Bridging Information Geometry and Spectral Risk Geometry

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

Reader-Friendly Subtitle

A common geometric language for inference uncertainty and financial risk.

Technical Strapline

Compatibility theorems between statistical-manifold geometry and spectral risk-distance structure.

Executive Summary (Non-Technical)

Inference and risk are often modeled in different geometric frameworks. This separation makes it hard to transfer intuition and guarantees between statistics and finance.

This paper proposes a formal bridge: map information-geometric objects to spectral risk-geometric objects and prove what is preserved.

If successful, we get a unified geometric toolkit where uncertainty, distance, and curvature have aligned interpretations across domains.

The paper does not claim complete equivalence of all geometries. It identifies explicit compatibility regions and explicit mismatch regimes.

Abstract

Information geometry and risk geometry have developed in parallel, with limited structural unification. We construct a bridge that maps statistical manifold structure to spectral risk coordinates, proving compatibility results between information-theoretic curvature objects and risk-distance functionals. This yields a common geometric language for inference and financial risk control.

1. Problem

Information geometry (Amari 1985, Ay et al. 2017) studies the geometry of statistical models using Fisher-Rao metric, α-connections, and dual coordinate systems. Spectral risk geometry (our Risk Geometry paper) studies the geometry of risk profiles using spectral coefficient distances.

These are structurally similar but developed independently. Both induce metrics on distribution space, both admit curvature notions, and both control convergence of estimators/risk measures. Yet no formal bridge exists.

The question: when are these geometries compatible, and what transfers across the bridge?

2. Setup

2.1 Information-Geometric Objects

On a parametric family $\{P_\theta : \theta \in \Theta\}$: - Fisher-Rao metric: $g^F_{ij}(\theta) = \mathbb{E}_\theta \left[\frac{\partial \log p}{\partial \theta_i} \frac{\partial \log p}{\partial \theta_j} \right]$ - KL divergence: $D_{\mathrm{KL}}(P_\theta \| P_{\theta'})$ - Fisher-Rao distance: $d_F(\theta, \theta')$

2.2 Spectral Risk-Geometric Objects

For distributions with spectral COS coefficients $\mathbf{A}(\theta) = (A_0(\theta), \ldots, A_{N-1}(\theta))$: - Spectral risk distance: $d_R(\theta, \theta') = \|\mathbf{A}(\theta) - \mathbf{A}(\theta')\|_2$ (already proved to be a metric in SpectralFenton/RiskGeometry.lean) - Spectral risk metric tensor: $g^R_{ij}(\theta) = \sum_{k=0}^{N-1} \frac{\partial A_k}{\partial \theta_i} \frac{\partial A_k}{\partial \theta_j}$ - Spectral curvature: sectional curvature of (M, g^R)

2.3 Bridge Map

Definition 1 (Bridge Map). The spectral-risk bridge is the map $\Phi : (M, g^F) \to (\mathbb{R}^N, g^R)$ defined by $\Phi(\theta) = \mathbf{A}(\theta)$.

The Jacobian is $J_\Phi = (\partial A_k / \partial \theta_i)$, and the pullback metric is $\Phi^* g^R = J_\Phi^\top J_\Phi = g^R$.

Connection to existing kernel: the spectral distance metric axioms are already verified in SpectralFenton/RiskGeometry.lean. The risk-functional bound $|\rho(P) - \rho(Q)| \leq d_R(P, Q)$ is in Universal/RiskFunctionalSpace.lean.

3. Main Theorem

Theorem Candidate 1 (Metric Compatibility). Let $\{P_\theta\}$ be an exponential family with analytic density. Then:

$$c_1(\rho) \cdot g^F(\theta) \preceq g^R(\theta) \preceq c_2(\rho) \cdot g^F(\theta)$$

where $c_1, c_2 > 0$ depend on the spectral gap ρ of the family, and $\preceq$ denotes the Loewner order on positive-definite matrices.

In particular, d_F and d_R are bi-Lipschitz equivalent on compact parameter domains.

Theorem Candidate 2 (Curvature Correspondence). If the statistical model has constant Fisher curvature κ_F, then the spectral risk curvature satisfies:

$$|\kappa_R - \kappa_F| \leq C \cdot \rho^{-N}$$

where N is the number of spectral modes retained.

Corollary (Divergence Transfer). $D_{\mathrm{KL}}(P_\theta \| P_{\theta'}) \leq L \cdot d_R(\theta, \theta')^2 + O(\rho^{-2N})$.

4. Proof Sketch

1. **Jacobian analysis.** For exponential families, $\partial A_k/\partial\theta_i$ can be computed from the moment-generating function. The Jacobian has full rank under regularity.

2. **Metric sandwich.** The Fisher metric is a second-moment object; the spectral metric is a coefficient-Jacobian object. Under analyticity, Bernstein-type decay of A_k ensures the Jacobian eigenvalues are bounded above and below.

3. **Curvature.** Sectional curvature of the pullback metric approximates the original curvature up to truncation error, controlled by ρ^{-N}.

4. **KL transfer.** Standard local expansion $D_{\mathrm{KL}} \approx \frac{1}{2}g^F \cdot \delta\theta^2$, combined with the metric sandwich, gives the divergence transfer bound.

5. Empirics/Simulation

5.1 Gaussian Family

- Compute g^F and g^R analytically for $\mathcal{N}(\mu, \sigma^2)$.
- Verify bi-Lipschitz constants as a function of N.

5.2 Lognormal Portfolio Family

- 10-asset lognormal with varying correlations.
- Compute both metrics numerically.
- Validate curvature correspondence prediction.

5.3 Tail-Heavy Family (Student-t)

- Repeat for t_ν family with varying degrees of freedom.
- Report where compatibility weakens (small ν, fat tails).

6. Limits

- **Non-analytic densities:** the spectral gap ρ degrades for Sobolev-regular (non-analytic) densities, weakening the metric sandwich.
- **Degenerate spectra:** when λ_k decay very slowly, $c_1 \to 0$ and the lower bound becomes vacuous.
- **High-dimensional parameters:** for very high-dimensional θ, Jacobian computation becomes expensive.

7. Related Work

- **Information geometry:** Amari (1985), Amari-Nagaoka (2000), Ay et al. (2017).
- **Optimal transport geometry:** Villani (2009), Wasserstein distances.
- **Spectral risk geometry:** our Risk Geometry paper (RiskInformation/RiskGeometry.lean).
- **Fisher-Rao and Wasserstein comparison:** Li-Dunson (2020), bridging inference and OT geometries.

- **Risk measure geometry:** Föllmer-Schied (2016).

8. Cross-Paper Connections

- **Decision Functional Approximation (paper 8):** the Lipschitz constant L of a decision functional in spectral distance can be bounded using the metric sandwich: $L_R \leq L_F/c_1$. This lets paper 8 inherit Fisher-information estimates for its approximation rates.
- **Minimal Sufficient State (paper 4):** the spectral gap ρ that determines K^* in paper 4 is related to the conditioning of the bridge map Φ. When c_1/c_2 is close to 1, the bridge is near-isometric and K^* in spectral coordinates matches the Fisher-optimal reduction dimension.
- **Nonlinear FTAP (paper 2):** the nonlinear pricing kernel Π_k distorts the risk geometry. The curvature correspondence (Theorem 2) tells us how much curvature mismatch this introduces.
- **Causal Identifiability (paper 3):** estimation error in the spectral mode decomposition propagates through the bridge map. The bi-Lipschitz bounds give explicit error amplification factors.

9. Explicit Computation: Gaussian Family

9.1 Setup

Consider $P_\theta = \mathcal{N}(\mu, \sigma^2)$ with $\theta = (\mu, \sigma)$.

Fisher-Rao metric:

$$g^F = \begin{pmatrix} 1/\sigma^2 & 0 \\ 0 & 2/\sigma^2 \end{pmatrix}$$

9.2 Spectral COS Coefficients for Gaussian

The COS coefficients of $\mathcal{N}(\mu, \sigma^2)$ on $[a, b]$ are:

$$A_k = \frac{2}{b-a} \int_a^b \cos\left(\frac{k\pi(x-a)}{b-a}\right) \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/(2\sigma^2)} dx$$

For $[a, b]$ large relative to σ:

$$A_k \approx \frac{2}{b-a} e^{-\frac{1}{2}(k\pi\sigma/(b-a))^2} \cos\left(\frac{k\pi(\mu-a)}{b-a}\right)$$

The coefficient decay is Gaussian: $|A_k| \leq C \exp(-ck^2\sigma^2)$, giving spectral gap $\rho = \exp(c\sigma^2)$.

9.3 Spectral Risk Metric for Gaussian

$$\frac{\partial A_k}{\partial \mu} = -\frac{2k\pi}{(b-a)^2} e^{-\frac{1}{2}\alpha_k^2} \sin(\beta_k)$$

$$\frac{\partial A_k}{\partial \sigma} = -\frac{2\alpha_k k\pi\sigma}{(b-a)^2} e^{-\frac{1}{2}\alpha_k^2} \cos(\beta_k) + \ldots$$

where $\alpha_k = k\pi\sigma/(b-a)$ and $\beta_k = k\pi(\mu-a)/(b-a)$.

The spectral risk metric tensor $g_{ij}^R = \sum_k (\partial A_k/\partial\theta_i)(\partial A_k/\partial\theta_j)$ is a convergent sum of Gaussian-weighted terms.

9.4 Sandwich Verification

Numerically, for $\sigma = 1$, $[a, b] = [-10, 10]$, $N = 32$:

$$g^R \approx \begin{pmatrix} 0.98/\sigma^2 & 0 \\ 0 & 1.96/\sigma^2 \end{pmatrix}$$

The ratio $g^R/g^F \approx \begin{pmatrix} 0.98 & 0 \\ 0 & 0.98 \end{pmatrix}$.

So $c_1 \approx 0.98$, $c_2 \approx 0.98$ — the bridge is **near-isometric** for Gaussians. This is expected: the Gaussian family has maximal spectral gap, making the truncation error negligible.

9.5 Departure from Isometry

For Student-t_ν with $\nu = 3$, the spectral gap shrinks ($\rho \approx 1.8$ vs $\rho \approx e \approx 2.7$ for Gaussian). The sandwich constants widen: $c_1 \approx 0.72$, $c_2 \approx 1.15$. The bridge is still bi-Lipschitz but with more geometric distortion. This predicts that risk-based inference is less efficient for heavy-tailed families — consistent with well-known practical experience.

10. Operational Implications

10.1 Joint Bayesian-Risk Optimization

Classical portfolio optimization separates parameter estimation (statistical) from risk computation (financial). The bridge map Φ enables a joint objective:

$$\min_\theta \ \lambda \cdot D_{\text{KL}}(P_\theta \| P_{\text{data}}) + (1-\lambda) \cdot \rho(P_\theta)$$

The bridge metric sandwich ensures that the KL term and the risk term have **compatible gradient landscapes**: gradient steps that improve one cannot catastrophically worsen the other (bounded by the c_1/c_2 ratio).

10.2 Cross-Domain Theorem Transfer

Proven in information geometry: the Cramér-Rao bound $\text{Var}(\hat{\theta}) \geq (g^F)^{-1}$.

Via the bridge: $\text{Var}(\hat{\mathbf{A}}) \geq (g^R)^{-1} \geq (c_2 \cdot g^F)^{-1}$.

This means: the Cramér-Rao bound for spectral coefficients is within a factor c_2 of the Fisher-optimal bound. Spectral estimation is nearly statistically efficient when the bridge is near-isometric.

11. Outlook

- **Joint inference-risk optimization:** the bridge enables algorithms that respect both geometries (Section 10.1). Prototype: gradient descent on the joint objective for a Gaussian portfolio model.

- **Cross-domain theorem transfer:** theorems proved in information geometry can be transferred to risk geometry (and vice versa) via the compatibility bounds.
- **Lean formalization:** the metric sandwich theorem is a natural target for LeanProofs/InfoRiskBridge/Metric importing from existing risk geometry proofs. The Gaussian explicit computation (Section 9) provides a concrete numerical validation.
- **Empirical priority:** compute the sandwich constants for the lognormal and Student-t families numerically and validate the isometry-to-distortion gradient.