two high-height datasets the strongest observed peaks agree with these frequencies to about 0.3 Fourier bins RMS, while their relative amplitudes closely follow the predicted arithmetic ranking. Separate holdout data reproduce the sub-bin agreement. Finally, a wide low-height block, which initially fails a fixed-height analysis, resolves into sharper lines when subdivided; the line centres then follow the predicted $1/\log T$ trajectories. The computations provide a direct numerical visualisation of arithmetic structure in finite-height pair correlation.

2020 Mathematics Subject Classification. Primary 11M26; Secondary 11M06, 15B52.

1 A conversation at tea

One of the most striking bridges between number theory and mathematical physics began with a short conversation. In April 1972 Hugh Montgomery was visiting the Institute for Advanced Study in Princeton and described to Freeman Dyson a result concerning correlations between the nontrivial zeros of the Riemann zeta function. Dyson immediately recognised the same correlation law from random-matrix theory, where it describes the eigenvalues of large random Hermitian or unitary matrices. Dyson wrote to Atle Selberg the next day with the relevant references. Montgomery later recorded the connection in his original paper, crediting Dyson with drawing his attention to the random matrix result [1, 2].

Write the nontrivial zeros, assuming the Riemann hypothesis for the discussion, as

$$\rho_j = \frac{1}{2} + i\gamma_j, \quad 0 < \gamma_1 < \gamma_2 < \cdots.$$

Near height T the mean density of zeros is approximately

$$\bar{\rho}(T) = \frac{1}{2\pi} \log \frac{T}{2\pi},$$

so adjacent gaps are unfolded by multiplying by this local density. The resulting mean spacing is approximately one.

Montgomery's conjecture says that the limiting two-point correlation of these unfolded zeros is

$$R_2(s) = 1 - \left(\frac{\sin \pi s}{\pi s} \right)^2. \quad (1)$$

Figure 1 is the plot that motivated the original numerical study [12]. Even for the first 10^5 zeros the empirical pair correlation already tracks the smooth curve in (1) remarkably well.




Figure 1: The empirical pair correlation for the first 10^5 Riemann zeros together with Montgomery's limiting curve $1 - (\sin \pi s / \pi s)^2$.

The usual plot stops after only a few mean spacings. Our question is different: what happens if the same construction is continued for hundreds of spacings, and then examined on a finer scale?

2 Building the pair correlation

It is useful first to see what is being accumulated. Let

$$\delta_j = (\gamma_{j+1} - \gamma_j) \bar{\rho}(\gamma_j)$$

be a normalised adjacent spacing. The separation between a zero and its k th successor is approximated by

$$s_{j,k} = \delta_j + \delta_{j+1} + \cdots + \delta_{j+k-1}.$$

For each k we histogram the values $s_{j,k}$ and then add the histograms as k increases. Figure 2 shows the process using the first 10^5 zeros. The individual separation distributions move progressively to the right, while their cumulative sum builds the pair-correlation curve.




Figure 2: How the empirical pair correlation is assembled. The left panels show the distributions of separations using (a) two adjacent zeros, (c) two through four adjacent zeros, and (e) two through nine adjacent zeros. The right panels (b), (d), and (f) show the corresponding cumulative curves. The red dotted curve is Montgomery’s limiting prediction. The panels are adapted from Figures 3–8 of the 2020 study [12].

For the long-range calculation we use lags $k = 1, \dots, 999$, histogram bin width 0.05, and hence a nominal range $0 < s < 1000$. The empirical curve is normalised by the number of spacings and the bin width. The zero data used in the original computations came from Odlyzko’s tables and the LMFDB; the principal blocks considered here include the first 10^5 and first 2×10^6 zeros, about 6.01×10^6 zeros near $T = 4.03646 \times 10^8$, and 7×10^6 zeros near $T = 3.0607946 \times 10^{10}$ [3, 4].

At range 1000 the curve looks almost flat after its initial rise, but it is not featureless. The 2020 calculations described a persistent “ripple” and a spectral peak near frequency one [12]. At the time the smaller neighbouring peaks had no interpretation. They are the main subject of the present article.

3 Arithmetic hiding near frequency one

To isolate the fine structure, we subtract the limiting GUE expression (1), remove a constant and linear trend, retain $1.875 \leq s \leq 998.125$, apply a Hann window (a smooth taper that suppresses Fourier leakage from the ends of the finite interval), and take a discrete Fourier transform. The corresponding frequency spacing is

$$\Delta f \simeq 0.0010037.$$

The spectrum is concentrated near $f = 1$, but at high zero heights it separates into a striking family of narrow lines.

The key to their interpretation lies in finite-height off-diagonal pair correlation. Bogomolny and Keating related the nonuniversal corrections to arithmetic data [5]. A later formulation by Keating and Smith rewrites the relevant off-diagonal contribution as an arithmetic double sum [6]. Schematically, with

$$\epsilon = \frac{s}{\bar{\rho}},$$

its oscillatory part contains terms of the form

$$e^{-2\pi i \bar{\rho} \epsilon} \left(\frac{n}{l}\right)^{-i\epsilon}, \quad (n, l) = 1, \quad (2)$$

with arithmetic weights involving the Möbius and Euler totient functions. Since $2\pi\bar{\rho} = \log(T/2\pi)$, equation (2) becomes

$$\exp \left[-2\pi i s \left(1 + \frac{\log(n/l)}{\log(T/2\pi)} \right) \right].$$

For an infinite correlation function these oscillatory terms give ideal spectral lines. Our empirical estimator is finite: truncating in s and multiplying by the Hann window convolves each ideal line with the transform of that window, broadening an isolated line without changing its theoretical centre; overlapping neighbouring lines may, however, shift the observed local maximum. Histogram binning limits the accessible frequency range. Thus the finite Fourier spectrum is predicted to show broadened peaks centred at

$$\boxed{f_{n,l}(T) = 1 + \frac{\log(n/l)}{\log(T/2\pi)}}. \quad (3)$$

Here $(n, l) = 1$, and the Möbius factor removes nonsquarefree n . A convenient leading ranking of line strengths is

$$W_{n,l} = \frac{n}{l} \left(\frac{\mu(n)}{\phi(n)} \right)^2. \quad (4)$$

We use (4) only to rank candidate lines; it is not treated as an exact amplitude prediction for the finite, windowed estimator. The frequency law has an immediate visual consequence. Ratios $n/l > 1$ give lines above one, ratios $n/l < 1$ give lines below one, and every line moves toward one as the height T increases.

Figure 3 superimposes a selection of the predicted rational lines on the empirical spectrum near $T = 4.03646 \times 10^8$. The labels are not fitted frequencies: each dashed line is fixed by (3).




Figure 3: Empirical Fourier spectrum near $T = 4.03646 \times 10^8$ with selected arithmetic frequencies from (3). Prominent peaks occur at ratios such as 2, 3, 6, 10 above one and at reciprocal or other rational ratios below one.

For example, with $L = \log(T/2\pi) = 17.97817\dots$, the predicted $n/l = 2$ line is

$$1 + \frac{\log 2}{L} = 1.038555,$$

while the observed peak is 1.038844. Similarly, $n/l = 3$ predicts 1.061108 versus 1.060925 observed, and $n/l = 10$ predicts 1.128077 versus 1.128174 observed. Below one, $n/l = 1/2$ predicts 0.961445 versus 0.961558 observed.

4 A quantitative line test

Visual agreement is suggestive but not enough, especially because rational numbers are dense. For the quantitative test we form candidates with $n, l \leq 200$, $(n, l) = 1$, $\mu(n) \neq 0$, and frequencies in $0.90 < f < 1.15$. We retain the 40 strongest candidates ranked by (4) and the 25 strongest observed local maxima. The assignment problem is simply to pair the 25 observed peaks with 25 of the 40 predicted locations, without using any predicted location twice, so that the total squared mismatch is as small as possible. We solve this standard one-to-one matching problem with the Hungarian algorithm. Restricting the candidate set in this way prevents arbitrary nearest-neighbour matching against an unlimited collection of rationals.

For the block near $T = 4.03646 \times 10^8$, the RMS frequency discrepancy over all 25 assigned peaks is

$$2.90 \times 10^{-4} = 0.289 \Delta f.$$

For the block near $T = 3.0607946 \times 10^{10}$ it is

$$3.06 \times 10^{-4} = 0.305 \Delta f.$$

Figure 4 plots observed against predicted frequencies. The identity line is shown for reference.




Figure 4: Observed and predicted frequencies for the strongest arithmetic lines. The deviations from the identity line are typically a small fraction of one Fourier bin.

As a conservative null check, we circularly shift the observed peak pattern through the band while preserving its internal spacings, giving $p \simeq 6.0 \times 10^{-5}$ and $p \simeq 2.14 \times 10^{-3}$ at the two heights. A second test with 100,000 random candidate sets likewise gives $p \lesssim 10^{-5}$. These are supporting checks; the more discriminating evidence below is the prescribed motion of the same lines with height.

The amplitudes provide an additional descriptive check. Their absolute normalisation is affected by the finite window and estimator, so we emphasise the relative ranking rather than an exact amplitude formula. For the matched lines, the Spearman correlation between observed amplitudes and the ranking in (4) is

$$\rho = 0.9915,$$

and the Pearson correlation of logarithmic amplitudes is about 0.996. The observed amplitudes show a roughly common attenuation relative to the leading estimate, with a typical ratio near 0.8. Because the lines were selected and ranked before this comparison, these correlations should be read as descriptive consistency checks rather than as independent significance tests. Figure 5 shows the relation.




Figure 5: Observed Fourier amplitudes compared with the leading arithmetic amplitude prediction. The principal result here is the strong relative ordering; the common attenuation in absolute amplitude has not been modelled.

We also varied the number of observed peaks, the frequency band, the Fourier window, and the number of retained theoretical candidates. Across 600 configurations (300 for each principal high-height dataset), the median RMS frequency error remained about 0.31 Fourier bins. The result is therefore not tied to a single window or cut-off choice.

5 Holdouts and a useful failure

A natural concern is retrospective tuning: perhaps the analysis choices were selected because they happened to work on the data that suggested the idea. After examining the two principal datasets, but before analysing the three holdout blocks, we fixed the specification just described—Hann window, $0.90 < f < 1.15$, 25 observed peaks and 40 theoretical candidates—and applied it without modification to the holdouts.

Two disjoint holdout blocks reproduce the result. A block near $T = 2.0396 \times 10^7$ gives an RMS error of 0.271 Fourier bins, while a block near $T = 1.02296 \times 10^8$ gives 0.251 bins. In each case all 25 assigned lines lie within about half a bin of their predicted positions, and the random-line test again gives $p < 10^{-5}$.

A third holdout initially appears to fail. It spans approximately $2.546 \times 10^6 < T < 4.646 \times 10^6$ and gives an RMS error of about 2.39 bins if represented by one effective height. Equation (3) explains why. Over so broad a fractional height interval the theoretical frequency itself moves appreciably. For $n/l = 2$, for example, the predicted line drifts by about 0.00239, or 2.38 of the original range-1000 Fourier bins, across the block. The failure of a fixed-height approximation is thus quantitatively of the expected size.

This suggests a further test rather than a reason to discard the block. We split the wide dataset into narrower contiguous pieces and recompute, in each piece, exactly the same range-1000 pair-correlation estimator, GUE subtraction, detrending and Hann-windowed Fourier transform used above. For each subblock, T is taken to be the geometric-mean zero height. The line centres should then move according to (3). Figure 6 shows the eight-block result. The

dashed theoretical curves contain no fitted trajectory parameters.




Figure 6: Exact range-1000 pair-correlation test in eight contiguous low-height blocks. Observed line centres are shown together with the predictions of (3). The apparent failure of the unsplit block resolves into systematic $1/\log T$ frequency motion.

As a separate reproducibility check, we reconstructed the eight-block spectra from the underlying zero ordinates and estimated each peak centre without using the theoretical frequency. We first located a discrete local maximum of the ordinary, unpadded Fourier transform, then applied three-point parabolic interpolation to that maximum and its two neighbouring Fourier samples. If A_{k-1}, A_k, A_{k+1} are the three amplitudes and Δf is the Fourier bin width, the sub-bin offset is

$$\delta = \frac{1}{2} \frac{A_{k-1} - A_{k+1}}{A_{k-1} - 2A_k + A_{k+1}}, \quad f_{\text{centre}} = f_k + \delta \Delta f.$$

Only after the empirical centres had been extracted were they compared with (3). The resulting eight-block RMS centre errors are 3.5×10^{-5} for $n/l = 2$, 4.0×10^{-5} for $2/3$, 4.2×10^{-5} for $3/2$, and 3.8×10^{-5} for 6 , corresponding to about 0.04 Fourier bins RMS. Thus the trajectory agreement survives a fully explicit, prediction-independent sub-bin centre estimator.

A four-block calculation gives similarly small centre errors for most lines, although the $n/l = 6$ feature is more strongly blended. Its median full width at half maximum decreases from 2.93×10^{-3} in four blocks to 2.10×10^{-3} in eight; most of the other lines remain near 2.0×10^{-3} in both splits, already close to the linewidth imposed by the finite s interval and Hann window. The identical-estimator test therefore strongly supports the predicted frequency motion, while the earlier apparent factor-of-two linewidth scaling is not present once the same range-1000 estimator is used throughout.

6 What the spectrum is telling us

The frequencies in (3) are not new. Bogomolny’s form-factor exposition displays the corresponding arithmetic delta-function structure near the Heisenberg time [7]; Berry and Keating survey the broader Riemann-zero/quantum-chaology connection [10]; and later work quantified finite-height corrections [8,9]. Very recent independent work by Murphy resolves the same pseudo-orbit locations with a band-limited variance statistic [11]. Here, by contrast, the Fourier transform of the empirical long-range pair-correlation curve directly resolves individual members as narrow lines, so their locations, relative strengths, and motion with height can be tested line by line.

This also clarifies the unexplained spectrum reported in 2020 [12]: the “frequency-one peak” contains arithmetic sidebands indexed by rational n/l . As T increases, $1/\log(T/2\pi)$ decreases and the lines move inward toward one, as seen in the separate height blocks.

This long-range phenomenon should be distinguished from the familiar local finite-size expansion of random-matrix statistics, which keeps unfolded separation fixed as $T \rightarrow \infty$; here separations extend to order 10^3 . The appropriate interpretation is the full finite-height arithmetic structure, not merely the first local $1/\log^2 T$ correction.

The common amplitude attenuation and still larger zero blocks remain useful follow-up tests, but neither affects the main frequency comparison.

7 Conclusion

Montgomery’s pair-correlation conjecture entered mathematics through an unexpected connection with random matrices. In long-range empirical zero data, the fine Fourier spectrum near frequency one directly resolves arithmetic lines at

$$f_{n,l}(T) = 1 + \frac{\log(n/l)}{\log(T/2\pi)}.$$

Across separate high- and intermediate-height blocks, the strongest lines agree with this prediction to a fraction of one Fourier bin and their relative amplitudes follow the expected arithmetic ranking. A wide low-height block provides a further check: after subdivision, the line centres trace the predicted $1/\log T$ trajectories under the same pair-correlation estimator. The result is a direct Fourier-space visualisation of known arithmetic information carried by the finite-height pair correlation of the Riemann zeros.

Declaration of generative AI use

ChatGPT was used during preparation of this manuscript for idea exploration, literature searching, coding assistance and computational checks. The author independently reviewed and verified the mathematical arguments, computations, references and takes full responsibility for the originality and integrity of the article.

References

- [1] H. L. Montgomery, The pair correlation of zeros of the zeta function, in *Analytic Number Theory*, Proc. Sympos. Pure Math. **24**, American Mathematical Society, 1973, 181–193.
- [2] K. D. Thomas, From prime numbers to nuclear physics and beyond, Institute for Advanced Study, 2013.
- [3] A. M. Odlyzko, Tables of zeros of the Riemann zeta function, online data tables.
- [4] The LMFDB Collaboration, The L-functions and Modular Forms Database (LMFDB).
- [5] E. B. Bogomolny and J. P. Keating, Gutzwiller’s trace formula and spectral statistics: beyond the diagonal approximation, *Phys. Rev. Lett.* **77** (1996), 1472–1475.
- [6] J. P. Keating and D. J. Smith, Twin prime correlations from the pair correlation of Riemann zeros, *J. Phys. A: Math. Theor.* **52** (2019), 365201.
- [7] E. Bogomolny, Spectral statistics, *Documenta Math.*, Extra Volume ICM 1998, Vol. III (1998), 99–108.
- [8] E. Bogomolny, O. Bohigas, P. Leboeuf, and A. G. Monastera, On the spacing distribution of the Riemann zeros: corrections to the asymptotic result, *J. Phys. A* **39** (2006), 10743–10754.
- [9] P. J. Forrester and A. Mays, Finite-size corrections in random matrix theory and Odlyzko’s dataset for the Riemann zeros, *Proc. R. Soc. A* **471** (2015), 20150436.
- [10] M. V. Berry and J. P. Keating, The Riemann zeros and eigenvalue asymptotics, *SIAM Review* **41** (1999), 236–266.

- [11] A. Murphy, Band-Limited Variance Tomography of the Riemann Zeros: Resolving Finite-Height Arithmetic Corrections at the Heisenberg Scale, draft v0.2, TOE-Share, 6 September 2026.
- [12] B. Scannell, Observations on the pair correlation of the Riemann zeros, viXra preprint, June 2020.