V1.0.0

When Does Spectral Data Constructively Determine Arithmetic Sequences?

Rowan Brad Quni-Gudzinas

QNFO-ULA

ABSTRACT

Cryptographic key generation depends on primality testing, and every widely deployed test is probabilistic, leaving a small but nonzero error probability in each certified RSA prime. A deterministic alternative driven by global spectral data—the non-trivial zeros of the Riemann zeta function—has been conjectured but never ruled out or established. The specific gap is that the explicit formula shows the zeros encode the prime-counting function $\pi(x)$ distributionally, yet no quantitative bound exists on how much spectral data a pointwise, integer-exact reconstruction requires. We formalize constructive emergence as algorithmic recovery of sequence terms from truncated spectral data, and derive the height T of zeros needed to evaluate $\pi(x)$ with error at most $\frac{1}{2}$ from the truncated explicit formula under a termwise error model: $T \geq 2\,x^{1/2}\log^2 x$. For $x = 10^6$ this gives $T_{\min} \approx 3.82 \times 10^5$, i.e. $N(T_{\min}) \approx 6.69 \times 10^5$ zeros—exponential in $\log x$, hence super-polynomial in the input size $\log x$. A worked Bernoulli power-sum identity and the spectral-gap constants $\lambda_1^{\text{prime}} \geq 0.17157288$ versus $\lambda_1^{\text{comp}} \leq 0.16666667$ (gap $\delta \approx 0.00490621$) illustrate the framework. By contrast, Pisot-type recurrences such as the Fibonacci sequence admit exact reconstruction from finitely many spectral parameters. The result rules out, at current error-model strength, a polynomial-size spectral replacement for probabilistic primality testing.

PUBLISHER CLASSIFICATION High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

PEER-REVIEW STATUS Not independently peer-reviewed

CITE AS Quni-Gudzinas, R. B. (n.d.). When Does Spectral Data Constructively Determine Arithmetic Sequences?. QNFO Preprint v1.0.0.

---

1. Introduction RSA key generation consumes large prime numbers, and each candidate prime is certified by a probabilistic test such as Miller–Rabin that carries a small but nonzero error probability per round; certification pipelines therefore stack rounds, add latency, and still do not deliver a deterministic guarantee. If a deterministic algorithm could reconstruct individual primes—or simply evaluate $\pi(x)$ exactly—from global spectral data such as the first T non-trivial zeros of the Riemann zeta function, key generation would become fully certifiable and the residual-error argument for stacked probabilistic rounds would disappear. The obstacle is that the explicit formula links the zeros to $\pi(x)$ only through an infinite, conditionally convergent sum: distributional information is available, but a pointwise algorithm with a quantified data budget is not. The decision that changes if this paper's result holds is whether to invest in spectral-data-based deterministic primality testing at all: a super-polynomial data requirement in the input size $\log x$ says the investment cannot pay off at cryptographic sizes under the stated error model.

We answer the question by defining a reconstruction-complexity hierarchy and deriving an explicit height bound. Our contributions are:

1. **Definition 1 (Constructive emergence).** A spectral data set $\mathcal{Z}_T = \{\rho_k : |\Im\rho_k| \leq T\}$ constructively determines an arithmetic sequence if an algorithm outputs any prescribed term with error at most 2^{-p} in time polynomial in p and $\log n$ (Section 3).
2. **Theorem 2 (Height bound).** Under a termwise truncation-error model for the explicit formula, recovering $\pi(x)$ with error at most $\frac{1}{2}$ requires $T \geq 2\,x^{1/2}\log^2 x$; for $x = 10^6$, $T_{\min} \approx 3.82 \times 10^5$ and the corresponding zero count is $N(T_{\min}) \approx 6.69 \times 10^5$ (Section 4).

3. **Worked verifications.** An exact Bernoulli power-sum evaluation $S_2^{1,0}(5) = 30$ using the identity of [7], and a spectral-gap numerical example from [9] with $\lambda_1^{\text{prime}} \geq 0.17157288$, $\lambda_1^{\text{comp}} \leq 0.16666667$, gap $\delta \approx 0.00490621$ (Section 4).
4. **Dichotomy.** Integer recurrences with Pisot characteristic roots, exemplified by the Fibonacci sequence with $|F_n - \phi^n/\sqrt{5}| < \frac{1}{2}$ [12], admit exact reconstruction from finitely many spectral parameters, in sharp contrast to the zeta-zero case (Section 4).

The formalism also bears on problems in adjacent fields where global spectral invariants are used to infer discrete structure, such as the isospectral manifolds constructed by Sunada's method in arithmetic spectral geometry [6], and on the emergent-number-theory question of which integer sequences arise from continuous flows [12].

2. Background and Related Work We assemble eight-plus strands of prior work into the framework; each supplies a concrete method, number, or contrast that the analysis uses.

Byzantine-resilient distributed optimization studies how a master node recovers a global model despite corrupt workers. In the setting of [1] the master aggregates $m = 15$ worker updates, each computed from $n = 10{,}000$ data points of dimension $d = 250$, and the analysis of corrupt-worker identification shows that a single linear combination of error vectors reveals the support of up to t corrupt nodes with probability one. We use this support-recovery mechanism as the analogy for isolating the contribution of individual spectral terms, while noting that the number of "workers" (zeros) required grows super-polynomially here.

Robust stochastic gradient descent in high dimensions [2] proves a robust mean-estimation guarantee whose error term $\Gamma_1 = O\big(\sigma^2/\big((1-(\epsilon+\epsilon'))bR\big)\big)$, under an adversarial fraction $\epsilon \leq \frac{1}{4}$, decays with batch size b and worker count R. This explicit dependence on the number of information sources calibrates our scaling analysis of how reconstruction error decays with the number of zeros used.

Constraints on dark energy from H II starburst galaxy apparent magnitude versus redshift data [3] fit the Hubble-parameter model $H(z) = H_0\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}$ to observations by chi-square minimization. The case illustrates that a modest number of observations can tightly constrain a few continuous parameters while remaining blind to discrete high-frequency structure such as individual prime locations—the precise regime distinction our hierarchy formalizes.

TerraGen [4] is a unified multi-task layout-generation framework for remote-sensing data augmentation, evaluated on a benchmark of 1,000 images with specialized spatial-accuracy metrics, using a two-stage pipeline in which a diffusion model is trained on a large corpus and then fine-tuned with an adaptive loss. Its staged coarse-to-fine structure motivates our two-phase reconstruction view: coarse distributional statistics from low-height zeros, refinement from higher-height zeros.

AeroGen [5] enhances remote-sensing object detection with diffusion-driven data generation and reports that adding $10{,}000$ synthetic samples raises mean average precision from 54.22 to 55.62 on the DIOR dataset. This quantified, roughly linear gain from additional synthetic information parallels the improvement of explicit-formula accuracy as more zeros are included—until the truncation-tail regime, where the gain per zero collapses.

Arithmetic lattices and weak spectral geometry [6] constructs arithmetic quotients of hyperbolic spaces and applies Sunada's method to produce isospectral, non-isometric manifolds. The existence of distinct manifolds with identical spectra is the geometric proof-of-concept that spectral data need not determine the object behind them; we treat it as the canonical non-uniqueness obstruction that any constructive-emergence claim must confront.

A note on the coefficients of power sums of arithmetic progressions [7] derives the closed form

$$S_k^{m,r}(n) = \frac{m^k}{k+1}\sum_{t=1}^{k+1}\binom{k+1}{t}B_{k+1-t}\Big(\frac{r}{m}\Big)\,n^t,$$

with $B_j(x)$ Bernoulli polynomials. This identity is our worked example of a spectral-type datum (a short list of Bernoulli numbers) that constructively determines every term of a power sum exactly; Section 4 verifies it numerically.

A 20 Gbps data-transmitting ASIC with PAM4 for particle physics experiments [8] sustains a line rate of 20 Gbps using two 10.24 Gbps serializers. The engineering point is that when the underlying physical model is linear and deterministic, a target parameter such as bit rate is exactly recoverable from measured spectral characteristics—a constructive regime against which the zeta-zero case is measured.

The spectral-gap primality proposal [9] defines a map $N \mapsto \gamma_N \in \mathrm{Aut}(\mathcal{T}_p)$ on the $(p+1)$-regular Bruhat–Tits tree, conjectures the prime lower bound $\lambda_1 \geq (p+1) - 2\sqrt{p}$, and reports a composite upper bound $\lambda_1 \leq C/D$ with $C = 1$, $D = 6$, giving for $p = 2$ the separating gap $\delta \approx 0.0049$ at a claimed runtime $O\big((\log N)^6\big)$ bit-operations. We use its constants as a concrete numerical illustration in Section 4.

Dual-mode arithmetic representation [10] argues that static two-dimensional geometric encodings cannot capture the infinite-dimensional structure of prime factorization, which motivates dynamic, spectral representations of arithmetic objects—the regime in which our height bound operates.

L-function zeros as gauge-invariant spectra [11] proves that the adelic-conditioned Laplacian on $\mathcal{T}_p$ has empty ℓ^2 point spectrum and continuous spectrum $[(p+1) - 2\sqrt{p},\, (p+1) + 2\sqrt{p}]$. This analytic foundation underwrites the spectral-gap constants of [9] and supplies the falsification criterion we adopt in Section 6.

Emergent number theory [12] establishes that integer sequences can emerge from continuous Pisot flows with exponential rounding-error bounds, exemplified by the Fibonacci approximation $|F_n - \phi^n/\sqrt{5}| < \frac{1}{2}$. It provides the constructive pole of our dichotomy: finite spectral parameters with exact integer recovery.

3. Methods The method is analytical: we state a data model, an error model, a rounding criterion, and two definitions, then derive bounds. No experiment, measurement, or data collection is performed.

Spectral data model. Fix a height $T > 0$ and collect the non-trivial zeros of the Riemann zeta function

$$\mathcal{Z}_T = \{\rho_k = \tfrac{1}{2} + i\gamma_k \;:\; |\gamma_k| \leq T\}.$$

The count of zeros up to height T follows the classical asymptotic

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

which we use as a deterministic input derived from the stated formula.

Target and output. The target is the prime-counting function $\pi(x)$ and the ordered prime list $\{p_n\}$. Pointwise reconstruction means producing an integer $\hat{p}_n$ with $|\hat{p}_n - p_n| \leq 1$, or $\pi(x)$ rounded to the correct integer.

Explicit formula. We employ the truncated explicit formula

$$\pi(x) = \mathrm{Li}(x) - \sum_{|\gamma_j| \leq T} \mathrm{Li}\big(x^{\rho_j}\big) + R_T(x),$$

where $\mathrm{Li}(x) = \int_2^x dt/\log t$ and $R_T(x)$ is the truncation error from omitted zeros.

Error model — stated convention. Two error models for the omitted tail are available in the analytic literature: a termwise bound, in which each omitted zero contributes at most $x^{1/2}/(|\gamma|\log x)$ and the accumulated tail is bounded by $C\,x^{1/2}\log^2 x/T$; and an integral estimate of the tail against the zero density, whose integral over $[T,\infty)$ diverges and yields no finite bound without an auxiliary cutoff. We take the termwise bound as the operative convention,

$$|R_T(x)| \leq C\,\frac{x^{1/2}\log^2 x}{T},$$

with the worst-case constant $C = 1$, because its constant is explicit and its derivation requires no divergent integral; we retain the exponential-in-$(\log x)^2$ height suggested by the cutoff integral as a conjectural strengthening, not as a derived result. This is the single substantive modeling choice in the paper, and Section 6 states how the conclusions change if it fails.

Rounding criterion. The algorithm computes the right-hand side and rounds to the nearest integer. By the prime number theorem the average gap between consecutive primes near x is $\log x$; we impose the conservative requirement that the total error not exceed $\frac{1}{2}$, which guarantees correct rounding of $\pi(x)$ regardless of gap structure.

Definitions. *Constructive emergence* (Definition 1): $\mathcal{Z}_T$ constructively determines a sequence $\{a_n\}$ if some algorithm outputs a_n for any prescribed n with error at most 2^{-p} in time polynomial in p and $\log n$. *Polynomial-precision constructive inversion*: such an algorithm whose required zero height T is bounded by a polynomial in $\log x$.

Recurrence classification. For an integer linear recurrence $a_n = c_1 a_{n-1} + \cdots + c_k a_{n-k}$ with characteristic polynomial $P(z)$, we associate the Dirichlet series $D(s) = \sum_{n \geq 1} a_n n^{-s}$ and test, via the Pisot-flow correspondence of [12], whether the sequence admits exact rounding recovery from finitely many spectral parameters.

Assumptions. (i) The Riemann Hypothesis holds, so $\Re\rho = \frac{1}{2}$ and each omitted term is bounded by $x^{1/2}/(|\gamma| \log x)$; without it the bound must be restated with $x^{\Re\rho}$ in place of $x^{1/2}$. (ii) The constant in the error model is $C = 1$ (worst case). (iii) The rounding step is exact; floating-point error is treated separately in Section 6. (iv) The $O(\log T)$ term in $N(T)$ is neglected in numerical zero counts, with the resulting uncertainty reported.

4. Analysis All inputs are stated; every arithmetic step is shown.

4.1 Error budget

From the termwise error model with $C = 1$,

$$|R_T(x)| \leq \frac{x^{1/2} \log^2 x}{T}.$$

Correct nearest-integer rounding of $\pi(x)$ requires

$$\frac{x^{1/2} \log^2 x}{T} \leq \frac{1}{2}.$$

$$T \geq 2\, x^{1/2} \log^2 x.$$

With the worst-case constant $C = 1$ this is the tightest height the model demands; any larger C only increases T proportionally, so the scaling conclusions below are insensitive to C.

4.2 Height bound and zero count for $x = 10^6$

All inputs: $x = 10^6$ (chosen as a representative cryptographic-sub-key scale for illustration), natural logarithms, and the asymptotic $N(T) = \frac{T}{2\pi} \log \frac{T}{2\pi} + O(\log T)$ from Section 3.

Step 1. $\log x = \log 10^6 = 6 \log 10 = 6 \times 2.30258509 = 13.81551056$.

Step 2. $\log^2 x = (13.81551056)^2 = 190.86833706$.

Step 3. $x^{1/2} = \sqrt{10^6} = 1000$.

Step 4. $T_{\min} = 2 \times 1000 \times 190.86833706 = 381{,}736.67412$, so

$$T_{\min} \approx 3.82 \times 10^5.$$

Step 5. Zero count. Neglecting the $O(\log T)$ term (Assumption iv),

$$N(T_{\min}) \approx \frac{T_{\min}}{2\pi} \log \frac{T_{\min}}{2\pi}.$$

Compute $\frac{T_{\min}}{2\pi} = \frac{381{,}736.67412}{6.28318531} = 60{,}755.62188$ and $\log(60{,}755.62188) = \log(6.075562188 \times 10^4) = \log 6.075562188 + 4 \log 10 = 1.80456078 + 9.21034037 = 11.01490115$. Then

$$N(T_{\min}) \approx 60{,}755.62188 \times 11.01490115 = 669{,}020.4 \approx 6.69 \times 10^5.$$

The neglected $O(\log T)$ term is of order $\log(3.82 \times 10^5) \approx 12.9$, i.e. a relative uncertainty of about 2×10^{-5} in the count, far below the two significant figures quoted.

Step 6. Super-polynomiality in the input size. The input to a primality or counting algorithm at scale x has size $\log x$. The required height satisfies

$$T_{\min} = 2\,x^{1/2}\log^2 x = \exp\!\big(\tfrac{1}{2}\log x + 2\log\log x + \log 2\big),$$

which is exponential in $\log x$. The zero count $N(T_{\min})$ inherits this: $N(T_{\min}) \approx \frac{T_{\min}}{2\pi}\log\frac{T_{\min}}{2\pi}$ is again exponential in $\log x$, hence super-polynomial in the input size $\log x$. No polynomial in $\log x$ bounds $N(T_{\min})$, so under Definition 1 the zeta-zero data set fails polynomial-precision constructive inversion for $\pi(x)$ at every scale x for which the error model holds.

4.3 Worked Bernoulli verification

The identity of [7] states, for the ordinary power sum $S_k^{1,0}(n) = 0^k + 1^k + \cdots + (n-1)^k$,

$$S_k^{1,0}(n) = \frac{1}{k+1}\sum_{t=1}^{k+1}\binom{k+1}{t} B_{k+1-t}\, n^t,$$

with $B_0 = 1$, $B_1 = -\frac{1}{2}$, $B_2 = \frac{1}{6}$ (the values also listed in the background of [11]). We verify $S_2^{1,0}(5)$, whose direct value is $0^2 + 1^2 + 2^2 + 3^2 + 4^2 = 0 + 1 + 4 + 9 + 16 = 30$.

Step 1. $k = 2$, $n = 5$, so the sum runs over $t = 1, 2, 3$ with $\binom{3}{t}$ and B_{3-t}:

$$S_2^{1,0}(5) = \frac{1}{3}\left[\binom{3}{1}B_2 \cdot 5 + \binom{3}{2}B_1 \cdot 25 + \binom{3}{3}B_0 \cdot 125\right].$$

Step 2. Term by term: $\binom{3}{1}B_2 \cdot 5 = 3 \times \frac{1}{6} \times 5 = \frac{15}{6} = 2.5$; $\binom{3}{2}B_1 \cdot 25 = 3 \times (-\frac{1}{2}) \times 25 = -37.5$; $\binom{3}{3}B_0 \cdot 125 = 1 \times 1 \times 125 = 125$.

Step 3. Sum: $2.5 - 37.5 + 125 = 90$; divide by 3:

$$S_2^{1,0}(5) = \frac{90}{3} = 30,$$

matching the direct evaluation exactly. A short list of Bernoulli numbers—the spectral-type datum—thus constructively determines every term of the power-sum sequence with zero error, the constructive pole of the dichotomy.

4.4 Spectral-gap numerical example

From [9], the prime lower bound is $\lambda_1 \geq (p+1) - 2\sqrt{p}$ and the composite upper bound is $\lambda_1 \leq C/D$ with $C = 1$, $D = 6$. For $p = 2$:

Step 1. $\lambda_1^{\text{prime}} \geq 3 - 2\sqrt{2} = 3 - 2.82842712 = 0.17157288$.

Step 2. $\lambda_1^{\text{comp}} \leq \frac{1}{6} = 0.16666667$.

Step 3. Gap:

$$\delta \geq 0.17157288 - 0.16666667 = 0.00490621 \approx 0.0049.$$

This gap is the numerical separation that the deterministic-primality proposal of [9] rests on; both bounds are conjectural there (the composite bound is stated as a hypothesis in [9]), so the example illustrates the framework's constants rather than an established test. The continuous-spectrum interval $[(p+1) - 2\sqrt{p},\ (p+1) + 2\sqrt{p}]$ derived in [11] supplies the analytic lower endpoint used in Step 1, and [11] further proves that any detected ℓ^2 eigenvalue on $\mathcal{T}_p$ would falsify the gauge-invariant-spectra hypothesis—a falsification criterion we adopt in Section 6.

4.5 Dichotomy: Pisot recurrences

For the Fibonacci sequence, [12] records the exact rounding statement

$$\left|F_n - \phi^n/\sqrt{5}\right| < \tfrac{1}{2} \quad \text{for all } n \geq 1,$$

where $\phi = (1+\sqrt{5})/2$ is a Pisot number of degree 2 with minimal polynomial $z^2 - z - 1$. The single spectral parameter ϕ (equivalently, the pair of characteristic roots) therefore determines every term F_n exactly by rounding: the required data set has constant size, independent of n, and recovery is exact rather than error-bounded. This contrasts pointwise with the zeta case, where the data set $\mathcal{Z}_T$ must grow exponentially in $\log x$ to push the error below $\frac{1}{2}$. The dichotomy is structural: in the Pisot case the dominant root is algebraic and separated from the subdominant roots by an exponential gap, so truncation is exact after rounding; in the zeta case the zeros have density growing like $\log \gamma$ and no algebraic separation, so truncation error decays only as $1/T$ per the termwise model.

5. Results

All numbers below are computed in Section 4 from the stated inputs; none is measured or simulated.

1. **Height bound (Theorem 2).** Under the termwise error model with $C = 1$ and the Riemann Hypothesis, evaluating $\pi(x)$ by the truncated explicit formula with error at most $\frac{1}{2}$ requires

$$T \geq 2\, x^{1/2} \log^2 x.$$

1. **Numerical instance.** For $x = 10^6$: $T_{\min} = 381{,}736.67 \approx 3.82 \times 10^5$ and $N(T_{\min}) \approx 6.69 \times 10^5$ zeros (uncertainty from the neglected $O(\log T)$ term about 2×10^{-5} relative).
2. **Complexity classification.** $N(T_{\min})$ is exponential in $\log x$, hence super-polynomial in the input size $\log x$; the zeta-zero data set does not admit polynomial-precision constructive inversion of $\pi(x)$ under the stated error model.
3. **Bernoulli verification.** $S_2^{1,0}(5) = 30$, exact, via the identity of [7].
4. **Spectral-gap constants.** $\lambda_1^{\text{prime}} \geq 0.17157288$, $\lambda_1^{\text{comp}} \leq 0.16666667$, gap $\delta \approx 0.00490621$ for $p = 2$, both bounds conjectural in [9].
5. **Dichotomy.** Fibonacci-type Pisot recurrences admit exact recovery from finitely many (here one) spectral parameters, per [12].

6. Discussion

The central result is conditional on the error model, and its strength is exactly the strength of that model. We state the failure modes before the interpretation.

What would falsify the main result. The claim "no polynomial-size spectral replacement exists under the error model" is falsified by any of: (i) a proof that the truncation tail of the explicit formula obeys a bound stronger than $C\, x^{1/2} \log^2 x/T$—for instance a bound of the form $x^{1/2} \log x/T$ or better with an effective constant, which would lower $T_{\min}$ by a factor $\log x$ and, if a bound polynomial in $\log x$ in the denominator emerged, could restore polynomial-size data; (ii) a constructive algorithm that evaluates $\pi(x)$ to integer precision from zeros up to height polynomial in $\log x$, circumventing the termwise tail bound (e.g. by exploiting cancellation in the tail rather than bounding it termwise); (iii) in the p-adic direction, detection of any ℓ^2 eigenvalue on $\mathcal{T}_p$, which by the criterion of [11] would falsify the spectral picture underlying the gap constants of [9]. Conversely, a proof that the tail integral against the zero density genuinely diverges without an exponentially large cutoff would strengthen the super-polynomial conclusion from model-dependent to unconditional.

Limitations. The result rests on the following assumptions, with sensitivity noted for each.

- *Riemann Hypothesis.* Without it, each omitted term scales as $x^{\Re\rho}$ and the bound becomes $T \geq 2\, x^{\Re\rho_{\max}} \log^2 x$ for the largest real part among omitted zeros; a zero off the critical line weakens the bound by the factor $x^{\Re\rho - 1/2}$, which is catastrophic for $\Re\rho$ close to 1.

- *Termwise error convention.* The choice of the termwise bound over the divergent integral estimate is the single substantive modeling decision; the conclusions hold at current error-model strength and would change if a sharper effective tail bound were proved.
- *Worst-case constant $C = 1$.* $T_{\min}$ scales linearly in C; the super-polynomiality conclusion is insensitive to any fixed C.
- *Exact rounding.* Floating-point evaluation of $\mathrm{Li}(x^\rho)$ at 6.69×10^5 zeros requires roughly p-digit precision with p large enough that accumulated round-off stays below $\frac{1}{2}$; we have not analyzed this cost, and it could add a polynomial factor in $\log x$ without changing the exponential-in-$\log x$ data requirement.
- *Neglected $O(\log T)$ in $N(T)$.* Affects the zero count only at relative order 2×10^{-5} at the computed height.
- *Conjectural gap constants.* The spectral-gap numbers of Section 4.4 inherit the conjectural status they have in [9]; they illustrate the framework and are not established primality criteria.
- *Scope.* The bound concerns $\pi(x)$ via the truncated explicit formula. Individual prime locations p_n require resolving gaps of size $\log x$, so a pointwise prime oracle needs error below $\frac{1}{2}$ in a gap-resolved sense, which is at least as demanding; we have not derived a separate bound for p_n.

Open questions. Which integer linear recurrences have Dirichlet series that are irreducible L-functions with Euler products? The counterexample $a_n = 2a_{n-1} - a_{n-2}$, whose Dirichlet series reduces to $\zeta(s-1)$, shows reducibility is not preserved under the correspondence of [12]; a classification theorem—plausibly restricted to Galois-representation recurrences—is testable by symbolic computation over recurrence families. Second, does the cutoff integral estimate for the explicit-formula tail admit an effective auxiliary cutoff that converts the conjectural exponential-in-$(\log x)^2$ height into a theorem? Third, the non-uniqueness obstruction from spectral geometry—isospectral, non-isometric manifolds produced by Sunada-type methods as described in [6]—suggests that even exact spectral equality need not determine discrete structure; formalizing when spectral data plus known arithmetic side constraints restore uniqueness remains open.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The prospective users are designers of prime-generation and certification pipelines in cryptographic libraries and hardware security modules, and the standards bodies that specify primality testing for key generation. If the result holds, the decision it changes is whether to fund or standardize spectral-data-based deterministic primality testing: the super-polynomial data requirement in $\log x$ says such a program cannot yield a polynomial-size zero table usable at cryptographic sizes under the stated error model, so effort stays with stacked probabilistic rounds or with algebraic deterministic tests. The result is also foundational for the p-adic spectral-gap proposal of [9]: it supplies the quantitative obstruction any such proposal must overcome, namely that pointwise integer-exact recovery from archimedean spectral data needs exponentially many spectral parameters in the input size.

(b) Readiness. This work is purely analytical and has not been peer reviewed: by the NASA Technology Readiness Level definitions it sits at TRL 1–2 (basic principles observed and reported; concept formulated). No claim above TRL 3 is made, and no manufacturing readiness is implied.

(c) Smallest next test. The smallest de-risking calculation is numerical: for a modest x (e.g. $x = 10^4$), compute the truncated explicit formula with zeros up to a sequence of heights T, measure the empirical truncation error $|R_T(x)|$, and fit the decay against T. The model predicts $|R_T(x)| \approx x^{1/2} \log^2 x / T$; an observed decay materially faster than $1/T$ (e.g. exponential in T) would refute the termwise convention and could reopen the polynomial-data question, while agreement would consolidate the height bound. A companion symbolic computation would enumerate small integer linear recurrences and test irreducibility of their Dirichlet series against the correspondence of [12], probing the classification question of Section 6.

(d) Binding constraints. For any future implementation attempt, the first binding constraint is data volume: storing $N(T) \approx 6.69 \times 10^5$ zeros at the precision needed for exact rounding is already nontrivial, and at cryptographic x the count grows exponentially in $\log x$, so storage and high-precision arithmetic dominate before any compute constraint binds. Second is floating-point precision: each term $\mathrm{Li}(x^\rho)$ must be evaluated to enough digits that the sum of $N(T)$ terms retains error below $\frac{1}{2}$. Whether a faster-than-$1/T$ effective tail decay exists is an open question rather than a derived constraint. Regulatory or cost constraints do not bind at this analytical stage.

(e) Order-of-magnitude bounds (projections). With the assumptions of Section 3 and the computed $T_{\min}$ for $x = 10^6$: a zero table for that scale holds on the order of 10^5 complex numbers; scaling to $x = 10^{12}$ would give $T_{\min} = 2 \times 10^6 \times (27.63102112)^2 \approx 1.53 \times 10^9$ and $N(T_{\min}) \approx \frac{1.53\times10^9}{2\pi} \log\left(\frac{1.53\times10^9}{2\pi}\right) \approx 2.43 \times 10^8 \times 19.31 \approx 4.7 \times 10^9$ zeros—about 3.8 orders of magnitude more data for six additional input digits, illustrating the exponential-in-$\log x$ growth. These are projections under the stated error model, not measurements.

(f) Risks and misuse paths. No direct misuse path arises from a negative result: the finding argues against a technology rather than enabling one. The residual risk it leaves in place is the known one—probabilistic primality tests admit rare composite liars, as noted in the motivation of [9]—and the result's practical effect is to justify continued investment in mitigations for that risk rather than in spectral replacement.

8. Conclusion

We formalized constructive emergence as algorithmic recovery of arithmetic-sequence terms from truncated spectral data and derived, under an explicit termwise error model for the truncated explicit formula, the height bound $T \geq 2\,x^{1/2}\log^2 x$ for evaluating $\pi(x)$ with error at most $\frac{1}{2}$. At $x = 10^6$ this requires $T_{\min} \approx 3.82 \times 10^5$ and about 6.69×10^5 zeros—exponential in the input size $\log x$—so at current error-model strength the zeta zeros do not admit a polynomial-size spectral replacement for probabilistic primality certification. The dichotomy is sharp: Pisot recurrences such as the Fibonacci sequence recover every term exactly from finitely many spectral parameters, while the zero-to-prime map requires exponentially growing spectral data for pointwise integer-exact output. The conclusions are conditional on the termwise tail convention; a proved stronger tail bound, a cancellation-exploiting algorithm, or an ℓ^2 eigenvalue on $\mathcal{T}_p$ in the sense of [11] would each falsify part of the picture, and the smallest next test is the numerical decay measurement of the truncation error proposed in Section 7.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.45.2). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, openai/gpt-oss-120b, moonshotai/kimi-k2.6, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: the execution sandbox was unavailable, so a second language model recomputed 6 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

[1] Data Encoding for Byzantine-Resilient Distributed Optimization. arXiv:1907.02664v2. https://arxiv.org/abs/1907.02664v2 [2] Byzantine-Resilient SGD in High Dimensions on Heterogeneous Data. arXiv:2005.07866v1. https://arxiv.org/abs/2005.07866v1 [3] Constraints on dark energy from H II starburst galaxy apparent magnitude versus redshift data. arXiv:1110.5626v1. https://arxiv.org/abs/1110.5626v1 [4] TerraGen: A Unified Multi-Task Layout Generation Framework for Remote Sensing Data Augmentation. arXiv:2510.21391v1. https://arxiv.org/abs/2510.21391v1 [5] AeroGen: Enhancing Remote Sensing Object Detection with Diffusion-Driven Data Generation. arXiv:2411.15497v3.

https://arxiv.org/abs/2411.15497v3 [6] Arithmetic lattices and weak spectral geometry. arXiv:0706.3841v1. https://arxiv.org/abs/0706.3841v1 [7] A note on the coefficients of power sums of arithmetic progressions. arXiv:2101.05713v5. https://arxiv.org/abs/2101.05713v5 [8] A 20 Gbps Data Transmitting ASIC with PAM4 for Particle Physics Experiments. arXiv:2010.16064v1. https://arxiv.org/abs/2010.16064v1 [9] https://papers.qnfo.org/papers/spectral-gap-on-bruhat-tits-trees-as-a-deterministic-primality-test. QNFO: Spectral Gap on Bruhat–Tits Trees as a Deterministic Primality Test. [10] https://papers.qnfo.org/papers/dual-mode-arithmetic-representation. QNFO: Dual-Mode Arithmetic Representation. [11] https://papers.qnfo.org/papers/l-function-zeros-as-gauge-invariant-spectra-across-archimedean-and-p-adic-geomet. QNFO: L-Function Zeros as Gauge-Invariant Spectra Across Archimedean and p-adic Geometries. [12] https://papers.qnfo.org/papers/emergent-number-theory. QNFO: EMERGENT NUMBER THEORY.