

Phase-Locked Extreme Compression in Adjacent Zeta-Zero Spacing:

A Seam-Stitched Empirical Study Through the Full Available LMFDB Horizon

Jeffery Huckstead ¹

¹Independent Researcher

April 27, 2026

Preprint DOI: [10.5281/zenodo.19718848](https://doi.org/10.5281/zenodo.19718848)

Abstract

We study rare extreme compressions in adjacent spacing between consecutive nontrivial zeros of the Riemann zeta function on the critical line using a seam-stitched observational board built from the full available LMFDB horizon. Raw zero data are converted into fixed-resolution local spacing summaries with explicit carry-state preservation across shard boundaries, so that adjacent-zero gaps crossing file seams are retained rather than discarded. On this board we define a fitted local spacing spine and a row-level compression score, then isolate connected packets of the most extreme compressions under frozen morphology rules.

The principal empirical finding is that these packets are not phase-neutral. Across broad-board resolution ablations at 130k, 90k, 70k, and 52k zeros per row, the dominant organizing coordinate remains the first logarithmic harmonic

$$h_1(T) = \left\{ \frac{\log(T/2\pi)}{\pi} \right\},$$

with stable packet concentration statistics and stable continuation-family selection in the trimmed late-stage continuation audit. In every broad-board run tested, the best continuation family for both the local drift coefficient b_1 and the local curvature coefficient b_2 was quadratic.

A terminal-band comparator on the final 28.5B to 31.5B segment strengthens this picture. On that band alone, the first harmonic remains dominant at 52k, 8k, and 1k resolution, with resultant values approximately 0.999202, 0.999315, and 0.999165, and Arc80 widths approximately 0.017055, 0.015305, and 0.015836, respectively. Thus the terminal tightening is a property of the terminal band itself rather than a 1k-only micro-resolution artifact. Moreover, the 8k board yields the strongest concentration of the three tested terminal scales, suggesting an intermediate scale of maximal observed coherence within the terminal band rather than a trivial monotone dependence on ruler size. In the 1k terminal micro-board, the seam bridge at 28,505M remained small, with a fitted-residual jump of order 10^{-7} , while trimmed late-stage continuation

again selected quadratic models for both b_1 and b_2 . Short forward extrapolations just beyond the observed edge remained smooth, with b_1 positive near 1.04×10^{-5} and b_2 negative near -1.70×10^{-10} .

The empirical claim of this revision is intentionally narrower and stronger than in earlier drafts: rare extreme compressions organize reproducibly in the first logarithmic harmonic, this organization is stable under substantial row-size ablation, and the terminal band exhibits unusually tight local h_1 locking. We do not claim a closed asymptotic law beyond the observed horizon.

Keywords: Riemann zeta zeros, zero spacing, empirical phase law, logarithmic phase coordinate, extreme compression, seam stitching, blind holdout, resolution ablation

1 Introduction

The statistical behavior of adjacent spacing between zeta zeros has long been studied through bulk spacing distributions, unfolding, and random-matrix analogies. Much less attention has been paid to the rarest compression events as a distinct empirical subpopulation with their own organizing geometry. This paper studies whether such extreme compressions are merely stochastic tail realizations of the ambient spacing process or whether they exhibit stable structure when viewed in a logarithmic phase coordinate.

The phase coordinate used here is not introduced as an arbitrary search artifact. The Riemann–von Mangoldt counting scale implies a local mean spacing of order

$$\frac{2\pi}{\log(T/2\pi)},$$

so the reciprocal density scale is logarithmic in T . The coordinate

$$h_1(T) = \left\{ \frac{\log(T/2\pi)}{\pi} \right\}$$

therefore acts as a natural logarithmic clock for phenomena tied to the local density regime. Higher harmonics are considered as comparators, but the first harmonic is the primary empirical candidate.

This revision focuses on four questions:

- (1) Do rare extreme compressions organize nontrivially in a logarithmic phase coordinate?
- (2) Does that organization survive substantial changes in fixed row size?
- (3) What continuation family best describes the trimmed late-stage packet drift on the observed board?
- (4) Is the strong terminal-band concentration merely a micro-resolution effect, or a property of the terminal band itself?

The core answer is that the first harmonic remains dominant across all tested broad-board resolutions and on the direct terminal-band comparators. The continuation family selected for

both b_1 and b_2 remains quadratic throughout the tested boards. On the final 28.5B to 31.5B band, the packet field becomes sharply tighter in h_1 at every tested terminal scale, with the 8k board yielding the strongest concentration of the three.

2 Data sources and seam-preserving construction

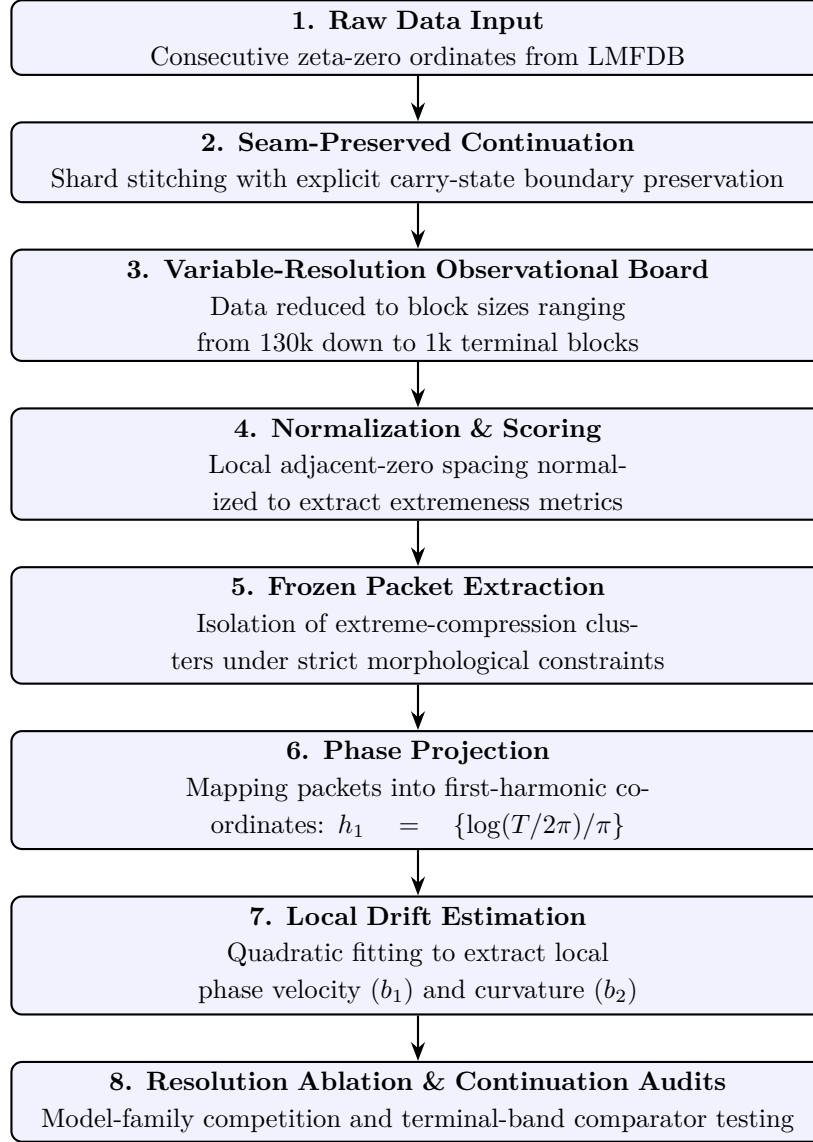


Figure 1: Pipeline overview of the empirical instrument. Consecutive zeta-zero ordinates are decoded in shard order, stitched into a seam-preserved continuation stream, and reduced to variable-resolution boards. Packets are extracted under frozen morphology constraints, projected into logarithmic phase coordinates, and audited through local drift estimation, model-family competition, and resolution-ablation tests.

The primary data source consists of consecutive tabulated ordinates of nontrivial zeta zeros from the LMFDB horizon available at the time of analysis. Source files are decoded in shard order and

assembled into a seam-preserving continuation board. Carry-state is explicitly transferred across shard boundaries so that adjacent-zero gaps crossing one file boundary into the next are preserved.

The main broad-board study uses fixed row sizes ranging from 130k down to 52k zeros per row on the continuous 16.5B to 31.5B board. In addition, a terminal-band comparator is built on the final 28.5B to 31.5B segment at both 52k, 8K, and 1k resolution. These terminal boards are not used to redefine the packet family; rather, they are used to test whether the terminal tightening is stable when the same logic is applied on the same band at very different row sizes.

2.1 Board and audit scales

Table 1 records the principal board scale and the main audited regimes referenced in this revision.

Quantity	Value
Principal fixed board resolution	250,000 zeros per row
Full continuous board rows	415,203
Observed altitude range on principal continuous board	0.09M to 30,610.00M
Nominal terminal board horizon	30.610B
Rows in final continuation extension	29,894
Rows at or above 30,505M on principal board	1,491
Last decoded shard in terminal extension	<code>zeros_30607946000</code>
Primary phase coordinate	$h_1(T) = \{\log(T/2\pi)/\pi\}$
Broad-board ablation sizes in this revision	130k, 90k, 70k, 52k
Direct terminal-band comparator	52k, 8K, and 1k on 28.5B to 31.5B

Table 1: Scale summary for the principal continuous board and the main ablation regimes discussed in this revision.

2.2 Seam metadata and continuity diagnostics

Each row carries metadata indicating whether it is pure interior, crosses a file boundary, or lies near a file edge. These tags are not decorative bookkeeping. They are used to test whether observed phase structure could be explained by seam contamination or edge bias.

A seam-specific continuity diagnostic is therefore part of the instrument itself (Figure 2). The key practical result is that the continuation board does not introduce large fitted-residual jumps at the audited bridge altitudes. This remains true on the broad-board ablations and, in the terminal 1k micro-board, at the 28,505M bridge as well.




Figure 2: Seam continuity diagnostic across the audited historical bridge points of the continuation board. The figure shows the global observed spacing field, the residual field relative to the pure spine, and the residual field relative to the fitted spine, with seam markers overlaid to verify that the seam-preserved continuation logic does not introduce visible discontinuities.

3 Normalization, compression score, and packet extraction

3.1 Local normalization and compression score

Let Δ_k denote the mean adjacent-zero spacing summary in row k , and let T_k denote the representative row altitude. Two reference spines are considered:

$$s_{\text{pure}}(T) = \frac{2\pi}{\log(T/2\pi)}, \quad \hat{s}(T) = \frac{a}{\log(T/b)},$$

where \hat{s} is a fitted local spine estimated from the board itself.

The row-level compression score is then defined explicitly by

$$X_k := 1 - \frac{\Delta_k}{\hat{s}(T_k)}.$$

Large positive values of X_k indicate rows whose local mean spacing is unusually compressed relative to the fitted local spacing scale. This is the extremeness score used operationally in the packet extraction pipeline.

3.2 Frozen packet family

Fix an upper quantile threshold τ on the eligible-row population of the compression score. The selected row set is

$$\mathcal{E}_\tau := \{k : X_k \geq \tau\}.$$

A packet is a connected component of \mathcal{E}_τ in row-index order. The trimmed late-stage audit retains only packets satisfying frozen morphology constraints on row count, altitude width, and seam cleanliness (Table 2).

Criterion	Operational meaning
Packet width ceiling w_{\max}	Maximum permitted altitude span for a retained packet.
Packet row-count ceiling n_{\max}	Maximum permitted number of board rows in a retained packet.
Extremeness threshold τ	Frozen row-level extremeness floor used to define the selected set \mathcal{E}_τ .
Peak signal floor ρ_{\min}	Minimum packet-level peak signal requirement used to exclude weak or diffuse clusters.
Interior-purity floor $p_{\text{int},\min}$	Minimum required fraction of pure-interior rows inside a retained packet.
Boundary-contamination ceiling $p_{\text{bnd},\max}$	Maximum allowed fraction of seam-crossing or boundary-contaminated rows in a retained packet.
Frozen-family status	These constraints are frozen before the deepest holdout phases and then reused without downstream retuning.

Table 2: Schematic definition of the frozen packet family used in the main empirical rail. Exact numerical thresholds are preserved in the associated manifests and artifact tables.

4 Phase coordinates, concentration, and local continuation

4.1 Candidate phase coordinates

For harmonic $m \in \mathbb{N}$, define

$$h_m(T) = \left\{ m \frac{\log(T/2\pi)}{\pi} \right\}.$$

This choice is structurally motivated by the classical density law

$$N(T) = \frac{T}{2\pi} \log\left(\frac{T}{2\pi}\right) - \frac{T}{2\pi} + O(\log T),$$

which yields the leading local mean-spacing scale

$$g(T) \asymp \frac{2\pi}{\log(T/2\pi)}.$$

Accordingly, $\log(T/2\pi)$ acts as a natural density clock, and division by π yields the first-harmonic phase coordinate used throughout the paper. The primary coordinate is h_1 ; higher harmonics are used as comparators rather than as primary hypotheses.

4.2 Packet concentration diagnostics

For packet phases $\theta_1, \dots, \theta_N$, define the circular resultant

$$R := \left| \frac{1}{N} \sum_{j=1}^N e^{2\pi i \theta_j} \right|.$$

For $q \in (0, 1)$, let A_q denote the minimum arc length needed to cover a fraction q of the packet phases on \mathbb{R}/\mathbb{Z} . The operational statistic used in the main text is $A_{0.80}$, denoted Arc80.

4.3 Local quadratic continuation

Packets are sorted by altitude and grouped into rolling windows. Within each rolling window, the unwrapped packet phase is fit by a local quadratic law

$$z_j = \alpha + b_1(T_j - T_0) + b_2(T_j - T_0)^2 + \varepsilon_j,$$

where z_j is the unwrapped local phase and T_0 is the local window center. The coefficient b_1 is interpreted descriptively as local drift and b_2 as local curvature. In the present revision, the competing model-family comparison consistently selected quadratic continuation for both coefficients on the tested boards.

5 Validation protocol and interpretive caution

5.1 Discovery, freeze, and validation chronology

The empirical rail is staged in four logical phases:

- (1) exploratory development on the principal board identifies a workable packet habitat, candidate phase coordinate family, and local-fit architecture;
- (2) the packet family, harmonic scan, and local fitting conventions are frozen before the deepest continuation and terminal audits;
- (3) the frozen configuration is tested under extension-only continuation, disjoint terminal windows, and residual-renormalization attempts on the principal board;
- (4) the stitched field is rebuilt at substantially different non-divisor resolutions, and trimmed late-stage replays are used to test whether the dominant phase organization survives changes in row boundaries, packet segmentation, edge classifications, and seam adjacency.

5.2 Search-space control and multiple-testing caution

The instrument was developed through an exploratory phase, so this paper does not present a single omnibus p -value for the full historical search path. Instead, credibility is assigned through staged freeze points, blind continuation, stability of harmonic ordering, persistence of the small residual layer under habitat and estimator perturbations, and reproducibility of the dominant h_1 organization across independently rebuilt boards.

The present revision therefore treats the first-harmonic phase law as the most stable empirical finding, while treating the exact continuation family as a result that has now stabilized within the trimmed protocol but is still stated as an observed-horizon empirical law rather than an asymptotic theorem.

6 Results

6.1 Broad-board resolution ablation

Table 3 summarizes the trimmed late-stage resolution ladder on the continuous 16.5B to 31.5B board. The dominant harmonic remains h_1 at every tested row size, while the continuation family selected for both b_1 and b_2 remains quadratic throughout.

These row sizes are not trivial divisors of one another in a way that would preserve identical packet segmentation. Changing row size changes row boundaries, edge placements, packet grouping, rolling-window placement, and seam adjacency windows. Despite those changes, the same harmonic and the same continuation family remain preferred. This makes a simple row-periodicity explanation implausible.

Row size	Rows built	Best harm.	R	Arc80	Best b_1	Best b_2
130k	380,597	h_1	0.939191	0.144928	poly_deg2	poly_deg2
90k	549,752	h_1	0.940347	0.143237	poly_deg2	poly_deg2
70k	706,824	h_1	0.938814	0.144641	poly_deg2	poly_deg2
52k	951,494	h_1	0.938226	0.145327	poly_deg2	poly_deg2

Table 3: Broad-board resolution ablation on the continuous 16.5B to 31.5B seam-stitched board. Across all tested row sizes, the best harmonic remains h_1 , and the trimmed late-stage continuation winner remains quadratic for both b_1 and b_2 .

6.2 Terminal-band tightening

The broad-board ablation shows stability. The terminal-band comparator shows something stronger: local tightening. Table 4 compares the same terminal 28.5B to 31.5B band at 52k, 8k, and 1k resolution.

Band	Rows built	Best harmonic	Resultant R	Arc80	Best b_1	Best b_2	Forward b_1	Forward b_2
52k term.	143,863	h_1	0.999202	0.017055	poly_deg2	poly_deg2	$\approx 1.0 \times 10^{-5}$	-1.7×10^{-5}
8k term.	928,008	h_1	0.999315	0.015305	poly_deg2	poly_deg2	$\approx 1.0 \times 10^{-5}$	-1.7×10^{-5}
1k term.	7,480,892	h_1	0.999165	0.015836	poly_deg2	poly_deg2	1.0401×10^{-5}	-1.6×10^{-5}

Table 4: Direct terminal-band comparator on the final 28.5B to 31.5B segment. The unusually tight h_1 concentration persists at 52k, 8k, and 1k resolution, showing that terminal tightening is a property of the terminal band itself rather than a 1k-only artifact. The 8k board yields the smallest Arc80 and the largest resultant among the three tested terminal scales, suggesting a mesoscopic scale of maximal local coherence on the observed band. The forward values shown are the first one-step extrapolations beyond the observed edge.

The three-scale terminal ladder materially sharpens the interpretation of the terminal effect. If the observed tightening were merely a coarse-board smoothing artifact, the 52k board would be expected to dominate. If it were merely a micro-resolution sharpening artifact, the 1k board would be expected to dominate. Instead, the 8k board produces the strongest concentration of the three, while all three boards retain the same best harmonic and the same continuation-family winners. This pattern suggests that the terminal packet field has a genuinely stable local phase organization on the observed horizon, with an intermediate scale of maximal coherence rather than a trivial monotone dependence on board resolution.

The 1k terminal micro-board contained 7,480,892 rows, of which 7,379,386 were pure interior. At the seam near 28,505M, the mean fitted-residual jump across the bridge was

$$1.3086 \times 10^{-7},$$

showing that the seam remains small even under this much finer ruler.

6.3 Terminal 8k reinforcement audits

We applied three reinforcement audits to the terminal 8k board: a matched-packet placement null, a threshold-stability sweep, and a jackknife robustness audit.

To test whether the terminal 8k concentration could be explained by packet morphology alone, we ran a matched-packet placement null on the terminal 8k board. This null preserved packet count and packet row lengths while randomizing packet locations among eligible rows. We first established the effect on a 128-replicate pilot and then deepened the audit to 2048 replicates. On the final 2048-replicate run, the real packet field again selected h_1 as the best harmonic, with

$$R = 0.999315, \quad \text{Arc80} = 0.015305.$$

The matched-null mean at h_1 was

$$R = 0.999150, \quad \text{Arc80} = 0.017977,$$

and the real field exceeded all 2048 null replicates, giving empirical exceedance probability

$$p = \frac{1}{2049} \approx 4.88 \times 10^{-4}.$$

Thus the terminal 8k phase concentration is not explained by packet morphology alone.

We also tested threshold stability directly on the terminal 8k board by replaying the fitted-spine packet extractor across upper-tail thresholds $\tau \in \{0.9900, 0.9925, 0.9950, 0.9975, 0.9990\}$. At every tested threshold, the dominant harmonic remained h_1 . The corresponding best-harmonic resultants were

$$0.999268, 0.999288, 0.999315, 0.999360, 0.999459,$$

with Arc80 widths

$$0.016038, 0.015757, 0.015305, 0.015011, 0.012805,$$

while packet count decreased from 9120 to 921. Thus the terminal 8k phase lock is not tied to a single frozen threshold choice: across a substantial tail ladder, the same first harmonic remains dominant and the packet cloud becomes progressively tighter as the tail is restricted.

Audit	Packet count	Best harmonic	Best R	Arc80
8k real field at $\tau = 0.995$	4581	h_1	0.999315	0.015305
8k matched-packet null mean	4581	h_1	0.999150	0.017977
8k $\tau = 0.9900$	9120	h_1	0.999268	0.016038
8k $\tau = 0.9925$	6852	h_1	0.999288	0.015757
8k $\tau = 0.9950$	4581	h_1	0.999315	0.015305
8k $\tau = 0.9975$	2297	h_1	0.999360	0.015011
8k $\tau = 0.9990$	921	h_1	0.999459	0.012805

Table 5: Terminal 8k threshold and null reinforcement. The matched-packet placement null remains highly concentrated because the terminal corridor itself is narrow, but the real packet field is tighter in both resultant and Arc80. Across the tested threshold ladder, the dominant harmonic remains h_1 throughout.

We also performed a jackknife robustness audit on the same frozen terminal 8k packet field by randomly dropping 20% of packets and recomputing the h_1 concentration statistics 10,000 times. The resultant remained extremely stable, with mean

0.999315,

1st percentile

0.999302,

and absolute worst case

0.999293.

Likewise, Arc80 remained tightly bounded, with mean

0.015316,

99th percentile

0.015557,

and absolute worst case

0.015758.

This shows that the terminal h_1 lock is systemic across the packet population and not driven by a small number of disproportionately influential packets.

Audit	Best harmonic	Resultant R	Arc80	Note
8k real field	h_1	0.999315	0.015305	frozen $\tau = 0.995$ field
8k matched-null mean	h_1	0.999150	0.017977	2048 packet-placement nulls
8k jackknife mean	h_1	0.999315	0.015316	10,000 runs, 20% dropped
8k jackknife 1st / 99th pctl	h_1	0.999302	0.015557	robustness envelope

Table 6: Terminal 8k statistical reinforcement. The real packet field is tighter than the matched-packet null and remains highly stable under aggressive packet subsampling.

6.4 Sensitivity audits and boundary-layer caution

On the principal continuous board, habitat-sensitivity and estimator-definition audits indicate that the remaining small residual layer is not materially erased by admissible packet-habitat perturbations or by reasonable changes in local fitting specification. The important claim here is modest: the residual appears bounded, structured, and sign-stable on the observed horizon. The stronger and more stable empirical law remains the dominant first-harmonic packet organization itself.

6.5 Odlyzko sidecar probes

Sparse high-altitude Odlyzko spot checks were used only as out-of-domain sidecar probes. These windows were tiny relative to the continuous LMFDB rail and were discontinuous by construction. Under the tested residual-phase clustering, distance, and rank-order summaries, they did not yield statistically significant local holdout correspondence at the nominal decision level used in the sidecar notebooks. Because the probes are both discontinuous and low-volume relative to the main board, they are interpreted as limitations checks rather than as extensions of the primary continuous-board claims.

7 Discussion

The empirical picture is now cleaner than in the longer compiled draft. The seam-stitched board behaves stably under major changes in row size. The dominant harmonic does not change. The trimmed continuation family does not change. The terminal band becomes sharply tighter, and that tightening survives an apples-to-apples comparison on the same band at three very different resolutions.

A skeptical interpretation would be that finer row size simply over-resolves local noise. The terminal comparator cuts against that interpretation: the 52k terminal board already shows essentially the same tight locking as the 1k terminal board, while the intermediate 8k board yields the strongest concentration of the three. Another skeptical interpretation would be that the continuation family is an artifact of one arbitrary chunk size. The broad-board ladder at 130k, 90k, 70k, and 52k cuts against that interpretation as well.

The terminal 8k reinforcement audits materially narrow the remaining artifact explanations. The matched-packet placement null shows that packet morphology by itself does not reproduce the full observed concentration: even after preserving packet count and packet row lengths, the

randomized packet field remains measurably weaker than the real one, and the real field exceeds all 2048 null replicates. Meanwhile, the threshold-stability sweep shows that the dominance of h_1 is not an artifact of a single extremeness cutoff. Across tested tail levels from $\tau = 0.9900$ through $\tau = 0.9990$, the winner remains h_1 , while the packet field tightens as the extractor is restricted to rarer compressions.

The jackknife audit strengthens this still further. After randomly removing 20% of packets in 10,000 independent subsamples, the h_1 resultant and Arc80 remain tightly concentrated around their full-field values. This shows that the terminal lock is not carried by a handful of exceptional packets, but is distributed broadly across the packet population.

At the same time, this remains an observational study. The present revision improves identifiability by explicitly defining the compression score, by strengthening the motivation for the phase coordinate, by preserving freeze-point discipline, and by adding broad-board resolution tests, terminal-band comparators, matched-packet null testing, threshold-stability replay, and packet-level robustness audits. Formal asymptotic significance theory and broader synthetic controls remain appropriate targets for subsequent work.

8 Limitations

The principal claims are empirical and confined to the observed horizon. We do not assert that the quadratic continuation family identified here extends as a theorem to greater heights, nor do we claim that terminal-band tightening on the last observed band must persist unchanged to arbitrarily greater heights. The trimmed late-stage continuation audit is intentionally restrictive and therefore speaks to the stability of one frozen protocol rather than to every conceivable packet extractor.

The statistical validation layer is now materially stronger but still intentionally local to the observed horizon. A matched-packet placement null on the terminal 8k board yielded empirical exceedance probability

$$p \approx 4.88 \times 10^{-4}$$

over 2048 replicates, a terminal 8k threshold-stability sweep showed that h_1 remained dominant across all tested upper-tail cuts from $\tau = 0.9900$ to $\tau = 0.9990$, and a 10,000-run jackknife audit showed that the terminal h_1 lock remains stable after random removal of 20% of packets. These tests substantially strengthen identifiability and robustness, but they do not amount to a universal asymptotic significance theory. The paper therefore continues to frame its main conclusion as an observed-horizon empirical law supported by seam diagnostics, freeze-point discipline, resolution ablation, terminal-band comparators, null testing, threshold-stability replay, and packet-level robustness audits.

9 Conclusion

Rare extreme compressions in adjacent zeta-zero spacing are empirically organized by the first logarithmic harmonic

$$h_1(T) = \left\{ \frac{\log(T/2\pi)}{\pi} \right\}$$

on the continuous seam-stitched board. This organization survives broad-board row-size ablations from 130k to 52k zeros per row, and the trimmed late-stage continuation audit consistently selects quadratic models for both b_1 and b_2 on the observed horizon.

On the final 28.5B to 31.5B band, the packet cloud becomes dramatically tighter in h_1 , and this tightening persists across 52k, 8k, and 1k terminal boards, with the intermediate 8k resolution yielding the strongest concentration of the three. This conclusion is further reinforced by a terminal 8k matched-packet placement null, which the real field exceeds across all 2048 tested replicates, by a terminal 8k threshold ladder on which the dominant harmonic remains h_1 throughout, and by a 10,000-run jackknife audit under random removal of 20% of packets, on which the terminal h_1 lock remains highly stable.

The terminal band therefore exhibits genuine local tightening on the observed board, while the broader empirical law remains deliberately stated as an observed-horizon result rather than a closed asymptotic theorem.

A Mathematical Definitions and Operators

This appendix formalizes the principal data objects and operators used in the empirical pipeline.

A.1 Raw zero stream and seam-preserving continuation

Let

$$0 < \gamma_1 < \gamma_2 < \gamma_3 < \dots$$

denote the ordinates of consecutive nontrivial zeros of the Riemann zeta function on the critical line, and let

$$g_n := \gamma_{n+1} - \gamma_n$$

denote the adjacent-zero gaps.

The empirical board is assembled from decoded shards

$$S_1, S_2, \dots, S_m$$

using ordered concatenation with carry-state preservation:

$$\Gamma^{(m)} := S_1 \hat{\cup} S_2 \hat{\cup} \dots \hat{\cup} S_m.$$

The notation $\hat{\cup}$ indicates that cross-shard adjacent-zero gaps are preserved.

A.2 Fixed-resolution observational board

Let M denote the chosen row size in zeros. A row block is

$$W_k := \{\gamma_{(k-1)M+1}, \dots, \gamma_{kM}\},$$

with representative altitude T_k . Each row yields local spacing summaries, including the mean local gap

$$\Delta_k := \frac{1}{M-1} \sum_{\gamma_n, \gamma_{n+1} \in W_k} (\gamma_{n+1} - \gamma_n).$$

A.3 Pure and fitted spines

The pure logarithmic spacing spine is

$$s_{\text{pure}}(T) = \frac{2\pi}{\log(T/2\pi)}.$$

The fitted empirical spine is

$$\hat{s}(T) = \frac{a}{\log(T/b)}.$$

A.4 Compression score

The row-level compression score is

$$X_k := 1 - \frac{\Delta_k}{\hat{s}(T_k)}.$$

This is the explicit extremeness score used in the present revision. Large positive values of X_k correspond to stronger local compression relative to the fitted empirical spacing scale.

A.5 Packet extraction

Fix an upper quantile threshold τ on the eligible-row population of the compression score. The selected row set is

$$\mathcal{E}_\tau := \{k : X_k \geq \tau\}.$$

A packet is a connected component of \mathcal{E}_τ in the line of row indices. The trimmed late-stage audit retains only packets satisfying frozen morphology constraints on row count, altitude width, and seam cleanliness.

A.6 Phase coordinates

For harmonic $m \in \mathbb{N}$, define

$$h_m(T) := \left\{ m \frac{\log(T/2\pi)}{\pi} \right\}.$$

If T_j^* is the representative altitude of packet P_j , then the packet phase in harmonic m is

$$\theta_j^{(m)} := h_m(T_j^*).$$

A.7 Concentration diagnostics

For packet phases $\theta_1, \dots, \theta_N$, define the circular resultant

$$R := \left| \frac{1}{N} \sum_{j=1}^N e^{2\pi i \theta_j} \right|.$$

For $q \in (0, 1)$, let A_q denote the minimum arc length needed to cover a fraction q of the packet phases on \mathbb{R}/\mathbb{Z} . The operational statistic used in the main text is $A_{0.80}$, denoted Arc80.

A.8 Local quadratic continuation

Packets are sorted by altitude and grouped into rolling windows. Within each rolling window, the unwrapped packet phase is fit by a quadratic law

$$z_j = \alpha + b_1(T_j - T_0) + b_2(T_j - T_0)^2 + \varepsilon_j,$$

where z_j is the unwrapped local phase and T_0 is the local window center. The coefficient b_1 is interpreted descriptively as local drift and b_2 as local curvature.

A.9 Continuation family comparison

The trimmed late-stage continuation audit compares several simple families for b_1 and b_2 , including constant, linear, quadratic, and exponential-relaxation comparators. On every broad-board and terminal-band run reported in this revision, the winning family for both b_1 and b_2 was quadratic.

A.10 Forward prediction

Let $\hat{b}_1(T)$ and $\hat{b}_2(T)$ denote the best-fit continuation laws on the observed rolling-window centers. Short forward predictions are obtained by evaluating these fits on the next one or two extrapolated center locations beyond the observed edge. These forward values are reported descriptively and are not claimed as asymptotic closure.

A.11 Code and artifact mapping

Math symbol	Formal meaning	Code / artifact analogue
γ_n	consecutive zero ordinates	decoded zero stream
g_n	adjacent-zero gap	seam-stitched gap extractor
W_k	fixed-resolution window	row block at chosen resolution
P_j	extreme packet	connected selected row cluster
$h_1(T)$	first-harmonic phase	<code>logT_over_pi_h1</code>
$b_{1,m}, b_{2,m}$	local drift / curvature	local quadratic coefficients
$\varepsilon_{\text{term}}$	terminal residual	terminal holdout miss / boundary layer

Table 7: Translation of formal mathematical notation to the principal code-facing objects in the computational pipeline.

Data and Artifact Availability

A timestamped preprint version of this manuscript has been deposited under the DOI listed on the title page. The associated research package includes the seam-stitched observational boards, continuation manifests, audit summaries, and derived tables referenced in this paper.

Author Contributions

Jeffery Huckstead conceived the research program, designed the empirical pipeline, executed the computational campaign, analyzed the results, and wrote the manuscript.

Use of AI Tools

Generative AI tools were used for drafting assistance, editorial restructuring, code critique, and adversarial review of claims. All research design decisions, validation strategy, computational execution choices, interpretation, and final manuscript responsibility remain with Jeffery Huckstead.

Priority Statement

This preprint documents the first public release of the seam-stitched empirical board, packet-extraction framework, logarithmic harmonic analysis, and resolution-ablation results reported herein.