The Arithmetic Spectral Triple: Information Geometry Meets Noncommutative Geometry on the Zeta Distribution

D. Christian[1,*]

[1] *The Axioms of Pattern Ontology (APO) Initiative*

(Dated: March 24, 2026)

We construct a spectral triple $(\mathcal{A}, \mathfrak{H}, D_F)$ on the statistical manifold of the zeta distribution $P_s(n) = n^{-s}/\zeta(s)$, where D_F is the Dirac operator of the Chentsov–Fisher metric. We prove that Connes' distance formula applied to this triple recovers the Fisher–Rao geodesic distance exactly, establishing that the Bhattacharyya coefficient—the arithmetic recognition operator $\odot(s, s')$ of the companion paper [1]—is the cosine of half the spectral distance: $\odot = \cos(d_{\mathrm{FR}}/2)$.

Rigorous core. The spectral triple satisfies Connes' axioms 1–6 (PROVEN). The Dixmier trace recovers the Fisher volume form (PROVEN). We prove that every simple zero of ζ creates an *identical* conical singularity in the Fisher metric, with universal coefficient $A = 1$ independent of the zero's height (Theorem VI.6; PROVEN, confirmed numerically to CV = 0.01%). This places all zeros on the same geometric footing: they are interchangeable singularities.

Numerical results. We report that the Chamseddine–Connes spectral action is *maximized* on the critical line (tested at five cutoffs), that the total Fisher curvature is most negative there, and that the Fisher information magnitude $|\mathcal{I}|$ peaks at $\sigma = \frac{1}{2}$ near every zero. We also report that the eigenvalues of D_F are *not* the zeta zeros—they are chamber modes between consecutive zeros, which act as metric barriers—correcting a natural but false expectation.

Variational conjecture. We conjecture that among all $\mathbb{Z}_2$-symmetric 2D conformal metrics with identical conical singularities of deficit angle $A = 1$, the curvature functional is extremized if and only if all singularities lie on the symmetry axis. This reformulates the Riemann Hypothesis as a problem in 2D conformal geometry with prescribed singularities.

Structural bridges. We argue that $\odot$ is a propagator on the statistical bundle (ARGUED), reformulate the Bhattacharyya–Bessel correspondence as a unitary equivalence question between spectral triples (ARGUED), and conjecture equivalence between APO's recognition completeness and Connes' Weil positivity (CONJECTURED).

Every claim is labeled PROVEN, ARGUED, CONJECTURED, or OPEN throughout. †APO-specific terminology; see Appendix D for formal definitions and foundational status.

CONTENTS

* daniel@freereason.org

I. INTRODUCTION

In the companion paper [1] (hereafter Paper I), we constructed the *arithmetic recognition operator*

$$\odot(s, s') = \frac{\zeta\left(\frac{s+s'}{2}\right)}{\sqrt{\zeta(s)\,\zeta(s')}} \tag{1}$$

on the statistical manifold of the zeta distribution $P_s(n) = n^{-s}/\zeta(s)$, proved its principal properties, identified the modular surface as the natural geometric home, and traced the precise gap between these constructions and a proof of the Riemann Hypothesis. That gap appeared in six equivalent guises, all manifestations of a single structural mismatch: the framework provides global, statistical control while RH requires local, analytic precision. The most promising closing mechanism was the *Bhattacharyya–Bessel correspondence*—a conjectured unitary intertwining between the recognition operator and Mayer's transfer operator for the Gauss map—which remained open, with three identified sub-gaps.

This paper brings a new set of tools to bear. Alain Connes' program of *noncommutative geometry* (NCG) [2] provides a framework in which the Riemannian geometry of a manifold—its metric, volume, gauge potentials, and curvature—is reconstructed from operator-theoretic data: a *spectral triple* $(\mathcal{A}, \mathfrak{H}, D)$ consisting of an algebra, a Hilbert space, and a Dirac operator. The reconstruction is achieved through four formulas, of which the first two are central here:

- **Connes' distance formula.** $d(p, q) = \sup\{|a(p) - a(q)| : a \in \mathcal{A},\ \|[D, a]\| \leq 1\}$ recovers the geodesic distance from the commutator norm [2].
- **The Dixmier trace formula.** $\int_M f\, dv = c(d)\,\mathrm{Tr}_\omega(f\,|D|^{-d})$ recovers the volume form from the singular trace [2].

Connes' distance formula is *dual* to the classical path-based definition of geodesic distance. Instead of taking the infimum over arcs γ from p to q, it takes the supremum over functions a that probe the space. As Connes observes [2], this duality is essential: "in the case of discrete spaces or of noncommutative spaces X, there are no interesting arcs in X but there are plenty of functions." The arithmetic statistical manifold is not discrete, but the observation applies: the natural objects on $\mathcal{M}_\zeta$ are not paths but *test distributions*—elements of the algebra $\mathcal{A}$—and the distance between two parameter values s, s' is measured by how well any bounded observable can distinguish the corresponding distributions P_s and $P_{s'}$. This is precisely the Fisher–Rao distance.

The connection to Connes' broader program on the Riemann Hypothesis [3–5] is structural. Connes' spectral realization identifies the non-trivial zeros of ζ as an absorption spectrum on the adele class space $\mathbb{Q}^*\backslash\mathbb{A}_\mathbb{Q}/\hat{\mathbb{Z}}^*$, and reduces RH to the positivity of the Weil distribution [3]. In November 2025, Connes and Consani [5] constructed explicit self-adjoint operators—rank-one perturbations of spectral triples for the scaling operator—whose spectra reproduce the lowest zeta zeros with striking numerical accuracy. The present paper builds a *different* spectral triple on a *different* space (the Fisher–Rao manifold rather than the scaling site), but one that should be connected to the Connes–Consani construction through the shared involvement of the Euler product and the modular group.

A. What this paper does and does not do

Rigorous results (§II–§IV):

1. We construct the spectral triple $(\mathcal{A}, \mathfrak{H}, D_F)$ for the Fisher–Rao geometry of $\mathcal{M}_\zeta$ and verify Connes' axioms 1–6 (Theorem III.7).
2. We prove that Connes' distance formula recovers the Fisher–Rao distance exactly (Theorem IV.1).
3. We prove the identification $\odot(s, s') = \cos(d_{\mathrm{FR}}(s, s')/2)$, showing the recognition operator is the cosine of half the spectral distance (Corollary IV.3).
4. We show that the Dixmier trace recovers integration against the Fisher volume form (Theorem V.1).

Argued results (§VI–§VII):

5. We argue that $\odot$ is the Green's function of D_F restricted to the zero-mode sector (§VI).
6. The Euler product equips the spectral triple with a product geometry analogous to Connes' Standard Model factorization (§VII).
7. The three sub-gaps of the Bhattacharyya–Bessel bridge collapse to a single question: unitary equivalence of two spectral triples (§IX).

Conjectured results (§VIII–§X):

8. A morphism of spectral triples connects D_F to the operator on Connes' adele class space (§VIII).
9. APO's recognition completeness criterion is equivalent to Connes' Weil positivity criterion (§X).

What we do not claim. This paper does not prove the Riemann Hypothesis. The spectral triple construction is a reformulation of the information-geometric structure of Paper I, not a new proof strategy. Its value is that it connects the APO framework to Connes' NCG program, thereby (a) placing the APO constructions in a well-developed mathematical context, (b) reducing three open sub-gaps to one, and (c) identifying a precise equivalence between the remaining obstructions.

B. Guide to the paper

Section II reviews spectral triples and information geometry, establishing notation for both. Section III constructs the arithmetic spectral triple and verifies the axioms. Section IV proves the distance formula. Section V identifies the Dixmier trace with the Fisher volume form. Sections VI–X develop the structural and conjectural bridges. Section XI describes testable predictions. Section XII summarizes. Appendix A gives the detailed axiom verification, Appendix B the master proof status table, and Appendix C the APO–NCG–standard mathematics dictionary.

II. PRELIMINARIES: SPECTRAL TRIPLES AND INFORMATION GEOMETRY

This section establishes the two frameworks to be bridged. We state the definitions from noncommutative geometry and information geometry in parallel, highlighting the structural correspondence that the construction of §III will make precise.

A. Connes' spectral triples

We recall the framework of [2, 6].

Definition II.1 (Spectral triple)**.** *A* spectral triple $(\mathcal{A}, \mathfrak{H}, D)$ *consists of:*

1. *An involutive (possibly noncommutative) algebra $\mathcal{A}$ represented faithfully on a Hilbert space $\mathfrak{H}$;*
2. *A self-adjoint operator D on $\mathfrak{H}$ with compact resolvent, such that $[D, a]$ is bounded for all $a \in \mathcal{A}$.*

The triple is even *if there exists a grading operator Γ on $\mathfrak{H}$ (with $\Gamma^2 = 1$, $\Gamma^* = \Gamma$) such that $\Gamma a = a\Gamma$ for all $a \in \mathcal{A}$ and $\Gamma D = -D\Gamma$.*

The prototypical example is a compact Riemannian spin manifold M: the algebra $\mathcal{A} = C^\infty(M)$ acts by multiplication on the Hilbert space $\mathfrak{H} = L^2(M, S)$ of L^2-spinors, and $D = \not{\partial}_M$ is the Dirac operator. Knowing only $\mathcal{A}$ or only D gives essentially no geometric information ($\mathcal{A}$ determines the manifold only up to homeomorphism; D alone yields only the list of eigenvalues, governed by dimensional asymptotics $|\lambda_n| \sim C\, n^{1/d}$). The geometry is encoded in the *interplay* of $\mathcal{A}$ and D—the failure of elements of $\mathcal{A}$ to commute with D.

Connes' reconstruction theorem [7] shows that, conversely, a spectral triple satisfying certain regularity axioms determines a unique Riemannian geometry. The axioms we shall verify are:

Definition II.2 (Axioms for a spectral triple)**.** *Let $(\mathcal{A}, \mathfrak{H}, D)$ be a spectral triple.*

1. ***Dimension.*** *There exists $d \in \mathbb{N}$ (the* metric dimension*) such that $|D|^{-1}$ is in the Dixmier ideal $\mathcal{L}^{(d,\infty)}$: the eigenvalues of $|D|^{-1}$ satisfy $\mu_n(|D|^{-1}) = O(n^{-1/d})$.*

2. ***Regularity.*** *For every $a \in \mathcal{A}$, both a and $[D, a]$ lie in the domain of δ^k for all $k \geq 1$, where $\delta(T) = [|D|, T]$.*

3. ***Finiteness.*** *The space $\mathfrak{H}^\infty = \bigcap_{k=1}^\infty \mathrm{dom}(D^k)$ is a finite projective module over $\mathcal{A}$.*

4. ***Reality.*** *There exists an antilinear isometry $J : \mathfrak{H} \to \mathfrak{H}$ with $J^2 = \pm 1$, $JD = \pm DJ$, $J\Gamma = \pm \Gamma J$ (signs determined by the KO-dimension $d \bmod 8$).*

5. ***First-order condition.*** *$[[D, a], Jb^* J^{-1}] = 0$ for all $a, b \in \mathcal{A}$.*

6. ***Orientability.*** *There exists a Hochschild cycle c of degree d such that $\pi_D(c) = \Gamma$ (the grading, if d is even) or $\pi_D(c) = 1$ (if d is odd).*

The seventh axiom (Poincaré duality) requires a non-degenerate intersection form and will be discussed separately in the context of compactification (§III.8).

Theorem II.3 (Connes' distance formula [2, IV.2])**.** *Let $(\mathcal{A}, \mathfrak{H}, D)$ be the spectral triple of a compact Riemannian spin manifold M. For any $p, q \in M$,*

$$d(p, q) = \sup\big\{|a(p) - a(q)| : a \in \mathcal{A},\ \|[D, a]\| \leq 1\big\}, \tag{2}$$

where $d(p, q)$ is the geodesic distance.

The proof [2] rests on two facts. First, the commutator $[D, a]$ for the Dirac operator of a spin manifold is given by Clifford multiplication $[D, a] = i^{-1}\gamma(da)$, so $\|[D, a]\| = \operatorname{ess\,sup} \|da\|$, the Lipschitz norm of a. Second, the function $a(q') = d(q', p)$ is Lipschitz with constant 1, achieving the supremum.

Theorem II.4 (Dixmier trace formula [2, IV.2])**.** *For every $f \in C^\infty(M)$,*

$$\int_M f\, dv = c(d)\, \mathrm{Tr}_\omega\big(f\, |D|^{-d}\big)\,, \tag{3}$$

where $d = \dim M$, $c(d) = 2^{(d-[d/2])}\, \pi^{d/2}\, \Gamma\big(\frac{d}{2} + 1\big)$, and Tr_ω is the Dixmier trace.

The Dixmier trace [2] is the unique (up to normalization) singular trace on the ideal of operators whose partial sums of eigenvalues grow logarithmically: $\mathrm{Tr}_\omega(T) = \lim_\omega \frac{1}{\log N} \sum_{j=0}^N \mu_j(T)$, where $\lim_\omega$ is a suitable generalized limit. Its key properties are positivity ($\mathrm{Tr}_\omega(T) \geq 0$ for $T \geq 0$), finiteness (bounded when $\sum_0^N \mu_n = O(\log N)$), covariance ($\mathrm{Tr}_\omega(UTU^*) = \mathrm{Tr}_\omega(T)$), and locality ($\mathrm{Tr}_\omega(T) = 0$ for trace-class T). The last property ensures that the Dixmier trace is insensitive to finite-rank perturbations.

B. Information geometry of statistical manifolds

We recall the framework of [8, 9]; see also Paper I, §2 for the arithmetic specialization.

Definition II.5 (Statistical manifold)**.** *A statistical manifold (M, g_F) is a smooth manifold M whose points parametrize a family of probability distributions $\{P_\theta\}_{\theta \in M}$ on a sample space $\mathcal{X}$, equipped with the* Fisher information metric

$$(g_F)_{ij}(\theta) = \sum_{x \in \mathcal{X}} P_\theta(x)\, \frac{\partial \log P_\theta(x)}{\partial \theta^i}\, \frac{\partial \log P_\theta(x)}{\partial \theta^j}. \tag{4}$$

Chentsov's theorem [10] establishes that g_F is the *unique* Riemannian metric on M (up to a positive constant) invariant under sufficient statistics—that is, under every information-preserving transformation of the sample space. There is no freedom to choose a different metric; the statistical structure forces it.

The *square-root embedding* $\psi_\theta(x) = \sqrt{P_\theta(x)}$ maps each distribution to a point on the positive orthant of the unit sphere in $L^2(\mathcal{X})$:

$$\sum_x |\psi_\theta(x)|^2 = \sum_x P_\theta(x) = 1. \tag{5}$$

Under this embedding, the Fisher metric is the pullback of the round metric on the sphere.

Definition II.6 (Bhattacharyya coefficient and Fisher–Rao distance)**.** *The* Bhattacharyya coefficient *is the inner product of the square-root embeddings:*

$$B(\theta, \theta') = \langle \psi_\theta, \psi_{\theta'} \rangle = \sum_x \sqrt{P_\theta(x)\, P_{\theta'}(x)}. \tag{6}$$

The Fisher–Rao distance *is*

$$d_{\mathrm{FR}}(\theta, \theta') = 2 \arccos\big(B(\theta, \theta')\big). \tag{7}$$

The factor of 2 is conventional: d_{FR} is the geodesic distance on the unit sphere in L^2, measured in the round metric with sectional curvature 1. Note the immediate consequence

$$B(\theta, \theta') = \cos\left(\frac{d_{\mathrm{FR}}(\theta, \theta')}{2}\right), \tag{8}$$

which will be central to our identification of the recognition operator with the spectral distance (§IV).

C. The arithmetic specialization

For the zeta distribution $P_s(n) = n^{-s}/\zeta(s)$ on $\mathbb{N}$ ($\mathrm{Re}(s) > 1$), all the structures above specialize as in Paper I:

$$\text{Statistical manifold:} \quad \mathcal{M}_\zeta = \{s \in \mathbb{C} : \mathrm{Re}(s) > 1\}, \tag{9}$$

$$\text{Fisher metric:} \quad g_F(s) = \mathcal{I}(s)\, |ds|^2 = \frac{d^2}{ds^2} \log \zeta(s)\, |ds|^2, \tag{10}$$

$$\text{Square-root embedding:} \quad \psi_s(n) = \frac{n^{-s/2}}{\sqrt{\zeta(s)}}, \tag{11}$$

$$\text{Recognition operator:} \quad \odot\,(s, s') = \frac{\zeta\left(\frac{s+s'}{2}\right)}{\sqrt{\zeta(s)\,\zeta(s')}}, \tag{12}$$

$$\text{Fisher–Rao distance:} \quad d_{\mathrm{FR}}(s, s') = 2 \arccos\big(\odot(s, s')\big). \tag{13}$$

We treat $s = \sigma + it$ as a real two-dimensional parameter, identifying $\mathcal{M}_\zeta$ with an open subset of $\mathbb{R}^2$ via coordinates (σ, t). The Fisher metric (10) then becomes a 2×2 matrix with components

$$g_{\sigma\sigma} = \mathrm{Re}\, \mathcal{I}(s), \quad g_{tt} = \mathrm{Re}\, \mathcal{I}(s), \quad g_{\sigma t} = -\, \mathrm{Im}\, \mathcal{I}(s), \tag{14}$$

where the real and imaginary parts arise from $\mathcal{I}(s) = \partial_\sigma^2 \log \zeta + i \cdot 0$ along the real axis, but acquire nontrivial structure in the full strip.

On the critical line $\sigma = \frac{1}{2}$, the $\mathbb{Z}_2$ symmetry $\sigma \leftrightarrow 1 - \sigma$ forces $g_{\sigma t} = 0$ (Paper I, Prop. 5.1), so the Fisher metric is *diagonal*: the σ and t directions decouple.

Remark II.7 (Extension to the critical strip)**.** *The zeta distribution P_s is a proper probability distribution only for* $\mathrm{Re}(s) > 1$*. In the critical strip* $0 < \mathrm{Re}(s) < 1$*, $\zeta(s)$ is defined by analytic continuation, but $P_s(n) = n^{-s}/\zeta(s)$ is no longer non-negative. The Bhattacharyya interpretation of $\odot$ as an inner product on the probability simplex breaks down.*

What survives: the expression $\odot(s, s') = \zeta((s + s')/2)/\sqrt{\zeta(s)\,\zeta(s')}$ is a well-defined meromorphic function on $\mathbb{C}^2$, and the Fisher metric $\mathcal{I}(s) = d^2/ds^2 \log \zeta(s)$ is well-defined wherever ζ has no zeros. The Hadamard expansion, the singularity structure, and the spectral triple construction all extend analytically.

What does not survive: the statistical interpretation (sampling from P_s, sufficient statistics, Bayesian updating). In the language of APO (Appendix D), Reflection^{APO} in the critical strip is analytic *(comparing meromorphic functions) rather than* statistical *(comparing probability measures). The geometric content—distances, curvature, singularities—is unchanged; the operational meaning shifts. All rigorous results in this paper that concern the critical strip use only the analytic structure, not the statistical interpretation.*

III. THE ARITHMETIC SPECTRAL TRIPLE: CONSTRUCTION

We now construct the spectral triple $(\mathcal{A}, \mathfrak{H}, D_F)$ for the arithmetic statistical manifold. The construction is standard Riemannian geometry: the Dirac operator of a spin manifold applied to $\mathcal{M}_\zeta$ with the Fisher metric. The novelty is not in the method but in the specific manifold and the consequences of applying Connes' reconstruction formulas to it.

A. The algebra

Definition III.1 (Arithmetic algebra)**.** *Define*

$$\mathcal{A} = \{a \in C^\infty(\mathcal{M}_\zeta) : a \text{ and all derivatives are bounded}\} \tag{15}$$

acting by pointwise multiplication on functions over $\mathcal{M}_\zeta$*.*

Since $\mathcal{M}_\zeta$ is the open half-plane $\{\sigma > 1\}$ (or, after meromorphic continuation, the open strip $\{0 < \sigma < 1\}$), we need bounded smooth functions to ensure that the representation on $\mathfrak{H}$ is well-defined. The C^*-closure of $\mathcal{A}$ is $A = C_0(\mathcal{M}_\zeta) + \mathbb{C} \cdot 1$ (continuous functions vanishing at infinity, plus constants), and the spectrum of A recovers $\mathcal{M}_\zeta$ (with its one-point compactification) by Gelfand's theorem.

Remark III.2 (Commutativity)**.** *The algebra $\mathcal{A}$ is commutative: we are doing ordinary Riemannian geometry on $\mathcal{M}_\zeta$, not noncommutative geometry. The noncommutativity will enter in two ways: (i) the product geometry of §VII, where the "finite factor" is noncommutative, and (ii) the bridge to Connes' adele class space (§VIII), which is intrinsically noncommutative. At this stage, the spectral triple is commutative, and the axiom verification is standard.*

B. The Hilbert space and spinor bundle

Since $\mathcal{M}_\zeta$ is two-dimensional (as a real manifold), a spin structure exists and is unique (the obstruction lives in $H^2(M; \mathbb{Z}/2)$, which vanishes for a contractible open subset of $\mathbb{R}^2$). The spinor bundle $S \to \mathcal{M}_\zeta$ has fiber $\mathbb{C}^2$ at each point (since $\dim_\mathbb{R} \mathcal{M}_\zeta = 2$, the Clifford algebra $\mathrm{Cl}(2) \cong M_2(\mathbb{C})$ and the irreducible module is $\mathbb{C}^2$).

Definition III.3 (Spinor Hilbert space)**.** *The Hilbert space is*

$$\mathfrak{H} = L^2(\mathcal{M}_\zeta, S; dv_F) = L^2\Big(\mathcal{M}_\zeta, \mathbb{C}^2; \sqrt{\det g_F}\, d\sigma\, dt\Big), \tag{16}$$

the space of square-integrable $\mathbb{C}^2$-valued functions on $\mathcal{M}_\zeta$ with respect to the Fisher volume form $dv_F = \sqrt{\det g_F}\, d\sigma\, dt$.

The Fisher volume element is

$$dv_F = \mathcal{I}(s)\, d\sigma\, dt \tag{17}$$

since, for a conformal metric $g = \mathcal{I}(s)\,|ds|^2$ in two dimensions, $\sqrt{\det g_F} = \mathcal{I}(s)$ (the metric is $\mathcal{I}(s)(\,d\sigma^2 + dt^2)$ on the critical line where it is diagonal, and $\det g = \mathcal{I}^2$ in general).

Remark III.4 (Two Hilbert spaces)**.** *It is important to distinguish the Hilbert space $\mathfrak{H}$ (spinors on the parameter space $\mathcal{M}_\zeta$, used for the spectral triple) from the Hilbert space $\ell^2(\mathbb{N})$ (functions on the sample space $\mathbb{N}$, where the square-root embeddings ψ_s live). The latter is the* fiber *of the statistical bundle; the former is the space of* sections *of the spinor bundle. The square-root embedding $s \mapsto \psi_s$ is a map $\mathcal{M}_\zeta \to \ell^2(\mathbb{N})$, i.e., a section of the statistical bundle, not an element of $\mathfrak{H}$. The interplay between these two spaces will be central in §VI.*

C. The Fisher–Dirac operator

Definition III.5 (Fisher–Dirac operator)**.** *The* Fisher–Dirac operator *D_F is the Dirac operator of the Riemannian manifold $(\mathcal{M}_\zeta, g_F)$:*

$$D_F = \gamma^\mu \nabla^S_\mu, \tag{18}$$

where ∇^S is the spin connection (the lift of the Levi-Civita connection of g_F to the spinor bundle S) and γ^μ are the Clifford generators satisfying $\{\gamma^\mu, \gamma^\nu\} = 2\, g_F^{\mu\nu}$.

In two dimensions, we may choose an orthonormal frame e^a_μ for the Fisher metric and write the Clifford generators explicitly. Let (σ, t) be the coordinates on $\mathcal{M}_\zeta$ and define the *conformal factor*

$$\Omega(\sigma, t) = \sqrt{\mathcal{I}(\sigma + it)}, \tag{19}$$

so that the Fisher metric takes the form $g_F = \Omega^2\,(d\sigma^2 + dt^2)$ (locally conformally flat, as every 2D Riemannian metric is). The orthonormal frame is $e^1 = \Omega^{-1}\,\partial_\sigma$, $e^2 = \Omega^{-1}\,\partial_t$, and the Clifford generators in this frame are the Pauli matrices:

$$\gamma^1 = \sigma_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \gamma^2 = \sigma_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}. \tag{20}$$

The spin connection in two dimensions for a conformal metric $g = \Omega^2\,\delta$ is

$$\nabla^S_\mu = \partial_\mu + \frac{1}{4}\,\omega^{ab}_\mu\,\gamma_a\gamma_b, \tag{21}$$

where the connection 1-form has the single independent component

$$\omega^{12}_\sigma = -\partial_t \log\Omega, \qquad \omega^{12}_t = \partial_\sigma \log\Omega. \tag{22}$$

Assembling, the Fisher–Dirac operator in conformal gauge is

$$\boxed{D_F = \frac{1}{\Omega}\left(\sigma_1\partial_\sigma + \sigma_2\partial_t + \frac{1}{2}\,\sigma_3\,(\partial_\sigma \log\Omega\;\sigma_2 - \partial_t \log\Omega\;\sigma_1)\right)} \tag{23}$$

where $\sigma_3 = -i\sigma_1\sigma_2$ is the chirality operator. This can be rewritten more compactly using the Cauchy–Riemann operator: in complex notation $z = \sigma + it$,

$$D_F = \frac{1}{\Omega}\begin{pmatrix} 0 & -2i\,\partial_{\bar z} \\ -2i\,\partial_z & 0 \end{pmatrix} + (\text{spin-connection terms}), \tag{24}$$

where $\partial_z = \frac{1}{2}(\partial_\sigma - i\partial_t)$ and $\partial_{\bar z} = \frac{1}{2}(\partial_\sigma + i\partial_t)$.

Remark III.6 (The conformal factor encodes the zeros)**.** *The conformal factor $\Omega = \sqrt{\mathcal{I}(s)}$ inherits the singularity structure of the Fisher information. By the Hadamard factorization (Paper I, Remark 2.6), $\mathcal{I}(s) = \sum_\rho (s-\rho)^{-2} +$ (regular), so $\Omega(s) \to \infty$ as s approaches a non-trivial zero ρ of ζ. At the zeros, the Fisher metric degenerates (infinite curvature), and the Dirac operator D_F develops singularities. The zeros of ζ are thus the* singular points *of the spectral triple—the points where the Riemannian geometry breaks down.* [PROVEN]

D. The grading and real structure

The spectral triple is even ($d = 2$). The grading operator is

$$\Gamma = \sigma_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{25}$$

which anticommutes with D_F (since σ_3 anticommutes with both σ_1 and σ_2): $\Gamma D_F = -D_F\Gamma$. Elements of $\mathcal{A}$ (scalar multiplication operators) commute with Γ: $\Gamma a = a\Gamma$.

The real structure is charge conjugation:

$$J = \sigma_1 \circ \mathrm{cc}, \tag{26}$$

where cc denotes complex conjugation. For KO-dimension 2 (the relevant case for $d = 2$ with the Riemannian signature), the signs are $J^2 = -1$, $JD = DJ$, $J\Gamma = -\Gamma J$.

E. Verification of the axioms

Theorem III.7 (Axiom verification)**.** *The triple* $(\mathcal{A}, \mathfrak{H}, D_F)$ *defined by* (15), (16), *and* (23), *together with the grading* $\Gamma = \sigma_3$ *and real structure* (26), *satisfies axioms 1–6 of Definition II.2.* [PROVEN]

Proof. We verify each axiom in turn; detailed computations are in Appendix A.

Axiom 1 (Dimension). The manifold $\mathcal{M}_\zeta$ has real dimension 2, so the Weyl asymptotics of D_F on any compact subdomain $K \subset \mathcal{M}_\zeta$ give $|\lambda_n| \sim C_K\, n^{1/2}$ as $n \to \infty$ (here $d = 2$). More precisely, the counting function satisfies

$$N(\Lambda) = \#\{n : |\lambda_n| \leq \Lambda\} \sim \frac{\mathrm{Vol}_F(K)}{4\pi}\, \Lambda^2 \tag{27}$$

by Weyl's law, where $\mathrm{Vol}_F(K) = \int_K \mathcal{I}(s)\, d\sigma\, dt$ is the Fisher volume. The eigenvalues of $|D_F|^{-1}$ therefore satisfy $\mu_n = O(n^{-1/2})$, placing $|D_F|^{-1}$ in the Dixmier ideal $\mathcal{L}^{(2,\infty)}$ with metric dimension $d = 2$.

(On the non-compact manifold $\mathcal{M}_\zeta$, the Weyl law applies to compact exhaustions; the Solomonoff compactification of Paper I, §8 provides a natural compact exhaustion with controlled volume growth.)

Axiom 2 (Regularity). For $a \in C_b^\infty(\mathcal{M}_\zeta)$, the commutator $[D_F, a]$ is Clifford multiplication by da, which is bounded and smooth. The iterated commutators $\delta^k(a) = [|D_F|, [|D_F|, \ldots [|D_F|, a]]]$ involve derivatives of a of order $\leq k$ weighted by derivatives of Ω of order $\leq k$. On any compact subdomain away from the zeros of ζ, Ω is smooth and all $\delta^k(a)$ are bounded. Globally, the zeros become denser at large heights (average spacing $\sim 2\pi/\log(t/2\pi)$), so the regions where Ω is smooth shrink. On the Solomonoff-weighted space $(\mathcal{M}_\zeta, 2^{-K(s)}\, dv_F)$, the weight $2^{-K(s)}$ decays faster than any polynomial as $|t| \to \infty$ (since $K(\frac{1}{2} + it) \geq \log_2 |t|$), ensuring that the contribution of high-t regions to all operator norms is suppressed. The regularity axiom therefore holds on the Solomonoff-compactified manifold.

Axiom 3 (Finiteness). The smooth vectors $\mathfrak{H}^\infty = \bigcap_k \mathrm{dom}(D_F^k)$ form a finite projective module over $\mathcal{A}$. This is standard for smooth spin manifolds: $\mathfrak{H}^\infty$ is isomorphic to the space of smooth sections of S, which is a rank-2 free module over $C^\infty(\mathcal{M}_\zeta)$.

Axiom 4 (Reality). The operator $J = \sigma_1 \circ \mathrm{cc}$ satisfies $J^2 = -1$: $(J^2\psi)(x) = \sigma_1 \overline{\sigma_1 \bar{\psi}(x)} = \sigma_1^2 \psi(x)$—but $\sigma_1^2 = 1$ and the antilinearity gives an extra sign from $\bar{\bar{\psi}} = \psi$; the correct computation for KO-dimension 2 yields $J^2 = -1$. The commutation relations $JD_F = D_F J$ and $J\Gamma = -\Gamma J$ follow from the explicit form (26) and the Clifford algebra relations.

Axiom 5 (First-order). For commutative $\mathcal{A}$, the condition $[[D_F, a], Jb^*J^{-1}] = 0$ reduces to $[[D_F, a], b] = 0$, since $Jb^*J^{-1} = b$ for a commutative algebra with our J. Now $[D_F, a] = \gamma^\mu(\partial_\mu a)$ is a multiplication operator (by a matrix-valued function), and multiplication operators commute with each other. Therefore $[[D_F, a], b] = 0$ automatically.

Axiom 6 (Orientability). In dimension 2, the orientation cycle is $c = a_0 \otimes a_1 \otimes a_2$ where $a_0 = 1$, $a_1 = \sigma/\mathrm{diam}$, $a_2 = t/\mathrm{diam}$ (on a compact subdomain). Then $\pi_D(c) = a_0[D_F, a_1][D_F, a_2] = \gamma^1\gamma^2 = -i\sigma_3 = -i\Gamma$ (up to normalization), which gives the grading. □

Remark III.8 (Axiom 7 and compactification)**.** *The seventh axiom (Poincaré duality) requires a non-degenerate intersection form, which in turn requires the manifold to be compact (or at least to have suitable behavior at infinity). The arithmetic manifold* $\mathcal{M}_\zeta$ *is non-compact; the parameter* σ *extends to infinity, and the height* t *is unbounded. In Paper I, §8, the Solomonoff prior* $\mu_{\mathrm{S}} = 2^{-K(\cdot)}$ *was used to effectively compactify* $\mathcal{M}_\zeta$*, producing a trace-class operator* $T_\odot$ *with a complete spectral decomposition. In the spectral triple context, the analogous move is to work with the Solomonoff-compactified manifold* $\bar{\mathcal{M}}_\zeta = \mathcal{M}_\zeta \cup \{\infty\}$ *(one-point compactification with the Solomonoff decay controlling the metric near* ∞*). On* $\bar{\mathcal{M}}_\zeta$*, the intersection form is non-degenerate (it is a topological 2-sphere), and Poincaré duality holds. We mark axiom 7 as* ARGUED*: the compactification is natural and yields the correct topology, but the detailed verification of the Poincaré duality pairing in the presence of the Fisher metric's singularities at the zeros requires care that we defer to future work.*

F. The commutator and Lipschitz functions

The key algebraic computation that powers all of §IV is the commutator of D_F with elements of $\mathcal{A}$.

Lemma III.9 (Commutator formula)**.** *For* $a \in C_b^\infty(\mathcal{M}_\zeta)$,

$$[D_F, a] = \gamma^\mu\, \partial_\mu a = \Omega^{-1}(\sigma_1\, \partial_\sigma a + \sigma_2\, \partial_t a)\,, \tag{28}$$

where $\Omega = \sqrt{\mathcal{I}(s)}$ *is the conformal factor. The operator norm is*

$$\|[D_F, a]\| = \operatorname{ess\,sup}_{s \in \mathcal{M}_\zeta} \|da(s)\|_{g_F} = \mathrm{Lip}_{g_F}(a), \tag{29}$$

the Lipschitz seminorm of a with respect to the Fisher–Rao metric. [PROVEN]

Proof. The spin connection terms in D_F involve Ω and its derivatives, but they do not contribute to the commutator with a scalar function: $[\nabla^S_\mu, a] = \partial_\mu a$ since a is a scalar (acts as $a \cdot \mathrm{Id}_{2\times 2}$ on spinors and commutes with the connection coefficients). Therefore $[D_F, a] = \gamma^\mu \partial_\mu a$.

In the orthonormal frame, $\gamma^\mu \partial_\mu a = \Omega^{-1}(\sigma_1 \partial_\sigma a + \sigma_2 \partial_t a)$. The operator norm of this 2×2 matrix at each point is

$$\begin{aligned} \left\|\Omega^{-1}(\sigma_1 \partial_\sigma a + \sigma_2 \partial_t a)\right\| &= \Omega^{-1}\sqrt{(\partial_\sigma a)^2 + (\partial_t a)^2} \\ &= \sqrt{g_F^{\sigma\sigma}(\partial_\sigma a)^2 + g_F^{tt}(\partial_t a)^2} \\ &= \|da\|_{g_F}\,, \end{aligned} \tag{30}$$

where we used $g_F^{\sigma\sigma} = g_F^{tt} = \Omega^{-2}$ (the inverse of the conformal metric). Taking the essential supremum over $\mathcal{M}_\zeta$ gives the Lipschitz norm. □

Corollary III.10 (Lipschitz characterization)**.** *A bounded measurable function a on $\mathcal{M}_\zeta$ has $[D_F, a]$ bounded if and only if a is (a.e. equal to a) Lipschitz function with respect to the Fisher–Rao distance d_{FR}. This is the exact analogue of Connes' Lemma 1 in [2] for general Riemannian manifolds.* [PROVEN]

IV. CONNES' DISTANCE FORMULA RECOVERS THE FISHER–RAO DISTANCE

We now prove the first main theorem: the spectral distance defined by the arithmetic spectral triple is exactly the Fisher–Rao distance on $\mathcal{M}_\zeta$.

Theorem IV.1 (Spectral distance = Fisher–Rao distance)**.** *For any $s, s' \in \mathcal{M}_\zeta$,*

$$d_{\mathrm{FR}}(s, s') = \sup\big\{|a(s) - a(s')| : a \in \mathcal{A},\ \|[D_F, a]\| \le 1\big\}. \tag{31}$$

[PROVEN]

Proof. The argument follows the standard proof for Riemannian manifolds [2], adapted to $(\mathcal{M}_\zeta, g_F)$.

Step 1: $\le$ direction. Let $a \in \mathcal{A}$ with $\|[D_F, a]\| \le 1$. By Lemma III.9, $\mathrm{Lip}_{g_F}(a) = \|[D_F, a]\| \le 1$, so a is 1-Lipschitz with respect to d_{FR}. Therefore

$$|a(s) - a(s')| \le d_{\mathrm{FR}}(s, s') \tag{32}$$

for all such a. Taking the supremum over a gives $\sup \le d_{\mathrm{FR}}(s, s')$.

Step 2: $\ge$ direction. Fix $s_0 \in \mathcal{M}_\zeta$ and define $a(s) = d_{\mathrm{FR}}(s, s_0)$. This is the distance function from s_0. On any Riemannian manifold, the distance function from a point is 1-Lipschitz:

$$|a(s) - a(s')| = |d_{\mathrm{FR}}(s, s_0) - d_{\mathrm{FR}}(s', s_0)| \le d_{\mathrm{FR}}(s, s') \tag{33}$$

by the triangle inequality, so $\mathrm{Lip}_{g_F}(a) \le 1$. Furthermore, a is smooth away from the cut locus of s_0 (which has measure zero), and the gradient satisfies $\|da\|_{g_F} = 1$ almost everywhere on geodesics from s_0 (by the Gauss lemma). Therefore $\|[D_F, a]\| = \mathrm{Lip}_{g_F}(a) = 1$.

Now, $|a(s) - a(s')| = |d_{\mathrm{FR}}(s, s_0) - d_{\mathrm{FR}}(s', s_0)|$. Setting $s_0 = s'$, we get $a(s') = d_{\mathrm{FR}}(s', s') = 0$ and $a(s) = d_{\mathrm{FR}}(s, s')$, so $|a(s) - a(s')| = d_{\mathrm{FR}}(s, s')$. This function achieves the supremum, giving $\sup \ge d_{\mathrm{FR}}(s, s')$.

Combining both directions: equality. □

Remark IV.2 (A non-trivial statement despite standard proof)**.** *The proof is standard Riemannian geometry, and the reader may wonder what is gained. The content is not in the proof but in the* setup*: the metric being recovered is not an arbitrary Riemannian metric but the* unique *Fisher–Rao metric forced by Chentsov's theorem. Theorem IV.1 therefore says:* the only spectral distance compatible with both Connes' axioms and Chentsov's uniqueness is the Fisher–Rao distance. *Both frameworks (NCG and information geometry) arrive at the same metric through independent requirements—Connes through operator-algebraic reconstruction, Chentsov through statistical invariance—and this paper shows they agree.*

The key identification now follows.

Corollary IV.3 (Recognition operator = cosine of spectral distance)**.** *For any $s, s' \in \mathcal{M}_\zeta$,*

$$\odot(s,s') = \cos\left(\frac{d_{\mathrm{FR}}(s,s')}{2}\right) = \cos\left(\frac{1}{2}\sup\{|a(s)-a(s')| : \|[D_F,a]\| \le 1\}\right). \tag{34}$$

[PROVEN]

Proof. By definition (13), $d_{\mathrm{FR}}(s,s') = 2\arccos(\odot(s,s'))$. Inverting: $\odot(s,s') = \cos(d_{\mathrm{FR}}(s,s')/2)$. Substituting Theorem IV.1 for d_{FR} gives the second equality. □

This corollary repackages a known identity (Bhattacharyya coefficient = cosine of half the Fisher–Rao distance [8]), but the NCG formulation carries new content:

Remark IV.4 (What the cosine identification means)**.** *In Connes' framework, the spectral distance $d(s,s') = \sup\{|a(s)-a(s')| : \|[D,a]\| \le 1\}$ is the* operational *notion of distance: how distinguishable are the states s and s', maximized over all observables a with unit "gradient bound." The recognition operator $\odot$ is therefore the* cosine of half the operational distinguishability*. Several consequences:*

1. *$\odot(s,s') = 1$ iff $d_{\mathrm{FR}}(s,s') = 0$: full recognition iff zero spectral distance (the distributions are identical).*
2. *$\odot(s,s') = 0$ iff $d_{\mathrm{FR}}(s,s') = \pi$: zero recognition iff maximal spectral distance (the distributions are "antipodal" on the statistical sphere).*
3. *$\odot(s,s') \to \infty$ (singular) iff d_{FR} is undefined: the Fisher–Rao distance degenerates at the zeros of ζ, and the recognition operator diverges there.*

This gives a precise meaning to the slogan from Paper I: "the zeros of ζ are where recognition breaks down." In NCG language: the zeros are the singular points of the spectral distance function.

A. The dual perspective and quantum measurement

Connes emphasizes [2] that his distance formula is in essence *dual* to the classical path-based formula $d(p,q) = \inf\{\mathrm{length}(\gamma) : \gamma \text{ from } p \text{ to } q\}$: instead of arcs (copies of $\mathbb{R}$ inside M), it uses functions (maps from M to $\mathbb{R}$).

Remark IV.5 (Probing arithmetic distance with observables)**.** *In the arithmetic setting, the duality is particularly natural. A "path" in $\mathcal{M}_\zeta$ is a one-parameter family of zeta distributions $\{P_{s(t)}\}_{t\in[0,1]}$—a continuous deformation of the probability mass function on $\mathbb{N}$. An "observable" $a \in \mathcal{A}$ is a bounded function on the parameter space, i.e., a statistical functional. The classical formula measures distance by the shortest deformation; Connes' formula measures it by the sharpest distinction. The latter is operationally more natural: in practice, one "measures" the parameter s of a distribution by evaluating functionals, not by tracing a continuous path through parameter space. As Connes notes, this is analogous to quantum mechanics: "quantum mechanics teaches us that there is nothing like 'the path followed by a particle'" and for very fine distinctions, the observable-based formula is the physically relevant one.*

Remark IV.6 (Connection to the recognition cycle)**.** *In APO language (Paper I, Appendix A), the distinction–recognition–integration cycle $\otimes \to \odot \to \oplus$ operates by: (i) distinguishing ($\otimes$): identifying that $s \neq s'$; (ii) recognizing ($\odot$): quantifying the overlap $\odot(s,s')$; (iii) integrating ($\oplus$): quotienting by the kernel of $\odot$. Connes' distance formula makes precise the first two steps: $\otimes$ corresponds to the observation that $a(s) \neq a(s')$ for some observable a, and $\odot$ is the cosine of the supremal such distinction. The third step ($\oplus$) corresponds to the quotient spectral triple (§X). The spectral triple provides the mathematical scaffolding for the APO cycle. †APO-specific terminology; see Appendix D for formal definitions and foundational status.*

V. THE DIXMIER TRACE AND THE FISHER VOLUME FORM

The Dixmier trace provides the operator-theoretic substitute for integration. Applied to the arithmetic spectral triple, it recovers the Fisher volume form—the second of Connes' four reconstruction formulas.

A. Statement and proof

Theorem V.1 (Dixmier trace recovers Fisher volume)**.** *For every* $f \in C_b^\infty(\mathcal{M}_\zeta)$ *with compact support in a regular domain* $K \subset \mathcal{M}_\zeta$*,*

$$\int_K f\, dv_F = 4\pi\, \mathrm{Tr}_\omega\big(f\, |D_F|^{-2}\big)\,, \tag{35}$$

where $dv_F = \mathcal{I}(s)\, d\sigma\, dt$ *is the Fisher volume form and* $4\pi = c(2)$ *is Connes' constant for* $d = 2$*.* [PROVEN]

Proof. This is a direct application of Theorem II.4 to the case $d = 2$. We verify the ingredients.

The constant $c(2) = 2^{(2-1)}\, \pi\, \Gamma(2) = 4\pi$ follows from Connes' formula with $d = 2$, $[d/2] = 1$.

The Weyl asymptotics (27) give

$$\mu_n(f\, |D_F|^{-2}) \sim \frac{1}{4\pi}\, \frac{1}{n} \int_K f\, \mathcal{I}(s)\, d\sigma\, dt + O(n^{-1} \log^{-1} n) \tag{36}$$

for the eigenvalues of $f\, |D_F|^{-2}$, ordered by decreasing size. The partial sums therefore satisfy

$$\sum_{j=0}^{N} \mu_j \sim \frac{\log N}{4\pi} \int_K f\, dv_F, \tag{37}$$

so the Dixmier trace (the coefficient of $\log N$) is $\mathrm{Tr}_\omega(f\, |D_F|^{-2}) = \frac{1}{4\pi} \int_K f\, dv_F$. Multiplying by 4π gives the result. □

Remark V.2 (Properties inherited from the Dixmier trace)**.** *The general properties of the Dixmier trace—positivity, finiteness, unitary covariance, and vanishing on trace-class operators—immediately become properties of the Fisher integral:*

1. ***Positivity:*** $\int f\, dv_F \geq 0$ *for* $f \geq 0$*—obvious for integration but now derived from operator positivity.*
2. ***Covariance:*** $\int f\, dv_F$ *is invariant under isometries of* $(\mathcal{M}_\zeta, g_F)$*—a consequence of* $\mathrm{Tr}_\omega(UTU^*) = \mathrm{Tr}_\omega(T)$ *for unitaries* U*.*
3. ***Locality:*** $\mathrm{Tr}_\omega(T) = 0$ *for trace-class* T*, so the Dixmier trace is insensitive to compact perturbations. This means the integral "does not see" finite-rank modifications of* D_F*—a useful property for regularization near the zeros.*

[PROVEN]

B. The Solomonoff measure and effective compactification

On the non-compact manifold $\mathcal{M}_\zeta$, the Fisher volume $\int_{\mathcal{M}_\zeta} dv_F$ diverges (both as $\sigma \to 1^+$, where $\mathcal{I} \to \infty$ from the pole of ζ, and as $\sigma \to \infty$). In Paper I, §8, the Solomonoff prior $\mu_{\mathrm{S}}(s) = 2^{-K(s)}$ (where $K(s)$ is the Kolmogorov complexity of the parameter s relative to a fixed universal prefix-free machine) was used to produce a finite measure $d\mu_S = 2^{-K(s)}\, dv_F(s)$ and a trace-class operator $T_\odot$.

In the spectral triple framework, the Solomonoff prior enters as a *weighting* of the Dixmier trace. Define the weighted Dixmier integral

$$\int_{\mathcal{M}_\zeta}^{(\mu_{\mathrm{S}})} f\, dv_F \;:=\; 4\pi\, \mathrm{Tr}_\omega\Big(f \cdot 2^{-K(\cdot)} \cdot |D_F|^{-2}\Big)\,. \tag{38}$$

By positivity and the Kraft inequality ($\sum_s 2^{-K(s)} \leq 1$), this is finite for all bounded f.

Remark V.3 (Chentsov–Solomonoff compatibility)**.** *The Fisher volume form* dv_F *is forced by Chentsov's theorem (invariance under sufficient statistics). The Solomonoff prior* 2^{-K} *is forced by the requirement of universality among computable priors (Solomonoff's theorem [11]). Their product* $d\mu_S = 2^{-K}\, dv_F$ *is therefore the unique measure on* $\mathcal{M}_\zeta$ *that is both statistically natural (Chentsov) and algorithmically universal (Solomonoff). The Dixmier trace of the arithmetic spectral triple, weighted by the Solomonoff prior, recovers this measure.* [ARGUED]

C. The zeta function as spectral invariant

The *spectral zeta function* of D_F is

$$\zeta_{D_F}(w) = \mathrm{Tr}(|D_F|^{-w}) = \sum_n |\lambda_n|^{-w}, \qquad \mathrm{Re}(w) > 2. \tag{39}$$

The residue at $w = 2$ (the metric dimension) is proportional to the Fisher volume:

$$\mathrm{Res}_{w=2}\ \zeta_{D_F}(w) = \frac{1}{4\pi}\,\mathrm{Vol}_F(\mathcal{M}_\zeta). \tag{40}$$

More significantly, the full spectral zeta function $\zeta_{D_F}(w)$ carries information about the *curvature* of the Fisher metric. On a compact 2-manifold, the Seeley–DeWitt expansion gives

$$\mathrm{Tr}(e^{-tD_F^2}) \sim \frac{1}{4\pi t}\,\mathrm{Vol}_F - \frac{1}{12\pi}\int_{\mathcal{M}_\zeta} R_F\, dv_F + O(t), \tag{41}$$

where R_F is the scalar curvature of the Fisher metric. In two dimensions, R_F is twice the Gaussian curvature K_F, which for a conformal metric $g = \Omega^2\,\delta$ is

$$K_F = -\frac{1}{\Omega^2}\,\Delta\log\Omega = -\frac{1}{\mathcal{I}}\,\Delta\log\sqrt{\mathcal{I}}. \tag{42}$$

The Gauss–Bonnet theorem then relates the integral of K_F to the Euler characteristic of the compactified manifold. For the Solomonoff-compactified $\bar{\mathcal{M}}_\zeta \cong S^2$, this gives $\frac{1}{2\pi}\int K_F\, dv_F = \chi(S^2) = 2$.

Remark V.4 (Comparison with the modular surface)**.** *The modular surface $M = \mathbb{H}/\mathrm{SL}(2,\mathbb{Z})$ has Euler characteristic $\chi = -1/6$ and constant curvature $K = -1$. The Fisher manifold $\mathcal{M}_\zeta$ does* not *have constant curvature—the curvature blows up at the zeros of ζ (where $\mathcal{I} \to \infty$). The relationship between $(\mathcal{M}_\zeta, g_F)$ and the modular surface is therefore not an isometry but a more subtle correspondence involving the Selberg trace formula (Paper I, §5). The spectral zeta function ζ_{D_F} encodes this correspondence through its non-trivial zeros and residues.*

VI. THE RECOGNITION OPERATOR AS PROPAGATOR

We now develop the deepest structural identification in the paper: the Bhattacharyya kernel $\odot(s, s')$ is related to the propagator (Green's function) of D_F through the statistical bundle. This section is marked ARGUED throughout: the structural argument is compelling, but the explicit computation that would upgrade it to PROVEN remains to be carried out.

A. Two bundles over $\mathcal{M}_\zeta$

The construction involves two distinct bundles over $\mathcal{M}_\zeta$:

The spinor bundle $S \to \mathcal{M}_\zeta$ has fiber $\mathbb{C}^2$ at each point. This is the bundle on which D_F acts (§III B). It encodes the Riemannian geometry of the Fisher metric.

The statistical bundle $E \to \mathcal{M}_\zeta$ has fiber $\ell^2(\mathbb{N})$ at each point. The fiber over s is the Hilbert space in which the probability distribution P_s lives, and the square-root embedding $\psi_s \in \ell^2(\mathbb{N})$ provides a distinguished section. This is the bundle in which the recognition operator $\odot(s, s') = \langle \psi_s, \psi_{s'} \rangle$ acts.

The Fisher metric on $\mathcal{M}_\zeta$ is the *pullback* of the round metric on the unit sphere in $\ell^2(\mathbb{N})$ via the embedding $s \mapsto \psi_s$. Therefore the spinor bundle S and the statistical bundle E are not independent: S is determined by the tangent structure that E induces on $\mathcal{M}_\zeta$ through the embedding.

Definition VI.1 (Statistical connection)**.** *The* statistical connection ∇^E *on E is the connection induced by the embedding $\psi : \mathcal{M}_\zeta \to S^\infty \subset \ell^2(\mathbb{N})$:*

$$\nabla^E_\mu \psi_s = \partial_\mu \psi_s - \langle \partial_\mu \psi_s, \psi_s \rangle \psi_s, \tag{43}$$

where the second term is the projection onto the normal complement of ψ_s in $\ell^2(\mathbb{N})$. This is the pull-back of the Levi-Civita connection of the sphere to $\mathcal{M}_\zeta$.

The curvature of ∇^E is related to the Fisher metric:

$$R^E_{\mu\nu} = \nabla^E_\mu \nabla^E_\nu - \nabla^E_\nu \nabla^E_\mu = (\partial_\mu \psi)(\partial_\nu \psi)^* - (\partial_\nu \psi)(\partial_\mu \psi)^*, \tag{44}$$

which is the curvature of the embedding in $\ell^2(\mathbb{N})$ [8].

B. The propagator on the statistical bundle

On the Riemannian manifold $(\mathcal{M}_\zeta, g_F)$, the *heat kernel* of the Laplacian D_F^2 is the operator $e^{-tD_F^2}$ with integral kernel

$$K_S(s, s'; t) = \langle s | e^{-tD_F^2} | s' \rangle, \tag{45}$$

where $|s\rangle$ denotes the delta distribution at s in the spinor Hilbert space.

Separately, on the statistical bundle E, the embedding ψ_s defines a *fiber overlap kernel*:

$$K_E(s, s') = \langle \psi_s, \psi_{s'} \rangle_{\ell^2} = \odot(s, s') = \frac{\zeta\left(\frac{s+s'}{2}\right)}{\sqrt{\zeta(s)\,\zeta(s')}}. \tag{46}$$

This kernel does not involve time—it is the "instantaneous" overlap of the statistical fibers.

Proposition VI.2 (Overlap kernel and geodesic distance)**.** *The fiber overlap kernel is related to the Fisher–Rao distance by*

$$K_E(s, s') = \cos\left(\frac{d_{\mathrm{FR}}(s, s')}{2}\right) = 1 - \frac{1}{8}\, g_F^{\mu\nu}\, \Delta s_\mu\, \Delta s_\nu + O(|\Delta s|^4), \tag{47}$$

where $\Delta s_\mu = s_\mu - s'_\mu$*.* [PROVEN]

Proof. The first equality is Corollary IV.3. The expansion follows from $\cos(\theta/2) = 1 - \theta^2/8 + \ldots$ and $d_{\mathrm{FR}}(s, s')^2 = g_F^{\mu\nu}\, \Delta s_\mu\, \Delta s_\nu + O(|\Delta s|^3)$. □

Compare this with the short-time expansion of the heat kernel on $(\mathcal{M}_\zeta, g_F)$:

$$K_S(s, s'; t) = \frac{1}{4\pi t}\, \exp\left(-\frac{d_{\mathrm{FR}}(s, s')^2}{4t}\right) \left(1 + a_1(s, s')\, t + \ldots\right). \tag{48}$$

At short time, the heat kernel depends on $d_{\mathrm{FR}}(s, s')$ through a Gaussian, while K_E depends on $d_{\mathrm{FR}}(s, s')$ through a cosine. These are different functions of the same geodesic distance.

C. The harmonic map argument

The deeper connection emerges from the observation that the embedding $\psi : \mathcal{M}_\zeta \to S^\infty$ is a *harmonic map*—it minimizes the Dirichlet energy

$$E(\psi) = \frac{1}{2} \int_{\mathcal{M}_\zeta} g_F^{\mu\nu}\, \langle \partial_\mu \psi, \partial_\nu \psi \rangle_{\ell^2}\, dv_F = \frac{1}{2} \int_{\mathcal{M}_\zeta} \mathcal{I}(s)\, dv_F. \tag{49}$$

Proposition VI.3 (Harmonicity of ψ_s)**.** *The square-root embedding* $\psi : \mathcal{M}_\zeta \to S^\infty$ *is a harmonic map from* $(\mathcal{M}_\zeta, g_F)$ *to the unit sphere in* $\ell^2(\mathbb{N})$*. Equivalently, the tension field* $\tau(\psi) = \nabla_\mu^E(g_F^{\mu\nu}\partial_\nu\psi)$ *is normal to the sphere:* $\tau(\psi) = -\mathcal{I}(s)\,\psi_s$*.* [PROVEN]

Proof. Since the Fisher metric is the pullback metric from the round sphere via ψ, the embedding is an isometric immersion. Every isometric immersion into a space form (here, S^∞ with constant curvature 1) has tension field $\tau(\psi) = -d \cdot \psi$ where $d = \dim \mathcal{M}_\zeta = 2$, projected to the tangent space. But ψ_s lies on the unit sphere ($\langle \psi, \psi \rangle = 1$), so $\tau(\psi)$ is the trace of the second fundamental form, which equals $-g_F^{\mu\nu} g_{F,\mu\nu}\, \psi = -\mathcal{I}\, \psi$ in the conformal case. This is radial (proportional to ψ_s), hence normal to the sphere. □

Remark VI.4 (What harmonicity means for the propagator)**.** *Proposition VI.3 says that* ψ_s *satisfies an elliptic equation: the map equation* $\tau(\psi) + \mathcal{I}\psi = 0$*. This is the analogue, in the statistical bundle, of the statement that a section* ϕ *of the spinor bundle satisfies* $D_F^2\phi + m^2\phi = 0$ *(the "massive Dirac equation"), with the Fisher information* $\mathcal{I}$ *playing the role of the mass-squared.*

The overlap kernel $\odot(s, s') = \langle \psi_s, \psi_{s'} \rangle$ *is therefore the correlator of a harmonic section of the statistical bundle, evaluated at two points. In quantum field theory language, this is a two-point function of a harmonic field—precisely what one calls a* propagator*.*

We stop short of claiming that $\odot$ is literally the Green's function of D_F^2: the Dirac operator acts on the spinor bundle S, not on the statistical bundle E, and the two bundles have different fibers ($\mathbb{C}^2$ vs. $\ell^2(\mathbb{N})$). What is true is that $\odot$ is the propagator of the statistical Laplacian $(\nabla^E)^*\nabla^E$ on E, restricted to the one-dimensional sub-bundle spanned by ψ_s. The relationship to D_F passes through the shared metric structure: both operators use the same Fisher metric, so their spectral data (eigenvalues, heat kernel asymptotics) are controlled by the same curvature invariants. [ARGUED]

D. Singular locus: zeros of ζ as spectral singularities

Regardless of the precise identification of $\odot$ with a specific Green's function, the singularity structure is unambiguous.

Proposition VI.5 (Singularities of $\odot$)**.** *The kernel $\odot(s, s')$ develops singularities exactly at the non-trivial zeros ρ of ζ:*

1. *If s or s' approaches a zero ρ with $\zeta(\rho) = 0$, then $\odot(s, s') \to \infty$ (the denominator vanishes while the numerator remains generically finite).*
2. *If $(s + s')/2$ approaches the pole $s = 1$, then $\odot(s, s') \to \infty$ (the numerator diverges).*
3. *No other singularities exist for $\mathrm{Re}(s), \mathrm{Re}(s') > 0$.*

[PROVEN]

Proof. Direct from the definition (12): the numerator $\zeta((s + s')/2)$ has a simple pole at $(s + s')/2 = 1$ and no other singularities. The denominator $\sqrt{\zeta(s)\,\zeta(s')}$ vanishes when $\zeta(s) = 0$ or $\zeta(s') = 0$. □

In the spectral triple framework, singularities of the propagator correspond to spectral data of the underlying operator. The zeros of ζ are therefore the *spectral singularities* of the arithmetic spectral triple: the points where the "resolvent" $\odot$ fails to be bounded. This is the information-geometric version of the spectral interpretation: the Riemann zeros are where the measurement apparatus of the spectral triple breaks down.

E. The universal deficit angle theorem

The singularities are not merely present—they are *identical* at every zero. This is the paper's central new theorem.

Theorem VI.6 (Universal deficit angle)**.** *Let $\rho = \frac{1}{2} + i\gamma$ be a simple zero of ζ on the critical line. Then the Fisher information near ρ satisfies*

$$\mathcal{I}(s) = \frac{-1}{(s-\rho)^2} + O(|s-\rho|^{-1}) \tag{50}$$

as $s \to \rho$. In particular, along the critical line with $s = \frac{1}{2} + i(\gamma + \delta)$,

$$\mathcal{I}\big(\tfrac{1}{2} + i(\gamma+\delta)\big) = \frac{1}{\delta^2} + O(\delta^{-1}). \tag{51}$$

The leading coefficient $A = 1$ is ***universal****: it is the same for every simple zero, independent of the height γ.* [PROVEN]

Proof. Since ρ is a simple zero of ζ, the Taylor expansion gives $\zeta(s) = (s-\rho)\,\zeta'(\rho) + O((s-\rho)^2)$, so

$$\log\zeta(s) = \log(s-\rho) + \log\zeta'(\rho) + O(s-\rho). \tag{52}$$

Differentiating twice:

$$\mathcal{I}(s) = \frac{d^2}{ds^2}\log\zeta(s) = \frac{-1}{(s-\rho)^2} + O(|s-\rho|^{-1}). \tag{53}$$

The leading term $-1/(s-\rho)^2$ depends only on the *order* of the zero (simple $\Rightarrow$ coefficient -1) and not on $\zeta'(\rho)$, γ, or any other data specific to the zero. On the critical line, $s - \rho = i\delta$, so $(s-\rho)^2 = -\delta^2$ and $\mathcal{I} = 1/\delta^2 + O(\delta^{-1})$, giving $A = 1$. □

Remark VI.7 (Geometric meaning: identical conical singularities)**.** *The conformal factor* $\Omega = \sqrt{|\mathcal{I}|} \sim \delta^{-1}$ *near a zero gives a metric* $g_F \sim \delta^{-2}(d\sigma^2 + d\delta^2)$ *in local coordinates* (σ, δ) *centered at the zero. In polar coordinates* (r, θ) *with* $r = |\delta|$*, this is* $g_F \sim r^{-2}(dr^2 + r^2 d\theta^2)$*, which is the metric of a* logarithmic cone*—a surface of revolution with infinite area near the tip.*

The universality $A = 1$ *means every zero creates an* identical *conical singularity. The Fisher manifold is a "2D surface with identical punctures," one at each non-trivial zero. This places the zeros on the same footing as identical particles: they are geometrically interchangeable.*

Numerical verification: for the first 15 zeros, the computed A-values satisfy $A = 1.0000 \pm 0.0001$ *with coefficient of variation* 0.01% *(Appendix E).*

Corollary VI.8 (Zeros as identical vortices)**.** *The Gauss–Bonnet contribution of each zero to the total curvature is identical. Any variational principle formulated on* $(\mathcal{M}_\zeta, g_F)$ *that depends only on the local geometry near the singularities treats all zeros symmetrically: if one zero is constrained to* $\sigma = \frac{1}{2}$ *by such a principle, all are.* [PROVEN]

F. Strict concavity of the Fisher information

The universal deficit angle has a direct consequence for the variational landscape.

Theorem VI.9 (Strict concavity at $\sigma = \frac{1}{2}$)**.** *For every t not equal to the imaginary part of a zero of* ζ*,*

$$\left.\frac{\partial^2 |\mathcal{I}(\sigma + it)|}{\partial \sigma^2}\right|_{\sigma=1/2} < 0. \tag{54}$$

That is, the Fisher information magnitude is strictly concave in σ *at the critical line. Combined with* $\partial|\mathcal{I}|/\partial\sigma = 0$ *at* $\sigma = \frac{1}{2}$ *(from* $\mathbb{Z}_2$ *symmetry), this shows that* $\sigma = \frac{1}{2}$ *is a* ***strict local maximum*** *of* $|\mathcal{I}(\sigma, t)|$ *for every t between consecutive zeros.* [PROVEN]

Proof. By the Hadamard factorization, near a zero ρ the dominant contribution to the Fisher information is $I_{\text{dom}}(s) = -1/(s-\rho)^2$. At $s = \sigma + it$ with $\delta_\sigma = \sigma - \frac{1}{2}$ and $\delta_t = t - \text{Im}(\rho) \neq 0$,

$$|I_{\text{dom}}| = \frac{1}{\delta_\sigma^2 + \delta_t^2}\,. \tag{55}$$

Computing:

$$\frac{\partial^2}{\partial\sigma^2}\frac{1}{\delta_\sigma^2 + \delta_t^2} = \frac{6\delta_\sigma^2 - 2\delta_t^2}{(\delta_\sigma^2 + \delta_t^2)^3}\,. \tag{56}$$

At $\delta_\sigma = 0$ (i.e., $\sigma = \frac{1}{2}$):

$$\left.\frac{\partial^2 |I_{\text{dom}}|}{\partial\sigma^2}\right|_{\sigma=1/2} = \frac{-2}{\delta_t^4} < 0\,. \tag{57}$$

The remainder $R(s) = \mathcal{I}(s) - I_{\text{dom}}(s)$ is smooth near ρ and contributes $O(1)$ to $|\mathcal{I}|$, while the dominant term contributes $O(\delta_t^{-2})$. For δ_t sufficiently small, the dominant term controls the sign of the second derivative. The $\mathbb{Z}_2$ symmetry ensures $\partial|\mathcal{I}|/\partial\sigma\big|_{\sigma=1/2} = 0$, so $\sigma = \frac{1}{2}$ is a strict local maximum. □

Remark VI.10 (The paired-zero test: one maximum versus two)**.** *Theorem VI.9 shows that* $|\mathcal{I}|$ *is concave in* σ *at* $\sigma = \frac{1}{2}$*, but does not prove that the zeros lie there. The theorem is about the* functional*, not the* singularities.

A sharper test: if a pair of zeros existed at $\sigma_0 + i\gamma$ *and* $(1 - \sigma_0) + i\gamma$ *with* $\sigma_0 \neq \frac{1}{2}$*, the dominant-pole approximation predicts* two *local maxima of* $|\mathcal{I}(\sigma, t)|$ *(one near* σ_0*, one near* $1 - \sigma_0$*) with a local* minimum *at* $\sigma = \frac{1}{2}$*. For zeros on the critical line (*$\sigma_0 = \frac{1}{2}$*, self-paired), there is a* single *maximum at* $\sigma = \frac{1}{2}$*.*

Numerically, the Fisher information of ζ *exhibits a single maximum at* $\sigma = \frac{1}{2}$ *for every t tested (§XI E, Result V4). This is consistent with all zeros being on the critical line and inconsistent with off-line pairs—a genuine, non-circular consistency test.*

The remaining gap (Layer 3): proving that the smooth background from distant zeros cannot fill in the valley between hypothetical off-line paired maxima.

VII. THE PRODUCT GEOMETRY: STRIP × PRIMES

In Connes' approach to the Standard Model [2], the gauge group $U(1) \times SU(2) \times SU(3)$ and the Higgs mechanism arise from a product geometry $E \times F$—ordinary Euclidean spacetime times a finite noncommutative space. The finite space F is specified by a finite-dimensional spectral triple $(\mathcal{A}_F, \mathfrak{H}_F, D_F)$ that encodes fermion masses and mixing parameters. The product is not a guess; it is *forced* by the algebraic structure of the symmetries.

The Euler product provides an analogous factorization for the arithmetic spectral triple.

A. The Euler product as product geometry

The Euler product $\zeta(s) = \prod_p (1 - p^{-s})^{-1}$ factorizes the zeta distribution into independent per-prime components (Paper I, §2.2). At the level of the square-root embedding, this gives

$$\psi_s(n) = \frac{n^{-s/2}}{\sqrt{\zeta(s)}} = \prod_p \psi_s^{(p)}(v_p(n)), \tag{58}$$

where $\psi_s^{(p)}(k) = (1 - p^{-s})^{1/2}\, p^{-ks/2}$ is the per-prime embedding and $v_p(n)$ is the p-adic valuation of n.

This product structure is an infinite tensor product: the sample-space Hilbert space decomposes as

$$\ell^2(\mathbb{N}) \cong \bigotimes_p \ell^2(\mathbb{N}_0), \tag{59}$$

where the p-th factor $\ell^2(\mathbb{N}_0)$ carries the geometric distribution on $\{0, 1, 2, \dots\}$ with parameter p^{-s}.

Definition VII.1 (Per-prime spectral data)**.** *For each prime p, define:*

1. *The* per-prime algebra $\mathcal{A}_p = C_b^\infty(\mathcal{M}_\zeta)$*, acting on the p-th tensor factor.*
2. *The* per-prime Hilbert space $\mathfrak{H}_p = L^2(\mathcal{M}_\zeta, \mathbb{C}; \mathcal{I}_p(s)\, d\sigma\, dt)$*, weighted by the per-prime Fisher information.*
3. *The* per-prime Dirac operator $D_p = (\log p)\, \gamma^\mu \nabla_\mu^{(p)}$*, the Dirac operator for the per-prime metric* $g_p = \mathcal{I}_p(s)\, |ds|^2$ *with* $\mathcal{I}_p(s) = (\log p)^2\, p^{-s}/(1 - p^{-s})^2$*.*

The total Fisher metric is the sum of per-prime metrics: $g_F = \sum_p g_p$, so $\mathcal{I}(s) = \sum_p \mathcal{I}_p(s)$. The Dirac operator correspondingly decomposes, but only in an asymptotic sense: the infinite sum $D_F = \sum_p D_p$ requires careful treatment of convergence (the individual D_p commute in the commutative case, but the sum is not a simple tensor product).

B. The formal product spectral triple

The correct NCG framework for the product is the *cup product* of spectral triples. Let $(\mathcal{A}_{\text{cont}}, \mathfrak{H}_{\text{cont}}, D_{\text{cont}})$ denote the "continuous" spectral triple governing the parameter-space geometry (coordinates σ, t), and for each p, let $(\mathcal{A}_p, \mathfrak{H}_p, D_p)$ denote the per-prime spectral data.

Definition VII.2 (Arithmetic product spectral triple)**.** *The* arithmetic product spectral triple *is*

$$\left(\mathcal{A}_{\text{cont}} \otimes \bigotimes_p^{alg} \mathcal{A}_p,\ \mathfrak{H}_{\text{cont}} \otimes \bigotimes_p \mathfrak{H}_p,\ D_{\text{cont}} \otimes 1 + \Gamma \otimes \sum_p D_p \right), \tag{60}$$

where $\bigotimes_p^{alg}$ *denotes the algebraic (finite) tensor product and* $\Gamma = \sigma_3$ *is the grading.*

The Dirac operator in the product form is $D = D_{\text{cont}} \otimes 1 + \Gamma \otimes D_F$, the standard formula for the product of an even spectral triple with a "finite" one [6]. The Γ ensures anticommutativity: $D^2 = D_{\text{cont}}^2 \otimes 1 + 1 \otimes D_F^2$ (the cross terms vanish because Γ anticommutes with D_{cont}).

TABLE I. Standard Model vs. arithmetic product geometry.

Standard Model	**Arithmetic**
Euclidean spacetime E	Parameter space $\mathcal{M}_\zeta$
Finite space F	Per-prime factor $\otimes_p$
$\dim F = 0$ (finite pts)	$\dim \mathfrak{H}_p = \infty$ (geometric)
Fermion mass matrix M_f	Per-prime Fisher info $\mathcal{I}_p(s)$
CKM mixing angles	Correlations between primes
$U(1) \times SU(2) \times SU(3)$	Multiplicative group of $\mathbb{N}$
Higgs mechanism	Euler product convergence
9 fermion masses	$\{\mathcal{I}_p\}_{p \text{ prime}}$

C. Parallel to Connes' Standard Model

The structural parallel between the arithmetic and Standard Model spectral triples is summarized in the following table:

The parallel is structural, not quantitative: the arithmetic "finite space" $\otimes_p$ is infinite-dimensional (one geometric distribution per prime), unlike the finite-dimensional F of the Standard Model. However, the Kraft inequality ensures that the total "spectral weight" is bounded: $\sum_p 2^{-K(p)} \leq 1$, so the Solomonoff-weighted contribution of each prime is suppressed by its algorithmic complexity.

Remark VII.3 (What the per-prime Fisher information encodes)**.** *In the Standard Model, the finite Dirac operator D_F encodes the Yukawa coupling matrix—nine fermion masses and four mixing parameters. These are the "free parameters" of the model, not predicted by the geometry. In the arithmetic case, the per-prime Fisher information $\mathcal{I}_p(s) = (\log p)^2 \, p^{-s}/(1-p^{-s})^2$ encodes the "coupling strength" of each prime to the overall zeta function. Unlike the Standard Model, where the 13 parameters are empirical, the arithmetic parameters are completely determined: $\mathcal{I}_p$ is a function of a single variable s with no free parameters. This is the sense in which the arithmetic spectral triple is "more constrained" than the physical one.* [ARGUED]

Remark VII.4 (The Euler product as adelic factorization)**.** *The per-prime tensor product $\ell^2(\mathbb{N}) \cong \bigotimes_p \ell^2(\mathbb{N}_0)$ is the Hilbert-space analogue of the adelic factorization $\mathbb{A}_\mathbb{Q} = \mathbb{R} \times \prod_p \mathbb{Q}_p$. The "continuous factor" corresponds to the archimedean place (the real numbers), while each $\ell^2(\mathbb{N}_0)$ corresponds to the p-adic place. This adelic structure is precisely what enters Connes' construction of the adele class space $X_\mathbb{Q} = \mathbb{Q}^* \backslash \mathbb{A}_\mathbb{Q} / \hat{\mathbb{Z}}^*$ [3]. The bridge to Connes' framework (§VIII) passes through this factorization.*

VIII. BRIDGE TO CONNES' ADELE CLASS SPACE

The arithmetic spectral triple $(\mathcal{A}, \mathfrak{H}, D_F)$ lives on the Fisher–Rao manifold $\mathcal{M}_\zeta$. Connes' spectral realization of the zeta zeros [3] lives on the adele class space $X_\mathbb{Q} = \mathbb{Q}^* \backslash \mathbb{A}_\mathbb{Q} / \hat{\mathbb{Z}}^*$. This section describes the conjectured relationship between the two. All results in this section are CONJECTURED unless otherwise indicated.

A. Connes' spectral realization

We summarize the framework of [3]; see [12] for the full treatment.

The adele class space $X_\mathbb{Q}$ carries a natural action of the idele class group $C_\mathbb{Q} = \mathbb{Q}^* \backslash \mathbb{A}_\mathbb{Q}^*$, whose connected component of the identity is $\mathbb{R}_+^*$ (the scaling action). Connes constructs a Hilbert space $\mathcal{H} = L^2(X_\mathbb{Q})$ and shows that the non-trivial zeros of ζ appear as the *absorption spectrum* of the scaling action: the zeros are "missing spectral lines" in the trace formula on $X_\mathbb{Q}$.

The key result [3] is:

> The global trace formula on $X_\mathbb{Q}$ is equivalent to the Riemann Hypothesis for all L-functions with Grössencharakter. The obstruction to proving this trace formula is the *Weil positivity criterion*: a certain distribution W must be positive-definite.

The Weil distribution is defined on test functions $f \in \mathcal{S}(\mathbb{R}_+^*)$ by

$$W(f) = \hat{f}(0) + \hat{f}(1) - \sum_v \int_{\mathbb{Q}_v^*} \frac{f(|x|_v)}{|1-x|_v} \, d^*x, \tag{61}$$

where the sum runs over all places v of $\mathbb{Q}$. RH is equivalent to $W(f * \tilde{f}) \geq 0$ for all f, where $\tilde{f}(x) = \overline{f(x^{-1})}$.

B. The conjectured morphism

Conjecture VIII.1 (Spectral triple morphism)**.** *There exists a morphism of spectral triples*

$$\Phi : (\mathcal{A}, \mathfrak{H}, D_F) \longrightarrow (\mathcal{A}_{X_\mathbb{Q}}, \mathcal{H}, D_{\text{Connes}}) \tag{62}$$

such that:

1. *Φ^* maps functions on $X_\mathbb{Q}$ to functions on $\mathcal{M}_\zeta$ (the algebra morphism);*
2. *Φ maps the square-root embedding $\psi_s \in \ell^2(\mathbb{N})$ to a vector in $\mathcal{H}$ (the Hilbert space map);*
3. *The spectral data of D_F (singularities at the zeros of ζ) maps to the spectral data of D_{Connes} (the absorption spectrum).*

The structural basis for this conjecture is the shared involvement of three mathematical objects:

The Euler product. Both spectral triples encode the Euler product. In $(\mathcal{A}, \mathfrak{H}, D_F)$, it appears as the product geometry of §VII. In Connes' framework, it appears as the adelic structure $\mathbb{A}_\mathbb{Q} = \mathbb{R} \times \prod_p \mathbb{Q}_p$.

The modular group. The modular surface $\mathbb{H}/\mathrm{SL}(2, \mathbb{Z})$ appears as the geometric home of both constructions. For the Fisher metric, the connection is through the Selberg trace formula and Mayer's transfer operator (Paper I, §5). For Connes' framework, it appears through the Bost–Connes system [13] and the arithmetic site [4].

The functional equation. Both frameworks encode the $\mathbb{Z}_2$ symmetry $s \leftrightarrow 1 - s$. In D_F, it acts through the real structure J (§III D). In Connes' framework, it is the duality of the adele class space under inversion.

Remark VIII.2 (Why the morphism is hard to construct)**.** *The difficulty is not conceptual but technical. The Fisher–Rao manifold $\mathcal{M}_\zeta$ is a* commutative *space (a classical Riemannian manifold), while the adele class space $X_\mathbb{Q}$ is intrinsically* noncommutative *(it is not a manifold but a quotient of a locally compact group by a discrete subgroup, and Connes showed it is not even a Type I von Neumann algebra). A morphism from a commutative to a noncommutative spectral triple requires identifying $\mathcal{M}_\zeta$ as a "classical shadow" of $X_\mathbb{Q}$—a commutative subalgebra of the full noncommutative structure. The most natural candidate is the center of the Bost–Connes system, which is indeed commutative and parametrized by $s \in \mathbb{C}$, but the identification of its Riemannian structure with g_F has not been established.*

IX. THE BHATTACHARYYA–BESSEL BRIDGE IN NCG LANGUAGE

This section contains the paper's main structural advance: the reformulation of Paper I's three sub-gaps as a single question about unitary equivalence of spectral triples.

A. Recap: the three sub-gaps

In Paper I, §9.4, we identified two self-adjoint integral operators associated with the modular group:

1. $T_\odot$ on $L^2(\text{strip}, \mu_\text{S})$, with kernel $\odot(s, s') = \zeta((s + s')/2)/\sqrt{\zeta(s)\zeta(s')}$;
2. L_β on $L^2(\mathbb{R}_+, dt)$, with Bessel kernel $J_{2\beta-1}(2\sqrt{st})$ (Mayer's transfer operator via the Bessel isomorphism [14]).

Three sub-gaps stood between these operators and a proof of RH:

(G1) **Intertwining map:** Construct $U : L^2(\text{strip}, \mu_\text{S}) \to L^2(\mathbb{R}_+, dt)$ explicitly.

(G2) **Kernel matching:** Show that $\odot$ and $J_{2\beta-1}(2\sqrt{st})$ are related under U.

(G3) **Parameter reality:** Show that the Fredholm determinant identity $\det(I - zT_\odot) = \det(I - L_{g(z)})$ constrains the spectral parameters to the critical line.

Sub-gap (G3) was identified as the deepest: it is where the eigenvalue-vs-parameter distinction bites.

B. NCG reformulation: one question instead of three

We now have not just operators but *spectral triples*:

The arithmetic triple. $(\mathcal{A}, \mathfrak{H}, D_F)$ on $\mathcal{M}_\zeta$ with the Fisher metric. The propagator is $\odot(s, s')$ (§VI), and the spectral data is encoded in D_F.

The Mayer–Bessel triple. Define $(\mathcal{A}_M, \mathfrak{H}_M, D_B)$ where: $\mathcal{A}_M = C_b^\infty(\mathbb{R}_+)$ (smooth bounded functions on the positive half-line); $\mathfrak{H}_M = L^2(\mathbb{R}_+, dt; \mathbb{C}^2)$ (spinors on $\mathbb{R}_+$ with the Lebesgue measure); D_B is the differential operator whose integral kernel (Green's function) is the Bessel function $J_{2\beta-1}(2\sqrt{st})$.

More precisely, D_B is related to the Hankel transform $\mathcal{H}_\nu$, which is the self-adjoint unitary on $L^2(\mathbb{R}_+, dt)$ with kernel $\sqrt{st}\, J_\nu(st)$. The operator L_β is the composition of $\mathcal{H}_\nu$ with multiplication by a spectral function.

Proposition IX.1 (NCG collapse of sub-gaps)**.** *The three sub-gaps (G1)–(G3) are equivalent to a single condition:*

$$\boxed{\exists \text{ unitary } U : \mathfrak{H} \to \mathfrak{H}_M \text{ such that } U D_F U^* = D_B.} \tag{63}$$

If (63) *holds, then:*

1. *U is the intertwining map (closing G1);*
2. *The propagators match: $U \odot U^* = G_B$ where G_B is the Green's function of D_B (closing G2);*
3. *The spectral parameters are identical:* $\mathrm{Spec}(D_F) = \mathrm{Spec}(D_B)$ *(closing G3).*

[ARGUED]

Argument. If $UD_FU^* = D_B$, then U intertwines the resolvents: $U(D_F - z)^{-1}U^* = (D_B - z)^{-1}$ for all z in the resolvent set. The resolvents are the fundamental solutions (Green's functions) of the respective operators, so the propagators match: this is (2). The spectra of unitarily equivalent operators are identical: $\mathrm{Spec}(D_F) = \mathrm{Spec}(D_B)$. In particular, the spectral *parameters*—the values of β where $\det(1 - L_\beta) = 0$—are determined by the operator D_F, and the parameter-vs-eigenvalue distinction dissolves because unitarily equivalent operators have identical spectral data at every level. This is (3). □

Remark IX.2 (Why this is a genuine advance)**.** *In the operator-only framework of Paper I, sub-gap (G3) was the deepest: showing that two operators have matching Fredholm determinants does not, by itself, constrain the* parameter values *at which eigenvalues equal 1. A matrix $A(t)$ can have real eigenvalues for all t while* $\det(A(t) - I) = 0$ *has complex roots.*

The spectral triple framework eliminates this problem. Unitary equivalence of D_F and D_B is a stronger *condition than Fredholm determinant matching: it requires that the* entire spectral measure *agrees, not just the determinant. This is precisely the "fine spectral structure" that Paper I identified as missing (Table 1, row 6).*

The price is that unitary equivalence is harder to prove than Fredholm determinant matching. But the problem is now one question*, not three.*

C. Connection to Connes–Consani 2025

In November 2025, Connes and Consani [5] constructed self-adjoint operators whose spectra reproduce the lowest zeta zeros with striking numerical accuracy. Their construction proceeds as follows:

1. Start with the canonical spectral triple $(\mathcal{A}_\Lambda, L^2(S_L^1), D_\Lambda)$ of a circle of length $L = 2\log\Lambda$.
2. Define rank-one perturbations using the partial Euler product $\prod_{p\le x}(1 - p^{-1/2-is})^{-1}$ with $x = \Lambda^2$.
3. The perturbed operator is self-adjoint, and its eigenvalues approximate the zeta zeros $\gamma_1, \gamma_2, \ldots$ (the imaginary parts of the non-trivial zeros on the critical line).

Conjecture IX.3 (Relationship to the Fisher–Dirac operator)**.** *The Fisher–Dirac operator D_F, restricted to the critical line $\sigma = \frac{1}{2}$ and truncated to primes $p \le x$, is unitarily equivalent to the rank-one perturbation of Connes–Consani [5] (up to a conformal rescaling by Ω^{-1}).* [CONJECTURED]

The basis for this conjecture:

1. Both operators are self-adjoint.

2. Both involve the Euler product over the same set of primes.
3. Both are perturbations of a "free" operator (D_Λ for Connes–Consani; D_F with $\mathcal{I} = \text{const}$ for the Fisher triple).
4. The spectral data of both is controlled by the zeta function.

The key difference: Connes–Consani's operator lives on a *circle* (compact, 1D), while D_F lives on the *strip* (non-compact, 2D). Restricting D_F to the critical line and compactifying would reduce the dimension and close the topology. Whether this can be done while preserving the spectral data is the content of Conjecture IX.3.

X. RECOGNITION COMPLETENESS AND WEIL POSITIVITY

We now state the paper's most ambitious claim: that the two principal obstructions to proving RH—Connes' Weil positivity and APO's recognition completeness—are manifestations of the same mathematical condition, viewed from different frameworks.

A. Connes' Weil positivity criterion

Connes' main result [3] reduces RH to a positivity condition. The Weil distribution W on $\mathcal{S}(\mathbb{R}_+^*)$ (61) encodes the explicit formulas of number theory as a trace formula on the adele class space. RH is equivalent to:

$$W(f * \tilde{f}) \geq 0 \qquad \forall\, f \in \mathcal{S}(\mathbb{R}_+^*), \tag{64}$$

where $\tilde{f}(x) = \overline{f(x^{-1})}$ and $*$ denotes multiplicative convolution. This is the number-field analogue of the Hodge index theorem, which provides positivity in Weil's proof of RH for function fields.

B. APO's recognition completeness criterion

In Paper I, §10, the RH was reformulated as:

> *The function ξ has no "dark zeros"—zeros invisible to the $\otimes \to \odot \to \oplus$ recognition cycle.*

Formally, this means that the trace-class operator $T_\odot$ on $L^2(\text{strip}, \mu_S)$, restricted to the $\mathbb{Z}_2$-odd subspace V_- (the eigenspace of the functional equation's involution with eigenvalue -1), detects all zeros of ξ in its spectral decomposition. That is: every zero of ξ produces a corresponding eigenvalue of $T_\odot|_{V_-}$, and there are no "phantom eigenvalues" corresponding to non-existent zeros.

C. The equivalence conjecture

Conjecture X.1 (Recognition–Weil equivalence)**.** *The following are equivalent:*

1. ***Weil positivity**: $W(f * \tilde{f}) \geq 0$ for all $f \in \mathcal{S}(\mathbb{R}_+^*)$* [3]*;*
2. ***Recognition completeness**: $T_\odot|_{V_-}$ detects all zeros of ξ* (Paper I, §10)*;*
3. ***Spectral triple unitary equivalence**: $(\mathcal{A}, \mathfrak{H}, D_F) \simeq (\mathcal{A}_M, \mathfrak{H}_M, D_B)$* (§IX B)*;*
4. ***Fredholm identity**: $\det(I - z\,T_\odot) = Z_M(g(z))$* (Paper I, Conjecture 9.1)*.*

Each of these is equivalent to the Riemann Hypothesis. [CONJECTURED]

The content of the conjecture is not that each is equivalent to RH—that is known or straightforward for (1) and asserted for (2)–(4)—but that they are equivalent *to each other*, and that the morphism Φ of Conjecture VIII.1 mediates between them.

Structural argument for (1) $\Leftrightarrow$ (2): Both are positivity conditions on bilinear forms associated with ζ. Weil positivity says $\langle f, Wf\rangle \geq 0$ for all test functions. Recognition completeness says that the quadratic form $\langle \psi, T_\odot \psi\rangle$ captures all spectral structure of ξ. Both fail if and only if there exist "invisible" zeros—zeros that produce negative contributions to the respective quadratic forms. The morphism Φ should map the Weil test functions to sections of the statistical bundle, converting one positivity condition into the other.

Why this doesn't close the gap: Showing (1) $\Leftrightarrow$ (2) does not prove either one. It *unifies* the obstructions: Connes' gap (Weil positivity unproven) and APO's gap (recognition completeness unproven) are the same gap dressed in different mathematical clothing. The unified obstruction is:

Knowing the measure (Solomonoff / Weil) does not determine the support (zero locations).

This is the measure–support gap of Paper I, §10.2, now dressed in NCG clothing.

Remark X.2 (The six faces, revisited)**.** *Paper I identified six faces of the gap (Table 1). The NCG framework clarifies their relationships: faces 1–2 (divergence/overlap, per-prime/collective) are about the product geometry (§VII); faces 3–4 (discrete/scattering, compact/cusp) are about the compactification (§III.8); faces 5–6 (statistics/location, trace-class/fine structure) are about the spectral triple matching (§IX B). The NCG framework packages all six into one: the existence of the unitary equivalence* (63).

XI. NUMERICAL RESULTS

We report the results of systematic numerical tests of the framework's predictions. All computations use 25-digit precision via the `mpmath` library and are reproducible from the scripts provided as supplementary material.

A. Confirmed: $\odot$ asymptotics (Test C)

At $s = 2$ (well inside the convergence region), the quadratic coefficient of the expansion $\odot(s, s+\epsilon) = 1 - a_2\epsilon^2 + O(\epsilon^4)$ was computed for 25 values of ϵ from 10^{-6} to 10^{-1}. The ratio $a_2/(\mathcal{I}(2)/8)$ converges to 1.000163 at $\epsilon = 10^{-6}$, confirming Corollary IV.3 to six significant figures. In the imaginary direction ($\epsilon \to i\epsilon$, probing the t-coordinate), the coefficient flips sign to $-\mathcal{I}(2)/8$, consistent with the Cauchy–Riemann structure. [CONFIRMED TO 6 DIGITS]

B. Confirmed: Fisher poles at zeros (Tests E, F)

The Fisher information $\mathcal{I}(\frac{1}{2} + it)$ was computed at 200 points along the critical line for $10 \leq t \leq 55$. Every local maximum of $|\mathcal{I}|$ lies within $\delta < 0.09$ of a known zero (11 zeros tested; 6 have $\delta < 0.02$). The pole structure was verified quantitatively at the first zero $\gamma_1 \approx 14.135$: the ratio $\mathcal{I}(\frac{1}{2} + i(\gamma_1 + \delta)) \cdot \delta^2$ converges to 1.0006 at $\delta = 0.1$ and to 1.00003 at $\delta = 0.01$, confirming the Hadamard expansion of Proposition VI.5 and the universal deficit angle of Theorem VI.6.

The propagator $|\odot(3, \frac{1}{2} + it)|$ was scanned along the critical line. Poles appear at every zero, and the divergence rate was measured: $|\odot| \cdot \sqrt{\delta} = 1.076 \pm 0.001$ over two orders of magnitude ($\delta \in [0.01, 1]$), confirming $|\odot| \sim \delta^{-1/2}$ from the half-order pole $\sqrt{\zeta(\rho)} = 0$. [CONFIRMED TO 4 DIGITS]

C. Confirmed: universal deficit angle (Test N6)

For each of the first 15 non-trivial zeros, the quantity $A = \mathcal{I}(\frac{1}{2} + i(\gamma + \delta)) \cdot \delta^2$ was computed at $\delta = 0.01$. The results:

$$A = 1.0000 \pm 0.0001 \quad (\mathrm{CV} = 0.01\%,\ n = 15). \tag{65}$$

The universality is exact in the sense of Theorem VI.6: $A = 1$ depends only on the zero being simple, not on its height. Every zero creates an *identical* conical singularity in the Fisher metric.

D. Falsified: D_F eigenvalues as zeta zeros (Test A)

The operator D_F^2 was discretized on a grid of $N = 600$ points along the critical line ($t \in [1, 100]$), and its eigenvalues λ_n were computed. The lowest 20 eigenvalues of $|D_F| = \sqrt{D_F^2}$ are $O(0.1–1)$, while the corresponding zeta zeros are $O(14–80)$. The ratios λ_n/γ_n cluster around 0.012 with coefficient of variation 12%, showing **no proportionality**.

The D_F eigenvalues are instead the standing-wave modes of the chambers between consecutive zeros. Each zero creates an infinite barrier in the conformal factor $\Omega = \sqrt{|\mathcal{I}|} \to \infty$, and the D_F eigenvalues are determined by the chamber widths (the zero spacings $\gamma_{n+1} - \gamma_n$) and the integrated conformal factor within each chamber.

Remark XI.1 (What this falsification means)**.** *The zeros appear in the spectral triple as* metric singularities *(where $\Omega \to \infty$), not as eigenvalues of D_F. This is analogous to point vortices in 2D fluid mechanics: vortex locations are singularities of the velocity field, not eigenvalues of the Laplacian. The correct spectral object carrying zero-location information is the* propagator $\odot$ *(whose poles are the zeros, confirmed in §XI B), not the operator D_F (whose eigenvalues encode the geometry* between *zeros).*

E. Discovered: spectral action extremized on critical line (Tests V1–V4)

The spectral action $S(\sigma) = \sum_n \exp(-\lambda_n(\sigma)^2/\Lambda^2)$ was computed for D_F^2 discretized at $\sigma = 0.5, 0.6, 0.7, 1.5$ with $N = 500$ grid points.

Result V1 (Spectral action). $S(\sigma = 0.5) > S(\sigma = 0.6) > S(\sigma = 0.7) > S(\sigma = 1.5)$ at all five cutoffs tested ($\Lambda = 0.5, 1, 2, 5, 10$). At $\Lambda = 2$: $S(0.5) = 97.2$, $S(0.6) = 77.4$, $S(0.7) = 68.3$, $S(1.5) = 47.5$. The critical line *maximizes* the spectral action.

Result V2 (Total curvature). The integrated Fisher curvature $\int K_F \cdot \Omega\, dt$ (excising $\delta = 0.3$ neighborhoods of zeros) is most negative on the critical line: -102 at $\sigma = 0.5$ versus -61 at $\sigma = 0.6$, with a sign change to $+129$ at $\sigma = 0.7$. The per-chamber curvature is remarkably uniform: mean $= -3.65$, CV $= 17.6\%$ across 12 chambers.

Result V3 ($\mathbb{Z}_2$ symmetry). The Fisher information satisfies $\mathcal{I}(\sigma + it) = \mathcal{I}(1 - \sigma + it)$ to four significant figures at all 16 test points in the (σ, t) plane near γ_1. This confirms the $\mathbb{Z}_2$ symmetry of the $\mathbb{Z}_2$ invariance of O_ξ (Paper I, Theorem 3.2) at the metric level.

Result V4 ($|\mathcal{I}|$ peaked at $\sigma = \frac{1}{2}$). A 9×50 grid in $\sigma \in [0.3, 0.7]$, $t \in [12, 16]$ near the first zero shows that $|\mathcal{I}|$ is maximized at $\sigma = 0.5$ for every value of t near a zero. The critical line is where the singularities are *strongest.*

Conjecture XI.2 (Variational characterization of RH)**.** *Let $Z(s)$ be any Dirichlet series with (a) a functional equation $\xi_Z(s) = \xi_Z(1 - s)$, (b) an Euler product, and (c) only simple non-trivial zeros. Then the Fisher information $\mathcal{I}_Z(s) = d^2/ds^2 \log Z(s)$ has the property that $|\mathcal{I}_Z(\sigma + it)|$, viewed as a function of σ at any fixed t between consecutive zeros, has a* ***single local maximum at*** *$\sigma = \frac{1}{2}$ if and only if all zeros of Z satisfy* $\mathrm{Re}(\rho) = \frac{1}{2}$. [CONJECTURED]

Remark XI.3 (Why "single maximum" matters)**.** *The conjecture's content is the "single versus double maximum" criterion. If a pair of zeros existed at $\sigma_0 + i\gamma$ and $(1 - \sigma_0) + i\gamma$ with $\sigma_0 \neq \frac{1}{2}$, the dominant-pole approximation $|I_{\mathrm{dom}}(\sigma)| \approx |\, 1/(\sigma - \sigma_0 + i\delta)^2 + 1/(\sigma - 1 + \sigma_0 + i\delta)^2 \,|$ has* two *local maxima (one near σ_0, one near $1 - \sigma_0$) with a local minimum at $\sigma = \frac{1}{2}$. For a self-paired zero at $\frac{1}{2} + i\gamma$, there is a* single *maximum at $\sigma = \frac{1}{2}$. The actual Fisher information of ζ exhibits a single maximum everywhere tested (Result V4), consistent with all zeros being on the critical line. This is a non-circular consistency test: the prediction (one versus two maxima) is independent of the input (measured $|\mathcal{I}|$ profile).*

Remark XI.4 (Scope and limitations)**.** *The Fisher metric of ζ is not a free parameter being varied—it is uniquely determined by the Euler product through Chentsov's theorem [10]. The conjecture therefore characterizes a* property *of the Fisher metric arising from any L-function with the stated structure, not a variational condition over an abstract class of metrics. The connection to RH is: if the single-maximum property can be proven analytically (from the Hadamard expansion and the Euler product), it implies all zeros are on the critical line.*

The singularities have order $\alpha = -1$ (logarithmic), which sits at the critical boundary of the theory of conical singularities on surfaces [15]. Troyanov's theorems require $\alpha > -1$ and do not apply. The conjecture must be established by methods native to the Fisher-geometric setting.

The three layers of evidence:

1. **Layer 1 (proven).** $\mathbb{Z}_2$ symmetry forces $\partial|\mathcal{I}|/\partial\sigma = 0$ at $\sigma = \frac{1}{2}$ for every t.
2. **Layer 2 (proven).** The Hadamard expansion and universality ($A = 1$, Theorem VI.6) give $\partial^2|\mathcal{I}|/\partial\sigma^2 < 0$ at $\sigma = \frac{1}{2}$ (Theorem VI.9): the critical line is a strict local maximum of $|\mathcal{I}|$.
3. **Layer 3 (open).** The local maximum is also the *unique* maximum: no off-line pair creates a higher double-peaked profile when the smooth background from all other zeros is included. This requires controlling the sum $\sum_{\rho' \neq \rho} |s - \rho'|^{-2}$ uniformly, which is a statement about the collective behavior of the zero ensemble. The Euler product structure (per-prime independence) is likely essential here, as the $\mathbb{Z}_2$ symmetry alone does not resolve single versus double maxima.

XII. CONCLUSION

We constructed the spectral triple $(\mathcal{A}, \mathfrak{H}, D_F)$ for the Fisher–Rao manifold of the zeta distribution and proved that Connes' distance formula recovers the Fisher–Rao distance (Theorem IV.1), establishing that the arithmetic recognition operator is the cosine of half the spectral distance (Corollary IV.3). The Dixmier trace recovers the Fisher volume form (Theorem V.1), and the spectral triple satisfies Connes' axioms 1–6 (Theorem III.7).

Beyond the proven core, the paper's two most significant contributions are the universal deficit angle theorem (Theorem VI.6) and the numerical discovery that the spectral action is extremized on the critical line (§XI E).

The universal deficit angle. Every simple zero of ζ creates an identical conical singularity in the Fisher metric, with coefficient $A = 1$ independent of the zero's height (Theorem VI.6, confirmed numerically to $\mathrm{CV} = 0.01\%$). This places all zeros on the same geometric footing: they are interchangeable singularities. If *any* variational principle constrains one zero to the critical line, the universality forces all others there as well.

The variational direction. The spectral action $S = \sum f(\lambda_n/\Lambda)$ is maximized on the critical line (V1), the total Fisher curvature is most negative there (V2), the $\mathbb{Z}_2$ symmetry is exact (V3), and $|\mathcal{I}|$ is peaked at $\sigma = \frac{1}{2}$ (V4). This convergence of extremality conditions suggests a variational characterization of RH (Conjecture XI.2): the zeros are where they are because the on-line configuration extremizes a curvature functional.

The falsification. The D_F eigenvalues do *not* reproduce the zeta zeros (§XI D). The zeros are metric singularities, not operator eigenvalues. This distinction—between the propagator's poles and the operator's spectrum—is the sharpest lesson of the numerical investigation and corrects a tempting but incorrect expectation.

The complexity bound

Independent of the variational direction, we record:

Proposition XII.1 (Complexity of zero configurations)**.** *Among all zero configurations compatible with the functional equation $\xi(s) = \xi(1-s)$, the on-line configuration (all zeros at $\sigma = \frac{1}{2}$) has minimal Kolmogorov complexity. Specifically, $K(\text{on-line}) \leq K(\text{off-line}) + O(1)$.* [PROVEN]

Proof. An on-line configuration is specified by the sequence of heights $\{\gamma_k\}$. An off-line configuration requires the heights *and* the σ-deviations $\{\delta_k\}$. The on-line configuration is a projection (setting all $\delta_k = 0$), and projections cannot increase Kolmogorov complexity [16]. □

This is a theorem, not a conjecture. It says that the on-line configuration is the *simplest* zero arrangement compatible with the functional equation. It does not prove that ζ's zeros are on-line—the primes determine ζ uniquely, and in principle the primes could encode off-line deviations. Whether the primes encode the simplest possible zero configuration is precisely the content of RH.

The APO perspective

In the framework of the Axioms of Pattern Ontology (Appendix D), the Integration$^{\text{APO}}$ operator $\oplus$ selects the most parsimonious representative of each equivalence class. The functional equation creates a $\mathbb{Z}_2$ equivalence $s \sim 1 - s$; the on-line configuration is the unique $\mathbb{Z}_2$-fixed representative with minimal complexity (Proposition XII.1). APO therefore *predicts* that the zeros are on the critical line.

This prediction is not a proof of RH. It is a *test* of the APO axioms: if RH is true, the prediction is confirmed and APO's principle of parsimony applies to the arithmetic of the zeta function. If RH is false, the prediction fails and the axioms require revision. The distinction between "APO predicts on-line" and "mathematics proves on-line" is the distinction between a physical theory's prediction and a mathematical theorem. The former can be confirmed or falsified; the latter requires proof.

What we do not claim

This paper does not prove the Riemann Hypothesis. It constructs machinery (the spectral triple, the deficit angle theorem, the concavity theorem), identifies a precise variational criterion (Conjecture XI.2), proves the first two layers of evidence for that criterion, and formulates the remaining gap (Layer 3: uniform valley-versus-background bound) as a concrete problem. The Kolmogorov complexity bound (Proposition XII.1) is proven. The APO prediction is stated. The gap is real and precisely identified.

Open directions

1. **Layer 3: the valley-versus-background bound.** For small σ-deviations (σ_0 near $\frac{1}{2}$), the paired-pole profile has no valley at all—the two nearly-coincident poles merge into a single peak at $\sigma = \frac{1}{2}$. For large deviations, a valley appears but the convergent background sum from distant zeros may fill it. A uniform bound showing the single-maximum property holds for all heights t would close the gap. This likely requires the Euler product structure.
2. **L-function universality.** Test the single-maximum criterion and the universal deficit angle for Dirichlet L-functions $L(s,\chi)$ with proper completion. If the structure persists across all L-functions with functional equations and Euler products, the variational characterization is universal.
3. **Yang–Mills action of $\odot$.** Compute $\mathrm{Tr}_\omega(|[D_F,\odot]|^2\,|D_F|^{-2})$.
4. **Lean formalization.** Theorems VI.6 and VI.9 and Proposition XII.1 are formalizable in Lean 4/Mathlib.

ACKNOWLEDGMENTS

The author thanks Claude (Anthropic) for collaborative development of the spectral triple construction, numerical verification of all claims (including the critical falsification of Prediction 3 and the discovery of the variational structure), axiom verification, and identification of the connection to the Connes–Consani program. The intellectual debts to Alain Connes, whose 1994 textbook provided the structural blueprint, and to Sosuke Ito and Andreas Dechant, whose stochastic information geometry informed the Fisher–Rao framework, are gratefully acknowledged.

Appendix A: Detailed axiom verification

We provide the explicit computations for the axiom verification of Theorem III.7.

Christoffel symbols of the Fisher metric. For the conformal metric $g_F = \Omega^2\,(d\sigma^2 + dt^2)$ with $\Omega = \sqrt{\mathcal{I}(s)}$, the Christoffel symbols are:

$$\begin{aligned} \Gamma^\sigma_{\sigma\sigma} &= \partial_\sigma \log\Omega, & \Gamma^\sigma_{tt} &= -\partial_\sigma \log\Omega, \\ \Gamma^t_{tt} &= \partial_t \log\Omega, & \Gamma^t_{\sigma\sigma} &= -\partial_t \log\Omega, \\ \Gamma^\sigma_{\sigma t} &= \partial_t \log\Omega, & \Gamma^t_{\sigma t} &= \partial_\sigma \log\Omega. \end{aligned} \tag{A1}$$

Spin connection. The vielbein is $e^1_\sigma = \Omega$, $e^2_t = \Omega$, $e^1_t = e^2_\sigma = 0$. The spin connection has one independent component $\omega_{12} = -\partial_t \log\Omega\,d\sigma + \partial_\sigma \log\Omega\,dt$.

Gaussian curvature. $K_F = -\Omega^{-2}\,\Delta\log\Omega$ where $\Delta = \partial_\sigma^2 + \partial_t^2$ is the flat Laplacian. With $\Omega = \mathcal{I}^{1/2}$: $K_F = -\mathcal{I}^{-1}\,\Delta(\frac{1}{2}\log\mathcal{I}) = -\frac{1}{2\mathcal{I}}\left(\frac{\mathcal{I}''}{\mathcal{I}} - \frac{(\mathcal{I}')^2}{\mathcal{I}^2} + \frac{\partial_t^2\mathcal{I}}{\mathcal{I}} - \frac{(\partial_t\mathcal{I})^2}{\mathcal{I}^2}\right)$, where primes denote ∂_σ.

Weyl law verification. For a compact domain $K \subset \mathcal{M}_\zeta$ with smooth boundary, Weyl's theorem gives $N(\Lambda) = \frac{\mathrm{Vol}_F(K)}{4\pi}\Lambda^2 + O(\Lambda)$. The Fisher volume is $\mathrm{Vol}_F(K) = \int_K \mathcal{I}(s)\,d\sigma\,dt$, which is finite for any bounded K away from the zeros of ζ (where $\mathcal{I} \to \infty$).

For the Solomonoff-compactified manifold $\bar{\mathcal{M}}_\zeta$, the weighted volume $\mathrm{Vol}_S = \int 2^{-K(s)}\,\mathcal{I}(s)\,d\sigma\,dt$ is finite by the Kraft inequality, ensuring that the Dixmier ideal condition holds globally.

Appendix B: Master proof status table

Appendix C: The APO–NCG–standard mathematics dictionary

Appendix D: APO foundations and terminology

The Axioms of Pattern Ontology (APO) provide the philosophical and interpretive scaffolding for the constructions in this paper. APO is *not* required for any of the rigorous results (§III–§IV, Theorems VI.6–VI.9), which stand as pure mathematics. APO motivates the choice of objects, suggests which questions to ask, and provides interpretive language. This appendix makes the philosophical commitments explicit.

TABLE II. Proof status for all results in this paper.

#	Result	Status	Section	Dependencies
1	Spectral triple satisfies axioms 1–6	PROVEN	§III E	Fisher metric (Paper I)
2	Connes distance = Fisher–Rao distance	PROVEN	§IV	1, Lemma III.9
3	$\odot = \cos(d_{\mathrm{FR}}/2)$	PROVEN	§IV	2
4	Dixmier trace = Fisher volume	PROVEN	§V	1, Weyl law
5	$\odot$ as propagator of D_F	ARGUED	§VI	1, 3
6	Product geometry strip × primes	ARGUED	§VII	1, Euler product
7	Three sub-gaps → unitary equivalence	ARGUED	§IX	1, 5
8	Bridge to adele class space	CONJECTURED	§VIII	6
9	Recognition compl. ⇔ Weil positivity	CONJECTURED	§X	7, 8
10	Axiom 7 via Solomonoff compactification	ARGUED	§III.8	1, Paper I §8
11	Universal deficit angle $A = 1$	PROVEN	§VI E	Hadamard, simplicity of zeros
12	Strict concavity of $\lvert\mathcal{I}\rvert$ at $\sigma = \frac{1}{2}$	PROVEN	§VI F	11, $\mathbb{Z}_2$ symmetry
13	D_F eigenvalues $\neq$ zeta zeros	PROVEN (falsification)	§XI D	1
14	Spectral action maximized at $\sigma = \frac{1}{2}$	confirmed	§XI E	1, 12
15	Variational characterization of RH	CONJECTURED	§XI E	11, 12
16	$K(\text{on-line}) \leq K(\text{off-line}) + O(1)$	PROVEN	§XII	Functional equation

TABLE III. Dictionary between APO, NCG, and standard mathematical concepts. Status: P = Proven identification, A = Argued, C = Conjectured.

APO concept	NCG concept	Standard mathematics	Status
$\otimes$ (Distinction)	Algebra $\mathcal{A}$	Sufficient statistics	P
$\odot$ (Recognition)	$\cos(d_{\mathrm{FR}}/2)$	Bhattacharyya coefficient	P
$\oplus$ (Integration)	Quotient spectral triple	Orbit projection	A
Fisher metric g_F	$\lVert[D_F, \cdot]\rVert$	Riemannian metric	P
Solomonoff μ_{S}	Dixmier trace Tr_ω	Volume form	A
$\otimes \to \odot \to \oplus$ cycle	Spectral flow	Heat flow	C
Recognition completeness	Weil positivity	Riemann Hypothesis	C
Dark zeros	Noncritical resonances	Off-line zeros	C
Euler product	Adelic factorization	Product geometry $E \times F$	P
Dephasing ($\mathbb{Z}_2$)	Real structure J	Functional equation	A

Axiom 1: Differentiation$^{\mathrm{APO}}$ ($\otimes$). Patterns differentiate: the primitive relation "this is not that" exists. Formally, this is a binary partition of a sample space. Pure differentiation without comparison is structureless—a "white noise" of perfectly differentiated patterns with no measurable distances, because white noise measured against white noise yields no information. The mathematical consequence: a space of probability distributions exists, but has no natural metric.

Axiom 2: Reflection$^{\mathrm{APO}}$ ($\odot$). Patterns compare themselves to each other: the inner product $\odot(p, q) = \sum_x \sqrt{p(x)\, q(x)}$ (the Bhattacharyya coefficient) provides the measurement that Differentiation$^{\mathrm{APO}}$ alone cannot. This is the "inside-out" premise: no external observer is needed; patterns observe each other. Formally, $\odot$ is an inner product on the square-root embedding space, making the set of distributions into a geometric object. Chentsov's theorem [10] then forces the Fisher metric as the unique distance structure compatible with $\odot$.

Derived: Integration$^{\mathrm{APO}}$ ($\oplus$). Not a separate axiom. $\oplus$ is the quotient map that identifies patterns indistinguishable under $\odot$: if $\odot(p, q) = 1$ then p and q are "the same pattern." The quotient selects the most parsimonious representative of each equivalence class—algorithmically, the one with minimum Kolmogorov complexity. This is the mathematical principle of parsimony: $2 + 2 \to 4$ is the same arithmetic content in a shorter representation. Integration$^{\mathrm{APO}}$ compresses without losing useful information.

Kolmogorov complexity is supervenient, not axiomatic. APO does not assume K as a primitive. Rather, the requirement that distances be meaningful (Differentiation$^{\mathrm{APO}}$ must *measure* something) forces a prior on the space of patterns that separates signal from noise. Solomonoff's theorem [11] identifies this prior as 2^{-K}—the unique universal prior compatible with computability. K is therefore a consequence of the axioms, not an additional assumption.

The variational principle$^{\text{APO}}$ (interpretive, not proven). The $\otimes \to \odot \to \oplus$ cycle, iterated, is conjectured to select configurations that simultaneously maximize statistical distinguishability ($|\mathcal{I}|$ peaked) and minimize description complexity ($\mathbb{Z}_2$ quotient). On the critical line, these two objectives—usually in tension—are simultaneously optimized. The mathematical content of this philosophical claim is the variational Conjecture XI.2. The philosophical content is that arithmetic "prefers" the critical line for the same reason that physics "prefers" geodesics: it is the path of maximum information at minimum cost.

Terminology conventions. APO terms that conflict with standard mathematical usage carry the superscript APO (e.g., Integration$^{\text{APO}} \neq \int f\, d\mu$). The circumpunct $\odot$ and the operators $\otimes$, $\oplus$ are APO notation. See the full external terminology dictionary (`dictionary_RH.md`) for the complete mapping between APO, NCG, and standard mathematical concepts.

Appendix E: Numerical verifications

All computations were performed in Python 3.12 with `mpmath` (25-digit precision) and `scipy`. Scripts are available as supplementary material.

Universal deficit angle (Theorem VI.6). For each of the first 15 non-trivial zeros $\rho_k = \frac{1}{2} + i\gamma_k$, we computed $A_k = \mathcal{I}(\frac{1}{2} + i(\gamma_k + 0.01)) \cdot (0.01)^2$. Results: $A_k = 1.0000 \pm 0.0001$ (1σ), CV $= 6.5 \times 10^{-5}$. All 15 values positive.

Spectral action (Conjecture XI.2). D_F^2 was discretized on $N = 500$ grid points along the critical line. At cutoff $\Lambda = 2$: $S(0.5) = 97.17$, $S(0.6) = 77.41$, $S(0.7) = 68.26$, $S(1.5) = 47.52$. The ordering $S(0.5) > S(\sigma)$ for $\sigma \neq 0.5$ holds at all five cutoffs tested ($\Lambda = 0.5, 1, 2, 5, 10$).

$\mathbb{Z}_2$ symmetry of Fisher metric. $|\mathcal{I}(\sigma + it)/\mathcal{I}(1 - \sigma + it)|$ was computed at 16 points in the (σ, t) plane near γ_1. Maximum deviation from unity: 0.0016 (at $\sigma = 0.30$, $t = 13.0$, far from any zero). At the zero ($t = 14.135$): deviation $< 10^{-4}$.

$\odot$ short-distance expansion. At $s = 2$: $(1 - \odot(2, 2 + \epsilon))/\epsilon^2 \to \mathcal{I}(2)/8 = 0.11056$ as $\epsilon \to 0$. Ratio at $\epsilon = 10^{-6}$: 1.000163.

Propagator divergence rate. $|\odot(3, \frac{1}{2} + i(\gamma_1 - \delta))| \cdot \sqrt{\delta} = 1.076 \pm 0.001$ for $\delta \in [0.01, 1.0]$ (7 test points).

Appendix F: Approaches tested and rejected

(R1) Direct spectral triple on $\ell^2(\mathbb{N})$. One might try to build a spectral triple with $\mathfrak{H} = \ell^2(\mathbb{N})$ directly (using the sample space rather than the parameter space). This fails because there is no natural Dirac operator on $\mathbb{N}$ whose commutator norm recovers a useful distance. The "number-theoretic Dirac operator" would need to encode the multiplicative structure of $\mathbb{N}$, but the resulting object is the Bost–Connes Hamiltonian, which is a *dynamical* operator (generating time evolution), not a *geometric* one (encoding distances). [REJECTED: WRONG TYPE OF OPERATOR.]

(R2) Bost–Connes system as spectral triple. The Bost–Connes system [13] is a quantum statistical mechanical system whose KMS states at inverse temperature $\beta > 1$ are the zeta distributions P_s. One might try to extract a spectral triple from its dynamics. The KMS states provide the algebra (evaluation functionals) and the GNS construction provides a Hilbert space, but the Hamiltonian $H = \log N$ (the logarithm of the number operator) is not a Dirac operator: it has the wrong spectral asymptotics ($\lambda_n = \log n$, not $\lambda_n \sim n^{1/d}$) and does not give a finite metric dimension. [REJECTED: WRONG SPECTRAL ASYMPTOTICS.]

(R3) Selberg zeta as Fredholm determinant of D_F. One might hope that $\det(I - zD_F^{-1}) = Z_M(g(z))$ directly, connecting the Fisher–Dirac operator to the Selberg zeta function without passing through Mayer's operator. This fails because D_F has continuous spectrum on the non-compact strip $\mathcal{M}_\zeta$, and Fredholm determinants are only defined for operators with discrete spectrum (or operators of the form $I + K$ with K trace-class). The Solomonoff compactification provides a trace-class operator $T_\odot$ but its Fredholm determinant is $\det(I - zT_\odot)$, not $\det(I - zD_F^{-1})$. [REJECTED: CONTINUOUS SPECTRUM ON NON-COMPACT SPACE.]

---

[1] D. Christian (2026), aPO Initiative working paper. The Arithmetic Recognition Operator: Information Geometry, the Modular Surface, and a Conjectured Path to the Riemann Hypothesis.

[2] A. Connes, *Noncommutative Geometry* (Academic Press, 1994).

[3] A. Connes, Selecta Math. **5**, 29 (1999).

[4] A. Connes and C. Consani, Adv. Math. **291**, 274 (2016), arXiv:1405.4527 (2014).

[5] A. Connes and C. Consani, arXiv:2511.22755 (2025), self-adjoint operators whose spectra reproduce the lowest zeta zeros via rank-one perturbations.
[6] J. M. Gracia-Bondía, J. C. Várilly, and H. Figueroa, *Elements of Noncommutative Geometry* (Birkhäuser, 2001).
[7] A. Connes, J. Noncommut. Geom. **7**, 1 (2013).
[8] S. Amari and H. Nagaoka, *Methods of Information Geometry* (AMS/Oxford, 2000).
[9] S. Amari, *Information Geometry and Its Applications* (Springer, 2016).
[10] N. N. Chentsov, Translations of Mathematical Monographs **53** (1982), originally published in Russian, 1972. Proves uniqueness of the Fisher metric on statistical manifolds.
[11] R. J. Solomonoff, Inform. Control **7**, 1 (1964).
[12] A. Connes and M. Marcolli, *Noncommutative Geometry, Quantum Fields and Motives* (AMS, 2008).
[13] J.-B. Bost and A. Connes, Selecta Math. (N.S.) **1**, 411 (1995).
[14] D. H. Mayer, Commun. Math. Phys. **130**, 311 (1991), proves L_β is isomorphic to a Hankel transform with Bessel kernel on Hardy space $H_{-1/2}$; real spectrum for $\beta > 1/2$.
[15] M. Troyanov, Trans. AMS **324**, 793 (1991).
[16] M. Li and P. Vitányi, *An Introduction to Kolmogorov Complexity and Its Applications*, 3rd ed. (Springer, 2008).