Riemann Ground State: Holographic Trace Formulas, Abel Finite Parts, and a Protocol Derivation of the Riemann Hypothesis in HPA–Omega

Haobo Ma (Auric)*
AELF PTE LTD.
#14-02, Marina Bay Financial Centre Tower 1, 8 Marina Blvd, Singapore 018981

December 26, 2025

Abstract

In the HPA–Omega (Holographic Polar Arithmetic / Omega axiom system) scan–readout paradigm, we present a closed protocol derivation chain that recasts the classical explicit formula for the Riemann zeta function as a *holographic trace formula.* Under finite-information readout and unitary scanning constraints, we characterize the Riemann Hypothesis (RH) as a *definability and boundedness requirement* for a canonical Abel-regularized trace along the path $r \uparrow 1$.

The core mechanism is simple and auditable: if there exists a trace identity embedding the spectral contributions of zeta zeros into a scan algebra, and if the resulting Abel trace is holomorphic on the unit disk $|r| < 1$ (as mandated by bounded unitary scan readout), then any zero off the critical line forces an *interior pole* of the spectral-side mode factor at $r = \mathrm{e}^{-(\rho-\frac{1}{2})}$ with $|r| < 1$. This contradicts holomorphy of the geometric side, hence all nontrivial zeros must lie on $\mathrm{Re}(s) = \frac{1}{2}$.

We conclude with reproducible numerics: (i) the star discrepancy of the golden-branch scan exhibits the expected logarithmic stability, and (ii) in a toy "zero-mode" signal model, an artificial shift $\mathrm{Re}(\rho) \neq \frac{1}{2}$ produces an Abel threshold blow-up and an energy explosion, illustrating the paper's polar rigidity mechanism.

Keywords: HPA; Omega axiom system; holographic trace formula; Abel finite part; star discrepancy; unitary scanning; Riemann hypothesis; prime spectrum.

Conventions. Unless otherwise stated, log denotes the natural logarithm. We use $t \in \mathbb{Z}_{\geq 0}$ for discrete protocol time. We identify the d-torus with $\mathbb{T}^d = (\mathbb{R}/\mathbb{Z})^d$ and use $[0, 1)^d$ as the standard fundamental domain when defining discrepancy. Throughout, "canonical Abel path" refers to the limit process $r \uparrow 1$ with $r \in (0, 1)$, applied to Abel generating functions that are defined (and holomorphic) for $|r| < 1$.

Contents

1 Introduction: from analytic infinity to protocol geometry . . . 4
1.1 Summary of the closed-layer derivation . . . 4
1.2 What this paper does and does not claim . . . 5

*Email: auric@aelf.io

2 Preliminaries: zeta, explicit formulas, and closed-layer scan objects **5**
2.1 The Riemann zeta function and zero notation 5
2.2 Explicit formulas as a trace template 5
2.3 A quantitative zero counting law 6
2.4 Scan–readout protocols and star discrepancy 6
2.5 Abel regularization and finite-part extraction 6

3 HPA–Omega axioms: finite information, unitary scanning, and readout rigidity **7**
3.1 Finite information and compact readout 7
3.2 Scan–projection readout 7
3.3 Unitary scan algebra and Weyl pairs 8
3.4 Abel-weighted traces are well-defined under boundedness 8

4 From symbolic orbits to "zeta": finite-state prototypes and infinite-state lift **8**
4.1 A finite-state prototype: the Zeckendorf shift and a dynamical zeta 9
4.2 Why an infinite-state lift is necessary 9
4.3 Primes as periodic orbit data (geometric-side motivation) 9

5 Holographic trace formula and polar rigidity: a protocol derivation of RH **9**
5.1 The Holographic Trace Formula assumption (HTF) 10
5.2 An internal pole obstruction for off-critical modes 10
5.3 Polar rigidity implies the critical line 11

6 Programmatic layer: ground state, energy, and discrepancy (not used in proofs) **12**
6.1 RH as "no exponential gain/loss" 12
6.2 Discrepancy as an auditable proxy for finite-resource "energy" 12

7 Reproducible numerics: discrepancy stability and "off-line zero" threshold blow-up **12**
7.1 Reproducibility protocol 13
7.2 Experiment A: one-dimensional golden scan star discrepancy 13
7.3 Experiment B: toy "zero-mode" signal and the Abel threshold 13

8 Conclusion: RH as a protocol-stable ground-state condition **15**

A Audit closure checklist (this version) **15**
A.1 Closed-layer inputs (assumptions) 15
A.2 Closed-layer derivations (theorems) 16
A.3 Programmatic layer (not used in proofs) 16

B Finite-state zeta prototypes and why infinite-state lift is unavoidable **16**
B.1 Finite-state dynamical zeta is rational 16
B.2 Nontrivial zeta zeros require an infinite-dimensional mechanism 16
B.3 Finite parts, Abel regularization, and zeta-regularized determinants (directional note) 16

C Orbit calculus, Abel means, and finite parts (rotation template) **17**
C.1 Abel generating functions are holomorphic on the unit disk 17
C.2 Abel means for uniquely ergodic rotations 17
C.3 Abel finite part as a constant term (uniqueness and a Fourier criterion) 18
C.4 Golden-branch star discrepancy bound via Denjoy–Koksma and Ostrowski 18

D Reproducible experiment code (pure Python) 19
D.1 Experiment A: exact one-dimensional star discrepancy for the golden scan 19
D.2 Experiment B: toy “zero-mode” signal, energy growth, and Abel thresholds . . . 21

1 Introduction: from analytic infinity to protocol geometry

The Riemann Hypothesis (RH) asserts that every nontrivial zero ρ of the Riemann zeta function satisfies

$$\mathrm{Re}(\rho) = \tfrac{1}{2}.$$

In its classical analytic number theory setting, RH is embedded in the machinery of analytic continuation, functional equations, and explicit estimates for prime-counting functions [1–3]. In the "number theory meets physics" tradition (e.g. the Hilbert–Pólya program and trace formula perspectives), the zeros behave like a frequency spectrum dual to prime periodic data.

This paper takes a different entry point. We treat the *finite-resource scan–readout process* as a first-class object rather than a late-stage regularization of an a priori infinite analytic formalism. Concretely, we adopt an audit-driven two-layer writing discipline:

Closed layer. Only protocol definitions, finite-resource objects, deterministic error certificates, and statements that can be proved within these notions are admitted as mathematical inputs.

Programmatic layer. Physical and teleological interpretations are allowed, but they are explicitly excluded from proofs and do not serve as theorem inputs.

Within the closed layer, our goal is a logically checkable implication chain that isolates where the nontrivial analytic difficulty resides and what is forced once a specific bridge is granted. The bridge is a structural assumption, formulated below as a *Holographic Trace Formula* (HTF), that embeds the classical explicit formula for ζ into a trace identity of a scan algebra. Under this assumption, RH becomes a rigidity statement about *Abel regularization along a canonical protocol path.*

1.1 Summary of the closed-layer derivation

We aim to establish the following conditional chain.

1. **Omega axioms (finite information + unitary scan + bounded readout).** Selected axioms of the HPA–Omega system imply that a broad class of Abel-weighted orbit traces define holomorphic functions on the open unit disk $|r| < 1$ (as power series with bounded coefficients), and that a finite-part extraction along the canonical real path $r \uparrow 1$ is an admissible operation in the closed layer (Section 3).
2. **Holographic Trace Formula assumption (HTF).** Assume the existence of a trace identity that is isomorphic, under the same regularization convention, to a Weil-type explicit formula: prime/periodic-orbit data form the geometric side and zeta zeros form the spectral side. Crucially, each zero contributes as an exponential mode $e^{(\rho-\frac{1}{2})t}$ in the log-time variable (Section 5).
3. **Internal pole obstruction for off-critical modes.** If $\rho = \beta + \mathrm{i}\gamma$ with $\beta > \frac{1}{2}$, then the corresponding mode factor
$$\frac{1}{1 - r\, e^{(\rho-\frac{1}{2})}}$$
has an interior pole at $r_\rho = e^{-(\beta-\frac{1}{2})}e^{-\mathrm{i}\gamma}$ with $|r_\rho| < 1$. Therefore any spectral-side expression that contains this mode with nonzero coefficient cannot extend holomorphically to the full unit disk, contradicting the Omega geometric-side holomorphy (Lemma 5.2).
4. **Polar rigidity implies the critical line.** Combining Omega's geometric-side holomorphy on $|r| < 1$ with HTF and the internal-pole obstruction rules out $\mathrm{Re}(\rho) \neq \frac{1}{2}$ and yields RH as a closed-layer consequence of (Omega + HTF) (Theorem 5.4).

1.2 What this paper does and does not claim

The argument does *not* claim an unconditional proof of RH in the standard axioms of analytic number theory. The nontrivial content is isolated into the explicit structural bridge HTF: constructing, in classical terms, an infinite-dimensional trace/transfer operator satisfying HTF is essentially the hard part of RH. What this paper provides is a clean protocol-level statement: *if* ζ admits an HTF embedding compatible with Omega's canonical Abel finite-part rule, *then* RH is forced by an analytic rigidity on the unit disk: bounded scan traces are holomorphic for $|r| < 1$, while any off-critical spectral mode would introduce an interior pole.

We emphasize that the core rigidity is analytic: once the geometric side is a holomorphic function on the unit disk (as mandated by bounded unitary scan readout) and the spectral side inserts each zero via a mode factor with pole at $r = \mathrm{e}^{-(\rho-\frac{1}{2})}$, any off-critical zero produces an interior pole and is thereby incompatible with the protocol trace identity.

2 Preliminaries: zeta, explicit formulas, and closed-layer scan objects

2.1 The Riemann zeta function and zero notation

We write

$$\zeta(s) = \sum_{n\geq 1} n^{-s} \qquad (\mathrm{Re}(s) > 1),$$

and use analytic continuation to extend ζ to $\mathbb{C} \setminus \{1\}$. Nontrivial zeros are denoted by $\rho = \beta + \mathrm{i}\gamma$ with $0 < \beta < 1$. The functional equation implies the usual symmetries: if ρ is a zero, so are $\overline{\rho}$ and $1 - \rho$ (ignoring the trivial zeros).

Completed zeta function. It is often convenient to work with the completed zeta function

$$\xi(s) := \tfrac{1}{2} s(s-1)\, \pi^{-s/2} \Gamma\big(\tfrac{s}{2}\big)\, \zeta(s),$$

which is entire and satisfies $\xi(s) = \xi(1-s)$ [1, 2]. The nontrivial zeros of ζ coincide with the zeros of ξ.

2.2 Explicit formulas as a trace template

Weil-type explicit formulas share a common shape: a *geometric side* supported on primes (or prime powers) equals a *spectral side* supported on the zeros ρ of ζ, up to controlled archimedean terms [1, 3]. Rather than reproducing technical hypotheses, we extract only the mode structure needed for the protocol argument.

Passing to log coordinates $t = \log x$, the term x^ρ becomes $\mathrm{e}^{\rho t}$. After centering at $\frac{1}{2}$ (the natural symmetry axis), the spectral contribution of a zero typically appears as an exponential mode of the form

$$\mathrm{e}^{(\rho-\frac{1}{2})t} = \mathrm{e}^{(\beta-\frac{1}{2})t}\, \mathrm{e}^{\mathrm{i}\gamma t}.$$

The closed-layer rigidity mechanism will use only this elementary growth/oscillation decomposition under Abel weights.

Proposition 2.1 (A standard explicit formula for $\psi(x)$ (quoted))**.** *Let $\psi(x) = \sum_{n\leq x} \Lambda(n)$ be the Chebyshev function and assume $x > 1$ is not a prime power. Then one has the explicit formula*

$$\psi(x) = x - \sum_{\rho} \frac{x^\rho}{\rho} - \log(2\pi) - \tfrac{1}{2} \log(1 - x^{-2}),$$

where the sum runs over the nontrivial zeros ρ of ζ (interpreted in the standard symmetric limiting sense).

Reference. See, e.g., [1, Ch. 14] or [2, Ch. 1]. □

The mode structure used in HTF is the direct consequence of the identity $x^\rho = e^{\rho \log x}$: in log-time $t = \log x$, each zero contributes a frequency γ and a growth rate β.

2.3 A quantitative zero counting law

Let $N(T)$ denote the number of nontrivial zeros $\rho = \beta + i\gamma$ with $0 < \gamma \le T$, counted with multiplicity. The Riemann–von Mangoldt formula gives the quantitative density of the spectral side:

$$N(T) = \frac{T}{2\pi} \log\left(\frac{T}{2\pi}\right) - \frac{T}{2\pi} + O(\log T) \qquad (T \to \infty).$$

We will not use this asymptotic in any proof, but it is useful for calibrating what kind of decay is needed for kernel-dependent coefficients $c_\rho(K)$ in order for spectral-side sums to define meromorphic functions on $|r| < 1$. Standard references include [1, 2].

2.4 Scan–readout protocols and star discrepancy

In the closed layer, a scan protocol and its readout are primary objects. A canonical example class is a Kronecker scan on the d-torus $\mathbb{T}^d = (\mathbb{R}/\mathbb{Z})^d$:

$$x_t = x_0 + t\alpha \pmod 1, \qquad t \in \mathbb{Z}_{\ge 0},$$

with irrational slope vector $\alpha \in \mathbb{R}^d$. Given a bounded kernel (readout observable) $f : [0,1)^d \to \mathbb{R}$, the finite-horizon readout is

$$\langle f \rangle_N = \frac{1}{N} \sum_{t=0}^{N-1} f(x_t),$$

a fully computable and auditable output at resource level N.

To certify uniformity of coverage in a deterministic way, we use the *star discrepancy* $D_N^*(P_N)$ of the point set $P_N = \{x_0, \dots, x_{N-1}\} \subset [0,1)^d$:

$$D_N^*(P_N) = \sup_{u \in [0,1]^d} \left| \frac{1}{N} \#(P_N \cap [0,u)) - \lambda([0,u)) \right|,$$

where $[0,u) = \prod_{j=1}^d [0, u_j)$ and λ is Lebesgue measure [4, 5].

In the one-dimensional golden-branch case $\alpha = \varphi^{-1}$ (with $\varphi = (1+\sqrt{5})/2$), an explicit logarithmic upper bound is available:

$$D_N^*(P_N) \;\le\; \frac{2(2 + \log_\varphi N)}{N}.$$

One convenient route to this explicit bound is via Denjoy–Koksma at convergent times together with the Ostrowski expansion of N (Appendix C; see also [4, 6, 7]). This certifies that, under a minimal-complexity irrational scan, finite-resource sampling error grows at most logarithmically (times $1/N$), in sharp contrast with the exponential growth induced by off-critical spectral modes.

2.5 Abel regularization and finite-part extraction

Abel regularization provides a protocol-native way to turn formal "infinite" orbit sums into auditable objects. Given a bounded function f along the scan orbit (x_t), define the Abel family

$$S_f(r) := \sum_{t=0}^{\infty} r^t f(x_t), \qquad 0 < r < 1.$$

For every $r < 1$, the series is absolutely convergent and computable. For uniquely ergodic scans (in particular, irrational translations on tori), Abel means converge to the space average for continuous observables; see, e.g., [8] and the Abel–Cesàro comparison in [9]. In the rotation template, one has the Abelian limit

$$\lim_{r\uparrow 1}(1-r)S_f(r) = \int_{[0,1)^d} f(x)\,\mathrm{d}x.$$

The *Abel finite part* (finite-part constant term) is defined, when the limit exists, by

$$\mathrm{FP}\sum_{t=0}^{\infty} f(x_t) \;:=\; \lim_{r\uparrow 1}\left(S_f(r) - \frac{\int f}{1-r}\right).$$

In the HPA–Omega formalism, the existence of this finite part along the canonical path $r \uparrow 1$ is not an afterthought: it is part of the admissibility conditions for using the corresponding infinite object in the closed layer. Appendix C records a precise asymptotic/finite-part template, including a Fourier-resolvent criterion that guarantees existence of the constant term for sufficiently regular readout kernels.

3 HPA–Omega axioms: finite information, unitary scanning, and readout rigidity

This section isolates the fragments of the Omega axiom system used by the closed-layer argument. The role of these axioms is to formalize a simple but strong constraint: *finite-resolution readout of a unitary scan on a compact state space cannot support exponential growth along the canonical Abel path.*

3.1 Finite information and compact readout

Axiom 3.1 (O2 (Finite information / compact readout))**.** *For any fixed finite resolution and finite observation horizon, effective observables are represented in a controlled (effectively finite) operator algebra on a compact state space. In particular, expectations and traces of bounded observables obey uniform boundedness constraints under admissible readout channels.*

In practice, O2 means that the geometric side of any protocol-defined trace built from bounded effect operators remains bounded whenever its defining series converges.

3.2 Scan–projection readout

Axiom 3.2 (O5 (Scan–projection readout))**.** *There exists a pointer observable V with spectral measure Π_V and a family of finite-resolution windows $\{w_k^{(\varepsilon)}\}_k$ such that the readout effects are*

$$E_k^{(\varepsilon)} = \int w_k^{(\varepsilon)}(x)\,\mathrm{d}\Pi_V(x).$$

Given an effective state ω_{eff}, the readout channel probabilities are

$$P_k^{(\varepsilon)} = \omega_{\mathrm{eff}}(E_k^{(\varepsilon)}).$$

This is a protocol-level encoding of the idea that continuous phases are only accessed through finite-resolution projections, producing finite symbol streams (e.g. Sturmian/Fibonacci words) rather than raw continuum data.

3.3 Unitary scan algebra and Weyl pairs

Axiom 3.3 (O6 (Unitary scan algebra / Weyl pair))**.** *The scan dynamics is generated by a unitary operator $U_{\rm scan}$. Together with the pointer operator V, it forms a Weyl pair*

$$U_{\rm scan}\, V = \mathrm{e}^{2\pi\mathrm{i}\alpha}\, V\, U_{\rm scan},$$

where α is irrational (the golden branch $\alpha = \varphi^{-1}$ is the minimal-complexity model case).

Unitarity implies that conjugation preserves operator norms:

$$\left\| U_{\rm scan}^{t} A U_{\rm scan}^{-t} \right\|_{\infty} = \|A\|_{\infty} \qquad (t \in \mathbb{Z}).$$

This yields a basic closed-layer fact: Abel-weighted protocol traces built from bounded single-step contributions are automatically well-defined for every $0 < r < 1$.

3.4 Abel-weighted traces are well-defined under boundedness

Proposition 3.4 (Abel definability for bounded scan contributions)**.** *Let A be a bounded observable and let $\omega_{\rm eff}$ be a bounded state functional. Define the Abel-weighted trace-like quantity*

$$T_A(r) := \sum_{t=0}^{\infty} r^t\, \omega_{\rm eff}(U_{\rm scan}^{t} A U_{\rm scan}^{-t}), \qquad |r| < 1.$$

Then $T_A(r)$ is absolutely convergent for every $|r| < 1$, defines a holomorphic function on the open unit disk, and satisfies the bound

$$|T_A(r)| \le \frac{\|A\|_{\infty}\, \|\omega_{\rm eff}\|}{1 - |r|}.$$

Proof. By boundedness of $\omega_{\rm eff}$ and unitarity,

$$\left| \omega_{\rm eff}(U_{\rm scan}^{t} A U_{\rm scan}^{-t}) \right| \le \|\omega_{\rm eff}\| \left\| U_{\rm scan}^{t} A U_{\rm scan}^{-t} \right\|_{\infty} = \|\omega_{\rm eff}\|\, \|A\|_{\infty}.$$

Hence for $|r| < 1$ the series is dominated by $\sum_{t \ge 0} |r|^t$ and converges absolutely. As a power series in r with radius of convergence at least 1, it defines a holomorphic function on the open unit disk. □

Proposition 3.4 is purely elementary, but it captures the essential closed-layer rigidity: within Omega, admissible geometric-side traces are defined on the entire interval $(0, 1)$. The additional Omega content is that a *canonical finite-part extraction along* $r \uparrow 1$ is treated as part of the protocol specification: only those traces for which the finite part exists on this path are admissible closed-layer objects.

4 From symbolic orbits to "zeta": finite-state prototypes and infinite-state lift

To connect scan orbits with zeta-like objects, it is useful to separate two layers: (i) a fully computable finite-state prototype where "trace $\leftrightarrow$ zeta" is completely explicit, and (ii) the infinite-state lift needed for a nontrivial zero spectrum.

4.1 A finite-state prototype: the Zeckendorf shift and a dynamical zeta

In HPA, the golden-branch readout naturally produces Fibonacci/Zeckendorf coding. A canonical symbolic dynamical system is the Zeckendorf shift: the subshift on $\{0,1\}^{\mathbb{Z}_{\geq 0}}$ forbidding the word "11". It is a subshift of finite type with adjacency matrix

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}.$$

Let $P_n = \operatorname{tr}(A^n)$ denote the number of period-n points. The Artin–Mazur dynamical zeta function is defined by [10, 11]

$$\zeta_\sigma(z) = \exp\left(\sum_{n \geq 1} \frac{P_n}{n} z^n\right).$$

For subshifts of finite type, one has the determinant formula

$$\zeta_\sigma(z) = \frac{1}{\det(I - zA)} = \frac{1}{1 - z - z^2}.$$

This illustrates the structural pattern relevant to our protocol viewpoint: *orbit traces* (here, $\operatorname{tr}(A^n)$) package into a zeta function via an exponential generating map or a determinant. However, because A is finite-dimensional, $\zeta_\sigma(z)$ is rational; its pole/zero structure is too coarse to model the subtle zeta-zero spectrum of $\zeta(s)$.

4.2 Why an infinite-state lift is necessary

Any finite-state automaton model yields a rational dynamical zeta function, hence admits only finitely many poles/zeros. To produce a genuine "zero spectrum" akin to that of the Riemann zeta function, one must pass to an infinite-dimensional limit (e.g. transfer operators, trace-class regularized determinants, or controlled refining partitions). This is precisely where the classical difficulty of RH resides.

Accordingly, in this paper we do not construct the required infinite-dimensional operator from scratch. Instead, we elevate the needed bridge to a clearly stated structural assumption (HTF in Section 5) and show that, once it exists and is compatible with Omega's regularization conventions, RH follows from an elementary analytic rigidity: off-critical modes force interior poles in the unit disk, which a bounded unitary scan trace cannot admit.

4.3 Primes as periodic orbit data (geometric-side motivation)

The HPA–Omega program also suggests a complementary geometric-side picture: after an adelic completion, primitive periodic orbit data of a scale flow can be matched with primes, with orbit length $\log p$. This aligns with Selberg/Connes-type spectral dualities where primes play the role of primitive closed geodesics and zeta zeros play the role of spectral frequencies [12]. In this paper, this picture serves only as motivation for treating the explicit formula as a trace identity; it is not used as an input to any closed-layer proof.

5 Holographic trace formula and polar rigidity: a protocol derivation of RH

We now state the structural bridge (HTF) and derive the core rigidity theorem. The proof uses only three ingredients: (i) Omega's geometric-side holomorphy of Abel traces on the unit disk, (ii) the HTF mode structure for zeros, and (iii) an elementary pole-location computation.

5.1 The Holographic Trace Formula assumption (HTF)

Assumption 5.1 (HTF (Holographic Trace Formula))**.** *There exists a family of controllable readout kernels (test objects) $\mathcal{K}$ and, for each $K \in \mathcal{K}$, an Abel-weighted scan trace function $T_K(r)$ constructed from the scan algebra (U_{scan}, V) and finite-resolution readout effects $\{E_k^{(\varepsilon)}\}$ such that:*

1. ***(Geometric side boundedness and canonical finite part).*** *For each fixed resolution ε and each $K \in \mathcal{K}$, the function $T_K(r)$ is given by a power series $\sum_{t \geq 0} a_t(K)\, r^t$ with bounded coefficients and therefore extends to a holomorphic function on the open unit disk $\{\,|r| < 1\}$ (cf. Proposition 3.4). Moreover, it admits a finite-part extraction along the canonical real path $r \uparrow 1$ in the sense of Section 2.5.*

2. ***(Explicit-formula isomorphism).*** *Under the same regularization convention, $T_K(r)$ admits a decomposition into a "prime/periodic-orbit" term and a "zero/spectral" term. On the spectral side, there exist complex coefficients $c_\rho(K)$ such that each nontrivial zero ρ contributes via an exponential mode sampled at integer times:*

$$S_\rho(r) := \sum_{t=0}^{\infty} r^t \mathrm{e}^{(\rho - \frac{1}{2})t},$$

   possibly after taking the same finite-part normalization used on the geometric side. Equivalently, $T_K(r)$ matches, kernel-by-kernel, a Weil-type explicit formula in which x^ρ appears as $\mathrm{e}^{\rho t}$ in log time. Whenever the above series converges (i.e. $|r\, \mathrm{e}^{\mathrm{Re}(\rho) - \frac{1}{2}}| < 1$), it sums to the meromorphic mode function

$$M_\rho(r) := \frac{1}{1 - r\, \mathrm{e}^{(\rho - \frac{1}{2})}}.$$

   HTF further assumes that the spectral-side expression

$$T_K^{\text{spec}}(r) := \sum_\rho c_\rho(K)\, M_\rho(r)$$

   is well-defined as a meromorphic function on $|r| < 1$ with poles only of the above mode form, in the sense that the series converges absolutely and uniformly on compact subsets of $\{|r| < 1\}$ that avoid the pole set. Finally, the HTF identity holds as an identity of meromorphic functions on $|r| < 1$ between $T_K(r)$ and its geometric+spectral decomposition (after the same canonical normalization on both sides).

3. ***(Detectability).*** *For every nontrivial zero ρ of ζ, there exists at least one kernel $K \in \mathcal{K}$ such that the residue of T_K^{spec} at the pole location $r = \mathrm{e}^{-(\rho - \frac{1}{2})}$ is nonzero.*

Closed-layer status. HTF is explicitly a structural assumption bridging scan algebra traces and the classical explicit formula. All subsequent closed-layer consequences will depend on HTF; no interpretation-layer content is used.

5.2 An internal pole obstruction for off-critical modes

Lemma 5.2 (Internal pole obstruction lemma)**.** *Let $\rho = \beta + \mathrm{i}\gamma$ with $\beta > \frac{1}{2}$. Then the meromorphic mode function*

$$M_\rho(r) = \frac{1}{1 - r\, \mathrm{e}^{(\rho - \frac{1}{2})}}$$

has a pole at

$$r_\rho = \mathrm{e}^{-(\beta-\frac{1}{2})}\,\mathrm{e}^{-\mathrm{i}\gamma}, \qquad |r_\rho| = \mathrm{e}^{-(\beta-\frac{1}{2})} < 1.$$

In particular, any spectral-side expression that contains a nonzero multiple of M_ρ cannot extend to a holomorphic function on the full unit disk.

Proof. The pole occurs where the denominator vanishes:

$$1 - r\,\mathrm{e}^{(\rho-\frac{1}{2})} = 0 \quad\Longleftrightarrow\quad r = \mathrm{e}^{-(\rho-\frac{1}{2})} = \mathrm{e}^{-(\beta-\frac{1}{2})}\mathrm{e}^{-\mathrm{i}\gamma}.$$

Its modulus is $\mathrm{e}^{-(\beta-\frac{1}{2})} < 1$ when $\beta > \frac{1}{2}$. ☐

Lemma 5.3 (Holomorphic–meromorphic incompatibility at an interior pole)**.** *Let $U \subset \mathbb{C}$ be open and let $r_0 \in U$. If F is holomorphic on U and G is meromorphic on U with a pole at r_0, then $F \not\equiv G$ on $U \setminus \{r_0\}$.*

Proof. Assume for contradiction that $F = G$ on $U \setminus \{r_0\}$. Choose a small disk $D \subset U$ centered at r_0. Since F is holomorphic on D, it has a Taylor series

$$F(r) = \sum_{k=0}^{\infty} b_k (r - r_0)^k \qquad (r \in D).$$

Since G is meromorphic on D with a pole at r_0, it has a Laurent expansion with a nontrivial principal part

$$G(r) = \sum_{k=-m}^{\infty} a_k (r - r_0)^k, \qquad a_{-m} \neq 0.$$

Equality on the punctured disk $D \setminus \{r_0\}$ forces equality of Laurent coefficients, hence $a_{-m} = 0$, a contradiction. ☐

5.3 Polar rigidity implies the critical line

Theorem 5.4 (Polar rigidity ⇒ critical line)**.** *Assume the Omega axioms O2/O5/O6 as used in Section 3 and assume HTF (Assumption 5.1). Then every nontrivial zero ρ of ζ satisfies* $\mathrm{Re}(\rho) = \frac{1}{2}$.

Proof. Fix a kernel $K \in \mathcal{K}$ and consider the corresponding Abel trace $T_K(r)$.

Step 1: geometric-side holomorphy on the unit disk. By construction in Omega, $T_K(r)$ is an Abel-weighted sum of bounded scan–readout contributions, hence a power series in r with radius of convergence at least 1. By Proposition 3.4, it defines a holomorphic function for every $|r| < 1$. Moreover, Omega requires that a canonical finite part along $r \uparrow 1$ exists for admissible traces (Assumption 5.1(1)).

Step 2: off-critical zeros force an internal pole. Suppose, for contradiction, that there exists a nontrivial zero $\rho = \beta + \mathrm{i}\gamma$ with $\beta > \frac{1}{2}$. By detectability (Assumption 5.1(3)), choose $K \in \mathcal{K}$ such that the spectral side has nonzero residue at the pole location $r_\rho = \mathrm{e}^{-(\rho-\frac{1}{2})}$. By Lemma 5.2, this location satisfies $|r_\rho| < 1$.

Step 3: contradiction with the Omega trace identity. HTF identifies the geometric-side holomorphic function $T_K(r)$ with the spectral-side meromorphic expression on $|r| < 1$. If an off-critical ρ were present with nonzero residue at r_ρ, the spectral side would have a pole at $r_\rho \in \{|r| < 1\}$, which is incompatible with holomorphy of $T_K(r)$ by Lemma 5.3. Therefore no off-critical ρ with $\mathrm{Re}(\rho) > \frac{1}{2}$ can be present.

Step 4: symmetry excludes $\beta < \frac{1}{2}$. Therefore no zero satisfies $\mathrm{Re}(\rho) > \frac{1}{2}$. By the functional equation symmetry $\rho \mapsto 1 - \rho$, this also excludes $\mathrm{Re}(\rho) < \frac{1}{2}$. Hence $\mathrm{Re}(\rho) = \frac{1}{2}$ for all nontrivial zeros. □

Corollary 5.5 (RH in the Omega+HTF protocol system)**.** *Within the protocol system defined by Omega (O2/O5/O6) plus HTF, the Riemann Hypothesis holds.*

Remark 5.6 (Closed-layer inputs)**.** *Theorem 5.4 depends only on: (i) Omega's geometric-side Abel definability and canonical finite-part admissibility, (ii) HTF's mode structure in which zeros contribute as* $\mathrm{e}^{(\rho-\frac{1}{2})t}$ *under Abel weights, and (iii) the interior-pole obstruction of Lemma 5.2 together with the holomorphic–meromorphic incompatibility Lemma 5.3. No programmatic interpretation (energy, teleology, cosmology) enters the proof.*

6 Programmatic layer: ground state, energy, and discrepancy (not used in proofs)

This section offers an interpretation-layer reading of the closed-layer rigidity. Nothing here is used as an input to Theorem 5.4.

6.1 RH as "no exponential gain/loss"

The mode decomposition

$$\mathrm{e}^{(\rho-\frac{1}{2})t} = \mathrm{e}^{(\beta-\frac{1}{2})t}\,\mathrm{e}^{\mathrm{i}\gamma t}$$

suggests a physical analogy: $\beta - \frac{1}{2}$ is a gain/loss rate and γ is a frequency. The critical line $\beta = \frac{1}{2}$ is precisely the locus of purely oscillatory modes, compatible with unitary scan evolution and bounded readout. From this point of view, Theorem 5.4 says that *protocol admissibility forces the zeta spectrum to be "ground-state stable"* in the sense of excluding exponential amplification under canonical Abel regularization.

6.2 Discrepancy as an auditable proxy for finite-resource "energy"

In HPA-style finite-resource analysis, deterministic error certificates often take the form

$$\left| \langle f \rangle_N - \int f \right| \;\leq\; \mathrm{Var}_{\mathrm{HK}}(f)\, D_N^*(P_N),$$

a Koksma–Hlawka type inequality controlling sampling error by the product of a kernel regularity measure and the star discrepancy of the scan orbit [4, 5]. The golden branch provides explicit logarithmic bounds on $D_N^*(P_N)$ (Section 2.4), suggesting a low-complexity "low-energy" scan family.

In this programmatic reading, discrepancy controls a provable resource cost for approximating regulated traces. Exponential modes induced by $\beta \neq \frac{1}{2}$ are precisely the kind of behavior that defeats such controlled, compact readout: they trigger internal Abel thresholds (Lemma 5.2) and hence force blow-ups in any reasonable finite-resource proxy. The toy numerics in Section 7 illustrate this contrast.

7 Reproducible numerics: discrepancy stability and "off-line zero" threshold blow-up

This section provides optional reproducible numerical illustrations consistent with the paper's closed-layer mechanism. No accepted mathematical fact in the main argument relies on computation: the rigidity theorem is purely analytic once HTF is assumed. The role of these examples

N	$D_N^*(P_N)$	bound $2(2+\log_\varphi N)/N$	ratio
1,000	1.347×10^{-3}	3.271×10^{-2}	0.0412
5,000	4.75×10^{-4}	7.88×10^{-3}	0.0603
10,000	2.908×10^{-4}	4.228×10^{-3}	0.0688
50,000	5.179×10^{-5}	9.794×10^{-4}	0.0529

Table 1: Exact one-dimensional star discrepancy for the golden-branch Kronecker sequence and the explicit logarithmic bound. The ratio column is D_N^*/bound.

is only to visualize the contrast between (i) logarithmically controlled scan irregularity and (ii) exponential-gain modes that destroy Abel admissibility.

7.1 Reproducibility protocol

All experiments are implemented in pure Python 3 with no third-party dependencies. To reproduce the tables, run the following scripts from this paper directory:

- `python3scripts/exp_golden_discrepancy.py`
- `python3scripts/exp_toy_zero_modes.py`

The scripts write LaTeX row files into `sections/generated/`, which are included below. All randomness is absent (deterministic computations only).

7.2 Experiment A: one-dimensional golden scan star discrepancy

We take $\alpha = \varphi^{-1}$ and the point set $P_N = \{\{t\alpha\}\}_{t=0}^{N-1} \subset [0,1)$. The script computes the one-dimensional star discrepancy $D_N^*(P_N)$ exactly from its definition and compares it with the explicit bound

$$D_N^*(P_N) \le \frac{2(2+\log_\varphi N)}{N}.$$

7.3 Experiment B: toy "zero-mode" signal and the Abel threshold

Let γ_k be the imaginary parts of the first K nontrivial zeros on the critical line (hard-coded constants in the script). We form a toy "critical-line" signal

$$e(t) = \sum_{k=1}^{K} \cos(\gamma_k t),$$

and then introduce an artificial real-part shift $\delta > 0$ on the first mode:

$$e_\delta(t) = \mathrm{e}^{\delta t} \cos(\gamma_1 t) + \sum_{k=2}^{K} \cos(\gamma_k t).$$

We monitor two proxies:

Discrete energy growth. Define the discrete L^2 energy up to horizon T by

$$E_\delta(T) = \sum_{t=0}^{T-1} |e_\delta(t)|^2.$$

Pure oscillatory signals ($\delta = 0$) exhibit at most linear growth, while $\delta > 0$ forces exponential growth in T.

T	δ	$E_\delta(T)$	$\max_{t<T} \lvert e_\delta(t)\rvert$
50	0.00	9.173×10^2	2×10^1
50	0.05	1.519×10^3	2×10^1
50	0.10	4.508×10^4	1.144×10^2
100	0.00	1.391×10^3	2×10^1
100	0.05	9.96×10^4	1.308×10^2
100	0.10	9.99×10^8	1.752×10^4
150	0.00	1.932×10^3	2×10^1
150	0.05	1.5×10^7	1.531×10^3
150	0.10	2.234×10^{13}	2.504×10^6

Table 2: Toy-mode energy growth. Exponential gain ($\delta > 0$) produces rapidly increasing energy and peak amplitude.

Lemma 7.1 (Closed-form growth of a single amplified mode)**.** *Let* $u_t = \mathrm{e}^{\delta t}\cos(\gamma t)$ *with* $\delta > 0$*. Then for every integer* $T \geq 1$*,*

$$\sum_{t=0}^{T-1} u_t^2 = \frac{1}{2}\sum_{t=0}^{T-1} \mathrm{e}^{2\delta t} \;+\; \frac{1}{2}\sum_{t=0}^{T-1} \mathrm{e}^{2\delta t}\cos(2\gamma t) = \frac{\mathrm{e}^{2\delta T}-1}{2(\mathrm{e}^{2\delta}-1)} \;+\; \frac{1}{2}\,\mathrm{Re}\left(\frac{1-(\mathrm{e}^{2\delta+2\mathrm{i}\gamma})^T}{1-\mathrm{e}^{2\delta+2\mathrm{i}\gamma}}\right).$$

In particular, the leading growth is $\Theta(\mathrm{e}^{2\delta T})$ *as* $T \to \infty$*.*

Proof. Use $\cos^2(\gamma t) = \frac{1}{2}(1+\cos(2\gamma t))$ and sum the resulting geometric series in complex form. □

Abel threshold. For integer time sampling, consider the finite-horizon Abel sum

$$S_\delta(r; T_{\max}) = \sum_{t=0}^{T_{\max}-1} r^t e_\delta(t).$$

The closed-layer prediction behind Lemma 5.2 is that the effective threshold is

$$r_0 = \mathrm{e}^{-\delta},$$

since $r^t \mathrm{e}^{\delta t} = (r\mathrm{e}^\delta)^t$. For the single complex mode $z_t := \mathrm{e}^{(\delta+\mathrm{i}\gamma_1)t}$ one has the exact geometric identity

$$\sum_{t=0}^{T_{\max}-1} r^t z_t = \sum_{t=0}^{T_{\max}-1} \left(r\,\mathrm{e}^{\delta+\mathrm{i}\gamma_1}\right)^t = \frac{1-\left(r\,\mathrm{e}^{\delta+\mathrm{i}\gamma_1}\right)^{T_{\max}}}{1-r\,\mathrm{e}^{\delta+\mathrm{i}\gamma_1}},$$

and for the real cosine mode $\mathrm{e}^{\delta t}\cos(\gamma_1 t) = \mathrm{Re}(z_t)$ one has

$$\sum_{t=0}^{T_{\max}-1} r^t \mathrm{e}^{\delta t}\cos(\gamma_1 t) = \mathrm{Re}\left(\frac{1-\left(r\,\mathrm{e}^{\delta+\mathrm{i}\gamma_1}\right)^{T_{\max}}}{1-r\,\mathrm{e}^{\delta+\mathrm{i}\gamma_1}}\right).$$

so the transition at $r\mathrm{e}^\delta = 1$ separates decay ($r < \mathrm{e}^{-\delta}$) from exponential growth in $T_{\max}$ ($r > \mathrm{e}^{-\delta}$), independent of any analytic continuation considerations. The script reports the magnitude $|S_\delta(r; T_{\max})|$ as r crosses r_0.

δ	r	$\lvert S_\delta(r;T_{\max})\rvert$
0.10	0.80	2.5675×10^{1}
0.10	0.90	2.6108×10^{1}
0.10	0.92	1.1153×10^{14}
0.10	0.94	5.359×10^{32}
0.10	0.96	1.0404×10^{51}

Table 3: Toy Abel threshold behavior at fixed $T_{\max}$. Values remain moderate for $r < \mathrm{e}^{-\delta}$ and grow rapidly once r crosses the predicted threshold.

8 Conclusion: RH as a protocol-stable ground-state condition

Within the HPA–Omega scan–readout paradigm, we formulated a closed derivation chain in which the classical explicit formula for ζ is treated as a holographic trace identity. Assuming a structural bridge (HTF) that embeds the spectral side into an Abel-regularized scan trace and assuming Omega's admissibility conditions (finite information, unitary scan, bounded readout, and canonical finite-part extraction along $r \uparrow 1$), we proved a simple rigidity theorem: any zero with $\mathrm{Re}(\rho) \neq \frac{1}{2}$ forces an internal Abel convergence barrier below 1, contradicting geometric-side definability on all of $(0,1)$. Thus RH becomes a protocol consequence of "polar rigidity" once the HTF bridge exists.

The closed-layer contribution of this paper is intentionally narrow: it clarifies which minimal structural inputs force RH and which part remains genuinely hard (constructing an HTF-compatible infinite-dimensional trace class). A natural next step is to make HTF constructive by exhibiting explicit transfer operators or trace identities compatible with the Omega finite-part convention, thereby compressing the protocol implication into a classical unconditional proof.

A Audit closure checklist (this version)

This appendix summarizes, in audit form, what is assumed in the closed layer and what is derived.

A.1 Closed-layer inputs (assumptions)

Omega axioms used.

- O2: finite information / compact readout (Axiom 3.1).
- O5: scan–projection readout (Axiom 3.2).
- O6: unitary scan algebra / Weyl pair (Axiom 3.3).
- Canonical Abel regularization and finite-part admissibility along $r \uparrow 1$ (Section 2.5; Appendix C records a sufficient existence criterion and uniqueness).

HTF structural bridge.

- Holographic Trace Formula assumption (Assumption 5.1): the Abel-weighted scan trace matches a Weil-type explicit formula under the same regularization, and zeros contribute as sampled exponential modes $\mathrm{e}^{(\rho-\frac{1}{2})t}$, equivalently via meromorphic mode factors $(1 - r\,\mathrm{e}^{(\rho-\frac{1}{2})})^{-1}$, in a way compatible with the geometric-side holomorphy on $|r| < 1$.

A.2 Closed-layer derivations (theorems)

- Lemma 5.2: if $\mathrm{Re}(\rho) > \frac{1}{2}$ then the mode factor $(1 - r\,\mathrm{e}^{(\rho-\frac{1}{2})})^{-1}$ has an interior pole at a point r_ρ with $|r_\rho| < 1$.
- Theorem 5.4: Omega geometric-side holomorphy on the unit disk $|r| < 1$ plus HTF mode structure implies $\mathrm{Re}(\rho) = \frac{1}{2}$ for all nontrivial zeros.

A.3 Programmatic layer (not used in proofs)

- "Energy" or "ground state" interpretations of the rigidity.
- Discrepancy-based intuitions beyond the deterministic, explicitly stated bounds.
- Cosmological, teleological, or narrative mappings of the protocol language.

B Finite-state zeta prototypes and why infinite-state lift is unavoidable

This appendix expands on the technical point used in Section 4.2: finite-state constructions yield rational zeta functions and thus cannot reproduce the nontrivial zero spectrum of $\zeta(s)$.

B.1 Finite-state dynamical zeta is rational

For a subshift of finite type with adjacency matrix A, the Artin–Mazur zeta function satisfies

$$\zeta_\sigma(z) = \frac{1}{\det(I - zA)}.$$

Since A is finite-dimensional, $\det(I - zA)$ is a polynomial and ζ_σ is rational. Therefore ζ_σ has only finitely many poles/zeros (counted with multiplicity), reflecting the finite-state nature of the underlying symbolic dynamics.

B.2 Nontrivial zeta zeros require an infinite-dimensional mechanism

The Riemann zeta function has infinitely many nontrivial zeros in the critical strip and they exhibit fine statistical structure. Any attempt to represent these zeros as a spectrum of an operator (Hilbert–Pólya type) or as poles/zeros of a trace-class determinant requires an infinite-dimensional setting: transfer operators, trace formulas, or controlled limits of refining partitions.

The HTF assumption in Section 5.1 is precisely a statement that such an infinite-dimensional mechanism exists and is compatible with Omega's protocol regularization conventions. The closed-layer rigidity then reduces RH to an elementary analytic constraint: the geometric side is holomorphic on the unit disk (bounded unitary scan), while any off-critical spectral mode would introduce an interior pole, which is forbidden for holomorphic traces.

B.3 Finite parts, Abel regularization, and zeta-regularized determinants (directional note)

In many contexts, finite-part extractions along canonical regulators can be related to zeta-regularized traces and determinants via Mellin transforms and Tauberian/Abelian comparisons. This suggests a plausible technical route to construct HTF: identify a transfer operator (or a scan-induced operator) whose regulated trace matches the explicit formula and whose determinant encodes the zeta function. This paper does not carry out that construction; it only records the precise closed-layer compatibility conditions that such a construction must satisfy to imply RH.

C Orbit calculus, Abel means, and finite parts (rotation template)

This appendix makes precise the closed-layer regularization conventions used in the main text. The simplest template is an irrational circle rotation (or, more generally, a torus translation), where Abel weights and finite-part extractions admit concrete analytic formulas.

C.1 Abel generating functions are holomorphic on the unit disk

Proposition C.1 (Holomorphy and boundedness of Abel generating functions)**.** *Let $(u_t)_{t\geq 0}$ be a bounded complex sequence: $|u_t| \leq M$. Define its Abel generating function*

$$S(r) := \sum_{t=0}^{\infty} u_t r^t.$$

Then $S(r)$ converges absolutely for every $|r| < 1$, defines a holomorphic function on the open unit disk, and satisfies

$$|S(r)| \leq \frac{M}{1 - |r|}.$$

Proof. Absolute convergence for $|r| < 1$ follows from domination by $\sum_{t\geq 0} M|r|^t$. Holomorphy is immediate since S is a power series with radius of convergence at least 1. The bound is the geometric-series estimate. □

C.2 Abel means for uniquely ergodic rotations

Let $X = \mathbb{R}/\mathbb{Z}$ and fix $\alpha \in (0,1) \setminus \mathbb{Q}$. Write $T_\alpha(x) = x + \alpha \pmod 1$.

Theorem C.2 (Orbit trace for an irrational rotation)**.** *If $f : X \to \mathbb{C}$ is continuous and $\alpha \notin \mathbb{Q}$, then the Cesàro orbit average*

$$A_N(f; x_0) := \frac{1}{N} \sum_{t=0}^{N-1} f\big(T_\alpha^t(x_0)\big)$$

converges as $N \to \infty$ to the Lebesgue integral $\int_X f \, \mathrm{d}x$, uniformly in x_0.

Reference. This is a standard consequence of unique ergodicity/equidistribution of irrational rotations; see, e.g., [4, 8, 13]. □

Theorem C.3 (Abel means converge to the space average)**.** *If $f : X \to \mathbb{C}$ is continuous and $\alpha \notin \mathbb{Q}$, then for every $x_0 \in X$,*

$$\lim_{r\uparrow 1} (1 - r) \sum_{t=0}^{\infty} r^t f\big(T_\alpha^t(x_0)\big) = \int_X f(x) \, \mathrm{d}x,$$

and the convergence is uniform in x_0.

Reference. One route is: unique ergodicity gives convergence of Cesàro means $A_N(f; x_0) \to \int f$ uniformly in x_0 (Theorem C.2), and Abel means preserve limits of Cesàro means (Abel–Cesàro comparison); see [8, 9]. □

C.3 Abel finite part as a constant term (uniqueness and a Fourier criterion)

Let $u_t := f(T_\alpha^t(x_0))$ and

$$S_f(r) := \sum_{t=0}^{\infty} u_t r^t, \qquad 0 < r < 1.$$

If f is bounded then $S_f(r)$ is well-defined for $0 < r < 1$. Following a standard Abel regularization convention [9], we treat the constant term in the asymptotic expansion of $S_f(r)$ as the finite part.

Definition C.4 (Abel finite part)**.** *Assume that as* $r \uparrow 1$ *the function* $S_f(r)$ *admits an asymptotic expansion*

$$S_f(r) = \sum_{j=1}^{J} \frac{c_j}{(1-r)^j} + C + o(1).$$

Then we define the Abel finite part by $\mathrm{FP}_\alpha(f) := C$*.*

Lemma C.5 (Uniqueness of the finite part)**.** *If* $S_f(r)$ *admits two asymptotic expansions of the form in Definition C.4 with constants* C *and* C'*, then* $C = C'$*.*

Proof. Subtract the two expansions; an induction on the highest pole order shows that all pole coefficients agree, hence the constant terms agree. □

For rotations, Abel sums admit an explicit Fourier–resolvent representation. Write Fourier coefficients of f as

$$\hat{f}(m) := \int_X f(x)\, \mathrm{e}^{-2\pi \mathrm{i} m x}\, \mathrm{d}x, \qquad m \in \mathbb{Z}.$$

Proposition C.6 (Fourier resolvent formula for Abel orbit sums)**.** *Assume* $\sum_{m\in\mathbb{Z}} |\hat{f}(m)| < \infty$*. Then for every* $|r| < 1$*,*

$$S_f(r) = \sum_{m\in\mathbb{Z}} \frac{\hat{f}(m)\, \mathrm{e}^{2\pi \mathrm{i} m x_0}}{1 - r\, \mathrm{e}^{2\pi \mathrm{i} m \alpha}}.$$

Proof. Under the absolute Fourier summability hypothesis, the Fourier series converges absolutely and uniformly, allowing termwise summation of the geometric series $\sum_{t\ge 0} (r\mathrm{e}^{2\pi \mathrm{i} m\alpha})^t$. □

Corollary C.7 (Simple pole and explicit finite part under a Diophantine/Fourier criterion)**.** *Assume the hypotheses of Proposition C.6 and in addition*

$$\sum_{m\in\mathbb{Z}\setminus\{0\}} \frac{|\hat{f}(m)|}{|1 - \mathrm{e}^{2\pi \mathrm{i} m\alpha}|} < \infty.$$

Then as $r \uparrow 1$*,*

$$S_f(r) = \frac{\hat{f}(0)}{1-r} \;+\; \sum_{m\in\mathbb{Z}\setminus\{0\}} \frac{\hat{f}(m)\, \mathrm{e}^{2\pi \mathrm{i} m x_0}}{1 - \mathrm{e}^{2\pi \mathrm{i} m\alpha}} \;+\; o(1).$$

In particular, $\mathrm{FP}_\alpha(f)$ *exists and equals the explicit constant term.*

C.4 Golden-branch star discrepancy bound via Denjoy–Koksma and Ostrowski

We record an explicit one-dimensional star-discrepancy bound used in Section 2.4. Let $P_N = \{\{x_0 + t\alpha\}\}_{t=0}^{N-1} \subset [0, 1)$.

Theorem C.8 (Ostrowski/Denjoy–Koksma bound for interval counts)**.** *Let $\alpha = [0; a_1, a_2, \ldots] \in (0,1)\setminus\mathbb{Q}$ and let q_k be the continued-fraction convergent denominators. Let $W \subset X$ be an interval and set $s_t = \mathbb{1}_W(x_0 + t\alpha)$ and $S_N = \sum_{t=0}^{N-1} s_t$. For every $N \geq 1$ write the Ostrowski expansion $N = \sum_{k=0}^{m} b_k q_k$ with admissible digits. Then*

$$\left|S_N - N\,\mu(W)\right| \;\leq\; 2\sum_{k=0}^{m} b_k,$$

uniformly in x_0 and W.

Reference. This is a standard consequence of Denjoy–Koksma at convergent times plus the Ostrowski block decomposition; see [6, 7, 14]. □

Corollary C.9 (Explicit bound for the golden branch)**.** *Let $\alpha = \varphi^{-1} = [0; 1, 1, 1, \ldots]$. Then for every $N \geq 1$,*

$$D_N^*(P_N) \leq \frac{2(2 + \log_\varphi N)}{N}.$$

Proof. For one-dimensional star discrepancy, it suffices to take $W = [0, u)$ in Theorem C.8 and divide by N, giving

$$D_N^*(P_N) \leq \frac{2}{N}\sum_{k=0}^{m} b_k.$$

For the golden branch all continued-fraction digits satisfy $a_k = 1$, hence all Ostrowski digits satisfy $b_k \in \{0, 1\}$, so $\sum_{k=0}^{m} b_k \leq m + 1$. Moreover, for the golden branch the convergent denominators are Fibonacci numbers $q_m = F_{m+1}$, and one has $q_m \geq \varphi^{m-1}$ for $m \geq 1$. Since $q_m \leq N$, it follows that $m \leq 1 + \log_\varphi N$, hence $m + 1 \leq 2 + \log_\varphi N$. Combining the bounds yields the stated inequality. □

D Reproducible experiment code (pure Python)

D.1 Experiment A: exact one-dimensional star discrepancy for the golden scan

```
#!/usr/bin/env python3
# -*- coding: utf-8 -*-
"""
Experiment A: exact 1D star discrepancy for the golden-branch Kronecker scan.

This script writes a LaTeX table row file into:
  sections/generated/golden_discrepancy_rows.tex

No third-party dependencies.
"""

from __future__ import annotations

import math
from pathlib import Path
from typing import List, Sequence


def golden_points(n: int) -> List[float]:
    phi = (1.0 + math.sqrt(5.0)) / 2.0
    alpha = 1.0 / phi  # golden branch
    return [(t * alpha) % 1.0 for t in range(n)]
```

```python
def star_discrepancy_1d(points: Sequence[float]) -> float:
    """
    Exact 1D star discrepancy:
      D_N^* = sup_{u in [0,1]} | (1/N) #{x_i < u} - u |
    computed from sorted points.
    """
    n = len(points)
    if n <= 0:
        return 0.0
    xs = sorted(points)
    inv_n = 1.0 / float(n)
    d_plus = 0.0
    d_minus = 0.0
    for i, x in enumerate(xs):
        a = (float(i + 1) * inv_n) - x
        b = x - (float(i) * inv_n)
        if a > d_plus:
            d_plus = a
        if b > d_minus:
            d_minus = b
    return max(d_plus, d_minus)


def bound_gold(n: int) -> float:
    phi = (1.0 + math.sqrt(5.0)) / 2.0
    return 2.0 * (2.0 + math.log(float(n), phi)) / float(n)


def fmt_sci_unsigned(x: float, sig: int = 4) -> str:
    """LaTeX scientific notation without sign, for nonnegative quantities."""
    if x <= 0.0:
        return "$0$"
    exp = int(math.floor(math.log10(x)))
    mant = x / (10.0 ** exp)
    mant = round(mant, max(sig - 1, 0))
    if mant >= 10.0:
        mant /= 10.0
        exp += 1
    mant_str = f"{mant:.{max(sig - 1, 0)}f}".rstrip("0").rstrip(".")
    return f"${mant_str}\\times 10^{{{exp}}}$"


def fmt_decimal(x: float, digits: int = 4) -> str:
    return f"${x:.{digits}f}$"


def write_rows(path: Path, lines: List[str]) -> None:
    path.parent.mkdir(parents=True, exist_ok=True)
    out = list(lines)
    if out:
        last = out[-1].rstrip()
        if last.endswith("\\\\"):
            last = last[:-2].rstrip()
        out[-1] = last
    path.write_text("\n".join(out).rstrip() + "\n", encoding="utf-8")
```

```
def main() -> None:
    root = Path(__file__).resolve().parent.parent
    gen = root / "sections" / "generated"

    ns = [1000, 5000, 10000, 50000]
    rows: List[str] = []
    for n in ns:
        pts = golden_points(n)
        dstar = star_discrepancy_1d(pts)
        bound = bound_gold(n)
        ratio = dstar / bound if bound > 0.0 else 0.0
        rows.append(
            f"{n:,} & {fmt_sci_unsigned(dstar)} & {fmt_sci_unsigned(bound)} &
            ↪  {fmt_decimal(ratio, digits=4)} \\\\"
        )

    write_rows(gen / "golden_discrepancy_rows.tex", rows)
    print(f"Wrote LaTeX rows into: {gen}")
    print("File: golden_discrepancy_rows.tex")


if __name__ == "__main__":
    main()
```

D.2 Experiment B: toy "zero-mode" signal, energy growth, and Abel thresholds

```
#!/usr/bin/env python3
# -*- coding: utf-8 -*-
"""
Experiment B: toy "zero-mode" signal, discrete energy growth, and Abel thresholds.

We build a signal from the imaginary parts of the first K nontrivial zeta zeros on
the critical line (hard-coded constants), then introduce an artificial real-part
shift delta>0 on a single mode and observe:
  (i) rapid energy growth in T, and
  (ii) an Abel threshold near r0 = exp(-delta).

This script writes LaTeX table row files into:
  sections/generated/toy_energy_rows.tex
  sections/generated/toy_abel_rows.tex

No third-party dependencies.
"""

from __future__ import annotations

import math
from pathlib import Path
from typing import List, Sequence, Tuple


# Imaginary parts gamma_k of the first 20 nontrivial zeros 1/2 + i*gamma_k.
# Values are standard numerical constants; high precision is not required here.
GAMMAS_20: List[float] = [
```

```
    14.134725141734693,
    21.022039638771555,
    25.010857580145689,
    30.424876125859513,
    32.935061587739190,
    37.586178158825671,
    40.918719012147495,
    43.327073280914999,
    48.005150881167160,
    49.773832477672302,
    52.970321477714461,
    56.446247697063395,
    59.347044002602353,
    60.831778524609810,
    65.112544048081607,
    67.079810529494174,
    69.546401711173979,
    72.067157674481908,
    75.704690699083933,
    77.144840068874805,
]


def signal(t: int, gammas: Sequence[float], delta: float = 0.0, idx: int = 0) ->
↪  float:
    """
    e_delta(t) = exp(delta*t)*cos(gamma_idx * t) + sum_{k != idx} cos(gamma_k * t).
    """
    s = 0.0
    for j, g in enumerate(gammas):
        amp = math.exp(delta * float(t)) if (delta != 0.0 and j == idx) else 1.0
        s += amp * math.cos(g * float(t))
    return s


def energy_discrete(T: int, gammas: Sequence[float], delta: float = 0.0) ->
↪  Tuple[float, float]:
    """
    Discrete energy proxy: E(T) = sum_{t=0}^{T-1} |e_delta(t)|^2.
    Returns (energy, max_abs).
    """
    e2 = 0.0
    m = 0.0
    for t in range(T):
        v = signal(t, gammas, delta=delta)
        av = abs(v)
        if av > m:
            m = av
        e2 += v * v
    return e2, m


def abel_partial_sum(Tmax: int, r: float, gammas: Sequence[float], delta: float = 0.0)
↪  -> float:
    """
    Finite-horizon Abel sum S_delta(r;Tmax) = sum_{t=0}^{Tmax-1} r^t e_delta(t).
    """
    s = 0.0
```

```
    wt = 1.0
    for t in range(Tmax):
        s += wt * signal(t, gammas, delta=delta)
        wt *= r
    return s


def fmt_sci_unsigned(x: float, sig: int = 4) -> str:
    """LaTeX scientific notation without sign, for nonnegative quantities."""
    x = float(x)
    if x <= 0.0:
        return "$0$"
    exp = int(math.floor(math.log10(x)))
    mant = x / (10.0 ** exp)
    mant = round(mant, max(sig - 1, 0))
    if mant >= 10.0:
        mant /= 10.0
        exp += 1
    mant_str = f"{mant:.{max(sig - 1, 0)}f}".rstrip("0").rstrip(".")
    return f"${mant_str}\\times 10^{{{exp}}}$"


def fmt_decimal(x: float, digits: int = 2) -> str:
    return f"${x:.{digits}f}$"


def write_rows(path: Path, lines: List[str]) -> None:
    path.parent.mkdir(parents=True, exist_ok=True)
    out = list(lines)
    if out:
        last = out[-1].rstrip()
        if last.endswith("\\\\"):
            last = last[:-2].rstrip()
        out[-1] = last
    path.write_text("\n".join(out).rstrip() + "\n", encoding="utf-8")


def main() -> None:
    root = Path(__file__).resolve().parent.parent
    gen = root / "sections" / "generated"

    gammas = GAMMAS_20
    Ts = [50, 100, 150]
    deltas = [0.00, 0.05, 0.10]

    energy_rows: List[str] = []
    for T in Ts:
        for delta in deltas:
            E, M = energy_discrete(T, gammas, delta=delta)
            energy_rows.append(
                f"{T:d} & {fmt_decimal(delta, digits=2)} & {fmt_sci_unsigned(E)} &
                ↪  {fmt_sci_unsigned(M)} \\\\"
            )
    write_rows(gen / "toy_energy_rows.tex", energy_rows)

    delta = 0.10
    Tmax = 2000
    rs = [0.80, 0.90, 0.92, 0.94, 0.96]
```

```
    abel_rows: List[str] = []
    for r in rs:
        S = abel_partial_sum(Tmax, r, gammas, delta=delta)
        abel_rows.append(
            f"{fmt_decimal(delta, digits=2)} & {fmt_decimal(r, digits=2)} &
             ↪  {fmt_sci_unsigned(abs(S), sig=5)} \\\\"
        )
    write_rows(gen / "toy_abel_rows.tex", abel_rows)

    print(f"Wrote LaTeX rows into: {gen}")
    print("Files: toy_energy_rows.tex, toy_abel_rows.tex")


if __name__ == "__main__":
    main()
```

References

[1] E. C. Titchmarsh. *The Theory of the Riemann Zeta-Function.* Oxford University Press, Oxford, 2 edition, 1986. Revised by D. R. Heath-Brown.

[2] Harold M. Edwards. *Riemann's Zeta Function.* Academic Press, New York, 1974.

[3] Henryk Iwaniec and Emmanuel Kowalski. *Analytic Number Theory*, volume 53 of *American Mathematical Society Colloquium Publications.* American Mathematical Society, Providence, RI, 2004.

[4] L. Kuipers and Harald Niederreiter. *Uniform Distribution of Sequences.* Wiley, New York, 1974.

[5] Michael Drmota and Robert F. Tichy. *Sequences, Discrepancies and Applications*, volume 1651 of *Lecture Notes in Mathematics.* Springer, Berlin, 1997.

[6] Arnaud Denjoy. Sur les courbes définies par les équations différentielles à la surface du tore. *Journal de Mathématiques Pures et Appliquées*, 11:333–375, 1932.

[7] J. F. Koksma. Ein mengentheoretischer satz über die gleichverteilung modulo eins. *Compositio Mathematica*, 2:250–258, 1935.

[8] Peter Walters. *An Introduction to Ergodic Theory.* Springer, New York, 1982.

[9] G. H. Hardy. *Divergent Series.* Oxford University Press, Oxford, 1949.

[10] Michael Artin and Barry Mazur. On periodic points. *Annals of Mathematics*, 81(1):82–99, 1965.

[11] David Ruelle. Zeta-functions for expanding maps and anosov flows. *Inventiones Mathematicae*, 34:231–242, 1976.

[12] Alain Connes. Trace formula in noncommutative geometry and the zeros of the riemann zeta function. *Selecta Mathematica (New Series)*, 5:29–106, 1999.

[13] Hermann Weyl. Über die gleichverteilung von zahlen mod. eins. *Mathematische Annalen*, 77:313–352, 1916.

[14] Anatole Katok and Boris Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, Cambridge, 1995.