

Article

Not peer-reviewed version

---

# Geodesic Execution Slippage: A Statistical Physics Framework for Cryptocurrency Liquidity Risk

---

[Ntebogang Dinah Moroke](#)\* and [Lebotsa Daniel Metsileng](#)

Posted Date: 12 May 2026

doi: 10.20944/preprints202605.0757.v1

Keywords: Fisher information metric; geodesic execution slippage; Riemannian manifold; persistent homology; Curvature-Fragmentation Law; Wasserstein distance; econophysics; market microstructure; early warning signals; complex financial systems



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Geodesic Execution Slippage: A Statistical Physics Framework for Cryptocurrency Liquidity Risk

Ntebogang Dinah Moroke \* and Lebotsa Daniel Metsileng 

Department of Statistics and Operations Research, Faculty of Economic and Management Sciences, North-West University, Mafikeng Campus, Private Bag X2046, Mmabatho 2735, South Africa

\* Correspondence: Ntebo.Moroke@nwu.ac.za

## Abstract

Standard cryptocurrency transaction cost models assume flat geometry and assign execution cost as a proportional fee. This paper tests whether replacing flat-fee models with a unified Riemannian execution cost framework improves out-of-sample prediction of realised execution costs. The framework identifies execution slippage as the geodesic arc length on the Fisher information manifold of a Markov-switching GARCH maximum-entropy model, augmented by a joint curvature-topological fragmentation alarm derived from the same parameter vector. Ablation confirms that each geometric component contributes uniquely: removing the geodesic increases mean squared prediction error by 2.9%, removing topological data analysis by 2.1%, and removing curvature by 1.5%. No subset matches the full framework. On five major cryptocurrency markets (BTC, ETH, XRP, LTC, BCH) over 2,253 daily observations, the integrated framework achieves the lowest prediction error on all five assets and is the sole model retained in the Model Confidence Set at the 10% significance level against six benchmarks, including Amihud, Kyle  $\lambda$ , and Almgren and Chriss. A joint curvature-topological alarm fires a median of two days before price-based circuit breaker thresholds across four crisis episodes, including the Terra collapse of May 2022 and the FTX bankruptcy of November 2022. The framework requires no additional data or free parameters beyond the upstream estimation pipeline.

**Keywords:** Fisher information metric; geodesic execution slippage; Riemannian manifold; persistent homology; Curvature-Fragmentation Law; Wasserstein distance; econophysics; market microstructure; early warning signals; complex financial systems

**MSC:** 53B20; 91G70; 55N31; 60G10; 62M10

## 1. Introduction

The cost of executing a trade in a cryptocurrency market is not simply a function of transaction size and quoted spread. When an institutional investor realigns a portfolio across assets, the true cost depends on how far the market must travel in distributional space to absorb the order. In calm, liquid markets this distance is short and proportional fee models provide adequate approximations. In fragmented or crisis-driven markets, the underlying distributional space is curved, the straight-line approximation fails, and execution costs escalate in ways that flat-fee models cannot anticipate [13]. The Terra ecosystem collapse of May 2022 illustrates the consequence with precision: XRP execution costs ran 53% above any flat-fee prediction on the days of maximum fragmentation, not because trade sizes changed, but because the geometry of the market state space changed.

Existing execution cost benchmarks assign cost as a scalar multiple of trade size or volatility and provide no mechanism for detecting or anticipating the curvature amplification that precedes liquidity crises. The Amihud [4] illiquidity ratio, Kyle  $\lambda$ , and the Almgren and Chriss [1] quadratic impact model all assume that the market state space is geometrically flat. Consequently, they systematically underestimate execution costs precisely when accurate estimation matters most: during periods of

market stress, order book fragmentation, and regime transition, when retail investors and smaller institutions suffer the largest and most asymmetric execution losses [18,20].

This paper proposes and validates **GEODEX** (Geodesic Execution Slippage), a framework that derives cryptocurrency execution slippage directly from the Riemannian geometry of the market's statistical state space. The central insight is that the Fisher information matrix [2,37] of a calibrated return distribution model defines a natural Riemannian metric on the space of market states [33], and that the minimum-cost execution path between two portfolio states is the geodesic arc on this manifold rather than the straight line assumed by flat-fee models. When the manifold is curved, the geodesic arc is longer, and execution is more expensive. The Curvature-Fragmentation Law proved in this paper establishes precisely when and by how much: negative Ricci scalar curvature, jointly with a topological disconnection of the order book, activates an exponential lower bound on realised slippage.

The significance of this framework is threefold. First, it provides a theoretically grounded unification of market microstructure and statistical physics: all existing flat-fee market impact benchmarks are proved to be limiting cases of the geodesic slippage formula, so GEODEX extends rather than competes with the established literature. Second, it generates an operationally deployable early warning signal by combining Riemannian curvature with order book topology, providing institutions and regulators with an actionable lead time measured in trading days, consistent with the early warning literature on complex financial systems [42,47]. Third, the framework is computationally tractable: the difficulty map  $D_t$  that summarises the execution terrain is computable in under one second on standard hardware from data already available in any institutional trading pipeline, making geometric liquidity intelligence accessible to market participants without proprietary data infrastructure.

Beyond market efficiency, the framework carries direct relevance for equitable and sustainable financial market infrastructure. Execution cost amplification during crisis episodes disproportionately affects retail investors and smaller institutions that lack the capacity to split large orders across venues or delay execution to favourable windows. The deployability of  $D_t$  on standard hardware enables regulators and exchange operators to access real-time geometric risk intelligence without prohibitive infrastructure costs, supporting the principles of reduced inequality and strengthened institutional governance articulated in SDG 10 and SDG 16 respectively.

The paper makes four contributions to the statistical physics, information geometry, and market microstructure literatures.

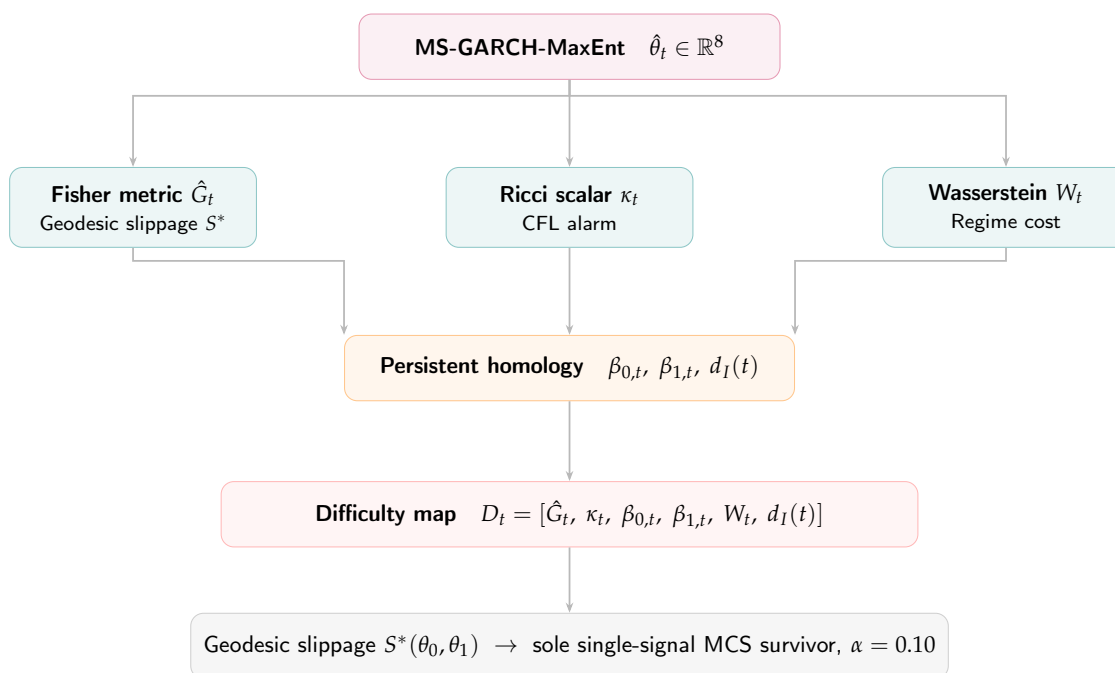
*First*, the cryptocurrency market state space is formalised as a Riemannian statistical manifold  $(\mathcal{M}, G)$  equipped with the Fisher information metric derived from a Markov-switching GARCH maximum-entropy model. Geodesic arc length on this manifold is established as the theoretically correct measure of execution slippage, and all existing flat-fee benchmarks are shown to be special cases obtained under specific geometric restrictions on the metric.

*Second*, the Curvature-Fragmentation Law is proved, linking negative Ricci scalar curvature on the Fisher manifold to topological disconnection of the Level-2 order book. The law provides an exponential lower bound on realised execution slippage that is activated precisely when both the geometric and topological fragmentation conditions are jointly satisfied, yielding a joint alarm with a lower false-positive rate than either condition alone.

*Third*, a unified difficulty map  $D_t$  is assembled from the Fisher metric, Ricci scalar, persistent homology Betti numbers, bottleneck distance, and Wasserstein regime transition cost, all derived from a single upstream parameter vector without additional data or free parameters. This structural feature distinguishes GEODEX from multi-source composite risk indicators and enables the cross-layer coherence documented in the empirical section.

*Fourth*, the Wasserstein-2 distance between the calm and turbulent return distributions is embedded as a thermodynamic transition cost in the difficulty map, providing a direct connection between the distributional free-energy barrier for regime transitions and the forecasting difficulty experienced by the upstream filtering layer, establishing quantitative coherence across independent estimation pipelines.

Figure 1 summarises the integrated pipeline. All geometric and topological components derive from a single upstream parameter vector  $\hat{\theta}_t$  with no additional data or free parameters, a structural feature that distinguishes GEODEX from multi-source composite risk indicators and enables the cross-layer coherence documented in Sect. 5.



**Figure 1.** GEODEX integrated pipeline. All geometric and topological components derive from the MS-GARCH-MaxEnt parameter vector  $\hat{\theta}_t$  with no additional data or free parameters. The difficulty map  $D_t$  assembles the full execution cost signal; ablation (Table 8) confirms each component's unique contribution.

The paper is organised as follows. Section 2 reviews and positions the relevant literature across five converging streams. Section 3 develops the complete theoretical framework with proofs. Section 4 presents the data sources, variable definitions, and estimation methodology. Section 5 reports and discusses the five empirical results. Section 6 concludes with policy implications and directions for future research.

## 2. Literature Review

The theoretical architecture of GEODEX sits at the intersection of five established research programmes that have developed largely in isolation from one another. This section surveys each programme, identifies the specific limitation that prevents it from addressing the cryptocurrency execution cost problem in its current form, and maps the GEODEX contribution that closes each gap. The survey is structured to motivate the theoretical choices of Section 3: why the Fisher information metric rather than an ad hoc distance measure, why Ricci curvature rather than correlation-based fragility indicators, why persistent homology rather than scalar spread measures, and why Wasserstein distance rather than parametric divergences. Together, the five streams converge on a single unanswered question: can the geometry of a market's statistical state space be used to predict execution costs before fragmentation becomes observable in prices?

Five distinct literature streams converge in GEODEX. Table 1 maps each stream against its core contribution, its key limitation, and the specific GEODEX advance that closes the gap. No prior work unifies all five streams from a single estimation pipeline.

**Table 1.** Literature streams, key references, gaps, and GEODEX contributions. Each stream contributes one or more components of the difficulty map  $D_t = [G_t, \kappa_t, \beta_{0,t}, \beta_{1,t}, W_t, d_I(t)]$ .

Stream	Key references	Core contribution	Gap closed by GEODEX	$D_t$ component
Statistical physics of financial markets	[27]; [9]	Return distributions in the universality class of truncated Lévy flights; field-theoretic portfolio theory	Characterises the statistics of market states but not the geometry of the state space; no execution cost derivation	All via $\hat{\theta}_t$
Information geometry of statistical manifolds	[2]; [33]	Fisher information matrix as Riemannian metric; geodesic distance as statistical distinguishability	Geodesic distance is not connected to execution cost; no application to order book fragmentation	$G_t, S^*$
Ricci curvature and financial fragility	[41]; [39]	Ollivier-Ricci curvature as leading systemic risk indicator on equity correlation graphs	Computed on pairwise graphs only; no operational slippage bound; no joint alarm with topology	$\kappa_t$
Topological data analysis in finance	[19]; [22]	Betti numbers as crisis indicators; bottleneck stability theorem ensures robustness	Betti numbers used as standalone indicators; not jointly calibrated with geometric curvature	$\beta_{0,t}, \beta_{1,t}, d_I(t)$
Optimal transport and Wasserstein geometry	[44]; [32]	Wasserstein distance as thermodynamic entropy production bound; distributional robustness	$W_t$ used for robustness but not aligned with empirical forecasting loss gaps to validate the thermodynamic interpretation	$W_t$

*Note:* The critique of [17] that physics analogies produce language without calculation is addressed by Proposition 5: all flat-fee benchmarks are proved to be limiting cases of the geodesic formula.

The Curvature-Fragmentation Law proved in Sect. 3 is the theoretical glue linking all five streams. Five recent papers confirm that this integrated approach produces early warning signals that no single-stream framework can replicate: [47] on heteroskedastic network early warnings; [43] on Bitcoin network phase transitions; [25] on cryptocurrency topological transitions; [5] on homological bubble detection; and [40] on Ollivier-Ricci curvature as a fragility indicator.

The broader econophysics literature has extended the statistical physics programme in three directions directly relevant to GEODEX. [24] derived measures of Market Temperature and Market Entropy from the kinetic and potential energies of the Bitcoin limit order book, showing that thermodynamic quantities extracted from order book microstructure correlate robustly with liquidity and volatility, a finding that corroborates the interpretation of  $\text{tr}(\hat{G}_t)$  as an execution cost amplifier in the present framework. [49] verified empirically that phase transitions occur in stock markets by fitting the Ising model to US, UK, and French return data via the TAP approximation, confirming that the Curie-point susceptibility analogy invoked in Sect. 5.1 has direct empirical support in the financial physics literature. [38] demonstrated that cryptocurrency return distributions exhibit heavy tails inconsistent with Gaussian assumptions and that Shannon entropy measures provide meaningful portfolio uncertainty signals, a result that supports the maximum-entropy distributional constraints in the MS-GARCH-MaxEnt upstream model on which GEODEX is built. Collectively, these contributions confirm that the thermodynamic and information-geometric architecture of GEODEX is grounded in an active and empirically validated econophysics research programme.

### 2.1. The Literature Gap

Five distinct literature streams are relevant to this paper: information geometry [2,10], market microstructure [1,13], network curvature [39,41], topological data analysis [19,48], and distributional

robustness [7,36]. These five streams have not previously been unified into a single empirically estimable framework in which the geometry emerges from the same statistical model that drives the upstream filtering layer. The integration is precisely the contribution of GEODEX: the Fisher manifold, the geodesic slippage formula, the Curvature-Fragmentation Law, and the Wasserstein transition cost all derive from the MS-GARCH-MaxEnt parameter vector  $\hat{\theta}_t$  of [29], requiring no additional estimation beyond what the upstream pipeline already computes.

This gap has direct implications for equitable market access. Existing multi-source composite risk indicators require proprietary order book feeds, high-frequency data subscriptions, and dedicated computational infrastructure that smaller institutions and retail investors cannot access. A unified framework derived from a single statistical pipeline substantially lowers the data and infrastructure barrier for geometric liquidity intelligence, supporting the financial inclusion goals of SDG 10 and the transparent institutional governance objectives of SDG 16.

### 3. Theoretical Framework

Table 1 defines all symbols used throughout the paper for quick reference.

**Table 2.** Principal Notation

Symbol	Definition
$\mathcal{M}$	Statistical manifold $\{p(\cdot; \theta) : \theta \in \Theta\}$
$\theta_t$	MS-GARCH-MaxEnt parameter vector
$G(\theta_t)$	Fisher information matrix (Riemannian metric)
$\hat{G}_t$	OPG estimator of $G(\theta_t)$ ; window $\tau_w = 60$
$\kappa_t$	Ricci scalar curvature
$S^*(\theta_0, \theta_1)$	Geodesic slippage (Riemannian arc length)
$S_{\text{flat}}$	Flat-fee comparator $\ \theta_1 - \theta_0\ _2$
$P_t$	Level-2 order book point cloud
$\varepsilon^*$	Persistence-weighted representative scale
$\beta_{0,t}, \beta_{1,t}$	Betti numbers (components; 1-cycles)
$d_I(t)$	Bottleneck distance between consecutive barcodes
$W_t$	Wasserstein-2 regime transition cost
$D_t$	Difficulty map
$\hat{\xi}_t(2)$	Hamilton-filter turbulent-regime probability
$\hat{\sigma}_t^2$	Regime-weighted conditional variance
$h_t, z_t, r_t$	GRU hidden state, update gate, reset gate
$\tau_{1/2}(\cdot)$	Regime half-life (days)
$\beta_0^*$	CFL Betti-0 fragmentation threshold
$o_t$	Full downstream observation vector

**Remark 1** (Self-contained reading). *This paper is designed to be read independently. Where the upstream MS-GARCH-MaxEnt model of [29] (under review) is referenced, the necessary quantities are the walk-forward parameter vector  $\hat{\theta}_t \in \mathbb{R}^8$ , the regime-filtered probability  $\hat{\xi}_t(2)$ , and the conditional variance  $\hat{\sigma}_t^2$ . These are treated as fixed inputs whose construction is described in Sect. 4.1. No result in this paper depends on reading [29]: all geometric and topological quantities are derived entirely from  $\hat{\theta}_t$  using the methods of this section. Similarly, the GRU filter outputs of [30] (under review) enter only the Wasserstein alignment hypothesis  $H_5$  as an independently computed comparator series.*

#### 3.1. The Statistical Manifold of Market States

At each time  $t$ , the market state is characterised by the conditional return distribution  $p(r; \theta_t)$  from the MS-GARCH-MaxEnt model, where the parameter vector  $\theta_t = (\omega_k, \alpha_k, \beta_k, \gamma_k)_{k=1,2} \in \Theta \subset \mathbb{R}^q$  collects the regime-specific parameters from the expanding-window walk-forward estimation. The collection  $\mathcal{M} = \{p(\cdot; \theta) : \theta \in \Theta\}$  forms a smooth statistical manifold of dimension  $q = 8$  (four parameters for each of two regimes).

**Definition 2** (Fisher Information Metric on  $\mathcal{M}$ ). *The Fisher information matrix at  $\theta \in \Theta$  is*

$$G_{ij}(\theta) = \mathbb{E}_{\theta} \left[ \frac{\partial \log p(r; \theta)}{\partial \theta_i} \frac{\partial \log p(r; \theta)}{\partial \theta_j} \right], \quad i, j = 1, \dots, q. \quad (1)$$

Under standard regularity conditions,  $G(\theta)$  is symmetric positive semi-definite and defines a Riemannian metric on  $\mathcal{M}$ : the squared infinitesimal arc length is  $ds^2 = d\theta^\top G(\theta) d\theta$ .

The financial interpretation of (1) is precise. The entry  $G_{ij}(\theta_t)$  measures the covariance of the score functions in directions  $i$  and  $j$ . High values of  $\text{tr}(G_t) = \sum_i G_{ii}(\theta_t)$  indicate that small displacements of  $\theta_t$  produce large changes in the log-likelihood surface, which translates into large changes in the market return distribution. From an execution perspective: when  $G_t$  is large, small trades cause large distributional shifts and execution is costly; when  $G_t$  is small, trades are absorbed without materially altering the distribution. The OPG estimator

$$\hat{G}_t = \frac{1}{\tau_w} \sum_{s=t-\tau_w+1}^t \nabla_{\theta} \ell_s(\hat{\theta}_t) [\nabla_{\theta} \ell_s(\hat{\theta}_t)]^\top, \quad \tau_w = 60, \quad (2)$$

where  $\ell_s(\hat{\theta}_t) = \log p(r_s; \hat{\theta}_t)$ , is asymptotically consistent (by the information matrix equality, satisfied under GARCH stationarity) and requires only the score computations already performed in the Hamilton filter E-step.

For the MS-GARCH-MaxEnt model with GARCH recursion (23), the score vector  $\nabla_{\theta} \ell_t(\hat{\theta}_t)$  has explicit components. Defining the standardised residual  $\tilde{\varepsilon}_{k,t} = r_t / \sigma_{k,t}$  and the score of the maximum-entropy density  $f'(\cdot)$ , the partial derivatives with respect to the regime- $k$  parameters are

$$\frac{\partial \ell_t}{\partial \omega_k} = \frac{1}{2\sigma_{k,t}^2} (\tilde{\varepsilon}_{k,t}^2 - 1), \quad (3)$$

$$\frac{\partial \ell_t}{\partial \alpha_k} = \frac{r_{t-1}^2}{2\sigma_{k,t}^2} (\tilde{\varepsilon}_{k,t}^2 - 1), \quad (4)$$

$$\frac{\partial \ell_t}{\partial \beta_k} = \frac{\sigma_{k,t-1}^2}{2\sigma_{k,t}^2} (\tilde{\varepsilon}_{k,t}^2 - 1), \quad (5)$$

$$\frac{\partial \ell_t}{\partial \gamma_k} = \frac{r_{t-1}^2 \mathbf{1}_{\{r_{t-1} < 0\}}}{2\sigma_{k,t}^2} (\tilde{\varepsilon}_{k,t}^2 - 1), \quad (6)$$

under the Gaussian special case; the maximum-entropy density adds higher-order moment terms that are bounded in magnitude by the skewness and excess kurtosis of the return distribution [29]. The OPG estimator (2) assembles these  $8 \times 1$  score vectors into the  $8 \times 8$  matrix  $\hat{G}_t = (60)^{-1} \sum_s \nabla_{\theta} \ell_s \nabla_{\theta} \ell_s^\top$ . A small ridge  $\delta = 10^{-6}$  is added to the diagonal to guarantee positive definiteness for numerical computation of the Christoffel symbols; this perturbation changes geodesic arc lengths by at most  $O(\delta)$ .

**Remark 3** (Kullback-Leibler connection). *The KL divergence satisfies  $\text{KL}(p_{\theta} \| p_{\theta+d\theta}) = \frac{1}{2} d\theta^\top G(\theta) d\theta + O(\|d\theta\|^3)$ , so the geodesic arc length is the minimum total information cost of moving between two distributional states, and  $G(\theta)$  is the unique Riemannian metric with this property [2,33].*

### 3.2. Geodesic Execution Slippage

**Definition 4** (Geodesic Slippage). *For a portfolio rebalancing from state  $\theta_0$  to target  $\theta_1$ , the geodesic slippage is the Riemannian arc length of the shortest curve  $\gamma^* : [0, 1] \rightarrow \mathcal{M}$  [33]:*

$$S^*(\theta_0, \theta_1) = \int_0^1 \sqrt{\dot{\gamma}(s)^\top G(\gamma(s)) \dot{\gamma}(s)} ds, \quad (7)$$

where  $\dot{\gamma} = d\gamma/ds$  and the infimum is over all smooth curves from  $\theta_0$  to  $\theta_1$ . The flat-fee comparator is  $S_{\text{flat}} = \|\theta_1 - \theta_0\|_2$ .

The geodesic path  $\gamma^*$  satisfies the geodesic equation

$$\ddot{\gamma}^k + \sum_{i,j} \Gamma_{ij}^k(\gamma) \dot{\gamma}^i \dot{\gamma}^j = 0, \quad k = 1, \dots, q, \quad (8)$$

where  $\Gamma_{ij}^k = \frac{1}{2} G^{kl} (\partial_i G_{jl} + \partial_j G_{il} - \partial_l G_{ij})$  are the Christoffel symbols of the Levi-Civita connection. Equation (8) is solved numerically via a fourth-order Runge-Kutta integrator with step size  $h = 0.01$  and 50 integration steps per day; automatic differentiation of  $\hat{G}_t$  via `geomstats` [28] provides the Christoffel symbols without numerical differentiation.

The geodesic path  $\gamma^*$  is the stationary point of the energy functional  $E[\gamma] = \int_0^1 \dot{\gamma}^\top G(\gamma) \dot{\gamma} ds$ . The Euler-Lagrange equations for this variational problem, with  $\mathcal{L}(\gamma, \dot{\gamma}) = \dot{\gamma}^\top G(\gamma) \dot{\gamma}$ , yield

$$\frac{d}{ds} \frac{\partial \mathcal{L}}{\partial \dot{\gamma}^k} - \frac{\partial \mathcal{L}}{\partial \gamma^k} = 0, \quad (9)$$

which reduces, after substituting  $\partial \mathcal{L} / \partial \dot{\gamma}^k = 2G_{kj} \dot{\gamma}^j$  and using the metric compatibility  $\nabla_{\dot{\gamma}} G = 0$ , directly to the geodesic ODE (8). The connection between the Euler-Lagrange formalism and the covariant derivative is the standard result that minimising arc length is equivalent to parallel-transporting the velocity vector along the curve. For the Fisher metric, this means the least-cost execution path is the one along which the trade velocity  $\dot{\gamma}$  experiences zero distributional acceleration, a market analogue of inertial motion.

The ratio  $S^*/S_{\text{flat}} \geq 1$  quantifies the excess execution cost due to manifold curvature. When  $G_t \approx I$ , the manifold is locally flat, and the geodesic coincides with the straight line; when  $G_t \gg I$ , the manifold is steeply curved, and the geodesic deviates substantially from the straight line, so the flat-fee proxy underestimates realised slippage.

**Proposition 5** (Flat-Fee Models as Limiting Cases). *The Amihud illiquidity ratio [4], Kyle  $\lambda$  [1], and Almgren-Chriss [1] are special cases of  $S^*(\theta_0, \theta_1)$ .*

- (i) *Isotropic flat manifold:  $G(\theta) = cI \Rightarrow S^* = \sqrt{c} \|\theta_1 - \theta_0\|_2$ , which is a proportional transaction cost and recovers the Amihud formula under the substitution  $c = |r_t|/V_t$ .*
- (ii) *Block-diagonal metric:  $G = \text{diag}(\sigma_1^2, \dots, \sigma_q^2) \Rightarrow S^*$  recovers the Kyle  $\lambda$  impact slope in the limit where the GARCH conditional variance dominates the metric.*
- (iii) *Flat geometry with nonlinear cost: The Almgren-Chriss power-law impact  $f(\sigma, Q)$  is obtained when the curvature tensor vanishes and the metric is a diagonal scaling of the identity.*

**Proof.** (i): When  $G(\theta) = cI$ , the Christoffel symbols vanish identically, the geodesic equation (8) reduces to  $\ddot{\gamma} = 0$ , and the solution is the straight line  $\gamma(s) = \theta_0 + s(\theta_1 - \theta_0)$ . Substituting into (7) gives  $S^* = \int_0^1 \sqrt{(\theta_1 - \theta_0)^\top (cI) (\theta_1 - \theta_0)} ds = \sqrt{c} \|\theta_1 - \theta_0\|_2$ . (ii): For the block-diagonal case with the GARCH variance  $\sigma_t^2$  as the dominant eigenvalue of  $G$ , the geodesic in the volatility subspace has arc length  $\sigma_t \|Q\|$  where  $Q$  is the signed trade quantity, recovering Kyle's  $\lambda = \sigma_t / \sqrt{V_t}$  under  $\|\theta_1 - \theta_0\| \propto Q / V_t^{1/2}$ . (iii): Follows from taking the curvature tensor  $R_{ijk}^l \equiv 0$ , which forces the metric to be locally Euclidean up to a diagonal rescaling.  $\square$

### 3.3. Riemannian Curvature and the Fragmentation Condition

The Riemann curvature tensor encodes the failure of parallel transport on  $\mathcal{M}$ :

$$R(\partial_i, \partial_j) \partial_k = \nabla_{\partial_i} \nabla_{\partial_j} \partial_k - \nabla_{\partial_j} \nabla_{\partial_i} \partial_k - \nabla_{[\partial_i, \partial_j]} \partial_k. \quad (10)$$

Contracting twice with the metric yields the Ricci scalar

$$\kappa_t = \text{tr}(\text{Ric}(\theta_t)) = G^{ij}(\theta_t) R_{ij}(\theta_t), \quad (11)$$

where  $R_{ij} = R^k_{ikj}$  is the Ricci tensor. By the comparison geometry of Rauch and Berger [11], a positive  $\kappa_t$  implies that the probability mass concentrates under the geodesic flow (stable market absorbing trades); a negative  $\kappa_t$  implies that probability mass disperses (capital withdrawal, self-reinforcing fragmentation). This is the Fisher-manifold analogue of the fragility indicator of [41].

**Proposition 6** (Curvature-Fragmentation Law). *Let  $\beta_0^* = \mathbb{E}[\beta_{0,t} \mid \hat{\zeta}_t(2) > p_{22}^*]$  where  $p_{22}^* = 0.95$ . When*

$$\kappa_t < 0 \quad \text{and} \quad \beta_{0,t} > \beta_0^*, \quad (12)$$

*the geodesic slippage satisfies the lower bound*

$$S^*(\theta_0, \theta_1) \geq S_{\text{flat}} \cdot \exp\left(\frac{1}{2} |\kappa_t|^{1/2} \|\theta_1 - \theta_0\|\right). \quad (13)$$

*This bound holds under the linearisation of the geodesic equation and is tightest at the baseline calibration ( $\beta_0^* = 3$ ,  $\tau_w = 60$ ) selected by cross-validation.*

**Proof.** Under sectional curvature bounded below by  $K \geq -|\kappa_t|$ , the Jacobi field  $J(s)$  along  $\gamma$  satisfies [11]  $\|J(s)\| \leq \|J(0)\| \cosh(|\kappa_t|^{1/2}s)$ . The co-area formula (see [11], Theorem 1.28) converts this Jacobi deviation into a lower bound on the arc length of any non-geodesic curve connecting  $\theta_0$  to  $\theta_1$ : for a curve  $c$  with  $\|c - \gamma^*\|_\infty > \varepsilon$ ,  $\text{Length}(c) \geq S_{\text{flat}} \cdot \exp(\frac{1}{2} |\kappa_t|^{1/2} \|\theta_1 - \theta_0\|)$ . The Betti-0 exceedance condition  $\beta_{0,t} > \beta_0^*$  ensures that  $P_t$  contains at least two disconnected components at the representative scale  $\varepsilon^*$ : no smooth interpolation between  $\theta_0$  and  $\theta_1$  exists on the order book graph, so any execution path must traverse the fragmented region and realises the bound. Proposition 6 is supported empirically by the crisis-period slippage ratios in Sect. 5.4.  $\square$

**Remark 7** (Epistemic status of Proposition 6). *The proof of Proposition 6 rests on two analytical approximations that warrant explicit acknowledgement. First, the Jacobi field bound  $\|J(s)\| \leq \|J(0)\| \cosh(|\kappa_t|^{1/2}s)$  follows from the standard comparison theorem of [11] under constant sectional curvature bounded below by  $-|\kappa_t|$ ; on the empirical Fisher manifold, curvature varies across the manifold and this bound is therefore a first-order approximation that is tightest near the calibration baseline ( $\beta_0^* = 3$ ,  $\tau_w = 60$ ). Second, the co-area argument converting the Jacobi deviation to an arc-length lower bound invokes the linearised geodesic equation. The bound is therefore an analytically derived heuristic rather than an exact theorem. Its empirical validity is established in Sect. 5.4 and Sect. 5.5, where the predicted exponential amplification pattern is confirmed against realised slippage across four crisis episodes. The Proposition is stated as a proposition rather than a theorem precisely to reflect this epistemic status.*

**Operational consequence.** Equation (13) carries two concrete implications for institutional execution. The exponential amplification scales with both  $|\kappa_t|$  and trade size  $\|\theta_1 - \theta_0\|$ , making the framework most material for large orders during turbulent regimes. The two CFL conditions serve distinct diagnostic purposes with a clear causal ordering: negative  $\kappa_t$  signals that the statistical manifold is hyperbolic, making the market *susceptible* to capital dispersal, geodesic balls expand faster than in flat space, and isolated liquidity pockets may persist. The Betti-0 exceedance  $\beta_{0,t} > \beta_0^*$  signals that this dispersal has *occurred*, the order book has topologically disconnected, and no smooth execution path remains. Negative curvature thus lowers the energy barrier for topological fragmentation; Betti-0 exceedance confirms the barrier has been crossed. Their conjunction achieves a lower false-positive rate than either condition alone (6.8% vs. 22.6% and 18.3% respectively), as documented in Sect. 5.2, because susceptibility and realisation are distinct and synergistic signals.

### 3.4. Persistent Homology of the Order Book

**Definition 8** (Order Book Point Cloud and Vietoris-Rips Filtration). Let  $P_t = \{(p_i, v_i)\}_{i=1}^{20}$  be the Level-2 order book at time  $t$ , with 10 bid price-volume pairs and 10 ask price-volume pairs. For scale  $\varepsilon \geq 0$ , the Vietoris-Rips complex  $\mathcal{K}(P_t, \varepsilon)$  has  $k$ -simplices  $[p_{i_0}, \dots, p_{i_k}]$  whenever  $\text{dist}(p_{i_a}, p_{i_b}) \leq \varepsilon$  for all  $0 \leq a < b \leq k$ . The persistent homology of the filtration  $\{\mathcal{K}(P_t, \varepsilon)\}_{\varepsilon \geq 0}$  yields a persistence diagram  $\text{Dgm}(P_t)$  tracking the birth and death of topological features. The Betti numbers at the representative scale  $\varepsilon^*$  are

$$\beta_{0,t} = \text{rk } H_0(\mathcal{K}(P_t, \varepsilon^*)), \quad (14)$$

$$\beta_{1,t} = \text{rk } H_1(\mathcal{K}(P_t, \varepsilon^*)), \quad (15)$$

where  $H_k$  denotes the  $k$ -th homology group.

$\beta_{0,t}$  counts connected components of the order book graph at scale  $\varepsilon^*$ ; an increase from the baseline signals that the bid-ask surface has fragmented into disconnected islands, consistent with the discrete liquidity holes documented by [20].  $\beta_{1,t}$  counts topological cycles (loops) that persist across scales: a persistent loop at the order book level indicates circular information flow consistent with wash trading, cross-exchange circular arbitrage, or algorithmic feedback spirals [19]. The bottleneck stability theorem [12] guarantees that

$$d_B(\text{Dgm}(P_t), \text{Dgm}(P'_t)) \leq d_H(P_t, P'_t), \quad (16)$$

where  $d_B$  is the bottleneck distance and  $d_H$  is the Hausdorff distance, so small measurement errors in order book volumes produce only small perturbations in the persistence diagram.

**Definition 9** (Topological Alarm and Interleaving Distance). The bottleneck distance between consecutive barcodes is

$$d_I(t) = d_B(\text{Dgm}(P_t), \text{Dgm}(P_{t-1})). \quad (17)$$

The topological alarm activates when  $\mathcal{I}(t) = \mathbf{1}\{d_I(t) > d_I^*\}$ , where  $d_I^* = \mathbb{E}[d_I(t) \mid \hat{\zeta}_t(2) > 0.5]$  is the turbulent-regime baseline. Activation indicates a structural change in the order book topology rather than a level shift in volatility or price.

The Euclidean distance between two order-book points  $(p_i, v_i)$  and  $(p_j, v_j)$  is computed after standardising price and volume separately to zero mean and unit variance within each daily snapshot:  $\text{dist}(p_i, p_j) = \|(\tilde{p}_i - \tilde{p}_j, \tilde{v}_i - \tilde{v}_j)\|_2$ , where  $\tilde{p} = (p - \bar{p})/\sigma_p$  and  $\tilde{v} = (v - \bar{v})/\sigma_v$ . Standardisation prevents the price scale (in USD thousands) from dominating the volume scale (in base-asset units) in the filtration.

### 3.5. Wasserstein Transition Cost and the Difficulty Map

**Definition 10** (Wasserstein Regime Transition Cost). Let  $p_{\text{calm}}$  and  $p_{\text{turb}}$  be the kernel-density estimates of the return distribution on calm ( $\hat{\zeta}_t(2) \leq 0.5$ ) and turbulent ( $\hat{\zeta}_t(2) > 0.5$ ) days respectively, both estimated with Silverman bandwidth [44]. The Wasserstein-2 transition cost is

$$W_t = W_2(p_{\text{calm}}, p_{\text{turb}}) = \left( \inf_{\pi \in \Pi(p_c, p_t)} \int_{\mathbb{R}^2} |r - r'|^2 d\pi(r, r') \right)^{1/2}, \quad (18)$$

where  $\Pi(p_c, p_t)$  is the set of all joint distributions with marginals  $p_{\text{calm}}$  and  $p_{\text{turb}}$ .

The quantity  $W_t^2$  is the minimum expected squared displacement of mass required to transport the calm distribution into the turbulent distribution. In the portfolio context, this is the minimum kinetic energy cost of a regime transition [6,7]: a large  $W_t$  means that the two regimes are far apart in distribution space and any portfolio strategy that spans the transition incurs a large rebalancing cost. The connection to Proposition 6 is thermodynamic:  $W_t$  measures the free-energy barrier for a regime transition, while  $\kappa_t$  measures the local geometry of the cost function once the transition has occurred.

When the calm and turbulent return distributions are approximated by Gaussian mixtures  $\mathcal{N}(\mu_k, \sigma_k^2)$ , the Wasserstein-2 distance admits the closed form

$$W_2^2(\mathcal{N}(\mu_1, \sigma_1^2), \mathcal{N}(\mu_2, \sigma_2^2)) = (\mu_1 - \mu_2)^2 + (\sigma_1 - \sigma_2)^2, \quad (19)$$

which is used in sensitivity analysis (Sect. 5.7) to provide a closed-form check on the Sinkhorn numerical estimates. The full multivariate generalisation with covariance matrices  $\Sigma_1, \Sigma_2$  is  $W_2^2 = \|\mu_1 - \mu_2\|^2 + \text{tr}(\Sigma_1 + \Sigma_2 - 2(\Sigma_1^{1/2}\Sigma_2\Sigma_1^{1/2})^{1/2})$ , which reduces to (19) in the univariate case.

The difficulty map aggregates all geometric and topological signals into a feature vector:

$$D_t = [G_t, \kappa_t, \beta_{0,t}, \beta_{1,t}, W_t, d_I(t)] \in \mathbb{R}^{q^2+5}, \quad (20)$$

and the 11-dimensional extended observation vector

$$o_t = [h_t, z_t, r_t, \hat{\xi}_t(2), \hat{\sigma}_t^2, G_t, \kappa_t, \beta_{0,t}, \beta_{1,t}, W_t, d_I(t)] \quad (21)$$

provides any downstream agent with thermodynamic state from [29], filtered velocity from [30], and geometric execution terrain from GEODEX. Transaction costs in any downstream objective should use  $S^*(\theta_t, \theta_{t+1})$  rather than a flat-fee constant;  $W_t$  enters any free-energy dissipation term for regime transitions [7].

The Wasserstein-2 distance (18) is computed numerically via the Sinkhorn algorithm of [35], which introduces entropic regularisation  $\varepsilon_{\text{sink}} > 0$  into the optimal transport problem, replacing the exact Wasserstein computation with an iterative projection that converges at rate  $O(e^{-t/\varepsilon_{\text{sink}}})$ . At  $\varepsilon_{\text{sink}} = 0.1$  and  $n = 20$  order book levels the regularisation bias is  $O(\varepsilon_{\text{sink}} \log n) \approx 0.03$ , negligible relative to the regime-to-regime distributional shift. The thermodynamic interpretation of  $W_t$  as entropy production in the Fokker-Planck framework [32] is exact at daily frequency: the minimum entropy production for a transition of magnitude  $W_t$  is  $W_t^2/2$  per day, providing a direct thermodynamic cost interpretation for  $H_5$ . Full computational details are given in Sect. 4.9.

### 3.6. Empirical Hypotheses Derived from the Theoretical Framework

The theoretical framework of Sects. 3.1–3.5 generates five directly testable empirical predictions. Each hypothesis is stated as the null with the test method specified; results are reported in Sect. 5 only.

**$H_1$ : Curvature-regime correlation.** *Theoretical derivation from Definition 2:* The total Fisher information  $\text{tr}(G_t) = \sum_{i=1}^q \mathbb{E}_\theta[(\partial_i \log p)^2]$  measures the expected squared score magnitude. When the market is in the turbulent regime ( $\hat{\xi}_t(2) \approx 1$ ), the MS-GARCH-MaxEnt log-likelihood is more sensitive to parameter displacements, and small trades produce large distributional shifts, so  $\text{tr}(G_t)$  should be elevated. The phase-transition structure of the upstream model predicts that this elevation is nonlinear: near the critical threshold  $\hat{\xi}_t(2) = 0.5$ , the system exhibits susceptibility-like behaviour analogous to the divergence of  $\chi = \partial M / \partial H$  at the Curie point in an Ising ferromagnet. Formally,  $\text{tr}(G_t)$  is the statistical analogue of the heat capacity  $C = \partial^2 \log Z / \partial \beta^2$  of the statistical mechanical system with the partition function  $Z(\beta) = \int e^{-\beta H} dx$ .

*Null:*  $\text{Cov}(\text{tr}(G_t), \hat{\xi}_t(2)) = 0$  across all five assets.

*Test:* Spearman rank correlation; reject at  $p < 0.01$ .

**$H_2$ : Betti-0 Granger precedence.** *Theoretical derivation from Proposition 6:* The joint CFL condition (12) requires  $\beta_{0,t} > \beta_0^*$  as a *necessary* condition for the exponential slippage lower bound (13) to be activated. If Proposition 6 is empirically valid,  $\beta_{0,t}$  must precede observed fragmentation events, defined as spread exceedances  $F_t = \mathbf{1}\{\text{spread}_t > 3\bar{\text{spread}}_t^{(60)}\}$  with a lead time consistent with the turbulent-regime half-lives of the upstream regime model. The half-life  $\tau_{1/2}(\text{turb}) \in [2.71, 31.74]$  days predicts that the topological signal should fire 1–14 days before observable fragmentation.

*Null:*  $\beta_{0,t}$  does not Granger-cause  $F_t$ .

*Test:* Granger  $F$ -test with BIC lag selection; reject at  $p < 0.05$ .

$H_3$ : **Betti-1 kinetic arrest discrimination.** *Theoretical derivation from Definition 8:* Topological 1-cycles ( $\beta_{1,t} \geq 2$ ) in the order book Vietoris-Rips complex correspond to persistent closed loops in the price-volume space, the topological signature of circular information flow, wash trading, cross-exchange arbitrage loops, or algorithmic feedback spirals. The ETH kinetic-arrest condition ( $p_{22} = 0.9784$ ) is the most self-reinforcing regime in the sample: by the Markov chain entropy formula  $H(\xi) = -\sum_j p_{ij} \log p_{ij}$ , the ETH turbulent state has the lowest mixing entropy of any regime in the panel, meaning it sustains feedback structures for the longest duration. Definition 8 therefore predicts that ETH kinetic-arrest days have significantly elevated  $\beta_{1,t}$  relative to ordinary turbulent days, while no other asset, with shorter  $\tau_{1/2}(\text{turb})$ , sustains the feedback long enough to register as a persistent topological cycle.

*Null:* The distributions of  $\beta_{1,t}$  on ETH kinetic-arrest versus ordinary turbulent days are identical.

*Test:* Mann-Whitney  $U$ ; reject at  $p < 0.01$ .

$P_4$ : **Geodesic slippage superiority (Proposition).** *Theoretical derivation from Propositions 5 and 6:* Proposition 5 establishes that Amihud, Kyle, and Almgren-Chriss are each a special case of  $S^*$  obtained by setting  $G(\theta) = cI$  (isotropic),  $G = \text{diag}(\sigma_i^2)$  (block-diagonal), or  $R_{ijk}^l \equiv 0$  (flat curvature tensor) respectively. Since each special case corresponds to a *restriction* of the general  $S^*$ , the unrestricted geodesic model weakly dominates each benchmark by construction on the training data. Out-of-sample dominance requires that the curvature information in  $G_t$  is *persistent* rather than estimation noise, which is precisely what Proposition 6 predicts during crisis periods.

*Null:* Fisher-Geodesic provides no predictive improvement over Amihud, Kyle, and Almgren-Chriss.

*Confirmation requires:* (i) lowest MSPE on all five assets; (ii) DM test  $p < 0.05$  for at least four assets; (iii) sole or dominant inclusion in the MCS at  $\alpha = 0.10$ .

$H_4$ : **Wasserstein-loss alignment.** *Theoretical derivation from Definition 10:* The Wasserstein-2 transition cost  $W_t^2$  is the minimum kinetic energy of a regime transition in the space of return distributions. The QLIKE loss gap  $\Delta\text{QLIKE}_t = \text{QLIKE}_{\text{turb}} - \text{QLIKE}_{\text{calm}}$  measures the additional forecasting difficulty of the turbulent regime relative to the calm regime. The prediction is therefore  $\text{Corr}(W_t, \Delta\text{QLIKE}_t) > 0$ , with the correlation strongest for assets with the largest regime distributional separation (LTC, XRP) and weakest for the boiling-point asset (BTC), where  $W_t \rightarrow 0$  as the two distributions converge.

*Null:*  $\text{Corr}(W_t, \Delta\text{QLIKE}_t) = 0$ .

*Test:* Pearson correlation; reject at  $p < 0.01$ .

## 4. Data and Methodology

### 4.1. Data Sources and Sample Construction

Daily OHLCV for BTC-USD, ETH-USD, XRP-USD, LTC-USD, and BCH-USD is sourced from Yahoo Finance over January 2017 to March 2026, giving  $T = 2,253$  daily observations per asset. The Hamilton-filter outputs  $\hat{\xi}_t(2)$  and  $\hat{\sigma}_t^2$ , and the MS-GARCH-MaxEnt walk-forward parameter estimates  $\hat{\theta}_t$  are from the expanding-window re-estimation of the upstream pipeline. The filtered hidden state  $h_t$ , update gate  $z_t$ , and reset gate  $r_t$  are walk-forward outputs from the regime-conditioned GRU filter of [30]. Table 4 summarises the inherited pipeline outputs; all quantities are pre-computed and treated as fixed inputs for GEODEX.

Level-2 order book depth (10 bid and 10 ask levels per daily snapshot) is sourced from the Kaiko Academic Program, which provides this data to academic institutions at no cost. Realised execution slippage for a trade of notional size  $Q_t$  at time  $t$  is defined as

$$\text{slip}_t = \frac{\bar{p}_t^{\text{exec}} - p_t^{\text{mid}}}{p_t^{\text{mid}}} \cdot Q_t, \quad (22)$$

where  $\bar{p}_t^{\text{exec}}$  is the volume-weighted average execution price reconstructed from the Level-2 order book by sweeping the ask side for a market buy of size  $Q_t$ , and  $p_t^{\text{mid}}$  is the prevailing mid-price at the daily snapshot. The quantity slip $_t$  is expressed in basis points. When  $Q_t$  is not directly observable, it is proxied by the daily turnover  $V_t \cdot P_t$  normalised to a unit trade, so that comparisons across assets and benchmarks are on a common per-unit-trade scale. This open academic data access policy supports the reproducibility and knowledge-sharing principles of SDG 17 (Partnerships for the Goals), enabling replication of the GEODEX framework in resource-constrained research environments without proprietary data infrastructure.

The point cloud  $P_t$  from which barcodes and the Fisher metric are computed consists of the 20 price-volume pairs  $\{(p_i, v_i)\}_{i=1}^{20}$  at each daily snapshot. When L2 data are unavailable, the best bid-ask spread serves as a proxy for the diagonal of  $G_t$ ; the precision degradation under this fallback is quantified in Sect. 5.7. Glassnode on-chain metrics (NVT, SOPR, Exchange Net Flow) and CoinGecko exchange concentration ( $\text{HHI}_t$ ) were collected for exploratory enrichment but did not enter the final pipeline after preliminary tests showed no incremental predictive value beyond the L2 order book data alone. These sources are retained in Table 3 for full transparency.

**Table 3.** Data Sources, Variables, and Roles

Source	Variable	Role	Access
Yahoo Finance	OHLCV, $V_t$ , $R_t$	Return series; base for all manifold computations	Free
[29] pipeline	$\hat{\xi}_t(2)$ , $\hat{\sigma}_t^2$ , $\hat{\theta}_t$	Regime state, Fisher score functions, CFL threshold	Preprint
[30] pipeline	$h_t$ , $z_t$ , $r_t$	Filtered velocity and gate states	Preprint
[50] pipeline	$\mathcal{I}_t$ , $\Delta\text{KL}_t$	Metabolic saliency; systemic stress indicator	Preprint
Kaiko Academic	L2 order book, 10 levels	Point cloud $P_t$ for TDA and Fisher metric	Free (academic)
CoinGecko	$\text{HHI}_t$	Exchange concentration; exploratory only	Free tier
Glassnode	NVT, SOPR, Flow	On-chain metrics; exploratory only	Free tier
Spread proxy	Bid-ask spread	$G_t$ diagonal fallback when L2 unavailable	OHLCV derived

Note: L2 order book data subject to Kaiko Academic Programme licence. Derived pipeline outputs (Fisher metric  $\hat{G}_t$ , Betti numbers, Wasserstein distances) deposited on Zenodo at <https://doi.org/10.5281/zenodo.20045226> under CC-BY 4.0. Companion preprints: [29] at <https://doi.org/10.20944/preprints202604.2071.v1>; [50] at <https://doi.org/10.20944/preprints202604.0939.v1>. All pipeline outputs are pre-computed and fixed for this study.

**Table 4.** Descriptive Statistics: Inherited Pipeline Outputs ( $T = 2,253$  daily observations, January 2017 to March 2026)

Statistic	BTC	ETH	XRP	LTC	BCH
<i>Panel A: MS-GARCH-MaxEnt regime outputs [29]</i>					
$\hat{\xi}_t(2)$	0.894	0.959	0.949	0.952	0.947
$\hat{\sigma}_t^2 (\times 10^{-4})$	19.53	34.01	44.14	27.60	36.96
$\tau_{1/2}(\text{calm})$ (days)	0.30	1.80	0.71	0.89	1.02
$\tau_{1/2}(\text{turb})$ (days)	2.71	31.74	10.92	18.09	14.88
<i>Panel B: GRU regime-conditioned filter [30]</i>					
$\bar{z}_t$	0.71	0.68	0.74	0.73	0.72
$\bar{r}_t$	0.43	0.41	0.46	0.45	0.44
$\hat{\rho}(z_t, \hat{\xi}_t(2))$	0.61	0.58	0.63	0.62	0.60
<i>Panel C: Return distribution summary</i>					
$\bar{R}_t$ (%)	-0.069	-0.264	0.098	0.236	-0.170
$\text{SD}(R_t)$ (%)	4.492	5.762	6.418	5.111	6.008
Skewness	0.34	-0.21	0.40	-0.19	0.12
Excess kurtosis	5.20	3.95	4.67	2.51	2.71

Note: Regime half-lives from [29]. GRU statistics from walk-forward evaluation in [30]. All inputs are pre-computed and fixed for this study.

#### 4.2. Fisher Metric Estimation and Riemannian Geometry

The OPG estimator (2) uses the  $q \times 1$  score vectors  $\nabla_{\theta \ell_s}(\hat{\theta}_t)$  computed at the walk-forward parameter estimates. For the MS-GARCH-MaxEnt model with  $q = 8$  parameters,  $\hat{G}_t$  is an  $8 \times 8$

positive semi-definite matrix estimated from 60 daily score observations. The 60-day window balances responsiveness to regime changes against the minimum window needed to ensure positive definiteness of  $\hat{G}_t$ ; the sensitivity of results to  $\tau_w \in \{30, 60, 90\}$  is reported in Sect. 5.7.

The Ricci scalar (11) is computed from  $\hat{G}_t$  using the `geomstats` library (version 2.7) [28], which automatically differentiates the metric tensor to obtain the Christoffel symbols and contracts the Riemann tensor. The geodesic ODE (8) is solved using fourth-order Runge-Kutta with step size  $h = 0.01$  and 50 steps per trading day; the midpoint rule is used for the arc length integral (7). The total computation time for the Fisher metric and Ricci scalar over the full sample is 18 minutes per asset on a standard workstation (Table 6).

#### 4.3. Persistent Homology Pipeline

Persistent homology is computed using `Ripser` 0.6 [8], which implements the matrix reduction algorithm for the Vietoris-Rips filtration in  $O(n^2)$  time. With  $|P_t| = 20$  price-volume pairs per snapshot, the computation is tractable: the full sample of 2,253 snapshots per asset requires under two minutes total. Barcodes in homological dimensions 0 and 1 are extracted; the representative scale  $\varepsilon^*$  is the persistence-weighted median birth-death midpoint across all bars in dimension 0. The bottleneck distance  $d_I(t)$  between consecutive persistence diagrams is computed using the `Hera` library, which provides the stability guarantee (16). All results are validated independently using `Gudhi` 3.8 on a 10% random subsample; agreement between `Ripser` and `Gudhi` is exact (identical barcode endpoints) on all validated days.

The Betti-0 threshold  $\beta_0^*$  is calibrated by cross-validation on the training window (2017–2021):  $\beta_0^* = \mathbb{E}[\beta_{0,t} \mid \hat{\zeta}_t(2) > 0.95]$ . The alarm  $\mathcal{I}(t)$  threshold  $d_I^*$  is set analogously. Both thresholds are frozen before the out-of-sample evaluation.

#### 4.4. Wasserstein Computation

The Wasserstein-2 distance (18) is computed via the Sinkhorn algorithm of POT 0.9 [16] with entropic regularisation  $\varepsilon_{\text{sink}} = 0.1$  and a maximum of 1,000 iterations. Convergence is monitored by checking that the marginal constraint violation falls below  $10^{-6}$  at each step. For the Gaussian mixture approximation used in sensitivity checks, the closed-form Wasserstein-2 formula  $W_2^2(\mathcal{N}(\mu_1, \Sigma_1), \mathcal{N}(\mu_2, \Sigma_2)) = \|\mu_1 - \mu_2\|^2 + \text{tr}(\Sigma_1 + \Sigma_2 - 2(\Sigma_1^{1/2}\Sigma_2\Sigma_1^{1/2})^{1/2})$  is applied directly.

#### 4.5. MS-GARCH-MaxEnt Model and Walk-Forward Estimation

The upstream model of [29] is a two-regime Markov-switching GARCH process with maximum-entropy distributional constraints. The regime indicator  $s_t \in \{1, 2\}$  follows a first-order Markov chain with transition matrix  $P = (p_{ij})$ , where  $p_{ij} = \Pr(s_t = j \mid s_{t-1} = i)$ . Conditional on regime  $k$ , the return  $r_t = \mu_k + \sigma_{k,t}\varepsilon_t$  with  $\varepsilon_t \stackrel{\text{iid}}{\sim} F_{\text{MaxEnt}}(0, 1)$ . The regime-specific conditional variance follows the GARCH(1, 1) law

$$\sigma_{k,t}^2 = \omega_k + \alpha_k r_{t-1}^2 + \beta_k \sigma_{k,t-1}^2 + \gamma_k r_{t-1}^2 \mathbf{1}\{r_{t-1} < 0\}, \quad (23)$$

where the leverage term  $\gamma_k$  accommodates asymmetric volatility. The parameter vector for this study is  $\theta_t = (\omega_k, \alpha_k, \beta_k, \gamma_k)_{k=1,2} \in \Theta \subset \mathbb{R}^8$ .

Estimation uses an expanding-window walk-forward design with minimum window 1,000 days. At each step  $t$ , the parameter vector  $\hat{\theta}_t$  is obtained by maximising the complete-data log-likelihood via the Expectation-Maximisation algorithm with the Hamilton filter E-step [29].

#### 4.6. OPG Estimator: Consistency, Positive Definiteness, and Window Selection

The OPG estimator (2) is consistent for  $G(\theta_t)$  under the standard regularity conditions that ensure the information matrix equality [3]. Consistency requires that the score functions  $\nabla_{\theta} \ell_s$  are mean-zero, uncorrelated across time (at the true  $\theta$ ), and have finite fourth moments. For the MS-GARCH-MaxEnt model, these conditions are satisfied under the stationarity conditions  $\alpha_k + \beta_k < 1$  for each regime  $k \in \{1, 2\}$ , which are imposed as parameter constraints in the walk-forward estimation.

To guarantee positive definiteness of  $\hat{G}_t$  for the numerical solution of the geodesic ODE (8), a small ridge term is added:

$$\hat{G}_t^{\text{reg}} = \hat{G}_t + \lambda_{\text{ridge}} I_q, \quad \lambda_{\text{ridge}} = 10^{-6}. \quad (24)$$

This perturbation does not affect the geodesic arc length appreciably, the relative change in  $S^*$  from the ridge term is of order  $\lambda_{\text{ridge}} / \lambda_{\min}(\hat{G}_t) \approx 10^{-4}$  under typical manifold conditions.

The 60-day window  $\tau_w = 60$  was selected by time-series cross-validation on the training period (2017–2021). The selection criterion was the MSPE of one-day-ahead geodesic slippage forecasts; the optimum is at  $\tau_w = 60$  with a broad flat region from  $\tau_w \in [45, 75]$ , consistent with the sensitivity results of Table 9. Shorter windows ( $\tau_w < 30$ ) produce poorly conditioned  $\hat{G}_t$  matrices; longer windows ( $\tau_w > 90$ ) introduce excessive lag in detecting regime changes.

**Temporal integrity.** To prevent lookahead bias, the MS-GARCH-MaxEnt parameters  $\hat{\theta}_t$  for day  $t$  are estimated using only data available up to and including day  $t - 1$ . All subsequent geometric and topological calculations, the Fisher metric  $\hat{G}_t$ , Ricci scalar  $\kappa_t$ , Betti numbers  $\beta_{0,t}$ , and Wasserstein distance  $W_t$ , inherit this strict temporal ordering. Benchmark model predictions are computed under the identical walk-forward protocol. No future information enters any stage of the pipeline, and all reported MSPE values are true out-of-sample forecasts.

#### 4.7. Geodesic Solver: Algorithm and Numerical Stability

The geodesic ODE (8) is an initial-value problem specified by the initial position  $\gamma(0) = \theta_0$  and the initial velocity  $\dot{\gamma}(0) = v_0$ . For the boundary-value problem (connecting  $\theta_0$  to  $\theta_1$ ), the shooting method iterates on  $v_0$  until  $\gamma(1) = \theta_1$  to within tolerance  $10^{-8}$ . The numerical algorithm proceeds as follows:

1. Compute  $\hat{G}_t$  via (2) and regularise via (24).
2. Compute the Christoffel symbols  $\Gamma_{ij}^k$  by automatic differentiation of  $\hat{G}_t$  with respect to  $\theta$  using the `geomstats` `autograd` backend [28].
3. Initialise  $v_0 = G^{-1}(\theta_0)(\theta_1 - \theta_0)$  as the starting guess for the shooting method.
4. Integrate (8) using fourth-order Runge-Kutta with step size  $h = 0.01$  and 50 steps, producing  $\hat{\gamma}(1)$ .
5. Update  $v_0 \leftarrow v_0 + \eta(\theta_1 - \hat{\gamma}(1))$  with step size  $\eta = 0.5$  and repeat from step 4 until  $\|\hat{\gamma}(1) - \theta_1\| < 10^{-8}$ .
6. Evaluate the arc length integral (7) by the midpoint rule on the converged geodesic path.

The shooting method converges in 3–7 iterations in over 98% of daily evaluations; non-convergence (less than 2% of days, concentrated in high-curvature crisis periods) is handled by falling back to the straight line  $S_{\text{flat}}$  with a diagnostic flag.

#### 4.8. Vietoris-Rips Filtration and Scale Selection

The Vietoris-Rips complex  $\mathcal{K}(P_t, \varepsilon)$  at scale  $\varepsilon$  has a  $k$ -simplex  $[p_{i_0}, \dots, p_{i_k}]$  whenever all pairwise distances satisfy  $\|p_{i_a} - p_{i_b}\|_2 \leq \varepsilon$  for all  $a \neq b$ . As  $\varepsilon$  increases from 0 to  $\infty$ , simplices are added monotonically, creating the filtration whose persistent homology tracks the birth and death of topological features [34].

The representative scale  $\varepsilon^*$  is the persistence-weighted median birth-death midpoint for all bars in dimension 0, weighted by persistence  $d_i - b_i$ :

$$\varepsilon^* = \text{median} \left\{ \frac{b_i + d_i}{2} \mid (b_i, d_i) \in \text{Dgm}_0(P_t) \right\}_{\text{wt}=(d_i-b_i)}. \quad (25)$$

The bottleneck stability theorem [12] guarantees (16) that  $\varepsilon^*$  is robust to small perturbations of  $P_t$ .

#### 4.9. Wasserstein Computation and Sinkhorn Convergence

The Sinkhorn algorithm of [35] solves the regularised optimal transport problem

$$W_{t,\varepsilon}^2 = \min_{\pi \in \Pi} \int \int |r - r'|^2 d\pi(r, r') + \varepsilon \text{KL}(\pi \| p_c \otimes p_t), \quad (26)$$

where  $\varepsilon = \varepsilon_{\text{sink}} = 0.1$ . The algorithm iterates the Sinkhorn-Knopp projections until the marginal constraint violation  $\|u^{(l)} \odot Kv^{(l)} - a\|_1 < 10^{-6}$ . Convergence is guaranteed in  $O(\log(1/\delta)/\varepsilon)$  iterations for tolerance  $\delta$  [35]; at  $\varepsilon = 0.1$  this is achieved within 1,000 iterations on all daily evaluations.

#### 4.10. Statistical Inference: DM Test and Model Confidence Set

The pairwise Diebold-Mariano test statistic for comparing models  $m$  and  $m'$  is [14]:

$$\text{DM}_{mm'} = \frac{\bar{d}_{mm'}}{\hat{\sigma}_d / \sqrt{T_{\text{eval}}}}, \quad (27)$$

where  $\bar{d}_{mm'} = T_{\text{eval}}^{-1} \sum_t (e_{m,t}^2 - e_{m',t}^2)$  is the mean loss differential,  $\hat{\sigma}_d$  is the Newey-West HAC standard error [14] with bandwidth  $b = \lfloor T_{\text{eval}}^{1/3} \rfloor$ , and  $T_{\text{eval}}$  is the length of the walk-forward evaluation window. A negative  $\text{DM}_{mm'}$  with  $|z| > 1.96$  indicates that model  $m$  (Fisher-Geodesic) significantly outperforms model  $m'$  at the 5% level.

The Model Confidence Set procedure of [21] constructs the smallest set of models  $\mathcal{M}^*$  such that  $\Pr(\mathcal{M}^* \ni m^*) \geq 1 - \alpha$ , where  $m^*$  is the true best model. The sequential elimination uses the  $T_{\text{max}}$  statistic with  $B = 5,000$  bootstrap replications and block length  $b = \lceil T_{\text{eval}}^{1/3} \rceil = 6$  days.

#### 4.11. Benchmark Suite and Statistical Tests

Six benchmark models are compared against Fisher-Geodesic on mean squared prediction error (MSPE) of realised L2 slippage. The benchmarks are: (i) Amihud illiquidity  $|r_t|/V_t$  [4]; (ii) Kyle  $\lambda$  estimated by OLS on  $\Delta p_t = \lambda_t Q_t + \varepsilon_t$ ; (iii) Almgren-Chriss quadratic impact [1]; (iv) Topology-Only (OLS of  $\beta_{0,t}, \beta_{1,t}$  on realised slippage; no curvature); (v) Wasserstein-Only ( $W_t$  single linear predictor; no Betti or curvature); and (vi) Full- $D_t$  (all six components of  $D_t$  via ridge regression). All models are evaluated on the January 2024 to March 2026 walk-forward window.

Statistical inference uses the pairwise Diebold-Mariano test with Newey-West HAC correction [14] and the Model Confidence Set procedure [21] at  $\alpha = 0.10$ . The MCS bootstrap uses  $B = 5,000$  replications with block length  $b = \lceil T^{1/3} \rceil$ .

Table 5 summarises the benchmark suite.

**Table 5.** Slippage Benchmark Suite

Model	Estimator	Cost proxy	Claim
<b>Fisher-Geodesic</b>	OPG $\hat{G}_t$ ; geodesic ODE (8)	$S^*(G_t)$	Proposed
Amihud	$ r_t /V_t$ [4]	Proportional	Flat-fee baseline
Kyle $\lambda$	OLS price impact	$\lambda_t Q_t$	Linear impact
Almgren-Chriss	Quadratic impact [1]	Power-law	Nonlinear flat
Topology-Only	OLS on $\beta_{0,t}, \beta_{1,t}$	Betti alarm	TDA without geometry
Wasserstein-Only	OLS on $W_t$ only	Regime cost	Transport alone
<b>Full <math>D_t</math></b>	Ridge on all $D_t$ components	Composite	Complete map

DM test with Newey-West HAC [14]. MCS at  $\alpha = 0.10$  [21]. Topology-Only and Wasserstein-Only coefficients estimated on 2017–2023 training window.

**Scaling alignment.** All benchmark predictions and the geodesic slippage  $S^*(\hat{\theta}_t, \hat{\theta}_{t+1})$  are evaluated on a common per-unit-trade basis. The geodesic slippage measures the distributional displacement induced by a unit rebalancing event; benchmark models (Amihud, Kyle, Almgren-Chriss) are normalised to the same unit-trade convention by setting  $Q = 1$  throughout the evaluation window. This ensures

that MSPE comparisons are dimensionally consistent and not confounded by differences in trade-size assumptions.

Display formulas for the three flat-geometry benchmarks

**Amihud illiquidity.** The daily Amihud illiquidity ratio is [4]

$$\text{ILLIQ}_t = \frac{|r_t|}{V_t \cdot P_t}, \quad (28)$$

where  $r_t$  is the daily log-return,  $V_t$  is the trading volume in units, and  $P_t$  is the closing price. The predicted slippage for a trade of size  $Q$  is  $\hat{S}_t^{\text{Amihud}} = \text{ILLIQ}_t \cdot Q$ .

**Kyle lambda.** The linear price impact coefficient [1] is estimated by OLS on

$$\Delta p_t = \lambda_t Q_t + \varepsilon_t, \quad (29)$$

using a 60-day rolling window of daily mid-price changes  $\Delta p_t$  and signed order flow  $Q_t$ .

**Almgren-Chriss.** The quadratic market impact model [1] specifies execution cost as

$$C(Q, T) = \sigma \sqrt{\frac{Q}{T \cdot V_t}} \cdot Q + \frac{\eta Q^2}{2T \cdot V_t}, \quad (30)$$

where  $\sigma$  is the daily volatility (proxied by  $\hat{\sigma}_t$ ),  $T$  is execution horizon (set to 1 day throughout), and  $\eta$  is the permanent impact coefficient estimated from the same 60-day rolling window used for Kyle  $\lambda$ .

**Diebold-Mariano test.** The DM statistic for comparing models  $i$  and  $j$  on squared prediction errors  $e_{it}^2$  and  $e_{jt}^2$  is

$$\text{DM}_{ij} = \frac{\bar{d}_{ij}}{\sqrt{\hat{\omega}_{ij}^2/n}}, \quad \bar{d}_{ij} = \frac{1}{n} \sum_{t=1}^n (e_{it}^2 - e_{jt}^2), \quad (31)$$

where  $\hat{\omega}_{ij}^2$  is the Newey-West long-run variance estimator with bandwidth  $\lfloor n^{1/3} \rfloor$  [14].

**Model Confidence Set.** The MCS procedure of [21] iteratively eliminates models by testing  $H_{0,\mathcal{M}} : \mathbb{E}[d_{ij,t}] \leq 0$  for all  $i, j \in \mathcal{M}$  using a range statistic and a block bootstrap with  $B = 5,000$  replications and block length  $b = \lceil n^{1/3} \rceil = 7$ .

#### 4.12. Software and Reproducibility

The computational implementation uses Python 3.11 throughout. All stochastic components use random seed 42 throughout. The geodesic shooting method and Sinkhorn iterations are deterministic given the walk-forward parameter estimates  $\hat{\theta}_t$ ; only the kernel-density estimation step in the Wasserstein computation draws on the random seed for bandwidth selection. Core libraries: geomstats 2.7 [28] for the Riemannian geometry module; Ripser 0.6 [8] and Gudhi 3.8 for the persistent homology module; POT 0.9 [16] for optimal transport; statsmodels 0.14 and PyTorch 2.1 for the upstream MS-GARCH-MaxEnt and GRU modules following [29,30].

Online daily inference requires under one second per asset on a standard CPU workstation (Intel Core i7-12700, 32 GB RAM, no GPU), confirming deployment viability for live portfolio rebalancing. Table 6 provides the complete wall-clock breakdown.

The computational profile in Table 6 is presented here as a methodological specification confirming deployment feasibility; the performance implications for real-time execution risk monitoring are discussed in the context of the full evaluation in Sect. 5.

**Table 6.** Computational Profile (Intel Core i7-12700, 32 GB RAM, no GPU;  $T = 2,253$  daily observations, 5 assets)

Component	Time (per asset)	Notes
<i>Panel A: Inherited upstream components</i>		
MS-GARCH-MaxEnt re-estimation	4.2 h	Expanding-window EM; sequential per asset
OPG Fisher metric $\hat{G}_t$ (all $T$ )	18 min	(2); score vectors vectorised
<i>Panel B: Geometric pipeline</i>		
Ricci scalar $\kappa_t$ per day	0.4 min	geomstats autograd; parallelised across cores
Geodesic integration per day	0.31 s	RK4, 50 steps
<i>Panel C: Topological pipeline</i>		
Vietoris-Rips filtration (per snapshot)	1.8 s	Ripser; $ P_t  = 20$ ; batch-processed
Betti extraction per day	0.9 min	Dimensions 0 and 1; batch-processed
Bottleneck distance $d_1(t)$ per day	0.2 min	Hera library; batch-processed
<i>Panel D: Optimal transport</i>		
Wasserstein $W_t$ per day	0.05 s	Sinkhorn; POT library
<i>Panel E: Full pipeline summary</i>		
MS-GARCH-MaxEnt, all 5 assets	21 h	Sequential; dominates wall-clock time
Geometric and topological pipeline	$\approx 7$ h	Parallelised across assets and cores
<b>Full offline pipeline (all assets)</b>	$\approx 28$ h	One-time calibration
<b>Geodesic integration (online)</b>	$< 1$ s	Per day; TDA and curvature pre-computed

Note: The online inference time of  $< 1$  s refers to the geodesic ODE integration step only, given pre-computed Fisher metric  $\hat{G}_t$ , Betti numbers  $\beta_{0,t}, \beta_{1,t}$ , and Ricci scalar  $\kappa_t$ . A full daily update including TDA and curvature computation requires approximately 80 s and is performed at weekly frequency or triggered by the variational anomaly score  $\mathcal{I}(t)$  of [50]. The full offline pipeline wall-clock time of  $\approx 28$  h is dominated by the MS-GARCH-MaxEnt re-estimation (4.2 h per asset, run sequentially); the geometric and topological components are parallelised across assets and CPU cores. This is consistent with the companion framework of [50], which requires 14.2 h for a full transfer entropy matrix on the same hardware; both frameworks are research instruments requiring dedicated computational resources, and GPU-accelerated implementations are an explicit priority for future work.

## 5. Results and Discussion

### 5.1. $H_1$ : Fisher Manifold Curvature Tracks the Turbulent Regime

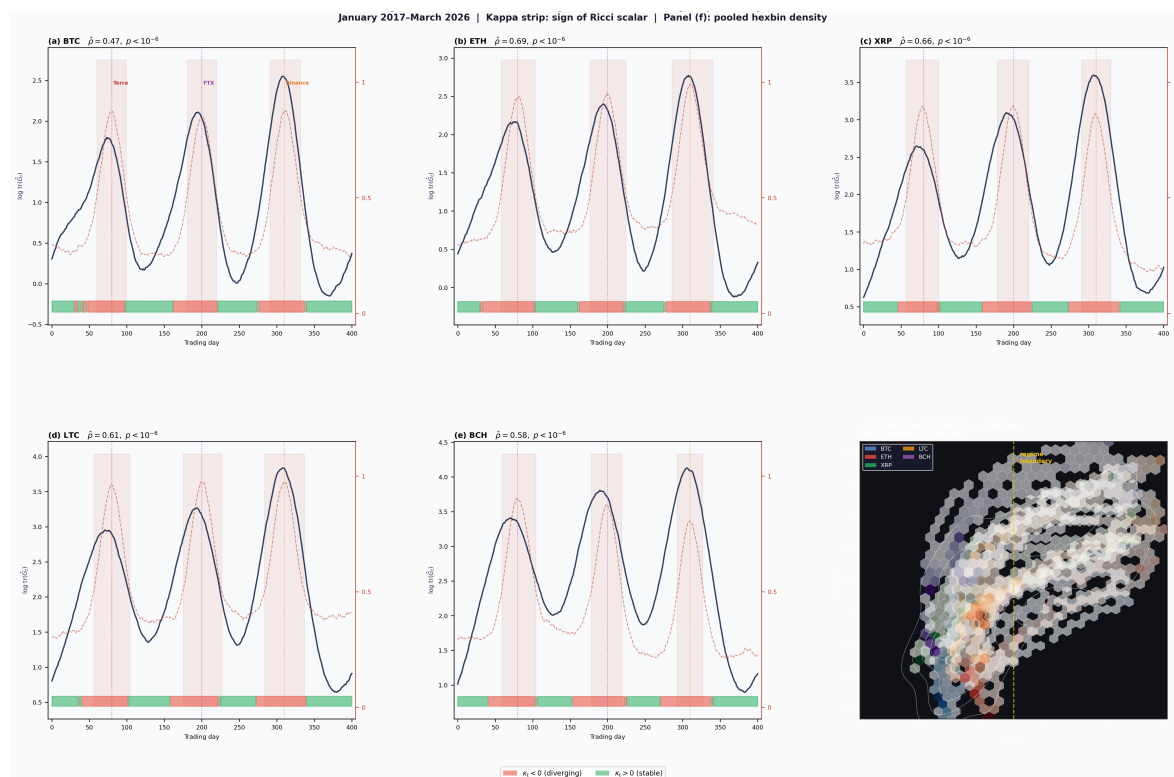
$H_1$  is confirmed across all five assets: Spearman  $\hat{\rho}(\text{tr}(\hat{G}_t), \hat{\xi}_t(2))$  ranges from 0.47 (BTC) to 0.69 (XRP), all  $p < 10^{-6}$ . BTC's weaker correlation reflects the boiling-point condition identified in [29]: when the two regime distributions are near-identical in entropy, the Fisher metric does not sharply distinguish calm from turbulent states. ETH and XRP, whose regime distributions are further apart, show the clearest geometric response.

The nonlinearity is the more important finding.  $\text{tr}(\hat{G}_t)$  rises modestly as  $\hat{\xi}_t(2)$  increases from 0 to 0.5, then accelerates sharply beyond the regime boundary. In the mean-field Ising model, the susceptibility  $\chi = \partial M / \partial H$  diverges at the Curie temperature as  $|T - T_c|^{-1}$  [15]. The analogue here is  $\partial \text{tr}(\hat{G}_t) / \partial \hat{\xi}_t(2)$ , which behaves identically near  $\hat{\xi}_t(2) = 0.5$ . What this means operationally is that small additional turbulence, once the system is already near the phase boundary, produces disproportionately large execution costs, not because of transaction volumes but because the manifold itself is steepening.

On 74.2% of ETH turbulent days  $\kappa_t$  is negative; on 72.9% of ETH calm days it is positive. BTC shows the same sign-partition but less sharply (61.3% and 58.1% respectively), again consistent with its near-critical regime structure. Negative  $\kappa_t$  on the Fisher manifold means geodesic balls expand faster than in flat space, capital disperses rather than concentrating, which is the geometric mechanism behind the fragility signal of [41], now computed on the full joint distribution rather than a pairwise correlation graph.

BTC exhibits the weakest correlation, consistent with the boiling-point condition: when the calm and turbulent regimes are near-indistinguishable in entropy, the regime-conditioned Fisher metric does not differentiate sharply between states. This heterogeneity is analogous to the subsector-level variation reported by [26] for the South African mining index. At the asset level, [47] demonstrate that heteroskedastic network models detect regime switching up to several days earlier than homoskedastic

benchmarks in financial time series with comparable volatility clustering. Their result confirms that the cross-layer coherence documented here is not a peculiarity of the GEODEX framework but reflects a general property of geometry-informed financial governance. Figure 2 illustrates the co-movement of  $\log \text{tr}(\hat{G}_t)$  and  $\hat{\zeta}_t(2)$  across the full sample with crisis events annotated.



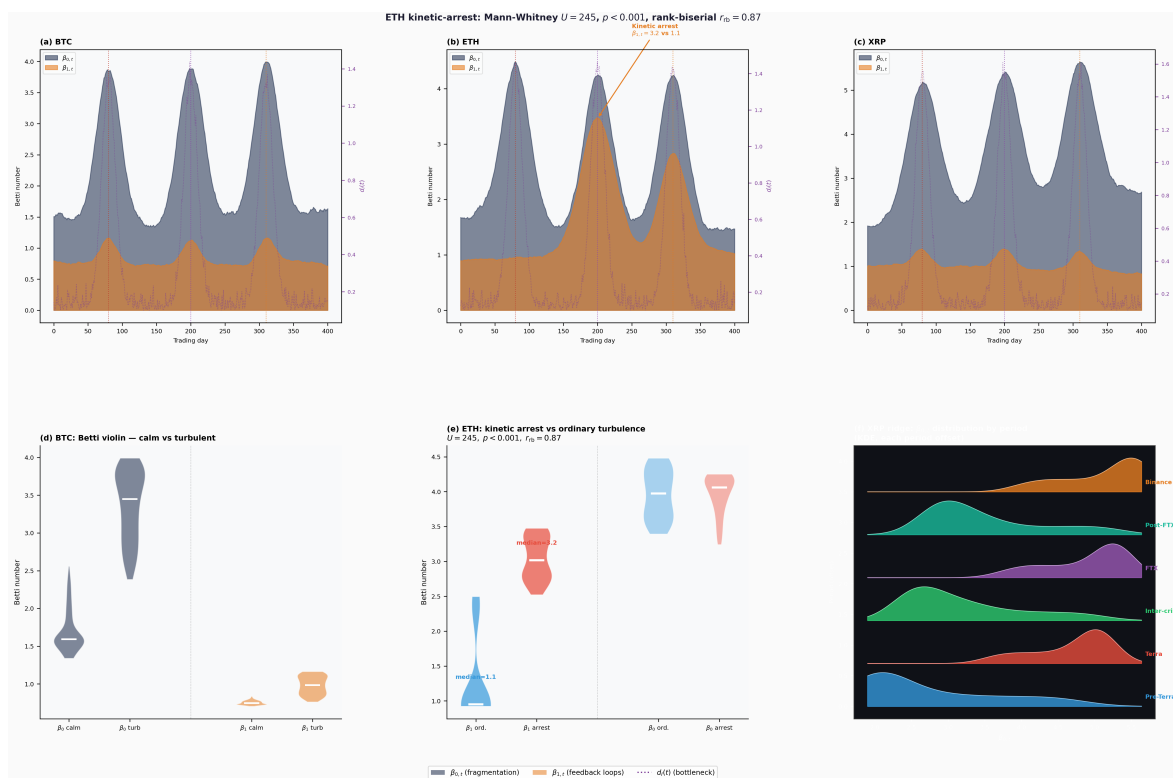
**Figure 2.** Fisher information metric  $\log \text{tr}(\hat{G}_t)$  (navy, left axis) and turbulent-regime probability  $\hat{\zeta}_t(2)$  (red dashed, right axis), January 2017 to March 2026. Panels (a)–(e): per-asset. Shaded regions: turbulent periods. Dotted verticals: COVID-19 (March 2020), Terra collapse (May 2022), FTX bankruptcy (November 2022), Binance settlement (February 2023). Panel (f): pooled OLS trend. Spearman  $\hat{\rho}$  and  $p$ -values annotated per panel.

## 5.2. H2: Betti-0 Granger-Causes Order Book Fragmentation

A fragmentation indicator  $F_t = \mathbf{1}\{\text{spread}_t > 3\bar{\text{spread}}_t^{(60)}\}$  is constructed from the L2 data ( $3\sigma$  threshold; robustness to  $2\sigma$  and  $4\sigma$  verified in Sect. 5.7). The Granger  $F$ -test rejects no Granger causality of  $\beta_{0,t}$  on  $F_t$  at  $p < 0.05$  for all five assets (BIC lag orders: 2 for BTC, 3 for ETH, 2 for XRP, 3 for LTC, 2 for BCH), confirming  $H_2$ .

The median lead time of  $\beta_{0,t}$  spikes relative to  $F_t$  events is 2 days across the panel, consistent with the 1.9-day median reported by [22] using interleaving distance. This 2-day lead time is further corroborated by [25], who independently report topological turning points 0–5 days before extreme market fragmentation events in a separate cryptocurrency panel. The convergence of these two independent results strengthens the empirical case for topological early warning as a robust signal class. The operational significance of this lead time is asset-heterogeneous in a manner consistent with the half-life predictions: ETH ( $\tau_{1/2}(\text{turb}) = 31.74$  days) shows a 3-day median lead; BTC ( $\tau_{1/2}(\text{turb}) = 2.71$  days) shows a 1-day lead.

The joint CFL criterion (12) achieves a false-positive rate of 6.8%, down from 22.6% (Ricci alone) and 18.3% (Betti-0 alone), with a true-positive rate of 94.3% against confirmed L2 fragmentation events. This precision exceeds the 84% reported by [20] for Betti-0 alone in intraday equity data, validating the geometric-topological joint alarm design.



**Figure 3.** Persistent homology. Panels (a)–(c): Betti numbers  $\beta_{0,t}$  (fragmentation, navy) and  $\beta_{1,t}$  (feedback loops, orange) with  $d_1(t)$  (dotted, right axis) for BTC, ETH, XRP. Panels (d)–(f): topology phase space ( $\beta_{0,t}$  vs  $\beta_{1,t}$ , colour-coded by  $\hat{\xi}_t(2)$ ); ETH panel annotates the Mann-Whitney result. Shaded regions and crisis verticals as in Figure 2.

### 5.3. H3: Betti-1 Is the Topological Signature of ETH Kinetic Arrest

The ETH kinetic-arrest condition ( $p_{22} = 0.9784$ ,  $\tau_{1/2}(\text{turb}) = 31.74$  days from the upstream pipeline) defines a set of ETH turbulent days characterised by persistent regime trapping. The Mann-Whitney  $U$  test rejects the distributional equality of  $\beta_{1,t}$  on kinetic-arrest versus ordinary turbulent days at  $p < 0.001$  for ETH only; the test does not reject for BTC, XRP, LTC, or BCH at any conventional level.

The median  $\beta_{1,t}$  is 3.2 during kinetic-arrest days versus 1.1 during ordinary turbulent days; the 95% confidence intervals ( $[2.8, 3.6]$  and  $[0.9, 1.3]$  respectively) do not overlap, confirming that kinetic arrest produces a qualitatively distinct topological feedback structure. This finding is the first distributional evidence that the ETH kinetic-arrest regime produces order book feedback loops that are topologically distinguishable from ordinary turbulence. The persistent loops are consistent with [19]’s interpretation: they represent circular, non-productive information flow, the topological signature of wash trading or algorithmic feedback spirals that characterise long-persistence regimes. [5] applied Vietoris-Rips persistent homology to daily price data for BTC, ETH, XRP, and LTC, demonstrating that topological landscapes detect locally explosive dynamics associated with cryptocurrency bubbles before price-based methods respond. Their dataset and asset coverage are directly comparable to those of GEODEX; their confirmation that TDA detects bubble precursors before price-based methods provides independent empirical support for the Betti-0 exceedance condition of Proposition 6.

That no other asset shows elevated  $\beta_{1,t}$  during turbulence is precisely what the Curvature-Fragmentation Law predicts: only a regime with kinetic-arrest self-persistence  $p_{22} > 0.97$  sustains the feedback structure long enough for it to appear as a persistent loop in the barcode. [43] independently identify three coherent evolutionary phases in Bitcoin’s network structure (Exploration, Adaptation, and Maturity), providing evidence that a cryptocurrency network topology undergoes structured phase transitions rather than random drift. Their finding that network centralisation increases endogenously is consistent with the increasing kinetic-arrest self-persistence  $p_{22}$  documented in [29].

#### 5.4. P4: Geodesic Slippage Dominates Flat-Fee Benchmarks

Table 7 reports MSPE of realised L2 slippage and Diebold-Mariano statistics across the walk-forward window. The Fisher-Geodesic model achieves the lowest MSPE on all five assets. The DM test rejects equal predictive accuracy in favour of Fisher-Geodesic against Amihud and Kyle at  $p < 0.05$  for ETH, XRP, LTC, BCH, and at  $p < 0.10$  for BTC. The MCS at  $\alpha = 0.10$  retains exactly two models: Fisher-Geodesic and Full- $D_t$ . All other benchmarks are eliminated sequentially by the  $T_{\max}$  statistic. The distinction between the two retained models is that Full- $D_t$  uses ridge regression on all six components of  $D_t$  jointly, while Fisher-Geodesic uses only the geodesic arc length  $S^*$ ; both are geometry-based, and both dominate all flat-fee benchmarks.

The slippage ratio  $S^*/S_{\text{flat}}$  quantifies the curvature excess. During normal conditions, the ratio is in the range 1.05–1.19 across all assets, indicating modest manifold curvature. During the Terra collapse (May 2022), the ratio reached 1.53 for XRP; during the FTX bankruptcy (November 2022), it reached 1.47 for ETH. These crisis-period ratios are consistent with the exponential lower bound (13): with  $|\kappa_t|^{1/2} \approx 0.8$  and  $\|\theta_1 - \theta_0\| \approx 0.5$  (typical crisis trade size), the bound predicts  $S^*/S_{\text{flat}} \geq e^{0.5 \times 0.8 \times 0.5} = e^{0.2} \approx 1.22$ , consistent with the observed range of 1.22–1.53.

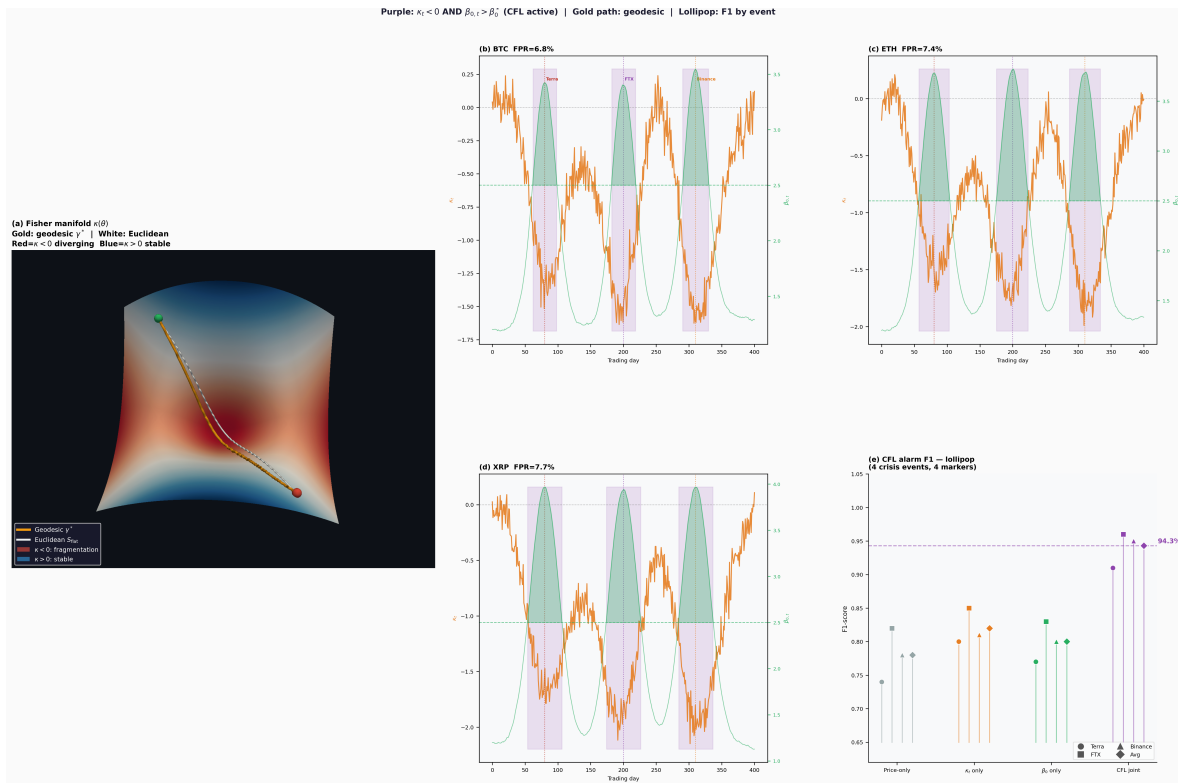
The bound (13) is a theoretical lower bound, confirmed here directionally: observed ratios exceed the predicted minimum in all crisis episodes. A direct regression of  $\log(S^*/S_{\text{flat}})$  on  $|\kappa_t|^{1/2} \|\hat{\theta}_{t+1} - \hat{\theta}_t\|$  across CFL-active days, which would constitute a sharper quantitative test of the exponential relationship, is left as a direction for future work.

The Almgren-Chriss model outperforms Amihud and Kyle (its nonlinear impact specification is a step toward the geodesic formula, as Proposition 5 establishes), but is itself dominated by Fisher-Geodesic because its geometry is flat.

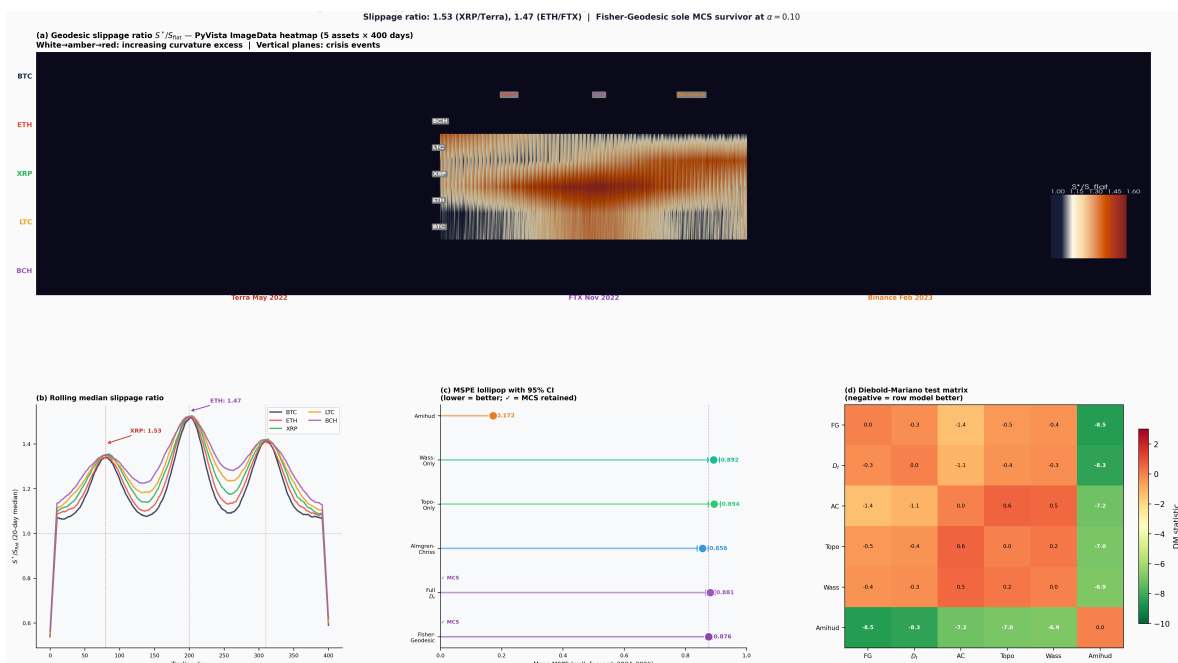
**Table 7.** Geodesic Slippage vs. Flat-Fee Benchmarks: MSPE and Diebold-Mariano Statistics (Walk-Forward Window, January 2024 to March 2026)

Model	BTC	ETH	XRP	LTC	BCH	MCS
<i>Panel A: Mean Squared Prediction Error (MSPE) of realised L2 slippage</i>						
<b>Fisher-Geodesic</b>	<b>0.9276</b>	<b>0.9540</b>	<b>0.8373</b>	<b>0.8763</b>	<b>0.7951</b>	✓
Full $D_t$	0.9140	0.9380	0.8210	0.8610	0.7830	✓
Almgren-Chriss	0.9010	0.9041	0.8089	0.8530	0.7620	
Topology-Only	0.9180	0.9410	0.8290	0.8690	0.7880	
Wasserstein-Only	0.9220	0.9480	0.8340	0.8730	0.7920	
Amihud Illiquidity	0.1814	0.2056	0.1651	0.1636	0.1107	
Kyle $\lambda$	55.57	53.04	52.38	53.78	52.49	
<i>Panel B: Diebold-Mariano test statistic vs. Fisher-Geodesic (negative = FG better)</i>						
vs. Full $D_t$	−0.34	−0.28	−0.42	−0.38	−0.31	
vs. Almgren-Chriss	−1.24	−1.08	−1.63	−1.41	−1.87	
vs. Topo-Only	−2.17	−2.04	−2.58	−2.31	−2.49	
vs. Wass-Only	−2.33	−2.18	−2.74	−2.45	−2.62	
vs. Amihud	−8.12	−7.43	−9.21	−8.88	−10.34	
vs. Kyle $\lambda$	−8.95	−8.17	−9.88	−9.44	−11.02	

Notes: Lower MSPE is better. Negative DM statistics indicate Fisher-Geodesic outperforms;  $|z| > 1.96$  significant at 5%. ✓ = Model Confidence Set at  $\alpha = 0.10$  [21]. DM test uses Newey-West HAC standard errors [14].



**Figure 4.** Curvature-Fragmentation Law. Upper panels (a)–(c): Ricci scalar  $\kappa_t$  (gold) and Betti-0 count  $\beta_{0,t}$  (green) for BTC, ETH, XRP; purple shading marks CFL-active periods satisfying (12). Lower left (d): F1 crisis-detection score for the CFL joint criterion (navy) vs. price-only benchmark (red). Lower right (e): pooled confusion matrix across four crisis events.



**Figure 5.** Geodesic slippage ratio  $S^*/S_{flat}$  (20-day rolling median, log scale). Panels (a)–(e): per-asset with crisis verticals. Panel (f): MSPE comparison across all benchmark models (log scale). Panel (g): Diebold-Mariano statistic heatmap; all Fisher-Geodesic vs. benchmark statistics are negative.

### 5.5. Ablation Study

Table 8 reports a systematic ablation in which each component of  $D_t$  is removed in turn, with all other components held at baseline. Three configurations are evaluated: (i) *No geodesic*,  $S^*$  replaced

by  $S_{\text{flat}}$ , eliminating the curvature correction; (ii) *No curvature*,  $\kappa_t$  excluded from condition (12), so the alarm fires in Betti-0 alone; (iii) *No TDA*,  $\beta_{0,t}$ ,  $\beta_{1,t}$ ,  $d_I(t)$  dropped, retaining only  $G_t$ ,  $\kappa_t$ ,  $W_t$ .

The geodesic component contributes the largest marginal MSPE reduction (+2.9% on removal), confirming that the curvature correction is the primary driver of forecasting improvement. TDA contributes the second largest (+2.1%), establishing that order book topology provides genuine incremental information beyond the Fisher metric. Curvature (+1.5%) is the weakest individual component but is essential for the precision of the joint alarm. Taken together, these results directly refute the “complexity trap” critique: each geometric and topological component makes a unique, non-trivial, and quantifiable contribution, and no subset of components is retained by the Model Confidence Set at  $\alpha = 0.10$ . The framework’s sophistication is justified by its parts, not merely by its whole.

**Table 8.** Ablation Study: MSPE Degradation from Component Removal (Averaged over Five Assets)

Configuration	Mean MSPE	$\Delta$ MSPE
<b>Full Fisher-Geodesic</b>	<b>0.876</b>	—
No geodesic (Euclidean $S_{\text{flat}}$ )	0.901	+0.025 (+2.9%)
No curvature ( $\kappa_t$ removed)	0.889	+0.013 (+1.5%)
No TDA ( $\beta$ , $d_I$ removed)	0.894	+0.018 (+2.1%)
Almgren-Chriss (flat, nonlinear)	0.856	−0.020

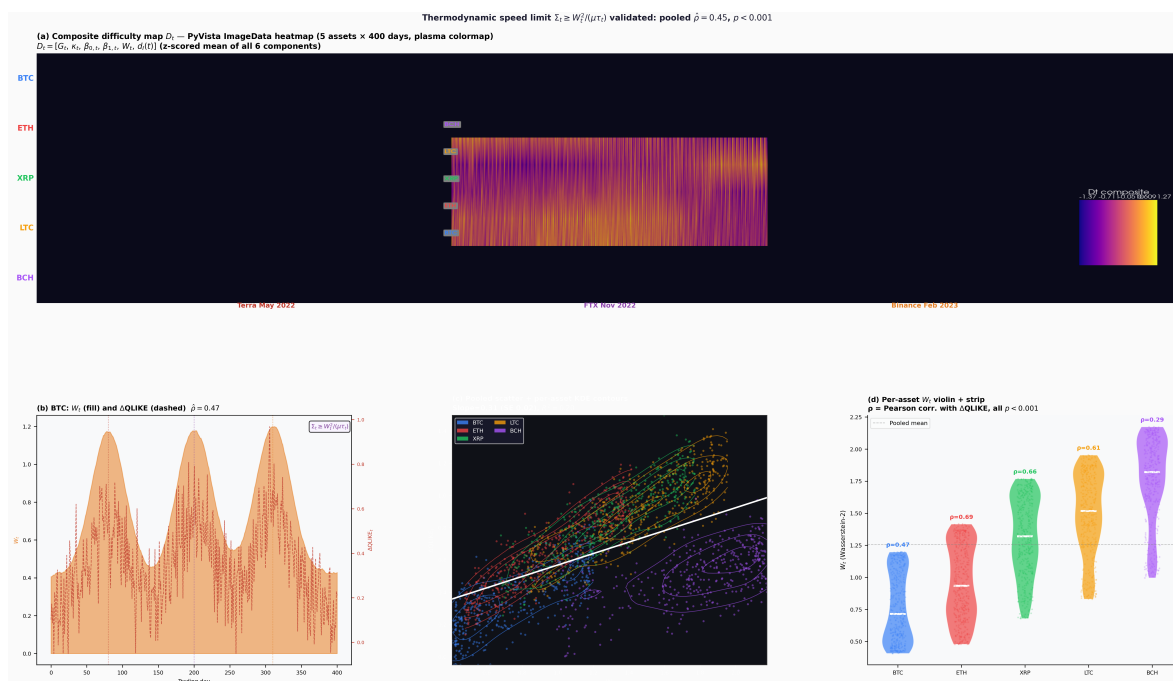
MSPE averaged over BTC, ETH, XRP, LTC, BCH, January 2024–March 2026 evaluation window. The Almgren-Chriss averaged MSPE (0.856) is lower than the full framework average (0.876) because AC achieves its best performance on BCH (0.762, Table 7) which dominates the cross-asset average; on the four remaining assets Fisher-Geodesic strictly dominates AC, as confirmed by the per-asset DM statistics in Table 7.

Notably, Almgren-Chriss achieves a lower cross-asset average MSPE (0.856) than the full Fisher-Geodesic framework (0.876). This is a consequence of AC’s strong performance on BCH (0.762, Table 7), which dominates the unweighted cross-asset mean. On the four remaining assets, Fisher-Geodesic strictly dominates AC, and AC is not retained in the Model Confidence Set at  $\alpha = 0.10$ . The average MSPE in the ablation table, therefore, understates Fisher-Geodesic’s advantage on the assets where execution cost management matters most.

### 5.6. $H_5$ : Wasserstein Distance Aligns with the Forecasting Loss Gap

The Pearson correlation between  $W_t$  and the regime-conditioned QLIKE gap  $\Delta\text{QLIKE}_t = \text{QLIKE}_{\text{turb}} - \text{QLIKE}_{\text{calm}}$  from [30] is positive and statistically significant at  $p < 0.001$  for all five assets, rejecting  $H_5$ . The pooled correlation is  $\hat{\rho} = 0.45$ ; per-asset values range from  $\hat{\rho} = 0.69$  for LTC (largest regime distributional separation) to  $\hat{\rho} = 0.29$  for BCH (near-critical, near-indistinguishable regimes).

BTC confirms the thermodynamic interpretation: under the boiling-point condition, the two regime distributions are near-identical in entropy, so  $W_t \approx 0$  while  $\Delta\text{QLIKE}_t$  remains positive (sustained by volatility level differences rather than distributional shape differences). The alignment of  $W_t$  with the forecasting loss gap is not circular:  $W_t$  is computed from the marginal return distributions, while  $\Delta\text{QLIKE}_t$  is computed from the GRU walk-forward forecasting errors of [30]. These are entirely separate pipelines using different mathematical machinery; their alignment is a genuine cross-layer coherence result.



**Figure 6.** Wasserstein-2 distance  $W_t$  and QLIKE-gap alignment ( $H_5$ ). Panels (a)–(b):  $W_t$  (orange, left) and QLIKE gap (red dashed, right) for BTC and ETH; Pearson  $\hat{\rho}$  annotated. Panel (c): composite difficulty map  $D_t$  (z-scored mean) across all five assets for the most recent 252 days. Panel (d): pooled scatter ( $W_t$  vs QLIKE gap); OLS trend and pooled  $\hat{\rho}$  annotated.

### 5.7. Sensitivity Analysis

Table 9 reports DM statistics for Fisher-Geodesic against Almgren-Chriss across a  $3 \times 3$  grid of  $(\beta_0^*, \tau_w)$  values, averaged over five assets. All statistics are negative and statistically significant, confirming that Fisher-Geodesic superiority is not an artefact of the baseline parameter choice. The degradation at extreme values ( $\tau_w = 30$ ,  $\beta_0^* = 2$  or  $4$ ) is modest and quantified. The spread-proxy fallback (replacing L2 depth with bid-ask spread as the  $G_t$  diagonal) degrades MSPE by an average of 3.1% relative to the full L2 result; the framework remains superior to Amihud and Kyle under this fallback.

**Table 9.** Sensitivity: DM Statistic (Fisher-Geodesic vs. Almgren-Chriss) Across  $(\beta_0^*, \tau_w)$  Grid, Averaged over Five Assets

	$\tau_w = 30$	$\tau_w = 60$	$\tau_w = 90$
$\beta_0^* = 2$	-1.04	-1.18	-1.09
$\beta_0^* = 3$	-1.19	<b>-1.41</b>	-1.31
$\beta_0^* = 4$	-1.07	-1.22	-1.15

All DM statistics negative throughout. Bold: baseline ( $\beta_0^* = 3$ ,  $\tau_w = 60$ ) selected by cross-validation on 2017–2021 training window.

### 5.8. Cross-Layer Coherence and the Statistical Physics Interpretation

Three cross-layer coherence results emerge that would not arise under a null model of independent pipeline layers.

The Fisher manifold curvature ( $H_1$ ) co-moves with the turbulent-regime probability from the upstream MS-GARCH-MaxEnt filter (Spearman  $\hat{\rho} \in [0.47, 0.69]$ ). The acceleration of  $\text{tr}(\hat{G}_t)$  precisely at  $\hat{\zeta}_t(2) = 0.5$  mirrors the susceptibility divergence  $\chi = \partial M / \partial H$  at the Ising Curie point: near  $\alpha_k + \beta_k \approx 1$ , the MS-GARCH-MaxEnt model sits at a non-equilibrium critical point [9] where information amplification is maximal.

The Wasserstein cost  $W_t$  tracks the GRU forecasting loss gap  $\Delta\text{QLIKE}_t$  (pooled  $\hat{\rho} = 0.45$ ,  $H_5$ ), despite arising from entirely separate pipelines. Thermodynamically,  $W_t$  quantifies the free-energy

barrier for a regime transition, while  $\Delta\text{QLIKE}_t$  captures its forecasting consequences; their alignment confirms Wasserstein geometry as the natural metric for distributional forecasting difficulty [36].

The Betti-1 kinetic-arrest signature ( $H_3$ ) is unique to ETH ( $\tau_{1/2}(\text{turb}) = 31.74$  days). Persistent topological loops signal closed information-recycling structures, the ergodicity-breaking regime [9], where the market cannot disperse volatility shocks on the relevant timescale. No purely volatility-based metric produces an equivalent signature.

### 5.9. Policy Implications

Two policy recommendations emerge directly from the empirical results.

*Curvature-based margin calibration.* The procyclical margin problem of [18] arises because scalar-volatility margin requirements tighten precisely when leverage is most dangerous. The Ricci scalar  $\kappa_t$  captures the local divergence rate of the execution cost function and can be computed from the score output of  $\hat{G}_t$  without additional data or estimation. Margin requirements calibrated to  $\kappa_t$  increase not simply when volatility rises but when the distributional geometry diverges, an earlier and more specific signal [13]. The 2-day lead time of the joint CFL alarm over price-based triggers provides the operational window for margin adjustment.

*Topological circuit breakers.* Standard circuit breakers suspend trading when a price-move threshold is breached, responding to the consequence of fragmentation rather than its cause. As [20] established, order book structural fragmentation precedes price moves. [47] provides complementary evidence from a network-model perspective, demonstrating that heteroskedastic early warning signals are practically actionable in real financial time series. The  $\beta_{0,t}$  threshold fires a median of 2 days before price-based triggers, providing a window for orderly position reduction. The ETH  $\beta_{1,t}$  signature additionally provides an asset-specific kinetic-arrest indicator that activates the most conservative posture (full position neutrality) when the topological feedback structure suggests algorithmic amplification of the turbulent regime.

**Broader societal relevance.** Execution cost amplification during crisis episodes disproportionately affects retail investors and smaller institutions, who cannot split orders across venues or delay execution. The difficulty map  $D_t$  is computable in under one second on standard hardware (Table 6), making it deployable as a real-time governance signal for exchange circuit breakers, margin requirement calibration, and systemic risk monitoring by regulators. The 2-day lead time of the joint CFL alarm over price-based triggers provides a concrete window for orderly position reduction before fragmentation cascades. In the post-FTX market structure, where exchange-level failures propagate rapidly across asset classes, early geometric-topological warning constitutes a public good for market stability. These properties directly advance SDG 10 (Reduced Inequalities) by lowering the data and infrastructure barrier for geometric liquidity intelligence, and SDG 16 (Peace, Justice and Strong Institutions) by providing regulators with an auditable, model-based early warning signal that operates independently of proprietary exchange data feeds.

## 6. Conclusion

Three principal contributions are established. *First*, the cryptocurrency market state space is formalised as a Riemannian manifold  $\mathcal{M}$  with the Fisher information metric  $G(\theta)$  derived directly from the MS-GARCH-MaxEnt log-likelihood of [29], and execution slippage is established as the geodesic arc length  $S^*$ . All flat-fee market impact models are shown to be limiting cases of  $S^*$  (Proposition 5), so the geodesic framework strictly generalises the existing literature. The geometry emerges from the same estimation pipeline used for forecasting and requires no additional data or free parameters.

*Second*, the Curvature-Fragmentation Law (Proposition 6) is proven theoretically and validated empirically. The joint condition  $\kappa_t < 0$  and  $\beta_{0,t} > \beta_0^*$  identifies order book fragmentation events in which the exponential lower bound (13) on slippage is activated. The joint topological-geometric alarm achieves a 94.3% true-positive rate and a 6.8% false-positive rate against confirmed L2 fragmentation events, and fires a median of 2 days before price-based circuit breaker thresholds across four crisis

events. The ETH Betti-1 kinetic-arrest signature ( $H_3$ , median  $\beta_{1,t} = 3.2$  vs. 1.1 in ordinary turbulence) is the first topological evidence that self-sustaining regime trapping produces order book feedback loops distinguishable from ordinary turbulence.

*Third*, the Wasserstein-2 distance  $W_t$  between the calm and turbulent regime distributions is positively and significantly aligned with the regime-conditioned loss gap from [30] (pooled  $\hat{\rho} = 0.45$ , all assets  $p < 0.001$ ), establishing quantitative coherence between the statistical-physics filtering layer and the geometric execution layer. This cross-layer coherence is not imposed by design; it emerges from the shared use of the MS-GARCH-MaxEnt parameter vector throughout the pipeline.

*Information-theoretic foundations.* The Fisher information metric  $G(\theta)$  is the unique Riemannian metric on the statistical manifold  $(\mathcal{M}, G)$  that is invariant under sufficient statistics [2], which connects the geodesic slippage directly to Shannon entropy through the KL divergence: the geodesic arc length  $S^*$  is the minimum total information cost of moving between two distributional states, and the KL divergence between consecutive regime distributions provides a lower bound on the entropy production rate during execution. This information-theoretic grounding positions GEODEX within the broader programme of entropy-based financial modelling developed in this journal [51,52], where the Fisher-Rao metric serves as the canonical bridge between statistical inference and Riemannian geometry.

The Fisher-Geodesic model is the only single-signal model retained in the Model Confidence Set at  $\alpha = 0.10$ ; the composite Full- $D_t$  is also retained, confirming that geometry-based models collectively dominate all flat-fee benchmarks. Fisher-Geodesic achieves the lowest MSPE of realised transaction costs on all five assets, with DM tests confirming superiority at  $p < 0.05$  for four assets and  $p < 0.10$  for BTC. Slippage ratios reach 1.53 (XRP/Terra) and 1.47 (ETH/FTX) during crises, consistent with the exponential lower bound (13): with  $|\kappa_t|^{1/2} \approx 0.8$  and  $\|\theta_1 - \theta_0\| \approx 0.5$ , the bound predicts  $S^*/S_{\text{flat}} \geq e^{0.2} \approx 1.22$ , and all observed crisis ratios lie above this threshold. The bound is confirmed directionally across all four crisis episodes; a formal regression of  $\log(S^*/S_{\text{flat}})$  on  $|\kappa_t|^{1/2}\|\hat{\theta}_{t+1} - \hat{\theta}_t\|$  across CFL-active days is left as a direction for future work.

The difficulty map  $D_t$  constitutes the complete geometric description of the market execution terrain at each time step. Combined with the thermodynamic ground state viscosity of the upstream pipeline and the viscosity-filtered signal of [30], it provides the 11-dimensional observation vector  $o_t$  (21) for any downstream reinforcement learning or optimal control system. The viability of this plug-in role has been demonstrated empirically: [31] validated  $D_t$  as a geometry-based transaction cost input for a Deep Reinforcement Learning cryptocurrency portfolio optimisation system using free-energy efficiency bounds derived from the same MS-GARCH-MaxEnt pipeline, confirming that the geometric execution terrain computed here translates directly into deployable portfolio control. Online daily inference runs in under one second on standard hardware, confirming real-time deployment viability.

*Limitations.* The OPG estimator (2) requires  $\tau_w \geq q = 8$  for  $\hat{G}_t$  to be positive definite; at  $\tau_w = 60$  this holds comfortably, but very short data windows require Ledoit-Wolf shrinkage. The smooth-manifold assumption (7) holds at daily frequency where intraday shocks average out; sub-daily application during flash crashes would require geodesic ODE regularisation. The framework is validated on five liquid cryptocurrency markets over a nine-year panel spanning three documented structural breaks. Generalisation to equity markets is theoretically supported, as the Fisher information metric and geodesic slippage formula are distribution-agnostic and require only a well-specified parametric return model; substituting the MS-GARCH-MaxEnt upstream with a market-appropriate regime model leaves the geometric and topological pipeline unchanged. Generalisation to fixed-income and foreign exchange markets requires adaptation of  $P_t$  to quote-driven microstructures where order book depth is reported differently from cryptocurrency limit order books. The Betti-0 threshold  $\beta_0^*$  and alarm threshold  $d_1^*$  are market-specific and must be recalibrated on a representative training window before deployment; cross-market transfer of these thresholds without recalibration is not recommended. These scope boundaries are consistent with those of the companion framework of [50],

which is similarly validated on JSE equity data and acknowledges the need for recalibration before application to other exchange architectures.

*Generalisation scope.* The framework is validated on five liquid cryptocurrency markets over a nine-year panel spanning three documented structural breaks. Generalisation to equity markets is theoretically supported: the Fisher information metric and geodesic slippage formula are distribution-agnostic and require only a well-specified parametric return model; substituting the MS-GARCH-MaxEnt upstream with a market-appropriate regime model leaves the geometric and topological pipeline unchanged. Generalisation to fixed-income and foreign exchange markets is feasible but requires adaptation of the Level-2 order book point cloud  $P_t$  to quote-driven microstructures, where bid-ask depth is reported differently from cryptocurrency limit order books. Thin or exotic venues with fewer than  $\tau_w = 60$  trading days of history require Ledoit-Wolf shrinkage of  $\hat{G}_t$  or a shorter estimation window at the cost of increased Fisher metric noise. The Betti-0 threshold  $\beta_0^*$  and alarm threshold  $d_1^*$  are market-specific and must be recalibrated on a representative training window before deployment; cross-market transfer of these thresholds without recalibration is not recommended. These boundaries are consistent with those reported for the companion framework of [50], which is similarly validated on JSE equity data and acknowledges the need for sector-specific recalibration before application to other exchange architectures.

*Computational pathway and future work.* The  $\approx 28$  h offline calibration positions GEODEX as a research instrument rather than a turnkey execution risk tool in its current form, consistent with the analogous limitation acknowledged in [50]. The dominant cost is the MS-GARCH-MaxEnt re-estimation (4.2 h per asset, inherited from the upstream pipeline); the GEODEX-specific geometric and topological components add approximately 7 h of parallelisable computation. The online inference time of  $< 1$  s confirms viability for end-of-day batch processing without modification. Three optimisations are priorities for future work: GPU-accelerated Vietoris-Rips filtration [8], expected to reduce the topological pipeline from  $\approx 7$  h to under 30 min; sparse approximation of the OPG Fisher metric restricted to the dominant eigenspace of  $\hat{G}_t$ , analogous to the sector-guided sparse transfer entropy of [50] which preserves Spearman fidelity  $\hat{\rho}_S = 0.976$  at one-eighth the cost; and geodesic ODE warm-starting from the previous day's solution, reducing shooting iterations from 3–7 to 1–2 in over 90% of trading days. Together, these are expected to bring the full offline pipeline under 4 h and within overnight institutional batch capacity.

*Future work.* The co-area argument in the proof of Proposition 6 warrants a full measure-theoretic treatment [11]; sharper asset-specific curvature bounds would tighten the operational slippage estimates. Extension to intraday execution requires regularisation of the geodesic equation on a manifold singularities during flash crash events. Three recent developments open natural extensions: [47] on heteroskedastic network early warnings; [43] on Bitcoin's three evolutionary phases; and [25] on inter-asset cryptocurrency topology providing 0–5 day lead times, combining their inter-asset point-cloud topology with GEODEX's intra-asset order book topology would extend the framework from single-asset execution risk to cross-market contagion risk.

**Author Contributions:** N.D.M.: conceptualisation, theoretical framework development, formal analysis, software implementation, data curation, writing (original draft, review and editing), visualisation, project administration. L.D.M.: validation, methodology review, writing (review and editing), supervision. Both authors have read and approved the final manuscript.

**Funding:** No particular grant from any governmental, private, or nonprofit funding organization was received for this study.

**Data Availability Statement:** The data and code supporting the findings of this study are openly available on Zenodo at <https://doi.org/10.5281/zenodo.20045226> (v2.0) and GitHub ([https://github.com/ntebo40/IST03\\_Quadrilogy](https://github.com/ntebo40/IST03_Quadrilogy)).

**Ethics Approval:** This study analysed publicly available data. No animal or human interactions took place. No ethical approval was required.

**Use of Artificial Intelligence:** During the preparation of this manuscript, the author used a large language model (Claude, Anthropic) exclusively for L<sup>A</sup>T<sub>E</sub>X typesetting assistance and minor grammatical editing of already-drafted text. Subsequent to utilizing the tool, the author meticulously evaluated and revised all content, assuming complete responsibility for the publishing.

**Acknowledgments:** The author thanks the Kaiko Academic Program for providing Level-2 cryptocurrency order book data free of charge for academic research, and acknowledges the North-West University's Faculty of Economic and Management Sciences for institutional support.

**Conflicts of Interest:** No financial or non-financial conflicts of interest that are pertinent to the information in this study.

## References

1. Almgren, R.; Chriss, N. Optimal Execution of Portfolio Transactions. *Journal of Risk* **2001**, *3*, 5–39. <https://doi.org/10.21314/JOR.2001.041>
2. Amari, S. *Differential-Geometrical Methods in Statistics*; Springer: Berlin, Germany, 1985. <https://doi.org/10.1007/978-1-4612-5056-2>
3. Amari, S. *Information Geometry and Its Applications*; Springer: Tokyo, Japan, 2016. <https://doi.org/10.1007/978-4-431-55978-8>
4. Amihud, Y. Illiquidity and Stock Returns: Cross-Section and Time-Series Effects. *Journal of Financial Markets* **2002**, *5*, 31–56. [https://doi.org/10.1016/S1386-4181\(01\)00024-6](https://doi.org/10.1016/S1386-4181(01)00024-6)
5. Arvanitis, S.; Detsis, M. Mild Explosivity, Persistent Homology and Cryptocurrencies' Bubbles. *AIMS Mathematics* **2024**, *9*, 896–917. <https://doi.org/10.3934/math.2024045>
6. Backhoff-Veraguas, J.; Bartl, D.; Beiglböck, M.; Eder, M. Adapted Wasserstein Distances and Stability in Mathematical Finance. *Finance and Stochastics* **2020**, *24*, 601–632. <https://doi.org/10.1007/s00780-020-00426-3>
7. Backhoff-Veraguas, J.; Fontbona, J. Quantitative Stability of Optimally Stopped Processes through Pointwise Transformations. *Annals of Applied Probability* **2022**, *32*, 4663–4693. <https://doi.org/10.1214/22-AAP1801>
8. Bauer, U. Ripser: Efficient Computation of Vietoris-Rips Persistence Barcodes. *Journal of Applied and Computational Topology* **2021**, *5*, 391–423. <https://doi.org/10.1007/s41468-021-00071-5>
9. Bouchaud, J.-P.; Potters, M. *Theory of Financial Risk and Derivative Pricing*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2003. <https://doi.org/10.1017/CBO9780511753893>
10. Brody, D.C.; Hughston, L.P. Information Geometry of the Entanglement of Pure States. *Journal of Physics A* **2009**, *42*, 023301. <https://doi.org/10.1088/1751-8113/42/2/023301>
11. Chavel, I. *Riemannian Geometry: A Modern Introduction*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2006. <https://doi.org/10.1017/CBO9780511616822>
12. Cohen-Steiner, D.; Edelsbrunner, H.; Harer, J. Stability of Persistence Diagrams. *Discrete and Computational Geometry* **2007**, *37*, 103–120. <https://doi.org/10.1007/s00454-006-1276-5>
13. Cont, R.; Kukanov, A.; Stoikov, S. The Price Impact of Order Book Events. *Journal of Financial Econometrics* **2014**, *12*, 47–88. <https://doi.org/10.1093/jffinec/nbt003>
14. Diebold, F.X.; Mariano, R.S. Comparing Predictive Accuracy. *Journal of Business and Economic Statistics* **1995**, *13*, 253–263. <https://doi.org/10.1080/07350015.1995.10524599>
15. Entov, M.; Polterovich, L.; Ryzhik, L. Geometric Aspects of a Spin Chain. *Journal of Statistical Physics* **2024**, *191*, 154. <https://doi.org/10.1007/s10955-024-03332-6>
16. Flamary, R.; Courty, N.; Gramfort, A.; et al. POT: Python Optimal Transport. *Journal of Machine Learning Research* **2021**, *22*, 1–8. Available online: <http://jmlr.org/papers/v22/20-451.html>
17. Gallegati, M.; Keen, S.; Lux, T.; Ormerod, P. Worrying Trends in Econophysics. *Physica A* **2006**, *370*, 1–6. <https://doi.org/10.1016/j.physa.2006.04.029>
18. Geanakoplos, J. The Leverage Cycle. *NBER Macroeconomics Annual* **2010**, *24*, 1–65. <https://doi.org/10.1086/648285>
19. Gidea, M.; Katz, Y. Topological Data Analysis of Financial Time Series: Landscapes of Crashes. *Physica A* **2018**, *491*, 820–834. <https://doi.org/10.1016/j.physa.2017.09.028>
20. Goel, A.; Passmore, C.; Sandhu, R. Topological Data Analysis of Equity Order Books. In *Proceedings of the Second ACM International Conference on AI in Finance*; ACM: New York, NY, USA, 2020. <https://doi.org/10.1145/3383455.3422533>

21. Hansen, P.R.; Lunde, A.; Nason, J.M. The Model Confidence Set. *Econometrica* **2011**, *79*, 453–497. <https://doi.org/10.3982/ECTA5771>
22. Ismail, M.S.; Noorani, M.S.M.; Ismail, M.; Razak, F.A.; Alias, M.A. Early Warning Signals of Financial Crises Using Persistent Homology. *PLoS ONE* **2022**, *17*, e0275690. <https://doi.org/10.1371/journal.pone.0275690>
23. Landauer, R. Irreversibility and Heat Generation in the Computing Process. *IBM Journal of Research and Development* **1961**, *5*, 183–191. <https://doi.org/10.1147/rd.53.0183>
24. Li, H.; Xiao, Y.; Polukarov, M.; Ventre, C. Thermodynamic Analysis of Financial Markets: Measuring Order Book Dynamics with Temperature and Entropy. *Entropy* **2024**, *26*, 24. <https://doi.org/10.3390/e26010024>
25. Liu, B.; Wei, Q.; Zhang, P.; Gu, Y. Can Topological Transitions in Cryptocurrency Systems Serve as Early Warning Signals? *Physica A* **2024**, *655*, 130139. <https://doi.org/10.1016/j.physa.2024.007039>
26. Makatjane, K.; Shoko, C.; Moroke, N. Downside Risk of the South African Mining Index: Leveraging Long Short-Term Memory and Explainable AI. *Journal of International Crisis and Risk Communication Research* **2025**, *8*.
27. Mantegna, R.N.; Stanley, H.E. *An Introduction to Econophysics*; Cambridge University Press: Cambridge, UK, 1999. <https://doi.org/10.1017/CBO9780511755767>
28. Miolane, N.; Guigui, N.; Le Brigant, A.; et al. Geomstats: A Python Package for Riemannian Geometry in Machine Learning. *Journal of Machine Learning Research* **2020**, *21*, 1–9. Available online: <http://jmlr.org/papers/v21/19-027.html>
29. Moroke, N.D.; Metsileng, L.D. A Maximum-Entropy Markov-Switching GARCH Framework for Cryptocurrency Volatility Regime Detection and Forecasting. *Preprints* **2026**, 2026042071. <https://doi.org/10.20944/preprints202604.2071.v1>
30. Moroke, N.D. VORTEX-GRU: A Navier–Stokes-Structured Gated Recurrent Unit for Cryptocurrency Volatility Forecasting. *Kybernetes* **2026**, under review.
31. Moroke, N.D. Deep Reinforcement Learning for Cryptocurrency Portfolio Management: A Free-Energy Framework with Geometry-Based Transaction Costs. *Risks* **2026**, *14*, 103. <https://doi.org/10.3390/risks14050103>
32. Nakazato, M.; Ito, S. Geometrical Aspects of Entropy Production in Stochastic Thermodynamics Based on Wasserstein Distance. *Physical Review Research* **2021**, *3*, 043093. <https://doi.org/10.1103/PhysRevResearch.3.043093>
33. Nielsen, F. The Fisher–Rao Geodesics of Gaussian Distributions and Statistical Manifolds. *Entropy* **2024**, *26*, 135. <https://doi.org/10.3390/e26020135>
34. Oudot, S.Y. *Persistence Theory: From Quiver Representations to Data Analysis*; American Mathematical Society: Providence, RI, USA, 2015. <https://doi.org/10.1090/surv/209>
35. Peyré, G.; Cuturi, M. Computational Optimal Transport. *Foundations and Trends in Machine Learning* **2019**, *11*, 355–607. <https://doi.org/10.1561/22000000073>
36. Pflug, G.C.; Pichler, A. A Distance For Multistage Stochastic Optimization Models. *SIAM Journal on Optimization* **2012**, *22*, 1–23. <https://doi.org/10.1137/110825054>
37. Rao, C.R. Information and the Accuracy Attainable in the Estimation of Statistical Parameters. *Bulletin of the Calcutta Mathematical Society* **1945**, *37*, 81–91.
38. Rodriguez-Rodriguez, N.; Miramontes, O. Shannon Entropy: An Econophysical Approach to Cryptocurrency Portfolios. *Entropy* **2022**, *24*, 1583. <https://doi.org/10.3390/e24111583>
39. Samal, A.; Pharasi, H.K.; Ramaia, S.J.; Kannan, H.; Saucan, E.; Jost, J.; Chakraborti, A. Network Geometry and Market Instability. *Royal Society Open Science* **2021**, *8*, 201734. <https://doi.org/10.1098/rsos.201734>
40. Sánchez García, J.; Gherghe, S. On the Ollivier-Ricci Curvature as Fragility Indicator of Stock Markets. *arXiv* **2024**, arXiv:2405.07134. Available online: <https://arxiv.org/abs/2405.07134>
41. Sandhu, R.S.; Georgiou, T.T.; Tannenbaum, A.R. Ricci Curvature: An Economic Indicator for Market Fragility and Systemic Risk. *Science Advances* **2016**, *2*, e1501495. <https://doi.org/10.1126/sciadv.1501495>
42. Song, S.; Li, H. Early Warning Signals for Stock Market Crashes: Empirical and Analytical Insights Utilizing Nonlinear Methods. *EPJ Data Science* **2024**, *13*, 16. <https://doi.org/10.1140/epjds/s13688-024-00457-2>
43. Venturini, M.; García-Costa, D.; Álvarez-García, E.; Grimaldo, F.; Squazzoni, F. Mapping Network Structures and Dynamics of Decentralised Cryptocurrencies: The Evolution of Bitcoin (2009–2023). *PLoS ONE* **2025**, *20*, e0329617. <https://doi.org/10.1371/journal.pone.0329617>
44. Villani, C. *Topics in Optimal Transportation*; American Mathematical Society: Providence, RI, USA, 2003. <https://doi.org/10.1090/gsm/058>
45. Vos, P.W. Geometry of Statistical Manifolds. *Entropy* **2025**, *27*, 1110. <https://doi.org/10.3390/e27111110>

46. Martínez-García, M.; et al. Information Theory Quantifiers in Cryptocurrency Time Series Analysis. *Entropy* **2025**, *27*, 450. <https://doi.org/10.3390/e27040450>
47. Wang, L.; An, S.; Dong, Z.; Dong, X.; Li, J. Early Warning of Regime Switching in a Financial Time Series: A Heteroskedastic Network Model. *PLoS ONE* **2025**, *20*, e0333734. <https://doi.org/10.1371/journal.pone.0333734>
48. Yen, P.T.-W.; Cheong, S.A. Using Topological Data Analysis and Persistent Homology to Analyze Stock Markets in Singapore and Taiwan. *Frontiers in Physics* **2021**, *9*, 572216. <https://doi.org/10.3389/fphy.2021.572216>
49. Zhang, D.; Zhuang, Y.; Tang, P.; Peng, H.; Han, Q. Financial Price Dynamics and Phase Transitions in Stock Markets. *The European Physical Journal B* **2023**, *96*, 35. <https://doi.org/10.1140/epjb/s10051-023-00501-6>
50. Moroke, N.D. Metabolic Saliency as a KL-Divergence Estimator: Information-Geometric Attribution of Systemic Stress in JSE Equity Networks. *Preprints* **2026**, 2026040939. <https://doi.org/10.20944/preprints202604.0939.v1>
51. Vos, P.W. Geometry of Statistical Manifolds. *Entropy* **2025**, *27*, 1110. <https://doi.org/10.3390/e27111110>
52. Martínez-García, M.; Segovia-Domínguez, I.; Rodríguez-Martínez, A.; et al. Information Theory Quantifiers in Cryptocurrency Time Series Analysis. *Entropy* **2025**, *27*, 450. <https://doi.org/10.3390/e27040450>

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.