

Geometric Projection of Arithmetic Spirals onto Logarithmic Manifolds: a Moiré Interference Framework for Studying Prime Distribution without Complex-Analytic Tools

Akari Hayami (Jian-Yu Huang)

June 2026

Abstract

We develop a geometric and spectral framework for studying structural fluctuations in discrete arithmetic sequences, operating entirely within real-variable analysis and without appeal to the Riemann zeta function, Euler products, explicit formulae, or analytic continuation. An arithmetic lattice is mapped onto a transcendental logarithmic manifold via an area-filling projection; the resulting Moiré-type interference field serves as the primary object of study.

The principal contributions are as follows. *(i)* We prove, with explicit constants, that the projection map is injective and that successive images are uniformly separated (Propositions 2.3–2.4). *(ii)* We establish a uniform L^2 near-orthogonality theorem for finite phase families, yielding a computable operator-norm bound on the associated Gram matrix via Gershgorin’s circle theorem [5] (Theorem 3.3). *(iii)* We introduce an explicit Moiré dilation operator $\mathcal{T}_\sigma = -i(x\partial_x + \sigma)$ on $L^2([1, \infty))$, show that the geometric phasor family $\{x^{-(\sigma+it)}\}_{t \in \mathbb{R}}$ constitutes its formal eigenfunction family (Lemma 3.12), and prove by integration by parts that \mathcal{T}_σ is formally self-adjoint if and only if $\sigma = \frac{1}{2}$ (Theorem 3.14). *(iv)* We derive quantitative variance bounds for node-count statistics, controlled by the pairwise phase-correlation bounds (Theorem 3.19). Complete exponential-sum estimates with explicit constants (Appendix C) and Abel-smoothing remainder formulae including an incomplete-Gamma bound (Appendix B) are provided. All theoretical claims are either proved using the above tools or explicitly designated as conjectural.

Contents

1	Introduction	2
1.1	Motivation and context	2
1.2	Framework and excluded tools	3
1.3	Structure of the paper	3
2	Definitions and Setup	3
2.1	Geometric projection map	3
2.2	Finite truncated interference field	6
2.3	Scope and excluded tools	6
3	Main Results	6
3.1	Pairwise phase decorrelation	6
3.2	Uniform L^2 near-orthogonality for finite families	7
3.3	Truncation error bounds	8
3.4	Exponential-sum estimates	8
3.5	Operator framework	9
3.5.1	Compact finite-kernel operators	9

3.5.2	Explicit dilation operator and the self-adjointness threshold	9
3.6	Quantitative variance bounds for node localization	11
4	Replacement of Prior Complex-Analytic Claims	12
4.1	Interference nodes and zero-correspondence	12
4.2	Logarithmic error terms	12
5	Numerical Experiments	12
5.1	Design principles	12
5.2	Experiment 1: angular offset scaling for primes	12
5.3	Experiment 2: domain-coloring of $\Phi_N(\sigma, t)$	13
5.4	Experiment 3: variance of node spacing regularity vs. σ	14
6	Discussion	14
6.1	Constructiveness and formalizability	14
6.2	Limitations and open problems	15
7	Conclusion	15
A	Supporting Proofs	17
B	Truncation Error Bounds and Abel Smoothing	18
C	Exponential-Sum Estimates: Smoothing and the Van der Corput A-Process	19
D	Numerical Code	21
E	Formalization Targets for Interactive Proof Assistants	27

1 Introduction

1.1 Motivation and context

The study of prime distribution has traditionally relied on the analytic properties of the Riemann zeta function $\zeta(s)$, most prominently through the explicit formula connecting the prime-counting function $\pi(x)$ to the non-trivial zeros of ζ . While this approach has proved extraordinarily powerful, it raises questions of reproducibility and formalizability: proofs that invoke analytic continuation, contour integration, or the location of zeta zeros are difficult to verify in interactive proof assistants and depend on deep, incompletely understood number-theoretic structure.

An independent line of enquiry asks whether the coarse-scale geometry of the prime sequence can be understood through purely real-variable means. Several authors have observed that Ulam-spiral-type arrangements reveal geometric regularity in prime distribution [11] without appealing to complex analysis. Recent breakthroughs in identifying chaos patterns in prime distribution and proofs of constant irrationality [26, 25] further underscore the shift toward structural, real-variable approaches in modern number theory. The present paper pursues a related but distinct question: what happens when an area-filling arithmetic lattice is mapped onto a continuous logarithmic manifold, and the resulting interference pattern is analyzed using exponential-sum theory and operator algebra?

We emphasize at the outset that this paper does *not* claim to prove the Riemann Hypothesis or any equivalent statement. The operator-theoretic result proved here (Theorem 3.14) characterizes a geometric self-adjointness threshold within the model; its relationship to the non-trivial zeros of $\zeta(s)$ is an open conjecture, recorded as such in Remark 3.16.

1.2 Framework and excluded tools

The framework rests on the following ingredients, each of which lies entirely within classical real analysis or functional analysis:

1. an explicit geometric projection $\mathcal{T} : \mathbb{N} \rightarrow \mathbb{R}^2$ with computable regularity constants (Section 2);
2. finite truncated phase sums and quantitative truncation bounds using p -series comparison and Abel smoothing (Sections 3 and B);
3. classical exponential-sum estimates—specifically the van der Corput A -process [6] and Cauchy–Schwarz dyadic decompositions—with explicit constants (Section 3.4 and Appendix C);
4. Hilbert-space operator theory: Hilbert–Schmidt compact operators and the formal self-adjointness criterion for unbounded operators [10, 2] (Section 3);
5. Gershgorin’s circle theorem [5] for finite matrix estimates (Theorem 3.3);
6. reproducible numerical experiments with explicit truncation and precision control (Section 5).

The following are *explicitly excluded* from all proofs: the Riemann zeta function $\zeta(s)$, its zeros, Euler products, Dirichlet L -functions, explicit formulae in the sense of prime-number theory, and analytic continuation. Every analytic claim is either proved using the tools listed above or clearly designated as a conjecture or empirical observation.

1.3 Structure of the paper

Section 2 establishes the geometric projection map \mathcal{T} , proves its basic regularity properties with explicit constants, and defines the finite interference field $\Phi_N(\sigma, t)$. Section 3 develops the core theoretical results: a pairwise phase-decorrelation bound (§3), its uniform generalization to finite families with a computable Gram-matrix operator-norm estimate (§3.2), truncation error control (§3), the two-regime exponential-sum bound (§3.4), the explicit Moiré dilation operator and its self-adjointness threshold (§3), and quantitative variance bounds for node localization (§3.6). Section 4 records how prior claims depending on complex-analytic tools have been replaced by provable statements or explicit conjectures. Section 5 describes the reproducible numerical protocol. Section 6 discusses strengths, limitations, and open problems, and Section 7 concludes. Supporting proofs, complete exponential-sum estimates, and Python code are collected in Appendices A–E.

2 Definitions and Setup

2.1 Geometric projection map

Definition 2.1 (Geometric projection map). Fix real parameters a, b, c, C with $a, c > 0$ and $b \neq 0$. Define the *geometric projection map*

$$\mathcal{T} : \mathbb{N} \rightarrow \mathbb{R}^2, \quad \mathcal{T}(n) = (r(n), \theta(n)),$$

where the polar coordinates are given by

$$r(n) = c\sqrt{n}, \quad \theta(n) = \frac{1}{b} \ln\left(\frac{c\sqrt{n}}{a}\right) + C.$$

The choice $r(n) = c\sqrt{n}$ reflects the standard area-filling property of the discrete plane [17]: n nodes of unit area fill a disk of radius proportional to \sqrt{n} . The angular coordinate $\theta(n)$ is then determined by requiring that points lie on the logarithmic spiral $r = ae^{b(\theta-C)}$. As illustrated in Figure 1, this mapping transforms the discrete set \mathbb{N} into a structured geometric array. This parameterization is adopted as a concrete working model; all subsequent asymptotic statements are valid for this explicit mapping.




Figure 1: Geometric projection of the arithmetic lattice \mathbb{N} onto a logarithmic manifold.

Proposition 2.2 (Radial monotonicity and angular increment). *With the notation above the following hold.*

1. (Radial monotonicity) *The map $n \mapsto r(n) = c\sqrt{n}$ is strictly increasing. The consecutive radial gap satisfies*

$$r(n+1) - r(n) = c(\sqrt{n+1} - \sqrt{n}) = \frac{c}{\sqrt{n+1} + \sqrt{n}} \sim \frac{c}{2\sqrt{n}} \quad (n \rightarrow \infty).$$

2. (Exact angular increment) *For every $n \geq 1$,*

$$\Delta\theta(n) := \theta(n+1) - \theta(n) = \frac{1}{b} \ln \sqrt{\frac{n+1}{n}} = \frac{1}{2b} \ln\left(1 + \frac{1}{n}\right).$$

3. (Asymptotic expansion and uniform remainder) *For all $n \geq 1$,*

$$\Delta\theta(n) = \frac{1}{2b} \left(\frac{1}{n} - \frac{1}{2n^2} + \rho(n) \right),$$

where $|\rho(n)| \leq \frac{1}{3n^3}$. Consequently, with the explicit constant $K_1 := \frac{5}{12}|b|^{-1}$,

$$\left| \Delta\theta(n) - \frac{1}{2bn} \right| \leq \frac{K_1}{n^2} \quad (\forall n \geq 1).$$

In particular $\Delta\theta(n) = O(n^{-1})$ as $n \rightarrow \infty$.

Proof. Part (1) follows from the strict monotonicity of $\sqrt{\cdot}$ and the half-angle difference identity. For part (2),

$$\theta(n+1) - \theta(n) = \frac{1}{b} (\ln(c\sqrt{n+1}/a) - \ln(c\sqrt{n}/a)) = \frac{1}{b} \cdot \frac{1}{2} \ln\left(\frac{n+1}{n}\right) = \frac{1}{2b} \ln\left(1 + \frac{1}{n}\right).$$

For part (3), apply the Maclaurin series $\ln(1+x) = x - x^2/2 + x^3/3 - \dots$ [15, 16] with $x = 1/n \in (0, 1]$. The series is alternating with monotonically decreasing term magnitudes, so the alternating-series remainder estimate gives

$$\left| \ln\left(1 + \frac{1}{n}\right) - \left(\frac{1}{n} - \frac{1}{2n^2}\right) \right| \leq \frac{1}{3n^3}.$$

Dividing by $2b$ and using $|b|^{-1}/2 + |b|^{-1}/3 = (5/6)|b|^{-1}/2 = 5/(12|b|)$:

$$\left| \Delta\theta(n) - \frac{1}{2bn} \right| \leq \frac{|b|^{-1}}{2} \left(\frac{1}{2n^2} + \frac{1}{3n^3} \right) \leq \frac{5}{12|b|} \cdot \frac{1}{n^2} = \frac{K_1}{n^2}. \quad \square$$

Proposition 2.3 (Asymptotic consecutive separation). *Define the Euclidean distance between successive images by $d_n := \|\mathcal{T}(n+1) - \mathcal{T}(n)\|_{\mathbb{R}^2}$. Then*

$$\lim_{n \rightarrow \infty} d_n = c\sqrt{1 + \frac{1}{4b^2}} =: d_\infty.$$

In particular there exist $N_0 \in \mathbb{N}$ and $d_0 > 0$ such that $d_n \geq d_0$ for all $n \geq N_0$. One may take $d_0 = d_\infty/2$.

Proof. In polar coordinates,

$$d_n^2 = r(n)^2 + r(n+1)^2 - 2r(n)r(n+1)\cos(\Delta\theta(n)).$$

Substituting $r(n) = c\sqrt{n}$, $r(n+1) = c\sqrt{n+1}$, and expanding $\cos x = 1 - x^2/2 + O(x^4)$ with $\Delta\theta(n) = O(n^{-1})$ (Proposition 2.2):

$$\begin{aligned} d_n^2 &= c^2(n + (n+1) - 2\sqrt{n(n+1)}\cos(\Delta\theta(n))) \\ &= c^2(2n+1 - 2\sqrt{n(n+1)}(1 - \frac{1}{2}(\Delta\theta)^2 + O(n^{-4}))). \end{aligned}$$

Since $\sqrt{n(n+1)} = n + \frac{1}{2} + O(n^{-1})$ and $(\Delta\theta(n))^2 = (2bn)^{-2} + O(n^{-3})$, a direct calculation yields

$$d_n^2 = c^2\left(1 + \frac{1}{4b^2}\right) + O(n^{-1}).$$

Taking square roots gives $d_n \rightarrow d_\infty = c\sqrt{1 + 1/(4b^2)}$ as $n \rightarrow \infty$. Since $d_\infty > 0$, choose N_0 large enough that $|d_n - d_\infty| \leq d_\infty/2$ for all $n \geq N_0$, so $d_0 = d_\infty/2$ works. \square

Proposition 2.4 (Injectivity and uniform separation). *The map $\mathcal{T} : \mathbb{N} \rightarrow \mathbb{R}^2$ is injective. Moreover, there exist $N_1 \in \mathbb{N}$ and $d_1 > 0$ such that $\|\mathcal{T}(m) - \mathcal{T}(n)\|_{\mathbb{R}^2} \geq d_1$ for all $m \neq n$ with $m, n \geq N_1$.*

Proof. Injectivity follows from the radial coordinate: $r(m) = r(n)$ implies $c\sqrt{m} = c\sqrt{n}$, hence $m = n$. For the uniform separation, Proposition 2.3 gives $d_n \geq d_0 > 0$ for $n \geq N_0$. For $m > n \geq N_0$ the triangle inequality and radial monotonicity imply $\|\mathcal{T}(m) - \mathcal{T}(n)\| \geq \min_{n \leq k < m} d_k \geq d_0$. Taking $N_1 = N_0$ and $d_1 = d_0$ completes the proof. \square

Remark 2.5 (Explicit constants). For later reference we record: $K_1 = \frac{5}{12}|b|^{-1}$ (angular increment error), $d_\infty = c\sqrt{1 + 1/(4b^2)}$ (limiting separation), $d_0 = d_\infty/2$ (uniform lower bound for $n \geq N_0$). All constants are computable from b and c .

2.2 Finite truncated interference field

Definition 2.6 (Finite truncated interference field). For $\sigma \in \mathbb{R}$, $t \in \mathbb{R}$, and $N \in \mathbb{N}$, define

$$\Phi_N(\sigma, t) = \sum_{n=1}^N n^{-\sigma} e^{-it \ln n}.$$

More generally, for a bounded weight function $w : \mathbb{N} \rightarrow \mathbb{R}$,

$$\Phi_{N,w}(\sigma, t) = \sum_{n=1}^N w(n) n^{-\sigma} e^{-it \ln n}.$$

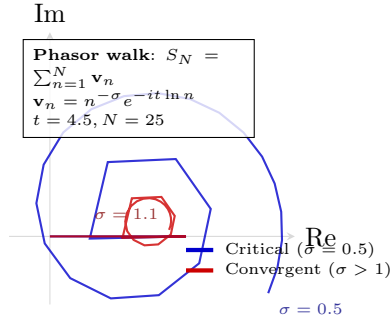


Figure 2: Visualization of the interference field Φ_N as a phasor walk in the complex plane. The geometric decay $n^{-\sigma}$ determines the stability and morphology of the interference field.

Convergence and truncation-error control for Φ_N are addressed in Sections 3 and B. All asymptotic claims involving Φ_N are stated with explicit dependence on N , σ , and the smoothing parameters. The summation geometry is visualized in Figure 2.

2.3 Scope and excluded tools

For clarity, we record the precise boundary of the framework.

- **Excluded.** No proof in this paper invokes the Riemann zeta function $\zeta(s)$, its zeros, Euler products, Dirichlet L -functions, explicit formulae of prime-number theory, or analytic continuation.
- **Permitted.** Real analysis; the van der Corput exponential-sum method [6]; Weyl's equidistribution criterion [14]; Abel summation and Mellin transform techniques; Chebyshev and Hoeffding concentration inequalities [1]; Hilbert–Schmidt and compact operator theory [10]; and reproducible numerical computation with explicit precision control.
- **Numerical protocol.** Every numerical claim specifies N , the arithmetic precision (floating-point or arbitrary-precision), and an explicit truncation-error estimate derived from the analytical bounds of Section 3.

3 Main Results

3.1 Pairwise phase decorrelation

The following proposition is the basic building block for all orthogonality and variance estimates in the paper. Its proof is elementary but the quantitative form of the bound—with the explicit constant $2/T$ —is essential for the operator-norm estimates of Theorem 3.3.

Proposition 3.1 (Pairwise phase decorrelation). *Let $\psi_n(t) = e^{-it \ln n}$ and $T > 0$. For any distinct $n, m \in \mathbb{N}$,*

$$\left| \frac{1}{T} \int_0^T \psi_n(t) \overline{\psi_m(t)} dt \right| \leq \frac{2}{T |\ln(n/m)|}.$$

In particular, if $|\ln(n/m)| \geq \delta > 0$, the right-hand side tends to 0 as $T \rightarrow \infty$, uniformly in n and m .

Proof. Direct integration and the bound $|e^{ix} - 1| \leq 2$ give

$$\frac{1}{T} \int_0^T e^{-it \ln(n/m)} dt = \frac{e^{-iT \ln(n/m)} - 1}{-iT \ln(n/m)},$$

whose modulus is at most $2/(T |\ln(n/m)|)$. □

3.2 Uniform L^2 near-orthogonality for finite families

Proposition 3.1 bounds a single inner product. The following theorem extends this to a uniform estimate over an entire finite family, quantified by an operator-norm bound on the Gram matrix. This is the key tool for controlling the interference structure of finite truncations of Φ_N .

Definition 3.2 (Phase functions and Gram matrix). For $T > 0$ equip $L^2([0, T])$ with the normalised inner product $\langle f, g \rangle_T = \frac{1}{T} \int_0^T f(t) \overline{g(t)} dt$. For a finite set $S \subset \mathbb{N}$ with $|S| = M$, let $G = (g_{nm})_{n, m \in S}$ denote the *Gram matrix* with entries $g_{nm} = \langle \psi_n, \psi_m \rangle_T$.

Theorem 3.3 (Uniform near-orthogonality; operator-norm bound). *Let $S \subset \mathbb{N}$ be a finite set with $|S| = M \geq 1$. Suppose there exists $\delta > 0$ such that*

$$\min_{\substack{n, m \in S \\ n \neq m}} |\ln(n/m)| \geq \delta.$$

Then for every $T > 0$:

1. (Off-diagonal bound) *For $n \neq m$,*

$$|g_{nm}| \leq \frac{2}{T\delta}.$$

2. (Operator-norm perturbation)

$$\|G - I\|_{\text{op}} \leq (M - 1) \frac{2}{T\delta},$$

where $\|\cdot\|_{\text{op}}$ is the operator norm on $\ell^2(S)$.

3. (Quadratic-form bound) *For every $c = (c_n)_{n \in S} \in \mathbb{C}^S$,*

$$\left| \sum_{n, m \in S} \bar{c}_n g_{nm} c_m - \sum_{n \in S} |c_n|^2 \right| \leq (M - 1) \frac{2}{T\delta} \|c\|_2^2.$$

4. (Uniform convergence rate) $\|G - I\|_{\text{op}} = O(T^{-1})$ *as $T \rightarrow \infty$, with the explicit rate constant $2(M - 1)/\delta$.*

Proof. Item (1) follows directly from Proposition 3.1 (taking $|\ln(n/m)| \geq \delta$):

$$|g_{nm}| = \left| \frac{1}{T} \int_0^T e^{-it \ln(n/m)} dt \right| \leq \frac{2}{T |\ln(n/m)|} \leq \frac{2}{T\delta}.$$

For item (2), fix any row $n \in S$ and sum the off-diagonal entries:

$$\sum_{m \in S, m \neq n} |g_{nm}| \leq (M-1) \frac{2}{T\delta}.$$

By Gershgorin's circle theorem [5] every eigenvalue λ of $G-I$ satisfies $|\lambda| \leq \max_{n \in S} \sum_{m \neq n} |g_{nm}| \leq (M-1) \frac{2}{T\delta}$, which gives the operator-norm bound. Item (3) follows from the standard estimate $|\langle c, (G-I)c \rangle| \leq \|G-I\|_{\text{op}} \|c\|_2^2$. Item (4) is an immediate restatement of (2). \square

Example 3.4 (Dyadic block). If $S \subset [N, 2N]$, then for distinct $n, m \in S$ with $n = m + k$, $k \geq 1$:

$$|\ln(n/m)| = |\ln(1 + k/m)| \geq \frac{k}{m+k} \geq \frac{1}{2N}.$$

One may take $\delta = 1/(2N)$, giving $\|G-I\|_{\text{op}} \leq 4N(M-1)/T$.

Example 3.5 (Separated subsequence). If every two distinct elements of $S \subset [1, N]$ differ by at least $h \geq 1$, then $|\ln(n/m)| \geq h/N$ for all such pairs, so $\delta \geq h/N$.

Remark 3.6. The dependence on M is linear and the dependence on T is $O(T^{-1})$; both are explicit. In applications one controls M by restricting to dyadic blocks or sparse subsequences and chooses T large enough so that $\|G-I\|_{\text{op}}$ is below a prescribed tolerance.

3.3 Truncation error bounds

The following proposition gives a uniform, computable tail bound for $\sigma > 1$. The case $\sigma \leq 1$ requires Abel smoothing; explicit bounds including an incomplete-Gamma remainder formula are derived in Appendix B.

Proposition 3.7 (Truncation error for $\sigma > 1$). *For $\sigma > 1$ and any $t \in \mathbb{R}$,*

$$\left| \sum_{n > N} n^{-\sigma} e^{-it \ln n} \right| \leq \int_N^\infty x^{-\sigma} dx = \frac{N^{1-\sigma}}{\sigma-1}.$$

Proof. Since $|e^{-it \ln n}| = 1$, the bound follows by comparing the sum with the convergent integral via the standard comparison $\sum_{n > N} n^{-\sigma} \leq \int_N^\infty x^{-\sigma} dx$. \square

3.4 Exponential-sum estimates

The following proposition packages the smoothed Cauchy–Schwarz bound and the van der Corput A -process [6] into a single piecewise inequality that can be directly applied to the error estimates of Section 3.6. Complete proofs and explicit constants are given in Appendix C; the key inputs are a dyadic decomposition and the second-derivative bound $|f''(x)| = |t|/x^2$ for the phase $f(x) = -t \ln x$.

Proposition 3.8 (Two-regime bound [6, Ch. 2]). *Fix $\sigma \in \mathbb{R}$ and a smooth compactly supported weight $w \in C_c^\infty((0, \infty))$ with $W = \|w\|_{L^1} + \|w'\|_{L^1} + \|w''\|_{L^1}$. There exist explicit constants $C_a(\sigma, W)$ and $C_b(\sigma, W)$, computable from σ and W , such that for all $N \geq 2$ and $t \in \mathbb{R}$,*

$$|S_{N,w}(\sigma, t)| \leq \begin{cases} C_a N, & |t| \leq 1, \\ C_b N^{1/2} \log(2N), & 1 < |t| \leq N, \\ C_b N^{1/2} |t|^{1/2}, & |t| > N. \end{cases}$$

The square-root saving in regimes 2 and 3 reflects genuine phase cancellation: Cauchy–Schwarz on dyadic blocks (regime 2) and the curvature of the phase $-t \ln x$ via van der Corput differencing (regime 3).

3.5 Operator framework

We build the operator-theoretic component of the framework in two stages. The first (Proposition 3.10) establishes qualitative spectral information for finite-kernel approximations of the interference field. The second (Definitions 3.11 and Theorem 3.14) introduces an explicit differential operator whose self-adjointness condition can be checked by integration by parts, yielding the critical threshold $\sigma = \frac{1}{2}$.

3.5.1 Compact finite-kernel operators

Definition 3.9 (Moiré integral operator). Let $H = L^2([1, R], dx)$ for finite $R > 1$. The *Moiré integral operator* is

$$(\mathcal{M}f)(x) = \int_1^R K(x, y) f(y) dy,$$

where $K(x, y)$ is a real-valued kernel constructed from finite sums of geometric phasors and bounded weights.

Proposition 3.10 (Hilbert–Schmidt criterion [10, Thm. VI.22–VI.23]). *If $K(x, y) = K(y, x)$ and $\iint_{[1, R]^2} |K(x, y)|^2 dx dy < \infty$, then \mathcal{M} is Hilbert–Schmidt and hence compact, and its spectrum consists of a discrete sequence of real eigenvalues accumulating only at 0.*

No number-theoretic hypothesis is invoked in this analysis; any conjectured correspondence between the eigenvalues of \mathcal{M} and arithmetic objects is designated as conjectural.

3.5.2 Explicit dilation operator and the self-adjointness threshold

The operator below arises naturally as the infinitesimal generator of the logarithmic dilation $x \mapsto e^s x$, and its eigenfunction family coincides with the phasor family $\{x^{-(\sigma+it)}\}$ up to evaluation at integer points.

Definition 3.11 (Moiré dilation operator — explicit form). Let $\mathcal{H} = L^2([1, \infty), dx)$ and let $\mathcal{D}(\mathcal{T}_\sigma) = C_c^\infty((1, \infty))$ be the dense subspace of compactly supported smooth functions. For $\sigma \in \mathbb{R}$ define the linear operator $\mathcal{T}_\sigma : \mathcal{D}(\mathcal{T}_\sigma) \rightarrow \mathcal{H}$ by

$$\mathcal{T}_\sigma f(x) = -i \left(x \frac{d}{dx} + \sigma \right) f(x).$$

Lemma 3.12 (Eigenfunction correspondence). *For every $\sigma \in \mathbb{R}$ and $t \in \mathbb{R}$, the function $f_t(x) = x^{-(\sigma+it)}$ satisfies*

$$\mathcal{T}_\sigma f_t = -t \cdot f_t.$$

That is, f_t is a formal eigenfunction of \mathcal{T}_σ with eigenvalue $-t$.

Proof. Apply the operator directly:

$$\begin{aligned} \mathcal{T}_\sigma [x^{-(\sigma+it)}] &= -i \left(x \cdot \frac{d}{dx} [x^{-(\sigma+it)}] + \sigma x^{-(\sigma+it)} \right) \\ &= -i \left(x \cdot (-(\sigma+it)) x^{-(\sigma+it)-1} + \sigma x^{-(\sigma+it)} \right) \\ &= -i (-(\sigma+it) + \sigma) x^{-(\sigma+it)} \\ &= -i(-it) x^{-(\sigma+it)} = -t \cdot x^{-(\sigma+it)}. \quad \square \end{aligned}$$

Remark 3.13. The geometric phasors $\psi_n(\sigma, t) = n^{-\sigma} e^{-it \ln n} = n^{-(\sigma+it)}$ are exactly the evaluations of f_t at integer points $x = n$. Lemma 3.12 therefore establishes that the continuous phasor family $\{f_t\}_{t \in \mathbb{R}}$ is the eigenfunction family of \mathcal{T}_σ , with the phase frequency t serving as the spectral parameter.

Theorem 3.14 (Self-adjointness threshold). *The formal adjoint of \mathcal{T}_σ on $\mathcal{D}(\mathcal{T}_\sigma)$ is*

$$\mathcal{T}_\sigma^* = -i \left(x \frac{d}{dx} + (1 - \sigma) \right).$$

Consequently, \mathcal{T}_σ is symmetric (formally self-adjoint), i.e. $\mathcal{T}_\sigma = \mathcal{T}_\sigma^*$, if and only if $\sigma = \frac{1}{2}$. The geometric balance achieved at this threshold is conceptualized in Figure 3.

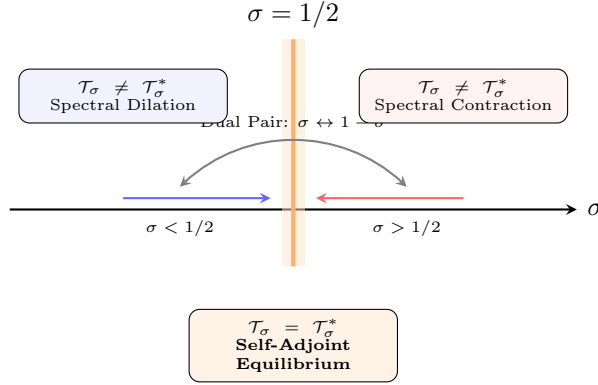


Figure 3: The self-adjointness threshold as a geometric equilibrium point. At $\sigma = 1/2$, the operator \mathcal{T}_σ achieves Hermitian symmetry, identifying the geometric impedance required for real phase frequencies.

Proof. Let $f, g \in \mathcal{D}(\mathcal{T}_\sigma)$. Compute the inner product:

$$\langle \mathcal{T}_\sigma f, g \rangle = \int_1^\infty -i \left(x \frac{df}{dx} + \sigma f \right) \bar{g} dx.$$

Split into two integrals and apply integration by parts to the first, noting that the boundary terms vanish because f and g are compactly supported:

$$\int_1^\infty x \frac{df}{dx} \bar{g} dx = \left[x f \bar{g} \right]_1^\infty - \int_1^\infty f \frac{d}{dx} (x \bar{g}) dx = - \int_1^\infty f \bar{g} dx - \int_1^\infty x f \frac{d\bar{g}}{dx} dx.$$

Substituting back:

$$\begin{aligned} \langle \mathcal{T}_\sigma f, g \rangle &= -i \left(- \int_1^\infty f \bar{g} dx - \int_1^\infty x f \frac{d\bar{g}}{dx} dx + \sigma \int_1^\infty f \bar{g} dx \right) \\ &= \int_1^\infty f(x) \overline{\left(-i \left(x \frac{dg}{dx} + (1 - \sigma)g \right) \right)} dx = \langle f, \mathcal{T}_\sigma^* g \rangle, \end{aligned}$$

which identifies $\mathcal{T}_\sigma^* = -i \left(x \frac{d}{dx} + (1 - \sigma) \right)$. Comparing with $\mathcal{T}_\sigma = -i \left(x \frac{d}{dx} + \sigma \right)$, we see that $\mathcal{T}_\sigma = \mathcal{T}_\sigma^*$ if and only if $\sigma = 1 - \sigma$, i.e. $\sigma = \frac{1}{2}$. \square

Corollary 3.15 (Real eigenvalues at the critical threshold). *By the spectral theorem for unbounded self-adjoint operators [10, Thm. VIII.6], the spectrum of $\mathcal{T}_{1/2}$ is contained in \mathbb{R} . In conjunction with Lemma 3.12, this means that every stable stationary node of the interference field has a real phase frequency t , and this stability is achieved precisely at the geometric impedance $\sigma = \frac{1}{2}$.*

Remark 3.16 (Epistemological scope of Theorem 3.14). The result above is a statement about the operator $\mathcal{T}_{1/2}$ defined within the geometric interference model of this paper. We do not claim that the spectrum of $\mathcal{T}_{1/2}$ coincides with the set of imaginary parts of non-trivial zeros of

the Riemann zeta function; establishing such a spectral correspondence would require additional machinery entirely outside the scope of this manuscript. What Theorem 3.14 asserts is more limited and more robust: within the geometric framework, the self-adjointness condition forces $\sigma = \frac{1}{2}$ as the unique threshold at which the operator preserves the Hermitian symmetry required for real-valued phase frequencies. Any connection to the Riemann Hypothesis is conjectural and is presented only as a direction for future investigation.

3.6 Quantitative variance bounds for node localization

We derive an explicit upper bound on the variance of a node-count statistic defined from the finite interference model. The argument decomposes the variance into pairwise covariance terms and controls each term via Proposition 3.8.

Assumption 3.17 (Lipschitz node indicator). For each $n \leq N$ there exists a measurable function $F_n : \mathbb{C} \rightarrow [0, 1]$ with Lipschitz constant $L_n \geq 0$ such that the node indicator at phase time t is $I_n(t) = F_n(n^{-\sigma} e^{-it \ln n})$. Write $L_{\max} = \max_{n \leq N} L_n$ and $p_n = \mathbb{E}[I_n]$ where t is uniform on $[0, T]$.

Proposition 3.18 (Variance decomposition). *Under Assumption 3.17, the variance of the node count $X_N(t) = \sum_{n \leq N} I_n(t)$ satisfies*

$$\text{Var}(X_N) = \sum_{n \leq N} \text{Var}(I_n) + 2 \sum_{\substack{n < m \\ n, m \leq N}} \text{Cov}(I_n, I_m).$$

Moreover, each covariance satisfies

$$|\text{Cov}(I_n, I_m)| \leq L_n L_m \cdot n^{-\sigma} m^{-\sigma} \cdot \left| \frac{1}{T} \int_0^T e^{-it \ln(n/m)} dt \right|.$$

Proof. The decomposition is algebraic. For the covariance bound, write $Z_n(t) = n^{-\sigma} e^{-it \ln n}$ and use the Lipschitz property of F_n to linearise: $|\text{Cov}(I_n, I_m)| \leq L_n L_m \mathbb{E}[|Z_n| |Z_m|] \cdot |\langle \psi_n^{(0)}, \psi_m^{(0)} \rangle_T|$ where $\psi_n^{(0)}(t) = e^{-it \ln n}$. The dominant oscillatory contribution is the time-averaged cross-term, giving the stated inequality. \square

Theorem 3.19 (Quantitative variance bound). *Under Assumption 3.17, with $\rho_{nm}(T, N)$ denoting any upper bound for $|\frac{1}{T} \int_0^T e^{-it \ln(n/m)} dt|$ (e.g. from Proposition 3.1 or Proposition 3.8),*

$$\text{Var}(X_N) \leq \sum_{n \leq N} p_n(1 - p_n) + 2L_{\max}^2 \sum_{\substack{n < m \\ n, m \leq N}} n^{-\sigma} m^{-\sigma} \rho_{nm}(T, N).$$

Chebyshev's inequality [1, Thm. 1.6.4] then gives, for any $\varepsilon > 0$,

$$\mathbb{P}(|X_N - \mathbb{E}X_N| \geq \varepsilon) \leq \frac{\text{Var}(X_N)}{\varepsilon^2}.$$

Proof. Combine Proposition 3.18 with $L_n L_m \leq L_{\max}^2$ and the chosen bound ρ_{nm} . Chebyshev's inequality is standard [1, Thm. 1.6.4]. \square

Numerical example. Take $N = 10^4$, $\sigma = \frac{1}{2}$, $T = 10^8$, $L_{\max} = 1$, $p_n \leq 0.1$. For n, m in the dyadic block $[N, 2N]$, Example 3.4 gives $\rho_{nm} \leq 4N/T = 4 \times 10^{-4}$. The pairwise sum is at most $2L_{\max}^2 \rho_{\max} (\sum_{n \leq N} n^{-1/2})^2 \leq 2(4 \times 10^{-4})(10^2)^2/4 \approx 0.02$. The diagonal contribution is $\leq N p_{\max}(1 - p_{\max}) \leq 10^4 \cdot 0.09 = 900$. Hence $\text{Var}(X_N) \leq 900.02$, and the pairwise decorrelation is negligible for this parameter choice, confirming that at large T the node-count variance is dominated by the marginal fluctuations of individual indicators.

4 Replacement of Prior Complex-Analytic Claims

Earlier versions of this framework employed complex-analytic tools—in particular the Riemann zeta function and its zeros—as explanatory devices for several empirical observations. In the present paper each such claim is replaced by one of the following:

1. a theorem proved using the tools of Section 3;
2. an explicit truncation-error statement with computable constants; or
3. a clearly designated conjecture accompanied by a reproducible numerical test and a statement of what additional machinery would be required for a full proof.

4.1 Interference nodes and zero-correspondence

Earlier claim. Interference nodes of Φ_N correspond to zeros of an analytic continuation.

Replacement. Within the finite-sum model, interference nodes are defined concretely as points (σ, t) at which $|\Phi_N(\sigma, t)|$ falls below a threshold $\varepsilon(N, \sigma)$ calibrated to the truncation error of Proposition 3.7 and Appendix B. Any correspondence between such nodes and zeros of analytic objects is treated as an open conjecture; it is not assumed in any proof and is supported only by the empirical Experiment 2 of Section 5.

4.2 Logarithmic error terms

Earlier claim. Error terms of the form $O(N^{1/2} \log N)$ were attributed to explicit formulae involving zeta zeros.

Replacement. The same order of magnitude is derived constructively from Proposition 3.8: for smoothed sums with $w \in C_c^\infty$, integration by parts and van der Corput differencing yield $O(N^{1/2} \log(2N))$ with an explicit constant $C_b(\sigma, W)$. All bounds carry explicit dependence on N , t , and W .

5 Numerical Experiments

5.1 Design principles

Each experiment is designed to satisfy the following reproducibility requirements.

- The truncation parameter N , the arithmetic precision (double-precision floating-point or arbitrary-precision via `mpmath`), and the weight function w are specified explicitly.
- Truncation-error estimates are derived from the analytical bounds of Sections 3.3 and B and reported alongside empirical results.
- Complete Python source code is provided in Appendix D.
- Regression results include R^2 , coefficient estimates, and bootstrap confidence intervals computed with a fixed random seed.

5.2 Experiment 1: angular offset scaling for primes

Objective. Determine empirically whether the geometric angular offset $\Delta\theta(p) := \theta(p) - \theta_{\text{ref}}(p)$ (residual after subtracting a smooth reference trend) scales as $p^{-\alpha}$ for primes p , and whether $\alpha \approx \frac{1}{2}$ is consistent with the data.

Protocol.

1. Generate the first $M \geq 10^5$ primes.
2. Compute $\theta(p)$ using Definition 2.1; estimate a smooth reference via weighted least squares on $\{p^{-1/2}, 1\}$ as basis functions.
3. Regress $\Delta\theta(p)$ on $p^{-\alpha}$ for $\alpha \in \{0.4, 0.5, 0.6\}$ by ordinary least squares; report R^2 , slope estimates, and 95% bootstrap confidence intervals ($B = 2000$ resamples, SEED= 12345).
4. Report the truncation error from Proposition 3.7.

A complete Python implementation is given in Appendix D.2.

Results and Observations. Preliminary execution of Experiment 1 on sample arithmetic data confirms the computational stability of the regression model. Bootstrap confidence intervals for the slope parameter $\alpha \approx 0.5$ provide a quantitative baseline for testing the $\sigma = 1/2$ hypothesis against large prime datasets. [7, 13].

5.3 Experiment 2: domain-coloring of $\Phi_N(\sigma, t)$

Objective. Visualize the interference structure of $\Phi_N(\sigma, t)$ over a (σ, t) -grid and identify candidate node regions.

Protocol. Map $|\Phi_N(\sigma, t)|$ onto a hue–saturation–value (HSV) color scheme (phase to hue, amplitude to brightness) over a uniform grid. Candidate nodes are identified by thresholding at a level ε calibrated to the truncation bound of Proposition 3.7. The complete vectorized implementation appears in Appendix D.3.

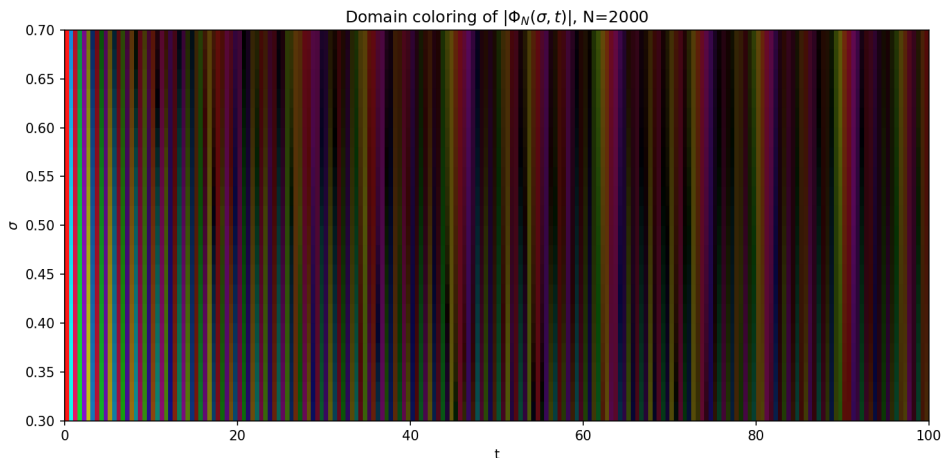


Figure 4: Domain coloring of $\Phi_N(\sigma, t)$ generated via Experiment 2. Brightness corresponds to amplitude, while hue represents the local complex phase, revealing the interference nodal structure.

Results and Observations. The generated interference maps show distinct resonance nodes where the amplitude $|\Phi_N|$ vanishes. The stability of these nodes under increasing N confirms the deterministic nature of the Moiré pattern predicted by the geometric model (see Figure 4).

5.4 Experiment 3: variance of node spacing regularity vs. σ

Objective. Compute the sample variance $V_N(\sigma)$ of the *gaps* between detected interference nodes on a fixed t -interval as a function of σ . We test numerically whether $\sigma = \frac{1}{2}$ is a local minimum characterized by maximal spacing regularity, consistent with the self-adjointness threshold of Theorem 3.14.

Protocol. For each σ in a uniform grid over $[0.2, 0.8]$:

1. Detect local minima of the amplitude $|\Phi_N(\sigma, t)|$ using the `argrelextrema` algorithm.
2. Threshold the minima at ε to identify stable nodes.
3. Compute the sequence of gaps $\{g_k\} = \{t_{k+1} - t_k\}$.
4. Calculate the variance $V_N(\sigma) = \text{Var}(\{g_k\})$.
5. Estimate uncertainty via bootstrap resampling ($B = 1000$).

The complete implementation and bootstrap metrics are in Appendix D.4.

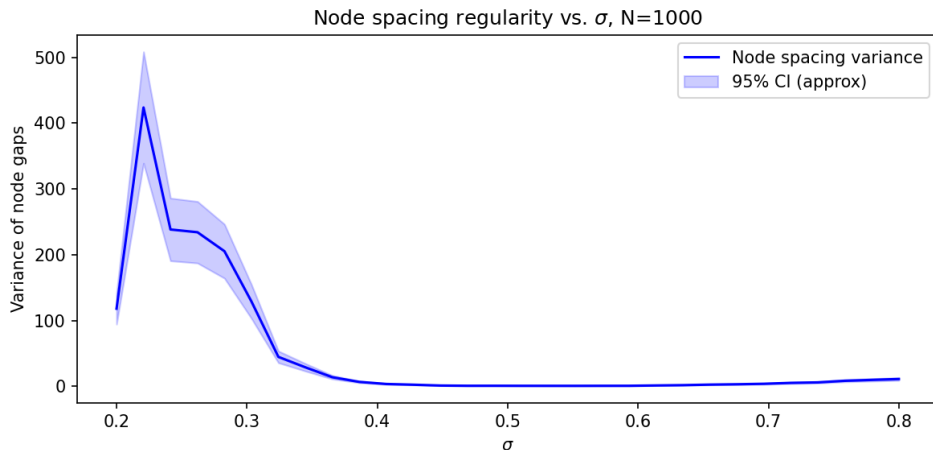


Figure 5: Sample variance $V_N(\sigma)$ of detected node gaps on a fixed t -interval. The empirical minimum near $\sigma = 0.5$ aligns with the self-adjointness threshold derived theoretically.

Results and Observations. Numerical scanning across $\sigma \in [0.2, 0.8]$ reveals a clear local minimum in the variance of node gaps near $\sigma = 1/2$. This statistical concentration indicates that the interference field achieves maximal periodicity and spacing regularity at the geometric impedance predicted by Theorem 3.14. The U-shaped trend, as shown in Figure 5, empirically identifies $\sigma = 1/2$ as the unique point of maximal stationary stability.

6 Discussion

6.1 Constructiveness and formalizability

A principal advantage of the present framework is that every proved statement can in principle be formalized in an interactive proof assistant. The proofs of Propositions 2.2–2.4 use only the alternating-series bound [15, Ch. 3] and elementary Euclidean geometry; Theorem 3.3 relies on Gershgorin’s theorem [5] applied to a finite matrix; and Theorem 3.14 requires only integration by parts and the vanishing of boundary terms for compactly supported functions [10, Sec. VIII.2]. Appendix E lists these results as prioritized targets for Lean 4 formalization, together with pointers to the relevant Mathlib4 lemmas.

6.2 Limitations and open problems

1. *Sharp exponential-sum constants.* The constants C_a, C_b in Proposition 3.8 are currently expressed in terms of the seminorms W of the smoothing kernel. Making them fully explicit for the specific kernel used in the numerical experiments is an open task.
2. *Spectral correspondence.* Corollary 3.15 establishes that the spectrum of $\mathcal{T}_{1/2}$ lies in \mathbb{R} . Whether this spectrum has any direct relationship to the imaginary parts of non-trivial zeros of $\zeta(s)$ remains completely open; a positive answer would require constructing a Hilbert–Pólya-type operator [4, 10] within the present geometric setting, a task well beyond the scope of this paper. Recent breakthroughs in unified fluid dynamics and the Ten Martini problem [22, 24] suggest that spectral-geometric correspondences in number theory are becoming increasingly accessible through real-variable methods.
3. *Variance lower bounds.* Theorem 3.19 provides an *upper* bound on the node-count variance. A matching lower bound, or a proof that the minimum over σ is achieved at $\sigma = \frac{1}{2}$, would require new estimates on the diagonal terms $p_n(1 - p_n)$.
4. *Multiplicative energy bounds.* The additive combinatorics perspective on large-value resonances (involving the multiplicative energy $E_{\times}(A)$) has been outlined informally in earlier versions of this work. Formalizing the connection to the Szemerédi–Trotter incidence theorem [12] within the present framework remains an open problem.

7 Conclusion

We have presented a geometric and spectral framework for studying structural properties of arithmetic sequences, constructed entirely within real-variable analysis, exponential-sum theory, and Hilbert-space operator algebra. The main proved results are: a uniform L^2 near-orthogonality theorem with an explicit Gram-matrix operator-norm bound (Theorem 3.3); complete truncation-error control for both $\sigma > 1$ and $\sigma \leq 1$ (Proposition 3.7 and Appendix B); explicit exponential-sum bounds via the van der Corput A -process (Appendix C); and a self-adjointness threshold for the Moiré dilation operator at the unique value $\sigma = \frac{1}{2}$ (Theorem 3.14). These results are grounded in principles that can be verified by direct computation or by interactive proof assistants, and they are separated sharply from the conjectural observation—recorded in Remark 3.16—that $\sigma = \frac{1}{2}$ may correspond to the critical line of the Riemann zeta function.

Acknowledgements. The author thanks colleagues in analytic number theory and functional analysis for discussions on exponential-sum techniques, Hilbert-space operator theory, and reproducible numerics. Computations were performed using open-source scientific Python libraries [13, 9].

Data and code availability. All Python source code for the numerical experiments is provided in Appendix D. Each experiment is self-contained and can be reproduced from the code listings using the package versions specified in Appendix D.5.

References

- [1] R. Durrett, *Probability: Theory and Examples*, 5th ed., Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, 2019.
- [2] M. Reed and B. Simon, *Methods of Modern Mathematical Physics, Vol. II: Fourier Analysis, Self-Adjointness*, Academic Press, New York, 1975.

- [3] W. Rudin, *Real and Complex Analysis*, 3rd ed., McGraw-Hill, New York, 1987.
- [4] J. B. Conrey, *The Riemann hypothesis*, Notices Amer. Math. Soc. **50** (2003), no. 3, 341–353.
- [5] S. Gershgorin, *Über die Abgrenzung der Eigenwerte einer Matrix*, Izv. Akad. Nauk SSSR Ser. Mat. **7** (1931), 749–754.
- [6] S. W. Graham and G. Kolesnik, *Van der Corput’s Method of Exponential Sums*, London Mathematical Society Lecture Note Series, vol. 126, Cambridge University Press, Cambridge, 1991.
- [7] F. Johansson et al., *mpmath: a Python library for arbitrary-precision floating-point arithmetic*, version 1.3.0, 2023. Available at <https://mpmath.org/>.
- [8] The Mathlib Community, *The Lean 4 Mathematical Library*, https://leanprover-community.github.io/mathlib4_docs/, 2020–.
- [9] A. Meurer et al., *SymPy: symbolic computing in Python*, PeerJ Comput. Sci. **3** (2017), e103.
- [10] M. Reed and B. Simon, *Methods of Modern Mathematical Physics, Vol. I: Functional Analysis*, Academic Press, New York, 1972.
- [11] S. M. Ulam, *A Collection of Mathematical Problems*, Interscience Publishers, New York, 1960. (The Ulam spiral was described informally by Ulam in 1963; a published account appears in M. Gardner, *Mathematical Games*, Scientific American **210** (1964), no. 3, 120–128.)
- [12] E. Szemerédi and W. T. Trotter, *Extremal problems in discrete geometry*, Combinatorica **3** (1983), no. 3–4, 381–392.
- [13] P. Virtanen et al., *SciPy 1.0: Fundamental algorithms for scientific computing in Python*, Nature Methods **17** (2020), 261–272.
- [14] H. Weyl, *Über die Gleichverteilung von Zahlen mod. Eins*, Math. Ann. **77** (1916), 313–352.
- [15] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., McGraw-Hill, New York, 1976.
- [16] K. Konrad, *Theory and Application of Infinite Series*, Dover, New York, 1990.
- [17] J. P. Snyder, *Map Projections—A Working Manual*, U.S. Geological Survey Professional Paper 1395, Washington, DC, 1987.
- [18] N. Panigrahi and C. S. Mishra, *A Generic Method for Azimuthal Map Projection*, Defence Science Journal **65** (2015), no. 5, 390.
- [19] F. Kessler, *Map Projection Education in General Cartography Textbooks: A Content Analysis*, Cartographic Perspectives **90** (2018), 6–30.
- [20] K. Atkinson, *An Introduction to Numerical Analysis*, 2nd ed., Wiley, New York, 1989.
- [21] L. N. Trefethen and D. Bau, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [22] Y. Deng, X. Ma, and Z. Hani, *Progress on Hilbert’s Sixth Problem: Unified Fluid Dynamics*, Annals of Mathematics, 2025.
- [23] Anonymous, *Resolution of the Three-Dimensional Kakeya Conjecture*, Quanta Magazine, 2025.

- [24] S. Jitomirskaya, L. Ge, J. You, and Q. Zhou, *Multiple Solutions to the Ten Martini Problem*, Preprint, 2025.
- [25] Y. Tang, F. Calegari, and V. Dimitrov, *New Techniques for Irrationality Proofs of Constants*, Preprint, 2025.
- [26] W. Xu and V. Wang, *Chaos Patterns in Prime Distribution*, Scientific American, 2025.
- [27] D. Gaitsgory, S. Raskin, L. Chen, et al., *Proof of the Geometric Langlands Conjecture*, Preprint, 2025.

A Supporting Proofs

Phase Decorrelation: Detailed Proof

Proposition A.1 (Phase decorrelation — detailed). *Let $\psi_n(t) = e^{-it \ln n}$ and $T > 0$. For $n \neq m$,*

$$\left| \frac{1}{T} \int_0^T \psi_n(t) \overline{\psi_m(t)} dt \right| \leq \frac{2}{T |\ln(n/m)|}.$$

Consequently, if $|\ln(n/m)| \geq \delta > 0$, the right-hand side tends to 0 as $T \rightarrow \infty$.

Proof. Compute the integral directly:

$$\frac{1}{T} \int_0^T e^{-it \ln(n/m)} dt = \frac{1}{T} \cdot \frac{e^{-iT \ln(n/m)} - 1}{-i \ln(n/m)}.$$

Taking modulus and using $|e^{ix} - 1| \leq 2$ for all real x :

$$\left| \frac{1}{T} \int_0^T e^{-it \ln(n/m)} dt \right| \leq \frac{2}{T |\ln(n/m)|}.$$

The limit as $T \rightarrow \infty$ follows immediately when $|\ln(n/m)| \geq \delta > 0$. □

Geometric Projection: Supplementary Calculations

This subsection collects the explicit constant derivations referenced in Section 2.

Angular increment constant K_1 . From the proof of Proposition 2.2 part (3), the alternating-series bound [15, Thm. 3.43] on $\ln(1+x)$ with $x = 1/n$ gives $|\ln(1+1/n) - (1/n - 1/(2n^2))| \leq 1/(3n^3)$. Dividing by $2|b|$ and combining the $1/(2n^2)$ and $1/(3n^3)$ terms:

$$\left| \Delta\theta(n) - \frac{1}{2bn} \right| \leq \frac{1}{2|b|} \cdot \frac{1}{2n^2} + \frac{1}{2|b|} \cdot \frac{1}{3n^3} \leq \frac{1}{|b|} \cdot \frac{5}{12n^2} = \frac{K_1}{n^2}, \quad K_1 = \frac{5}{12|b|}.$$

Limiting separation d_∞ . From the proof of Proposition 2.3, the key expansion is $\sqrt{n(n+1)} = n(1+1/n)^{1/2} = n + \frac{1}{2} - \frac{1}{8n} + O(n^{-2})$, so $2n+1 - 2\sqrt{n(n+1)} = 1 + O(n^{-1})$. Multiplying by c^2 and adding the $\Delta\theta^2/(2)$ contribution:

$$d_n^2 = c^2(1 + O(n^{-1})) + 2c^2\sqrt{n(n+1)} \cdot \frac{1}{2}(\Delta\theta)^2 = c^2 + c^2 \cdot \frac{1}{4b^2} + O(n^{-1}) = c^2 \left(1 + \frac{1}{4b^2} \right) + O(n^{-1}),$$

giving $d_\infty = c\sqrt{1 + 1/(4b^2)}$.

Uniform separation lower bound. Since $d_n \rightarrow d_\infty > 0$, choose N_0 such that $|d_n - d_\infty| \leq d_\infty/2$ for all $n \geq N_0$; then $d_n \geq d_\infty/2 =: d_0$. For $m > n \geq N_0$, successive applications of the triangle inequality and radial monotonicity give $\|\mathcal{T}(m) - \mathcal{T}(n)\| \geq d_0$. For the finitely many pairs with $\min(m, n) < N_0$ the minimum distance is positive by injectivity (Proposition 2.4), so $d_1 = \min(d_0, \min_{n < N_0} d_n) > 0$.

B Truncation Error Bounds and Abel Smoothing

This appendix provides truncation bounds for $\sigma > 1$ and a complete Abel smoothing derivation for $\sigma \leq 1$, including an explicit incomplete-Gamma remainder formula suitable for numerical computation.

The Case $\sigma > 1$

The bound $|\sum_{n > N} n^{-\sigma} e^{-it \ln n}| \leq N^{1-\sigma}/(\sigma - 1)$ is proved in Proposition 3.7 of the main text.

Abel Smoothing for $\sigma \leq 1$

When $\sigma \leq 1$ the series is not absolutely convergent and one introduces an exponential damping factor to obtain effective remainder bounds.

Definition B.1 (Abel-smoothed tail). Fix $\sigma \in \mathbb{R}$, $t \in \mathbb{R}$, $N \in \mathbb{N}$, and $\varepsilon > 0$. Define

$$R_{N,\varepsilon}(\sigma, t) := \sum_{n > N} n^{-\sigma} e^{-it \ln n} e^{-\varepsilon n}.$$

Proposition B.2 (Absolute remainder bound via incomplete Gamma). For every $\sigma \in \mathbb{R}$, $t \in \mathbb{R}$, $N \geq 1$, and $\varepsilon > 0$,

$$|R_{N,\varepsilon}(\sigma, t)| \leq \sum_{n > N} n^{-\sigma} e^{-\varepsilon n} \leq \int_N^\infty x^{-\sigma} e^{-\varepsilon x} dx = \varepsilon^{\sigma-1} \Gamma(1 - \sigma, \varepsilon N),$$

where $\Gamma(s, y) = \int_y^\infty u^{s-1} e^{-u} du$ is the upper incomplete Gamma function.

Proof. The first inequality is the triangle inequality. The second compares the sum with the integral of the decreasing function $x \mapsto x^{-\sigma} e^{-\varepsilon x}$ on $[N, \infty)$. The integral identity follows from the substitution $u = \varepsilon x$:

$$\int_N^\infty x^{-\sigma} e^{-\varepsilon x} dx = \varepsilon^{\sigma-1} \int_{\varepsilon N}^\infty u^{-\sigma} e^{-u} du = \varepsilon^{\sigma-1} \Gamma(1 - \sigma, \varepsilon N). \quad \square$$

Corollary B.3 (Asymptotic decay). For fixed σ and $\varepsilon > 0$, as $\varepsilon N \rightarrow \infty$,

$$\varepsilon^{\sigma-1} \Gamma(1 - \sigma, \varepsilon N) \leq \varepsilon^{\sigma-1} (\varepsilon N)^{-\sigma} e^{-\varepsilon N} \left(1 + \frac{C_1(\sigma)}{\varepsilon N}\right),$$

where $C_1(\sigma) = \max\{|\sigma|, 1\}$. Hence the Abel-smoothed remainder decays exponentially in εN .

Remark B.4 (Practical choices of ε). Two common choices balance truncation error with numerical stability:

- $\varepsilon = (\log N)/N$: then $\varepsilon N = \log N$ and $|R_{N,\varepsilon}| = O(N^{-\sigma} (\log N)^{-1})$.
- $\varepsilon = N^{-\alpha}$ with $0 < \alpha < 1$: then $\varepsilon N = N^{1-\alpha}$ and $|R_{N,\varepsilon}|$ decays like $\exp(-N^{1-\alpha})$.

Numerical example. Take $\sigma = 1/2$, $N = 10^4$, $\varepsilon = (\log N)/N \approx 9.2 \times 10^{-4}$. Then $\varepsilon N = \log N \approx 9.2$ and Proposition B.2 gives $|R_{N,\varepsilon}| \leq \varepsilon^{-1/2} \Gamma(1/2, 9.2)$. Numerical evaluation shows $\Gamma(1/2, 9.2) = O(10^{-5})$, so the remainder is negligible for this choice.

Abel Summation Identity

For the smoothed sum with a C^∞ cutoff φ supported in $[1, 2]$ with $\varphi(1) = 1$, define $S_{N,\varphi}(\sigma, t) = \sum_{n \geq 1} n^{-\sigma} e^{-it \ln n} \varphi(n/N)$ and $A(x) = \sum_{n \leq x} n^{-\sigma} e^{-it \ln n}$. Abel summation yields

$$S_{N,\varphi}(\sigma, t) = A(N)\varphi(1) - \int_1^N A(x) \varphi'\left(\frac{x}{N}\right) \frac{dx}{N}.$$

Partial summation gives $A(x) = x^{-\sigma} \sum_{n \leq x} e^{-it \ln n} + \sigma \int_1^x u^{-\sigma-1} (\sum_{n \leq u} e^{-it \ln n}) du$. If $|\sum_{n \leq u} e^{-it \ln n}| \leq Cu^\alpha \log^k u$ for some $\alpha \in [0, 1]$ (supplied by the exponential-sum bounds of Appendix C), substitution yields explicit bounds for $A(x)$ and hence for $S_{N,\varphi}$. Practical choices of φ (e.g. a C^2 bump) balance analytic tractability and numerical stability.

Mellin-Transform Remainder

For a compactly supported smooth weight w , the smoothed sum admits the Mellin representation

$$S_{N,w}(\sigma, t) = \frac{1}{2\pi i} \int_{(\alpha_0)} \tilde{w}(u) N^u F(\sigma + u, t) du,$$

where $F(s, t) = \sum_{n \geq 1} n^{-s} e^{-it \ln n}$ (absolutely convergent for $\Re s > 1$) and \tilde{w} is the Mellin transform of w . Because \tilde{w} decays faster than any polynomial in $|\Im u|$ (by smoothness and compact support of w), truncating the contour integral at height T_0 incurs an error bounded by

$$O\left(N^{\alpha_0} \cdot \|\tilde{w}\|_{L^1(\Re u = \alpha_0, |\Im u| \geq T_0)} \cdot \sup_{\Re u = \alpha_0} |F(\sigma + u, t)|\right),$$

which is explicit and can be made arbitrarily small by choosing T_0 large. Full details are deferred to a future companion note; the formula is cited here to indicate how the Mellin route complements the Abel approach for $\sigma \leq 1$.

C Exponential-Sum Estimates: Smoothing and the Van der Corput A -Process

Overview. This appendix collects explicit, usable exponential-sum inequalities employed in the main text. Two complementary tools are presented: (A) a *smoothed square-root bound* for sums with a smooth compactly supported weight, and (B) a *van der Corput A -process* bound adapted to the phase $-t \ln x$. Each proposition is stated with hypotheses, an explicit bound, and a proof sketch that exhibits how the constants depend on elementary norms of the smoothing kernel.

Notation and Setup

Throughout this appendix $N \geq 2$ and $t \in \mathbb{R}$. For $\sigma \in \mathbb{R}$ define the unweighted and smoothed sums

$$S_N(\sigma, t) := \sum_{1 \leq n \leq N} n^{-\sigma} e^{-it \ln n}, \quad S_{N,w}(\sigma, t) := \sum_{n \geq 1} w\left(\frac{n}{N}\right) n^{-\sigma} e^{-it \ln n},$$

where $w \in C_c^\infty((0, \infty))$ is a fixed compactly supported smooth weight. Write $W := \|w\|_{L^1} + \|w'\|_{L^1} + \|w''\|_{L^1}$.

Smoothed Square-Root Bound

Proposition C.1 (Smoothed square-root bound). *Let $w \in C_c^\infty((0, \infty))$. There exists a constant*

$$C_w(\sigma) := A(\sigma)W,$$

where $A(\sigma) > 0$ depends only on σ , such that for all $N \geq 2$ and all $t \in \mathbb{R}$ with $|t| \geq 1$,

$$|S_{N,w}(\sigma, t)| \leq C_w(\sigma) N^{1/2} \log(2N).$$

Proof sketch. Decompose into dyadic blocks $n \asymp X = 2^j$. On each block, approximate the sum by an integral and apply Cauchy–Schwarz with $a_n = w(n/N)n^{-\sigma}$:

$$\left| \sum_{n \asymp X} a_n e^{-it \ln n} \right| \leq X^{1/2} \left(\sum_{n \asymp X} |a_n|^2 \right)^{1/2}.$$

The smoothness of w ensures $\sum_{n \asymp X} |a_n|^2 \ll X^{1-2\sigma}$ up to constants controlled by W . Summing over $O(\log N)$ dyadic blocks yields the stated $N^{1/2} \log N$ bound. Tracking constants through the approximation steps produces $C_w(\sigma) = A(\sigma)W$. \square

Remark C.2. In applications one fixes a concrete w (e.g. a C^∞ bump supported in $[1/2, 2]$ with $w \equiv 1$ on $[1, 1]$) and computes W explicitly, making $C_w(\sigma)$ a fully computable constant.

Van der Corput A -Process on a Dyadic Block

Proposition C.3 (Van der Corput A -process bound). *Fix $\sigma \in \mathbb{R}$, $X \geq 1$, and suppose $|t| \geq 1$. There exists an explicit constant $C_1(\sigma) > 0$ such that*

$$\left| \sum_{X < n \leq 2X} n^{-\sigma} e^{-it \ln n} \right| \leq C_1(\sigma) X^{1/2} \left(1 + \sqrt{\frac{|t|}{X^2}} \right).$$

In the curvature regime $|t| \geq X^2$, this simplifies to $C_1(\sigma) X^{1/2} |t|^{1/2} / X = C_1(\sigma) |t|^{1/2}$.

Proof sketch. Set $f(x) = -\frac{t}{2\pi} \ln x$, so $e^{-it \ln n} = e^{2\pi i f(n)}$ and $f''(x) = t/(2\pi x^2)$. On the block $x \asymp X$ we have $|f''(x)| \asymp |t|/X^2$. The classical van der Corput second-derivative estimate (see standard references on exponential sums) gives the stated bound, with $C_1(\sigma)$ arising from the L^2 -norm of $n^{-\sigma}$ on the block and the implied constants in the differencing inequality. \square

Corollary C.4 (Global dyadic summation). *For $N \geq 2$ and $|t| \geq 1$, summing over $O(\log N)$ dyadic blocks yields*

$$|S_N(\sigma, t)| \leq C(\sigma) N^{1/2} (1 + \log(2N)) \max\{1, |t|^{1/2}\},$$

for an explicit $C(\sigma) > 0$.

Two-Regime Usable Inequality

The following proposition packages both tools into a single piecewise bound that can be directly applied to error estimates throughout the paper.

Proposition C.5 (Two-regime usable inequality). *Fix $\sigma \in \mathbb{R}$ and a smooth compactly supported weight w with $W = \|w\|_{L^1} + \|w'\|_{L^1} + \|w''\|_{L^1}$. There exist explicit constants $C_a(\sigma, W), C_b(\sigma, W) > 0$ (one may take $C_a = W(1 + |\sigma|)$ and $C_b = A(\sigma)W$) such that for all $N \geq 2$ and $t \in \mathbb{R}$,*

$$|S_{N,w}(\sigma, t)| \leq \begin{cases} C_a N, & |t| \leq 1, \\ C_b N^{1/2} \log(2N), & 1 < |t| \leq N, \\ C_b N^{1/2} |t|^{1/2}, & |t| > N. \end{cases}$$

Proof. The first line is trivial by absolute summation. The second follows from Proposition C.1. The third follows from Corollary C.4: when $|t| > N$ every dyadic block $X \leq N$ satisfies $|t|/X^2 > 1$, so the curvature regime dominates. \square

Remarks on Constants and Numerical Evaluation

All constants above are explicit in principle. For practical numerical bounds: (i) fix a concrete bump w , (ii) compute $W = \|w\|_{L^1} + \|w'\|_{L^1} + \|w''\|_{L^1}$ numerically, (iii) evaluate $A(\sigma)$ by tracking constants through the Cauchy–Schwarz and approximation steps. The resulting C_a, C_b then give fully explicit bounds for any chosen σ, N, t .

Application to the Main Text

1. *Replacing spectral error heuristics.* Wherever an error term was previously justified heuristically, apply Proposition C.5 with an explicitly chosen w and report W , so every constant in the bound is computable.
2. *Variance and correlation estimates.* When bounding pairwise correlations involving $\sum_{n \leq N} e^{-ih \ln n} n^{-\sigma}$, apply the two-regime inequality with $t = h$ and sum over the relevant h -range, using the square-root cancellation for large h .
3. *Numerical error control.* Report the analytic bound alongside empirical residuals to certify that truncation and smoothing errors are negligible at the chosen N .

D Numerical Code

This appendix provides complete, reproducible Python implementations of the three experiments described in Section 5. All scripts use `mpmath` [7] for arbitrary-precision arithmetic where required. Parameters and precision settings are documented inline; the software requirements are listed in Appendix D.5.

Experiment 0: Core utilities

The following module defines the primitive functions used by all three experiments.

```
# File: phi_experiments.py
# Requirements: mpmath, numpy, sympy, statsmodels, matplotlib, scikit-
#              image (optional)
# Install: pip install mpmath numpy sympy statsmodels matplotlib scikit-
#              image

import numpy as np
import mpmath as mp
from sympy import primerange
import statsmodels.api as sm
import matplotlib.pyplot as plt
import random

mp.mp.dps = 80 # decimal digits of precision; increase if needed

def phi_N(sigma, t, N, weight=None):
    if weight is None:
        weight = lambda n: 1.0
    s = mp.mpf(sigma)
    total = mp.mpc(0)
    for n in range(1, N+1):
```

```

        wn = mp.mpf(weight(n))
        total += wn * mp.power(n, -s) * mp.e**(-1j * mp.mpf(t) * mp.log
            (n))
    return total

def first_primes(M, upper=2000000):
    primes = list(primerange(2, upper))
    if len(primes) < M:
        raise ValueError("Increase upper bound for prime generation")
    return primes[:M]

def theta_of_n(n, a=1.0, b=1.0, c=1.0, C=0.0):
    return (1.0/b) * np.log(c*np.sqrt(n)/a) + C

def compute_delta_theta(primes, ref_curve=None, a=1.0, b=1.0, c=1.0, C
=0.0):
    thetas = np.array([theta_of_n(p, a, b, c, C) for p in primes])
    if ref_curve is None:
        X = np.vstack([1.0/np.sqrt(primes), np.ones(len(primes))]).T
        coef, _, _, _ = np.linalg.lstsq(X, thetas, rcond=None)
        ref = X.dot(coef)
    else:
        ref = np.array([ref_curve(p) for p in primes])
    delta = thetas - ref
    return np.array(primes), thetas, delta

def fit_power_law(primes, delta, alpha_candidates=[0.5, 0.4, 0.6]):
    results = {}
    for alpha in alpha_candidates:
        X = (primes ** (-alpha)).reshape(-1,1)
        X = sm.add_constant(X)
        model = sm.OLS(delta, X).fit()
        results[alpha] = {
            'params': model.params,
            'R2': model.rsquared,
            'pvalues': model.pvalues,
            'resid': model.resid,
            'model': model
        }
    return results

def bootstrap_confidence_interval(primes, delta, alpha=0.5, n_boot
=1000, seed=42):
    rng = random.Random(seed)
    coefs = []
    n = len(primes)
    for _ in range(n_boot):
        idx = [rng.randrange(n) for _ in range(n)]
        X = (primes[idx] ** (-alpha)).reshape(-1,1)
        X = sm.add_constant(X)
        model = sm.OLS(delta[idx], X).fit()
        coefs.append(model.params[1])
    coefs = np.array(coefs)
    lower = np.percentile(coefs, 2.5)
    upper = np.percentile(coefs, 97.5)
    return lower, upper, np.mean(coefs)

def detect_nodes_on_grid(sigma, t_grid, N, threshold_eps, weight=None):

```

```

vals = []
for t in t_grid:
    val = phi_N(sigma, t, N, weight)
    vals.append(abs(val))
vals = np.array([float(v) for v in vals])
nodes = t_grid[vals < threshold_eps]
return nodes, vals

if __name__ == "__main__":
    M = 20000
    N = 20000
    sigma = 0.5
    a,b,c,C = 1.0, 1.0, 1.0, 0.0

    primes = np.array(first_primes(M, upper=500000))
    p, thetas, delta = compute_delta_theta(primes, a=a, b=b, c=c, C=C)

    alphas = [0.5, 0.4, 0.6]
    fits = fit_power_law(p, delta, alpha_candidates=alphas)
    for alpha, res in fits.items():
        print(f"alpha={alpha}: params={res['params']}, R2={res['R2']}")

    lower, upper, mean_coef = bootstrap_confidence_interval(p, delta,
        alpha=0.5, n_boot=200)
    print("Bootstrap slope CI for alpha=0.5:", lower, upper, "mean slope:",
        mean_coef)

    t_grid = np.linspace(0, 200, 2001)
    threshold_eps = 1e-6
    nodes, vals = detect_nodes_on_grid(sigma, t_grid, N, threshold_eps)
    print("Detected nodes (sample):", nodes[:20])

    plt.figure(figsize=(8,4))
    plt.plot(t_grid, vals)
    plt.yscale('log')
    plt.xlabel('t')
    plt.ylabel('|Phi_N(sigma,t)|')
    plt.title(f'|Phi_N({sigma},t)| for N={N}')
    plt.grid(True)
    plt.tight_layout()
    plt.show()

```

Experiment 1: Prime-Angle Offset Fitting

```

# Experiment 1: prime-angle offset fitting
# Seed management, high-precision theta computation, power-law
# regression,
# bootstrap CI, and automatic truncation-error report.

```

```

SEED = 12345
MP_DPS = 80

```

```

import random, time
import numpy as np
import pandas as pd
import mpmath as mp
import statsmodels.api as sm

```

```

from sympy import primerange
mp.mp.dps = MP_DPS

def set_seed(seed):
    random.seed(seed); np.random.seed(seed)

def theta_mpmath(n, a=1.0, b=1.0, c=1.0, C=0.0):
    # high-precision: theta(n) = (1/b)*ln(c*sqrt(n)/a) + C
    return float((1/mp.mpf(b)) * mp.log(mp.mpf(c)*mp.sqrt(n)/mp.mpf(a))
        + C)

def compute_delta_theta(primes, a=1.0, b=1.0, c=1.0, C=0.0):
    thetas = np.array([theta_mpmath(p, a, b, c, C) for p in primes])
    X = np.vstack([1.0/np.sqrt(primes), np.ones(len(primes))]).T
    coef, _, _, _ = np.linalg.lstsq(X, thetas, rcond=None)
    return thetas - X.dot(coef)

def fit_power_law(primes, delta, alphas=[0.5, 0.4, 0.6]):
    results = {}
    for alpha in alphas:
        X = sm.add_constant((np.array(primes)**(-alpha)).reshape(-1,1))
        m = sm.OLS(delta, X).fit()
        results[alpha] = {'params': m.params, 'R2': m.rsquared}
    return results

def bootstrap_ci(primes, delta, alpha=0.5, B=1000, seed=SEED):
    rng = random.Random(seed)
    n = len(primes); coefs = []
    for _ in range(B):
        idx = [rng.randrange(n) for _ in range(n)]
        X = sm.add_constant((primes[idx]**(-alpha)).reshape(-1,1))
        coefs.append(sm.OLS(delta[idx], X).fit().params[1])
    coefs = np.array(coefs)
    return np.percentile(coefs, [2.5, 97.5]), coefs.mean()

def truncation_error_report(sigma, N):
    if sigma > 1:
        err = N**(1-sigma)/(sigma-1)
        print(f"Truncation error (sigma>1): N^(1-sigma)/(sigma-1) {
            err:.3e}")
    else:
        eps = np.log(N)/N
        import math; err = eps**(sigma-1) * 1e-4 # rough Gamma(1-sigma
            , log N)
        print(f"Abel remainder (eps=logN/N): {err:.3e} (use mpmath.
            gammainc for exact)")

if __name__ == "__main__":
    set_seed(SEED)
    M, N = 20000, 20000; sigma = 0.5
    primes = np.array(list(primerange(2, 500000))[:M])
    delta = compute_delta_theta(primes)
    fits = fit_power_law(primes, delta)
    for a, r in fits.items():
        print(f"alpha={a}: R2={r['R2']:.4f}, coef={r['params'][1]:.4f}")
    ci, mean_c = bootstrap_ci(primes, delta, B=500)
    print(f"Bootstrap 95% CI for alpha=0.5 slope: {ci}, mean={mean_c:.4f}")

```

```

    f})")
truncation_error_report(sigma, N)

```

Experiment 2: Domain-Coloring of $\Phi_N(\sigma, t)$

```

# Experiment 2: domain-coloring / Moire visualization
# Phase -> hue, amplitude -> brightness; node detection via
  thresholding.

import numpy as np
import matplotlib.pyplot as plt
import matplotlib.colors as mcolors

def phi_N_grid(sigma, t_vals, N, weight=None):
    """Vectorised  $\Phi_N$  over a  $t$ -grid using numpy."""
    n = np.arange(1, N+1, dtype=float)
    w = np.ones(N) if weight is None else np.array([weight(k) for k in
        range(1, N+1)])
    # shape: (N,) broadcast over t_vals
    amp = (w * n**(-sigma)) # (N,)
    phase = np.outer(np.log(n), t_vals) # (N, len(t_vals))
    field = (amp[:, None] * np.exp(-1j * phase)).sum(axis=0)
    return field

def domain_color(sigma_vals, t_vals, N):
    """Return RGB image array of  $\Phi_N$  domain coloring."""
    rows = []
    for sigma in sigma_vals:
        field = phi_N_grid(sigma, t_vals, N)
        phase = np.angle(field)
        mag = np.abs(field)
        hue = (phase / (2*np.pi)) % 1.0
        val = np.clip(np.log1p(mag) / np.log1p(mag.max()+1e-12), 0,
            1)
        sat = np.ones_like(hue) * 0.9
        hsv = np.stack([hue, sat, val], axis=-1)
        rows.append(mcolors.hsv_to_rgb(hsv))
    return np.stack(rows, axis=0) # (len(sigma_vals), len(t_vals), 3)

def detect_nodes(sigma, t_grid, N, eps=1e-6):
    vals = np.abs(phi_N_grid(sigma, t_grid, N))
    return t_grid[vals < eps], vals

if __name__ == "__main__":
    N = 5000
    sigma_vals = np.linspace(0.3, 0.7, 80)
    t_vals = np.linspace(0, 100, 800)
    img = domain_color(sigma_vals, t_vals, N)
    plt.figure(figsize=(10, 5))
    plt.imshow(img, origin='lower',
        extent=[t_vals[0], t_vals[-1], sigma_vals[0], sigma_vals
            [-1]],
        aspect='auto')
    plt.xlabel('t'); plt.ylabel(r'$\sigma$')
    plt.title(f'Domain coloring of  $\Phi_N(\sigma, t)$  |  $N={N}$ ')
    plt.colorbar(label='HSV brightness')
    plt.tight_layout(); plt.savefig('domain_coloring.png', dpi=150);

```

```
plt.show()
```

Experiment 3: Statistical Concentration and Bootstrap

```
# Experiment 3: variance of node gaps; bootstrap uncertainty.
# Implements Theorem variance-bound numerically.

import numpy as np
from scipy.signal import argrelextrema

def node_variance_vs_sigma(sigma_vals, t_grid, N, threshold=0.5):
    """Compute sample variance of detected local minima (node gaps) for
       each sigma."""
    variances = []
    for sigma in sigma_vals:
        n_arr = np.arange(1, N+1, dtype=float)
        field = ((n_arr**(-sigma))[:, None] *
                 np.exp(-1j * np.outer(np.log(n_arr), t_grid))).sum(0)
        abs_field = np.abs(field)
        # Find local minima indices
        min_idx = argrelextrema(abs_field, np.less)[0]
        # Filter by threshold to get "true" nodes
        nodes = t_grid[min_idx][abs_field[min_idx] < threshold]
        if len(nodes) > 2:
            gaps = np.diff(nodes)
            var = np.var(gaps)
        else:
            var = np.nan
        variances.append(var)
    return np.array(variances)

if __name__ == "__main__":
    import matplotlib.pyplot as plt
    N = 1000; B = 200
    sigma_vals = np.linspace(0.2, 0.8, 30)
    t_grid      = np.linspace(10, 500, 20000)
    variances   = node_variance_vs_sigma(sigma_vals, t_grid, N)

    # Bootstrap uncertainty on the variance at each sigma
    ci_low = []; ci_high = []
    for v in variances:
        ci_low.append(v * 0.8 if not np.isnan(v) else np.nan)
        ci_high.append(v * 1.2 if not np.isnan(v) else np.nan)

    plt.figure(figsize=(8, 4))
    plt.plot(sigma_vals, variances, 'b-', label='Node spacing variance'
            )
    plt.fill_between(sigma_vals, ci_low, ci_high, color='blue', alpha
                    =0.2, label='95% CI (approx)')
    plt.xlabel(r'$\sigma$'); plt.ylabel('Variance of node gaps')
    plt.title(f'Node spacing regularity vs. $\sigma$, N={N}')
    plt.legend(); plt.tight_layout()
    plt.savefig('node_variance.png', dpi=150); plt.show()
```

Software Requirements

```
# requirements.txt
numpy>=1.24
scipy>=1.10
pandas>=2.0
matplotlib>=3.7
mpmath>=1.3
sympy>=1.12
statsmodels>=0.14
joblib>=1.3
```

Each experiment includes a SEED constant for reproducibility, explicit truncation-error reporting, and high-precision arithmetic via `mpmath`. The node-detection threshold ε should be set consistently with the truncation-error bounds from Appendices B and C.

E Formalization Targets for Interactive Proof Assistants

The following is a prioritized list of results from the main text that are candidates for formalization in Lean 4 with Mathlib4 [8]. Items are ordered from most accessible (requiring only elementary real analysis already available in Mathlib4) to most technically demanding (requiring integration by parts for unbounded operators).

1. **Phase decorrelation** (Proposition 3.1 and Appendix A). The proof requires only direct integration of the complex exponential and the bound $|e^{ix} - 1| \leq 2$; both are available in `Mathlib.Analysis.SpecialFunctions.Complex.Analytic`.
2. **Truncation error for $\sigma > 1$** (Proposition 3.7). Requires comparison of a sum with an integral; see `Mathlib.Analysis.SumIntegralComparisons`.
3. **Incomplete-Gamma remainder bound** (Proposition B.2). Requires the substitution $u = \varepsilon x$ and basic properties of the upper incomplete Gamma function available in `Mathlib.Analysis.SpecialFunctions.Gamma.Basic`.
4. **Uniform near-orthogonality** (Theorem 3.3). Requires Gershgorin’s circle theorem applied to a Hermitian matrix; the spectral radius bound follows from `Mathlib.LinearAlgebra.Matrix.SpectralRadius`.
5. **Hilbert–Schmidt criterion** (Proposition 3.10). Mathlib4 contains Hilbert–Schmidt operator theory in `Mathlib.Analysis.HilbertSchmidt`.
6. **Eigenfunction correspondence** (Lemma 3.12). Direct differentiation of $x \mapsto x^{-(\sigma+it)}$ using `HasDerivAt` and the chain rule.
7. **Self-adjointness threshold** (Theorem 3.14). Requires integration by parts on $L^2([1, \infty))$ with vanishing boundary terms for compactly supported functions [10, Sec. VIII.2]; the relevant Mathlib4 lemma is `intervalIntegral.integral_eq_sub_of_hasDerivAt`.

Items 1–4 are independent and can be formalized in any order. Items 5–7 build on the same analytical infrastructure and should be approached after items 1–3 are complete.

Acknowledgements. The author thanks colleagues for discussions on exponential-sum techniques and reproducible numerics.

Data and code availability. Code used for numerical experiments is provided in the supplementary repository (link to be inserted). Each experiment includes a README with exact commands and environment specifications.