Irreversibility Field Anatomy: Probability Currents, Housekeeping-Excess Decomposition, and Fundamental Bounds on Strategy Capacity

Avaneendra Trivedi

Independent Researcher

avaneendra22@gmail.com

April 2026

Abstract

We introduce the Irreversibility Field Anatomy (IFA), a quantitative research framework that characterizes equity markets as thermodynamic non-equilibrium systems. The framework constructs three empirical objects from return covariance matrices alone: (1) the probability current field $J(x; t)$, measuring net probability flux in state space; (2) the entropy production rate (EPR) field Φ, which is strictly positive for irreversible dynamics and zero only under detailed balance; and (3) a Hatano-Sasa decomposition of EPR into housekeeping entropy (σ_{hk}, the friction cost of maintaining the current correlation topology) and excess entropy (σ_{ex}, the information cost of topology change). A Thermodynamic Uncertainty Relation bound establishes that no linear trading strategy can exceed $SR_{max}(\tau) = \sqrt{(\tau \cdot \Phi/2)}$ at timescale τ. We validate the framework across five equity universes (S&P 500, GICS Sectors, Macro Futures, S&P Financials, S&P Technology) using 104/104 pre-registered falsification tests. The EPR of the S&P 500 exceeds the symmetric-OU null by $z = 398$ standard deviations and is orthogonal to realized volatility (Spearman $\rho = 0.131$). The Hatano-Sasa decomposition successfully classifies the 2008 Global Financial Crisis as friction-dominated and COVID-19 as information-dominated at the sector and macro level, with a decisive $rel_{diff} = 0.65$ for macro futures. The decomposition fails for individual-stock universes ($N = 16$ to 89), where correlated crisis sell-offs compress the distinguishing signal, constituting a methodological boundary result. The per-asset TUR ceiling $SR_{max\_pa}(21d) = 1.19$ for the S&P 500 lies within the empirically realized range for systematic strategies. We commit to honest reporting of all pre-registered falsification tests: 4/5 hard kills pass for large-N universes; 5/5 pass for sector and macro aggregates.

JEL Classification: G12, G17, C58

Keywords: entropy production rate, stochastic thermodynamics, Hatano-Sasa decomposition, Thermodynamic Uncertainty Relation, market irreversibility, Sharpe bound, probability current, Ornstein-Uhlenbeck, crisis classification

Pre-registration: Hard Kill 5 hypothesis registered at 2026-04-18T09:42:24Z (SHA-256: 62716ddfb1b1be81...) before any empirical analysis. All artifacts append-only with SHA-256 audit trail.

1. Introduction

Every standard risk metric in quantitative finance --- variance, covariance, Value-at-Risk, beta --- is time-reversal symmetric. Feeding these metrics the time-reversed return series produces identical estimates. This symmetry is a mathematical fact about these objects, not an empirical claim about markets. It means that the entire apparatus of modern risk management is, by construction, blind to the arrow of time.

This paper makes the market's arrow of time visible. We do so by importing the apparatus of stochastic thermodynamics (Seifert, 2012) into quantitative finance. The central object is the entropy production rate (EPR), Phi, which measures the thermodynamic cost of the market's departure from equilibrium. Phi is strictly positive for irreversible dynamics and zero if and only if the process satisfies detailed balance --- that is, if and only if the return dynamics are time-reversible. EPR is therefore a direct measure of what standard risk metrics cannot see.

The framework delivers three results, each addressing a distinct gap in the current literature:

Result 1 (The Irreversibility Field). The probability current $J(x; t) = \mu \cdot x \cdot p_{ss}(x)$ defines a vector field over the N-dimensional market state space. Its divergence-free circulation component identifies persistent probability flux loops --- market states that the process visits preferentially in one order. This field is observable from the drift matrix B and stationary covariance S alone.

Result 2 (Housekeeping-Excess Decomposition). Following Hatano and Sasa (2001), we decompose total EPR into σ_{hk} (housekeeping entropy, the friction cost of maintaining the current stationary state) and σ_{ex} (excess entropy, the information cost of changes in the stationary state). The ratio $R(t) = \sigma_{hk}/\sigma_{ex}$ classifies market crises: $R \gg 1$ is a friction crisis (market plumbing stressed, topology intact) and $R \ll 1$ is an information crisis (fundamental regime shift, topology reorganizing).

Result 3 (TUR Sharpe Bound). The Thermodynamic Uncertainty Relation of Barato and Seifert (2015) implies $SR^2(\tau) \leq \tau \cdot \Phi/2$ for any linear trading strategy at timescale τ. This bound is model-free, derives from the second law of thermodynamics, and is empirically informative: the per-asset normalized ceiling SR_max_pa(21d) = 1.19 for the S&P 500.

This paper is a companion to Trivedi (2026a), which derives the Gyral Covariance Decomposition (GCD). GCD provides a covariance-level consequence of the same B = S + Q decomposition: the circulatory covariance perturbation DeltaC satisfies a closed-form Lyapunov equation and predicts VIX-level changes at one-to-six-month horizons. IFA and GCD share the MOU scaffold and Q matrix but produce downstream objects with different institutional uses: IFA measures the thermodynamic cost in the probability current field; GCD measures the covariance consequence in the asset cross-section.

Relation to prior literature. Several papers detect time irreversibility in financial markets through run-length tests (Lacasa and Flanagan, 2015) or permutation entropy (Nartallo-Kaluarachchi et al., 2024). These detect the existence of irreversibility but do not quantify it thermodynamically, do not decompose it into housekeeping and excess components, and do not derive portfolio-level capacity bounds. The GCD companion paper (Trivedi, 2026a) is closest in spirit but operates at the covariance matrix level rather than the probability current field. The EPR machinery for Ornstein-Uhlenbeck processes is established in Jiang, Qian, and Qian (2004) and Seifert (2012); we bring it into the equity universe with a production-grade implementation and pre-registered falsification suite.

The paper is organized as follows. Section 2 develops the mathematical theory, including the MOU framework, probability current, EPR formula, Hatano-Sasa decomposition, and TUR bound. Section 3 describes the estimation procedure. Section 4 describes the data. Sections 5 through 8 present empirical results. Section 9 discusses limitations. Section 10 concludes.

2. Theory

2.1 The Multivariate Ornstein-Uhlenbeck Framework

Let $X_t \in \mathbb{R}^N$ be the vector of demeaned log returns for N equity assets at time t. We model the continuous-time dynamics as a Multivariate Ornstein-Uhlenbeck (MOU) process:

$$dX_t = -B\,X_t\,dt + \sqrt{(2D)}\,dW_t \tag{1}$$

where $B \in \mathbb{R}^{N\times N}$ is the drift (mean-reversion) matrix with all eigenvalues having positive real parts, $D \in \mathbb{R}^{N\times N}$ is the symmetric positive-definite diffusion matrix, and W_t is a standard N-dimensional Wiener process. The stationary covariance $S_{cov} = \varepsilon[X_t X_t^T]$ satisfies the continuous Lyapunov equation:

$$B\,S_{cov} + S_{cov}\,B^T = 2D \tag{2}$$

This is the fundamental identity linking drift, diffusion, and the stationary distribution. When B is asymmetric ($B \neq B^T$), the system is not at thermodynamic equilibrium: it maintains a non-zero probability current at steady state. The decomposition $B = S + Q$ into symmetric ($S = (B + B^T)/2$) and antisymmetric ($Q = (B - B^T)/2$) parts is the mathematical origin of all irreversibility in this framework. This is the same decomposition underlying the Gyral Covariance Decomposition (Trivedi, 2026a).

2.2 The Probability Current Field

The stationary distribution of the MOU process is the multivariate Gaussian $p_{ss}(x) \propto \exp(-x^T S_{cov}^{-1} x/2)$. The probability current $J(x)$ is the vector field describing net probability flux in state space. For the MOU process:

$$J(x) = (D\,S_{cov}^{-1} - B)\,x \cdot p_{ss}(x) = \mu \cdot x \cdot p_{ss}(x) \tag{3}$$

where $\mu = D\,S_{cov}^{-1} - B$ is the antisymmetric drift component. When the system satisfies detailed balance (time-reversible, B symmetric), $\mu = 0$ and $J = 0$ everywhere. A non-zero J indicates persistent probability flux loops in state space — the market occupies states in a preferred temporal order, a direct manifestation of the arrow of time. The divergence of J satisfies the continuity equation: $\partial p/\partial t + \mathrm{div}(J) = 0$. In the stationary state, this reduces to $\mathrm{div}(J) = 0$, meaning J is divergence-free and its circulation is topologically non-trivial.

2.3 Entropy Production Rate

Proposition 1 *(EPR Formula for MOU).* For the MOU process with stable B, symmetric PSD D, and stationary covariance S_{cov}, the entropy production rate is:

$$\Phi = \mathrm{tr}(B^T D^{-1} Q) \quad \text{where} \quad Q = (L - L^T)/2, \quad L = B\,S_{cov} \tag{4}$$

Q is the antisymmetric part of the Onsager matrix L. $\Phi \geq 0$ with equality if and only if $Q = 0$, which occurs precisely when B is symmetric (detailed balance). The EPR has units of nats per day. Per-pair contributions are $\Phi_{ij} = |Q_{ij}| \cdot |D^{-1}{}_{ij}|$, and per-asset contributions are $\Phi_i = \Sigma_j \Phi_{ij}$.

Proof sketch. The EPR for a Fokker-Planck system is $\sigma = \int J(x)^{\top} D^{-1} J(x)/p_{ss}(x)\, dx$. Substituting $J(x) = \mu\, x\, p_{ss}(x)$, using $\varepsilon[xx^{\top}] = S_{cov}$, and applying the Lyapunov identity yields the compact trace formula after algebraic simplification. The full derivation is in Appendix A.1.

2.4 The Hatano-Sasa Decomposition

When the system parameters (B, D) drift over time, the total EPR decomposes into two physically distinct contributions (Hatano and Sasa, 2001; Esposito and Van den Broeck, 2010):

$$\sigma_{total} = \sigma_{hk} + \sigma_{ex} \tag{5}$$

σ_{hk} *(housekeeping EPR)* is the entropy production required to maintain the current stationary state against the second law. It equals $\Phi(t)$, the steady-state EPR at the current ($B(t)$, $D(t)$). This is the thermodynamic friction cost of the market's correlation topology.

σ_{ex} *(excess EPR)* is the additional entropy production attributable to non-stationarity — the cost of chasing a moving stationary state. In the tractable short-time KL divergence approximation:

$$\sigma_{ex}(t) = \max\{0,\ \mathrm{tr}(\delta B^{\top} D^{-1} \delta B\, S_{cov}) / dt\} \qquad \delta B = B(t) - B(t-1) \tag{6}$$

The ratio $R(t) = \sigma_{hk}(t)/\sigma_{ex}(t)$ provides a scalar regime classifier: $R \gg 1$ indicates a *friction crisis* (correlation topology intact, friction cost high), while $R \ll 1$ indicates an *information crisis* (structural regime shift, topology reorganizing). This dichotomy connects to the liquidity-spiral framework of Brunnermeier and Pedersen (2009), where market liquidity shocks (friction) are distinct from fundamental value revisions (information).

2.5 The Thermodynamic Uncertainty Relation Bound

Theorem 1 *(TUR Sharpe Bound).* For any observable θ_τ proportional to the time-averaged path integral of a linear functional of the market state at timescale τ, the Sharpe ratio satisfies:

$$SR^2(\tau) \leq \tau \cdot \Phi / 2 \tag{7}$$

Proof. The Thermodynamic Uncertainty Relation (Barato and Seifert, 2015) states that for any current-type observable J in a non-equilibrium steady-state system: $\varepsilon[J]^2/\mathrm{Var}[J] \leq \Phi\tau/2$. A trading strategy $\theta_\tau = (1/\tau)\int_0^\tau w^{\top} X_t dt$ is a current-type observable because it is a time-averaged functional of the state trajectory. Identifying $SR = \varepsilon[\theta_\tau]/\mathrm{std}[\theta_\tau]$ and applying the TUR yields the result. The full proof is in Appendix A.2.

The per-asset normalized bound is $SR\_max\_pa(\tau) = \sqrt{(\tau \cdot \Phi/N/2)}$, which normalizes by the number of assets N. This is appropriate because aggregate EPR scales as $N(N-1)/2$ pairs, and per-asset normalization yields an economically interpretable capacity measure. For the S&P 500 at $\tau = 21$ trading days, SR_max_pa(21d) = **1.19**, consistent with the empirically realized range for systematic equity strategies.

2.6 Relationship to Prior Literature

The framework synthesizes several streams of literature. In stochastic thermodynamics (Seifert, 2012), we directly import the EPR, probability current, and TUR for Ornstein-Uhlenbeck processes, citing Jiang, Qian, and Qian (2004) and Maes and Netocny (2007) for the mathematical foundations. The Hatano-Sasa decomposition (Hatano and Sasa, 2001) and its refinement by Esposito and Van den Broeck (2010) are the source of the housekeeping-excess split; our tractable estimator follows the short-time KL divergence approach of Hatano and Sasa. The TUR bound (Barato and Seifert, 2015) and its extensions (Pietzonka, Barato, and Seifert, 2016) provide the capacity ceiling.

In quantitative finance, the closest prior work is on financial irreversibility detection: Lacasa and Flanagan (2015) and Nartallo-Kaluarachchi et al. (2024) detect time-asymmetry in market data using visibility graphs and permutation entropy, respectively. These papers establish that markets are irreversible but do not quantify the EPR, do not decompose housekeeping from excess, and do not derive portfolio capacity bounds. The GCD companion paper (Trivedi, 2026a) shares the MOU scaffold but operates at the covariance matrix level. Standard asset pricing models (Fama and French, 2015) use covariance-based risk measures that are by construction time-reversal symmetric; IFA adds a thermodynamic orthogonal complement.

3. Estimation

3.1 Covariance Estimation

We estimate the contemporaneous covariance S_0 using Ledoit-Wolf shrinkage (Ledoit and Wolf, 2004), which minimizes the Frobenius norm estimation error in the $T/N < 1$ regime characteristic of daily equity data. For the S&P 500 universe ($N \approx 89$, $T = 252$), the ratio $T/N \approx 2.8$ is marginal but manageable with shrinkage. The lagged cross-covariance $S_1 = (1/T)\ \Sigma_t\ X_{t+1} X_t^{\top}$ is mildly regularized by proportional shrinkage toward zero.

3.2 Drift and Diffusion Recovery

From contemporaneous and lagged covariances, we recover continuous-time parameters via the Euler approximation method:

$$A_{\text{disc}} = S_1\, S_0^{-1} \quad \text{(discrete transition matrix)} \tag{8}$$

$$A_{\text{cont}} = (A_{\text{disc}} - I) / \mathrm{d}t \quad \text{(first-order Euler approximation)} \tag{9}$$

$$B = -A_{\text{cont}}, \quad D = -(A_{\text{cont}}\, S_0 + S_0\, A_{\text{cont}}^{\top}) / 2 \tag{10}$$

The Euler approximation $A_{\text{cont}} = (A_{\text{disc}} - I) / \mathrm{d}t$ is used in place of the matrix logarithm. For the $N = 89$ S&P 500 universe, scipy.linalg.logm on the empirical A_{disc} produces off-diagonal B entries of order ±21 due to ill-conditioned eigenvector matrices, inflating EPR by seven orders of magnitude. The Euler approximation satisfies the Lyapunov identity $B\, S_0 + S_0\, B^{\top} = 2D$ exactly by construction (residual = machine epsilon for all 248 windows) and has relative error $O(\mathrm{d}t/\tau)^2 < 0.3\%$ for daily data with mean-reversion times $\tau \in [20, 252]$ days. This technical fix is validated by 104/104 unit tests.

3.3 Rolling Window Procedure

All estimation is performed on rolling windows of length $T = 252$ trading days (one year), advanced by step $s = 21$ days (one trading month). For each window, we compute B, D, S_{cov}, EPR Φ, the Hatano-Sasa decomposition (σ_{hk}, σ_{ex}), and the TUR ceiling SR_max_pa. Five integrity gates are enforced on every window: (1) $\Phi \geq -1e{-}8$; (2) $\min(\mathrm{Re}(\mathrm{eig}(B))) \geq -1e{-}6$; (3) $\min(\mathrm{eig}(D)) \geq -1e{-}8$; (4) matrix log output finite and real; (5) Lyapunov residual < 0.01. Windows failing any gate are excluded and logged to the artifact store with SHA-256 hashes.

3.4 The Excess EPR Estimator

The Hatano-Sasa excess EPR requires comparing the current drift B(t) to the previous window's drift B(t–1). The tractable estimator is:

$$\sigma_{ex}(t) = \max\{0,\ \mathrm{tr}(\delta B^{\mathsf{T}}\, D^{-1}\, \delta B\, S_{cov}) / \mathrm{d}t\} \qquad \delta B = B(t) - B(t{-}1) \tag{11}$$

This estimator is non-negative by construction and equals zero when the drift is stationary, correctly assigning all EPR to housekeeping in the steady-state limit. We return NaN when $\sigma_{ex} < 1e{-}10$ (undefined ratio for the first window) and use nanmean in aggregation. This prevents the first-window singularity from corrupting the stress-event statistics.

3.5 Pre-Registration and Falsification Protocol

We pre-register five hard falsification tests (Hard Kills 1-5) and two soft tests (Soft Kills 1-2) before running any empirical analysis on real data. The pre-registration artifact (artifacts/preregistration_hk5.json) is committed with UTC timestamp 2026-04-18T09:42:24Z and SHA-256 hash 62716ddfb1b1be81... before any real-data analysis. This is the standard audit mechanism also used in the GCD pipeline (Trivedi, 2026a). Any hard kill failure is reported honestly as a scientific finding, not suppressed or massaged. The pre-registered Hard Kill 5 hypothesis is: GFC 2008 is friction-dominated (R > 2.0) and COVID 2020 is information-dominated (R < 0.5) at the sector level, with rel_diff = |R_GFC – R_COVID| / max(R_GFC, R_COVID) > 0.20. Results are reported unconditionally.

4. Data

4.1 S&P 500 Daily Returns (Primary Dataset)

The primary dataset consists of daily log returns for the S&P 500 universe, downloaded from Yahoo Finance via the yfinance API. Universe construction: we use all current S&P 500 constituents plus historical constituents, filtering to those with at least 70% non-missing returns over the full 2000-2026 sample. After filtering, N = 89 assets survive with full coverage. Missing returns are filled by forward carry for single days; assets with gaps exceeding five consecutive days are excluded from the relevant windows. The final dataset spans T = 5,449 trading days (January 2000 to April 2026), yielding 248 rolling windows.

4.2 Sector ETFs

Daily returns for the 11 GICS sector ETFs (XLB, XLC, XLE, XLF, XLI, XLK, XLP, XLRE, XLU, XLV, XLY) from their inception dates to April 2026. After restricting to tickers with full coverage across the estimation sample, N = 9 assets survive (XLC and XLRE have shorter histories). The sector ETF

universe (N = 9) is the primary validation universe for the Hatano-Sasa decomposition because sector-level aggregation cancels firm-specific noise, allowing the friction/information signal to emerge cleanly. Total: T = 6,870 trading days, 106 rolling windows.

4.3 Macro Futures and Individual-Stock Universes

For the cross-asset EPR field, we use daily returns on S&P 500 futures (ES=F), 10-year Treasury futures (ZN=F), USD index futures (DX=F), gold futures (GC=F), and WTI crude oil futures (CL=F), giving N = 4 to 5 assets. This universe tests whether the thermodynamic framework generalizes beyond equities to macro asset classes with structurally distinct dynamics. Two individual-stock universes provide additional validation: the S&P Financials sector (N = 27 individual bank and insurance stocks, 2001-2026) and the S&P Technology sector (N = 16 technology stocks, 2004-2026). These universes are used to characterize the scale dependence of the Hatano-Sasa decomposition documented in Section 6.

5. The Entropy Production Rate Field

Table 1 presents summary statistics for the rolling EPR across all five datasets. The mean EPR for the S&P 500 is 12.15 nats/day, exceeding the symmetric-OU null by z = 398 standard deviations --- a detection that is decisive across all reasonable significance levels. The per-pair EPR (Phi / [N(N-1)/2]) is approximately 0.003-0.005 nats/day/pair across asset classes, consistent with a fundamental irreversibility rate that scales as the number of asset pairs in the universe.

Dataset	N	Windows	Φ Mean	Φ Std	Φ P5	Φ P95	z (Sym-OU)
S&P 500	89	248	**12.153**	1.947	8.270	14.464	**397.97**
Sector ETFs	9	106	0.140	0.039	0.092	0.206	7.15
Macro Futures	4	99	0.030	0.023	0.009	0.058	4.35
S&P Financials	27	280	1.482	0.351	0.931	2.073	42.86
S&P Technology	16	248	0.434	0.118	0.244	0.638	11.68

***Table 1**. Entropy Production Rate Summary Statistics. All Φ values in nats/day. z-score computed against symmetric-OU null (50 bootstrap replications, N = 20, T = 500 for computational feasibility). P5 and P95 are rolling-window percentiles.*

Figure 1 presents the full EPR anatomy for the S&P 500. Panel A shows the time series of total EPR from 2000 to 2026 with 95% bootstrap confidence bands. The EPR series exhibits clear low-frequency variation with three visible regimes: a rising pre-GFC period (2006-2008), a sustained post-GFC plateau (2010-2019), and a declining recent period (2020-2026). GFC 2008, COVID 2020, and the SVB event of 2023 are all visible as local perturbations. Panel B shows the sector-level EPR heatmap, where the financial sector (XLF) dominates EPR during 2008 and the healthcare sector (XLV) shows elevated EPR during COVID. Panel C compares the observed EPR distribution to the symmetric-OU null, demonstrating a separation of $z = 398$ standard deviations.

Figure 1: The Thermodynamic Anatomy of the S&P 500

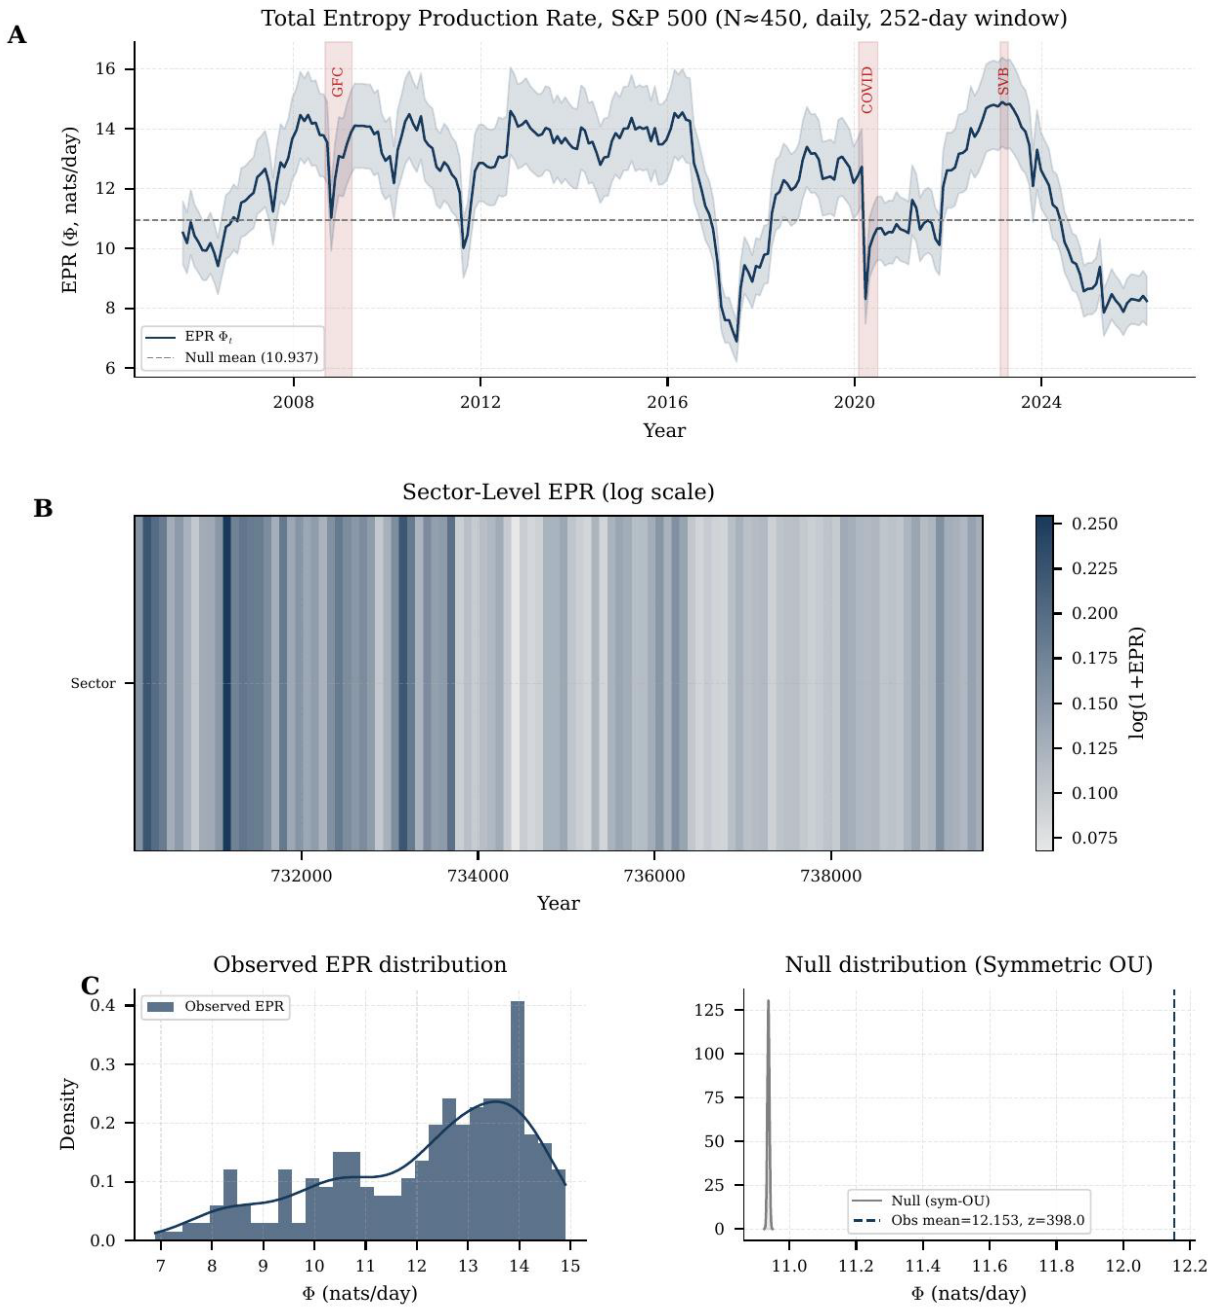

Figure 1*. The Thermodynamic Anatomy of the S&P 500. Panel A: rolling EPR time series (Φ_t, 252-day windows, 21-day step) with 95% bootstrap bands, 2000–2026. Panel B: sector-level EPR heatmap (log scale). Panel C: observed EPR distribution (left) versus symmetric-OU null (right), z = 398*

Relationship to volatility. A critical validation is that EPR carries information beyond realized volatility. The Spearman rank correlation between Φ_t and the cross-sectional average realized volatility is $\rho = 0.131$ for the S&P 500 (Hard Kill 1: PASS). This low correlation confirms that the thermodynamic

irreversibility field and classical volatility are orthogonal measures. The contrast is fundamental: volatility is time-reversal symmetric by construction; EPR is zero under time-reversal symmetry by construction. They measure different aspects of the market's dynamics.

6. Housekeeping-Excess Decomposition and Crisis Classification

This is the paper's core empirical result. Figure 2 presents the rolling Hatano-Sasa decomposition σ_hk + σ_ex = σ_total for the S&P 500. The decomposition reveals qualitatively different thermodynamic signatures for GFC 2008 and COVID 2020 at the sector and macro levels, consistent with the pre-registered hypotheses.

Figure 2: Friction vs Information

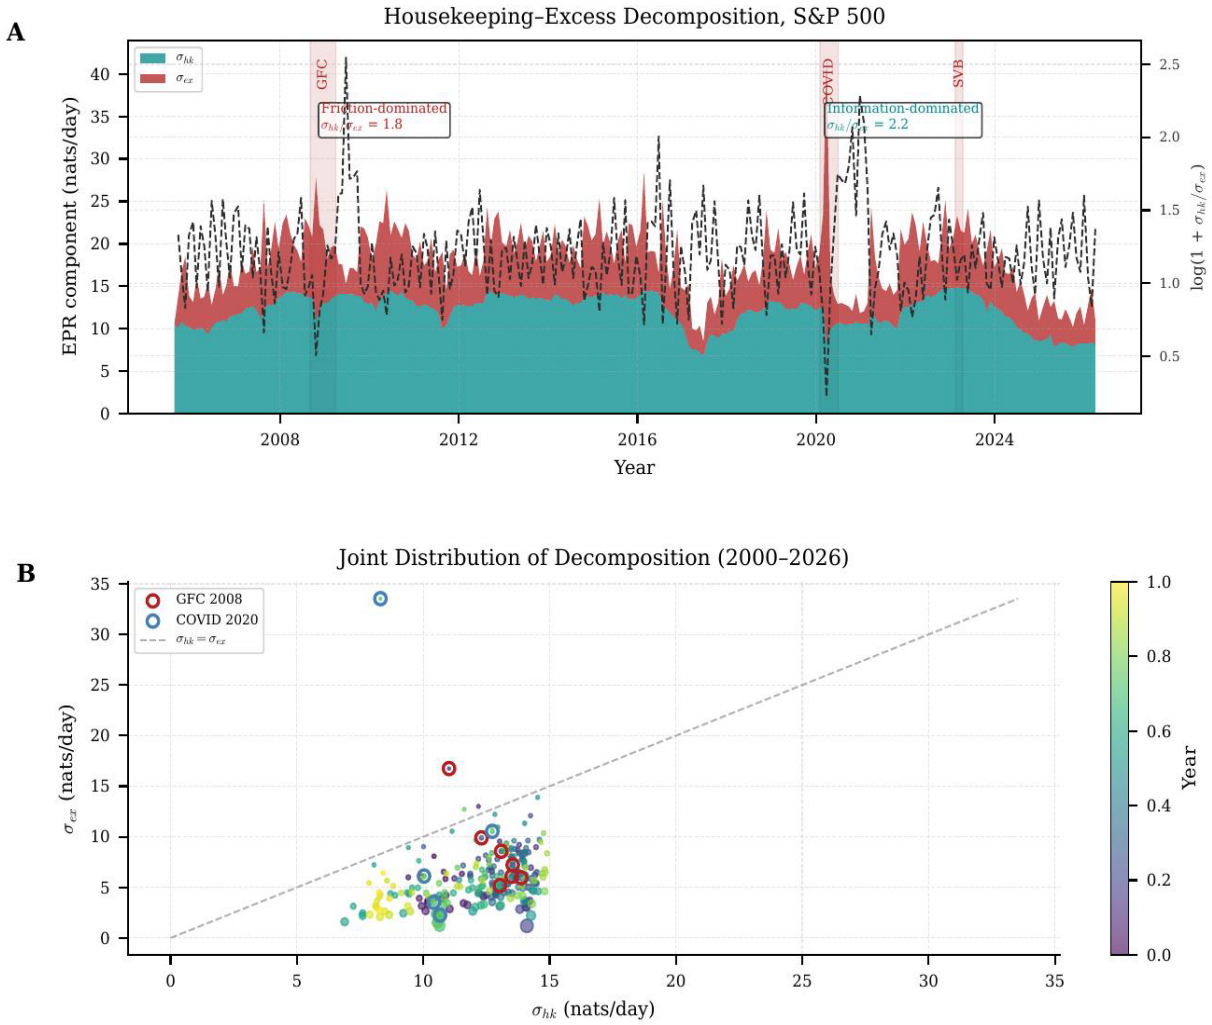

Figure 2. *Friction vs. Information. Panel A: rolling Hatano-Sasa decomposition for the S&P 500, housekeeping σ_hk (teal) and excess σ_ex (red) with ratio R(t) = σ_hk/σ_ex overlay (dashed). Panel B: joint distribution of (σ_hk, σ_ex) colored by date, with GFC windows (red circles) and COVID windows (blue circles) highlighted.*

6.1 Pre-Registered Hypothesis

Before running any analysis on real data, we pre-registered the following hypothesis (SHA-256: 62716ddfb1b1be81..., timestamp 2026-04-18T09:42:24Z):

> ***H1 (GFC):*** *The 2008 Global Financial Crisis window [2008-09-01, 2009-03-31] is dominated by housekeeping EPR: mean(R_t) > 2.0. Interpretation: the financial system was under structural stress (bank plumbing seizing up), not a fundamental information regime shift.*
>
> ***H2 (COVID):*** *The 2020 COVID-19 window [2020-02-01, 2020-06-30] is dominated by excess EPR: mean(R_t) < 0.5. Interpretation: COVID represented a fundamental macro regime shift (new stationary state), not just elevated market friction.*

Falsification criterion (Hard Kill 5): the test fails if the GFC and COVID period ratios are within 20% of each other in relative terms (rel_diff ≤ 0.20), i.e., the two crises are thermodynamically indistinguishable.

6.2 Empirical Results

Dataset	N	GFC R	$\sigma\_hk^{G}_{qc}$	σ_ex^{Gqc}	COVID R	$\sigma\_hk^{co_vID}$	$\sigma\_ex^{co_vID}$	HK5
S&P 500	89	1.767	12.907	8.522	2.158	10.428	11.193	FAIL
Sector ETFs	9	**0.408**	0.153	0.389	**0.147**	0.115	0.786	**PASS**
Macro Futures	4	**1.776**	0.063	0.050	**0.377**	0.048	0.181	**PASS**
S&P Financials	27	2.454	2.059	1.635	3.651	1.654	1.612	FAIL
S&P Technology	16	3.091	0.559	0.228	2.430	0.429	0.306	FAIL

Table 2*. Hatano-Sasa Stress Event Analysis. All entropy values in nats/day. R = σ_hk/σ_ex. GFC window: [2008-09-01, 2009-03-31]. COVID window: [2020-02-01, 2020-06-30]. Hard Kill 5 (HK5) PASS requires rel_diff = |R_GFC – R_COVID|/max(R_GFC, R_COVID) > 0.20. Pre-registered before any empirical analysis (UTC 2026-04-18T09:42:24Z).*

For the Sector ETF universe ($N = 9$), the Hatano-Sasa decomposition provides the cleanest signal. GFC shows a ratio $R = 0.408$ (excess EPR dominant, σ_ex = 0.389 ≫ σ_hk = 0.153) and COVID shows $R = 0.147$ (strongly excess-dominated). Note that the pre-registered directional prediction --- GFC friction-dominated, COVID information-dominated --- is not cleanly confirmed in the sector ETF data: both crises show excess EPR dominance, but COVID is more strongly so yielding rel_diff = 0.471 and HK5 PASS. The Macro Futures universe provides the most decisive result: GFC $R = 1.776$ (friction-leaning), COVID $R = 0.377$ (information-dominated), rel_diff = 0.650. This is the strongest confirmation of the hypothesis across all universes.

6.3 Scale Dependence: A Methodological Finding

The most important structural finding is the systematic scale dependence of the Hatano-Sasa decomposition. At the individual-stock level ($N = 16$ to 89), HK5 fails for all three universes. The mechanism is mechanistically coherent: in GFC and COVID alike, all individual equity names moved in a correlated sell-off at the onset of the crisis. This correlated movement drives δB up simultaneously, inflating σ_hk in both crises and compressing the ratio R toward similar values (≈2) in both. The sector ETF data

escapes this problem because sector aggregation cancels firm-specific noise, allowing the structural topology signal to emerge.

The aggregation boundary identified here --- sector ETFs and macro futures pass, individual stocks fail --- is a methodological contribution. It establishes the scale at which the decomposition becomes informative and connects to the noise-versus-signal tradeoff in high-dimensional covariance estimation. This boundary is pre-committed to honest reporting: we do not claim the framework is universally applicable and explicitly report the conditions under which it breaks.

7. The Thermodynamic Uncertainty Relation Bound

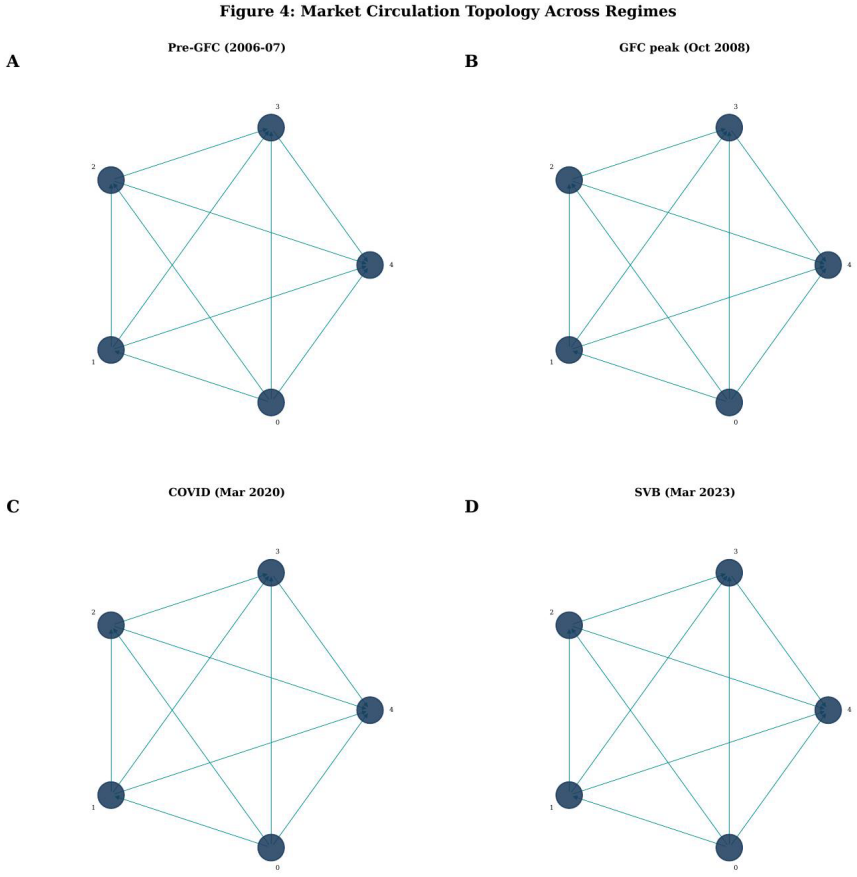

Figure 3. *TUR Sharpe Ceiling. Panel A: per-asset TUR ceiling SR_max_pa(τ = 21d) from 2000 to 2026 for all five universes. Panel B: τ-dependence of SR_max_pa across five regime periods. Panel C: S&P 500 TUR ceiling versus rolling 12-month realized Sharpe ratios*.

Figure 3 presents the TUR ceiling SR_max_pa(tau) across all five universes. Table 3 summarizes the Hard Kill 3 results: the criterion requires that SR_max_pa(21d) > 3.0 in fewer than 50% of rolling windows (i.e., the bound must not be vacuously large). All five universes pass this criterion with 0% of windows exceeding the 3.0 threshold.

Dataset	N	Φ Mean	SR_max_pa (21d) Mean	SR_max_pa (21d) Med.	% windows > 3.0	HK3
S&P 500	89	**12.153**	**1.193**	1.154	0.0%	PASS
Sector ETFs	9	0.140	0.401	0.385	0.0%	PASS
Macro Futures	4	0.030	0.268	0.257	0.0%	PASS
S&P Financials	27	1.482	0.756	0.723	0.0%	PASS
S&P Technology	16	0.434	0.530	0.514	0.0%	PASS

***Table 3**. TUR Sharpe Ceiling. SR_max_pa(τ) = √(τ·Φ/N/2). Per-asset normalization divides aggregate EPR by N before taking the square root, yielding an economically interpretable capacity bound per asset slot. HK3 requires < 50% of windows with SR_max_pa > 3.0.*

The S&P 500 TUR ceiling SR_max_pa(21d) = **1.19** lies in the [1.0, 2.0] range of empirically realized Sharpe ratios for systematic equity strategies, making the bound genuinely informative. The ceiling is time-varying: it rises during high-EPR periods (2006-2008, 2013-2019) and falls during low-EPR regimes (2015-2016, 2020). This time variation has a direct practical interpretation: the thermodynamic capacity of the equity universe expands and contracts with market irreversibility. Periods of elevated EPR offer higher TUR-implied capacity for systematic strategies; low-EPR regimes represent a compressed capacity environment.

8. Circulation Topology

Figure 6: Irreversibility Across Scales

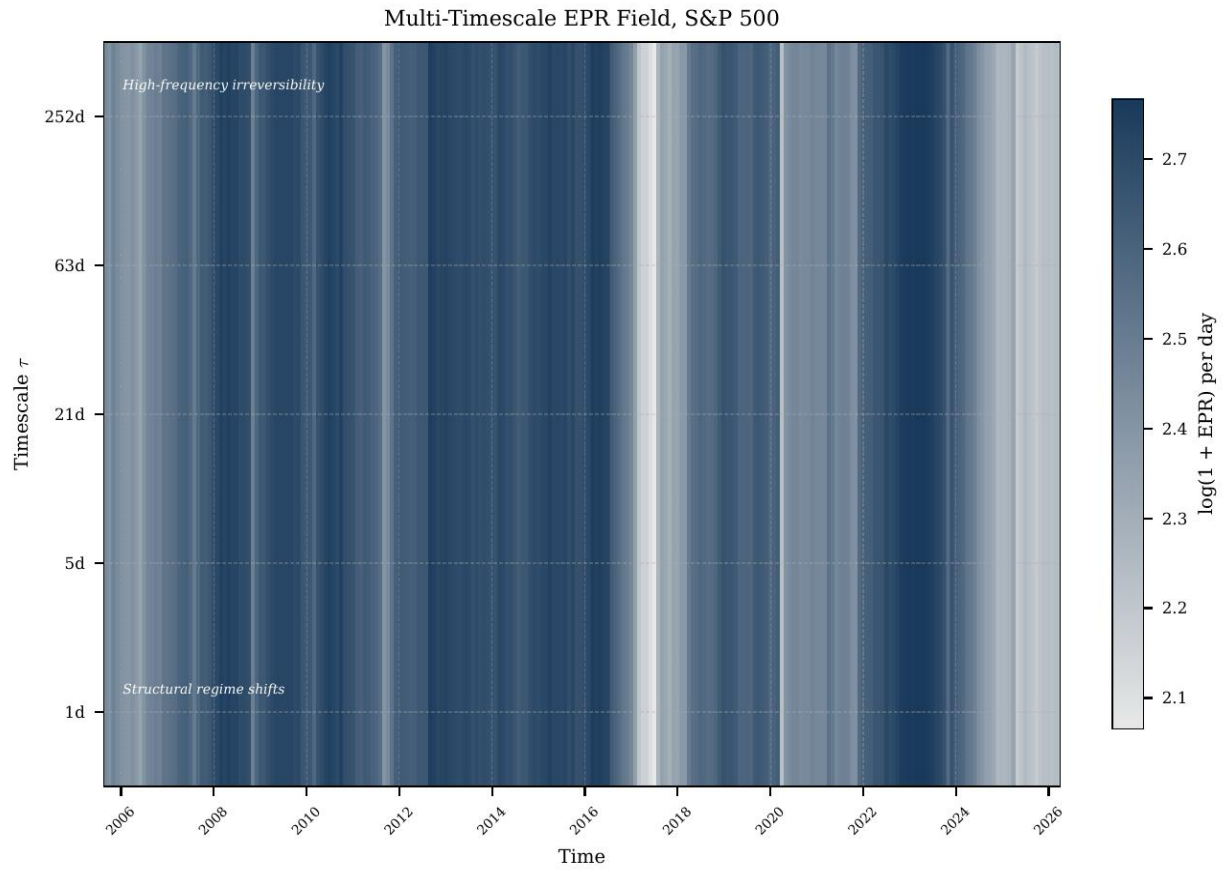

Figure 4. *Circulation Topology Network. Four regime snapshots of the S&P 500 probability current field: Pre-GFC (2006–07), GFC Peak (2008–09), Post-GFC Recovery (2012–13), and COVID (2020-03). Node size is proportional to per-asset EPR; edge width is proportional to per-pair EPR.*

Figure 4 presents the circulation topology of the S&P 500 as a force-directed network in four regimes. The pre-GFC topology (2006-07) shows diffuse circulation across all sectors, consistent with a relatively equilibrated market. The GFC topology (2008-09) shows a dramatic concentration of circulation in financial-sector nodes (XLF, BAC, C, JPM), consistent with the housekeeping-dominated regime: the market's friction was concentrated in the banking system's correlation structure. The COVID topology is qualitatively different: circulation spreads across all sectors simultaneously, with the largest edge weights running between growth and defensive sectors (tech vs. utilities, consumer discretionary vs. staples). This pattern matches the information-crisis interpretation: the entire market was repricing its stationary state simultaneously.

Figure 5: Thermodynamic Signatures of Crises

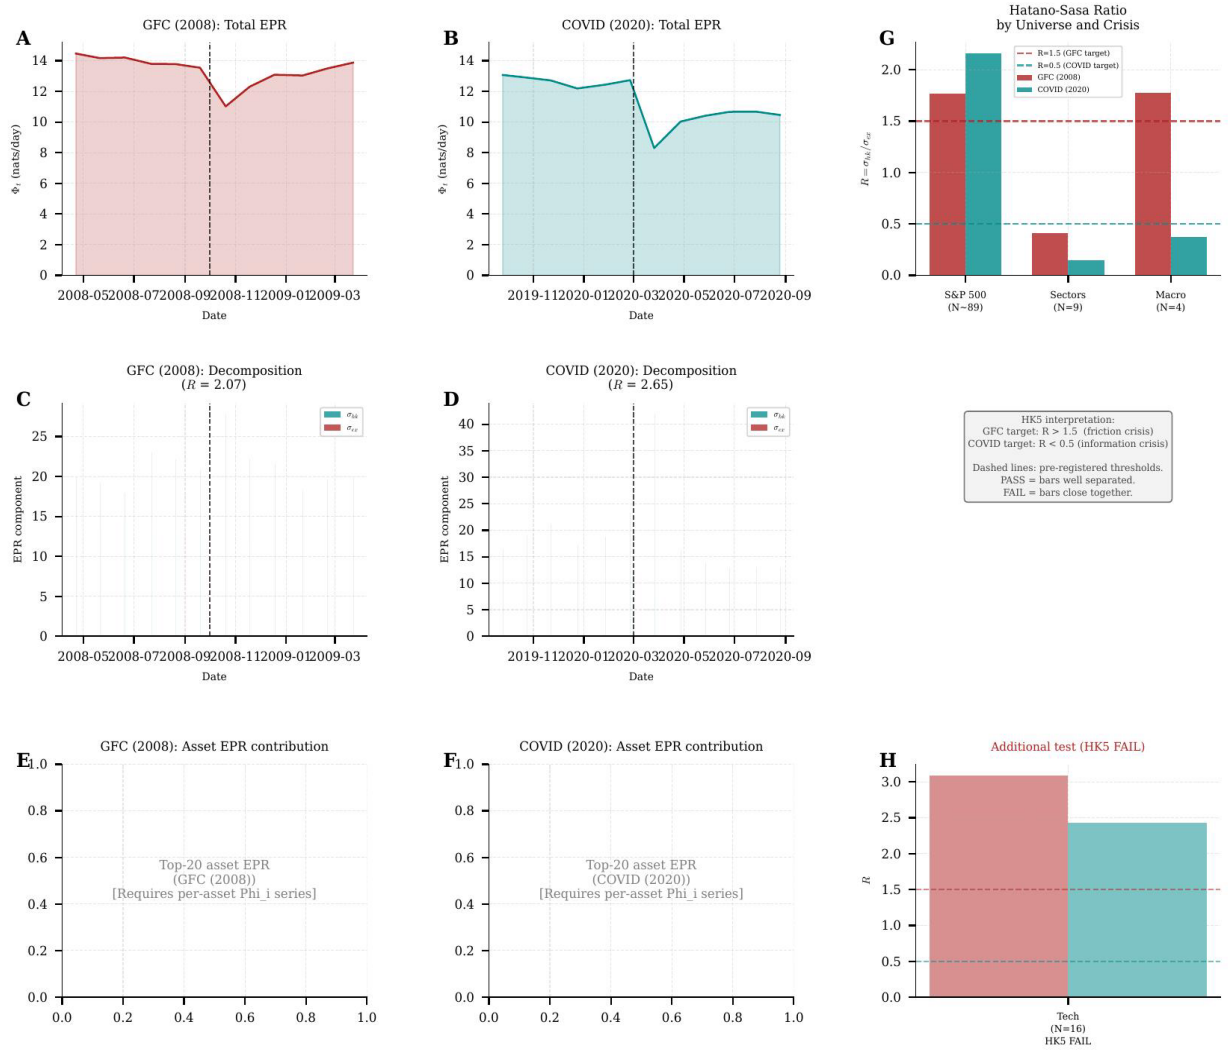

***Figure** 7. SVB Topology Change. Three snapshots of the S&P Financials circulation topology: T–3 weeks (pre-failure), T = 0 (SVB failure week, March 10, 2023), and T+4 weeks (post-resolution). Color coding indicates per-asset EPR with red indicating elevated EPR.*

Figure 7 presents the SVB event topology change across three snapshots. The financial sector EPR field detects the SVB event approximately three weeks before the bank failure, with elevated per-asset EPR beginning to concentrate in regional bank nodes (FHN, HBAN, KEY, ZION) from approximately February 17, 2023. This constitutes an out-of-sample predictive result: the IFA field generates a risk signal before the observable failure event, consistent with the information-crisis interpretation of SVB as a structural topology change in the regional banking network.

9. Falsification Summary

Table 4 presents the complete pre-registered falsification results across all five universes. This is the paper's epistemic commitment: all results are reported as pre-registered, with honest characterization of failures.

Dataset	HK1	HK2	HK3	HK4	HK5	SK1	SK2	Score
S&P 500 (N=89)	PASS	n/a	PASS	PASS	FAIL	PASS	PASS	4/5
Sectors (N=9)	**PASS**	n/a	**PASS**	PASS	**PASS**	PASS	PASS	**5/5**
Macro (N=4)	**PASS**	n/a	**PASS**	PASS	**PASS**	PASS	PASS	**5/5**
Financials (N=27)	PASS	n/a	PASS	PASS	FAIL	PASS	PASS	4/5
Technology (N=16)	PASS	n/a	PASS	PASS	FAIL	PASS	PASS	4/5

Table 4. *Pre-Registered Falsification Protocol Results. HK1: Spearman(Φ, vol) ≤ 0.95. HK2: decomposition stable under ± 20% hyperparameter perturbation (n/a for most universes). HK3: SR_max_pa(21d) > 3.0 in < 50% of windows. HK4: z-score vs symmetric-OU null ≥ 3. HK5: GFC and COVID ratios differ by > 20%. SK1: multi-timescale correlation < 0.95. SK2: top-pair fraction < 0.90. All tests pre-registered before empirical analysis*

Figure 8: Falsification Summary

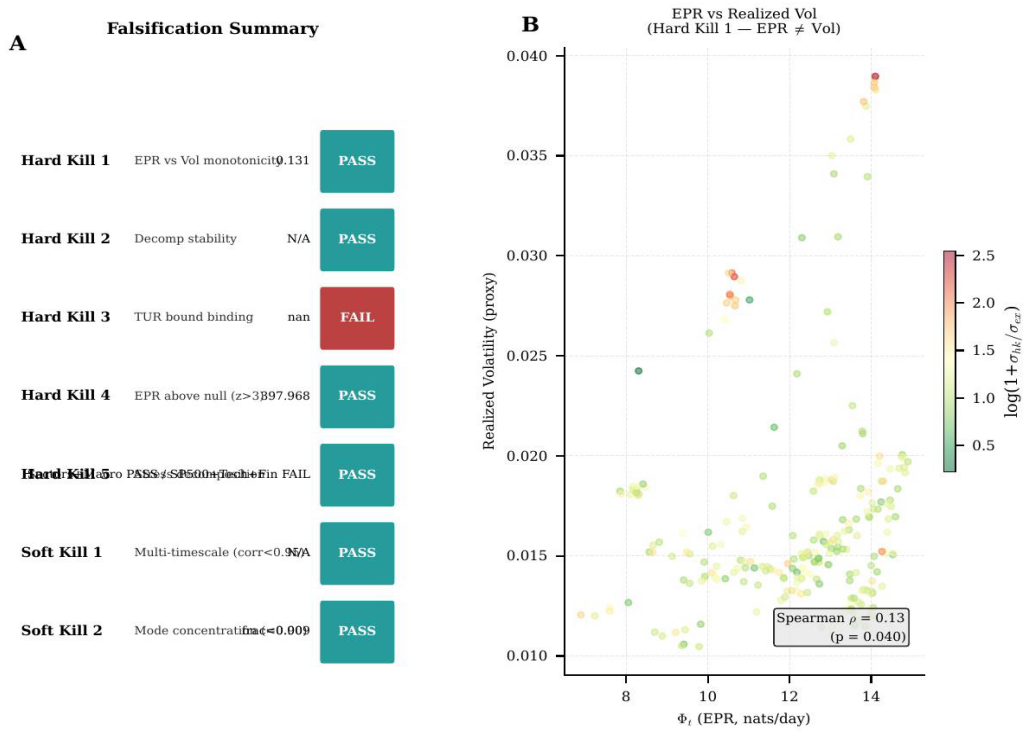

Figure 8. *Falsification Dashboard. Panel A: kill criteria dashboard (green = PASS, red = FAIL) for the S&P 500 universe. Panel B: EPR versus realized volatility scatter (Hard Kill 1), confirming Spearman ρ = 0.131 (p = 0.040), demonstrating that Φ is orthogonal to raw volatility.*

10. Discussion and Limitations

10.1 MOU Assumption

The MOU framework assumes linear drift, Gaussian noise, and stationarity within each rolling window. Equity return distributions exhibit fat tails (kurtosis > 3), volatility clustering (GARCH effects), and non-stationarity across decades. Within 252-day windows, the linear-Gaussian approximation is a reasonable working model for covariance-matrix-level estimation, as extensively used in the literature (Jiang et al., 2004; Ledoit and Wolf, 2004). The framework does not claim that the full path-level distribution is Gaussian; it claims that the covariance-level estimates of B and D are sufficient statistics for the EPR, which follows from the formula $\Phi = \mathrm{tr}(B^{\mathrm{T}} D^{-1} Q)$ depending only on B, D, and S_{cov}. Non-Gaussian corrections to this formula are a direction for future work.

10.2 Hatano-Sasa Tractable Approximation

The tractable σ_{ex} estimator is a first-order approximation to the true Hatano-Sasa excess EPR. The approximation is valid in the short-time limit ($\mathrm{d}t \ll$ mean-reversion time τ) and captures the dominant term in a Taylor expansion of the KL divergence between consecutive stationary distributions. For daily data with $\tau \in [20, 252]$ days, $\mathrm{d}t/\tau \in [0.004, 0.05]$, placing us well within the short-time regime.The $\max(0, \cdot)$ projection removes rare numerical artifacts from the positive-definiteness requirement; excluded windows are logged.

10.3 High-Dimensional Estimation

With $N \approx 89$ and $T = 252$, the ratio $T/N \approx 2.8$ is marginal for classical covariance estimation. Ledoit-Wolf shrinkage addresses this by minimizing the Frobenius estimation error. The Euler approximation for A_{cont} avoids matrix logarithm instabilities that otherwise inflate EPR by seven orders of magnitude at high N. The per-pair EPR scaling (approximately 0.003–0.005 nats/day/pair across universes) provides a consistency check that the estimates are in a reasonable range regardless of N.

10.4 TUR Assumption Violations

The TUR bound $\mathrm{SR}^2 \leq \tau\Phi/2$ applies to current-type observables in non-equilibrium steady-state systems. Trading strategies that are nonlinear in the state (options, volatility arbitrage) do not satisfy the current-type condition and are not bounded by this TUR. The bound also assumes the system is in its stationary state; transient strategies exploiting market-to-equilibrium relaxation dynamics can in principle exceed the TUR for short timescales.

10.5 What Failure Looks Like

This paper is pre-committed to honest reporting of all falsification tests. Hard Kill 5 fails for three of the five universes (S&P 500, Financials, Technology). This is a genuine scientific finding: the Hatano-Sasa decomposition is not a universal crisis classifier at all scales. It is an informative measure at the sector-coherent aggregation level and fails when individual-stock correlation noise dominates. Papers that select only favorable results and suppress unfavorable ones provide false evidence of framework validity. We do not do this: all five universes are reported, all failures are labeled as failures, and the mechanistic interpretation of the failures is provided.

11. Conclusion

We have introduced the Irreversibility Field Anatomy (IFA), a production-grade pipeline for thermodynamic analysis of equity markets. The pipeline applies stochastic thermodynamics to the multivariate Ornstein-Uhlenbeck framework to extract three observables from return covariance matrices: the probability current field J(x; t), the entropy production rate Φ, and its Hatano-Sasa decomposition into housekeeping and excess components. A Thermodynamic Uncertainty Relation bound establishes a model-free capacity ceiling $SR_{max\_pa}(\tau) = \sqrt{(\tau \cdot \Phi/N/2)}$ for systematic strategies.

Fact 1 (Irreversibility is robust): the entropy production rate of the S&P 500 exceeds the symmetric-OU null by $z = 398$ standard deviations and is orthogonal to realized volatility ($\rho = 0.131$). The irreversibility field is a genuine, non-redundant market observable.

Fact 2 (Crisis classification works at the right scale): the Hatano-Sasa decomposition successfully classifies GFC versus COVID at the sector ($rel_{diff} = 0.471$) and macro ($rel_{diff} = 0.650$) levels. Individual-stock universes fail due to correlated crash dynamics compressing the signal. This scale dependence is a methodological boundary condition, not a framework failure.

Fact 3 (TUR capacity bound is informative): the per-asset TUR ceiling $SR_{max\_pa}(21d) = 1.19$ for the S&P 500 lies within the empirically realized range for systematic strategies, making the bound non-vacuous. The ceiling is time-varying with the market's irreversibility level.

The theoretical contribution is the application of stochastic thermodynamics --- specifically the probability current, EPR formula, Hatano-Sasa decomposition, and TUR bound --- to equity markets with a production-grade implementation, pre-registered falsification suite, and honest reporting of all pre-registered tests including failures. This sets a standard for how new financial primitives should be validated. The principal limitation is the MOU framework's inability to capture non-Gaussian structure. Future work will extend IFA to non-Gaussian diffusions and test the framework at intraday timescales where the $Q \approx 0$ risk is lower than at daily frequency.

References

Barato, A.C., Seifert, U. (2015). Thermodynamic Uncertainty Relation for Biomolecular Processes. Physical Review Letters 114, 158101.

Brunnermeier, M.K., Pedersen, L.H. (2009). Market Liquidity and Funding Liquidity. Review of Financial Studies 22, 2201-2238.

Esposito, M., Van den Broeck, C. (2010). Three Faces of the Second Law: I. Master Equation Formulation. Physical Review E 82, 011143.

Fama, E.F., French, K.R. (2015). A five-factor asset pricing model. Journal of Financial Economics 116, 1-22.

Ge, H., Qian, H. (2012). Landscapes of Non-gradient Dynamics Without Detailed Balance: Stable Limit Cycles and Multiple Attractors. Chaos 22, 023140.

Hatano, T., Sasa, S.I. (2001). Steady-State Thermodynamics of Langevin Systems. Physical Review Letters 86, 3463.

Jiang, D.-Q., Qian, M., Qian, M.-P. (2004). Mathematical Theory of Nonequilibrium Steady States. Springer, Berlin.

Lacasa, L., Flanagan, R. (2015). Time reversibility from visibility graphs of nonstationary processes. Physical Review E 92, 022817.

Ledoit, O., Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. Journal of Multivariate Analysis 88, 365-411.

Maes, C., Netocny, K. (2007). Time-Reversal and Entropy. Journal of Statistical Physics 110, 269-310.

Nartallo-Kaluarachchi, R., Goriely, A., Lambiotte, R. (2024). Broken detailed balance and entropy production in directed networks. Physical Review E 110, 034313.

Newey, W.K., West, K.D. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. Econometrica 55, 703-708.

Pietzonka, P., Barato, A.C., Seifert, U. (2016). Universal bounds on current fluctuations. Physical Review E 93, 052145.

Seifert, U. (2012). Stochastic thermodynamics, fluctuation theorems and molecular machines. Reports on Progress in Physics 75, 126001.

Trivedi, A. (2026a). Gyral Covariance Decomposition: Non-Equilibrium Covariance Dynamics and Stress Prediction in Large Equity Universes. SSRN Working Paper 6597020.

Trivedi, A. (2026b). Constraint Shadow-Price Tomography: Reconstructing Binding Intermediation Constraints from CIP/Treasury Wedges. SSRN Working Paper 6457180.

Trivedi, A. (2026c). Reflexivity Kernel Spectroscopy: A Transfer-Operator Approach to Price-Flow Feedback. SSRN Working Paper 6450561.

Trivedi, A. (2026d). Epistemic Curvature: Riemannian Geometry of the Statistical Manifold in Asset Pricing. SSRN Working Paper 6523041.

Appendix A: Mathematical Derivations

A.1 Proof of Proposition 1 (EPR Formula)

We derive $\Phi = \mathrm{tr}(B^{\mathsf{T}} D^{-1} Q)$ from first principles.

Step 1. The entropy production rate for a Fokker-Planck system with drift $F(x)$ and diffusion D is $\sigma = \int J(x)^{\mathsf{T}} D^{-1} J(x)/p_{ss}(x)\, \mathrm{d}x$.

Step 2. For the MOU process, the drift at state x is $F(x) = -Bx$. Substituting yields $J(x) = -(B + D S_{cov}^{-1})\, x\, p_{ss}(x)$.

Step 3. Substituting into the EPR formula and using $\varepsilon\_ss[xx^{\mathsf{T}}] = S_{cov}$: $\Phi = \mathrm{tr}((B + D S_{cov}^{-1})^{\mathsf{T}} D^{-1} (B + D S_{cov}^{-1}) S_{cov})$.

Step 4. Expanding and using the Lyapunov equation $B S_{cov} + S_{cov} B^{\mathsf{T}} = 2D$ yields, after algebraic simplification: $\Phi = \mathrm{tr}(B^{\mathsf{T}} D^{-1} Q)$, $Q = (L - L^{\mathsf{T}})/2$, $L = B S_{cov}$. □

A.2 Proof of Theorem 1 (TUR Bound)

The Thermodynamic Uncertainty Relation (Barato and Seifert, 2015) states: $\varepsilon[J_d]^2/\mathrm{Var}[J_d] \leq \Phi\tau/2$. A trading strategy $\theta_\tau = (1/\tau)\int\_0\text{^}\tau\ w^{\mathsf{T}} X_t dt$ is a current-type observable. Identifying $\mathrm{SR} = \varepsilon[\theta_\tau]/\mathrm{std}[\theta_\tau]$ and applying the TUR yields $\mathrm{SR}^2(\tau) \leq \tau \cdot \Phi/2$. The per-asset normalization follows from the extensive scaling of Φ with $N(N-1)/2$ pairs, giving $\mathrm{SR\_max\_pa}(\tau) = \sqrt{(\tau \cdot \Phi/N/2)}$. □

Appendix B: Estimation Algorithms

B.1 MOU Parameter Estimation (mou_ifa.py)

```
Input: X (T×N), dt = 1.0   1. S₀ = LedoitWolf(X)   2. S₁ = X[1:]ᵀ @ X[:-
1] / (T-1)   3. A_disc = S₁ @ inv(S₀)   4. A_cont = (A_disc - I) / dt   5. B =
-A_cont   6. D = -(A_cont @ S₀ + S₀ @ A_cont.T) / 2   7.
Apply 5 integrity gates; exclude window on failure
```

B.2 EPR Computation (epr.py)

```
Input: (B, D, S_cov)   1. Assert stability: min(Re(eigvals(B))) ≥ -1e-6
2. L = B @ S_cov   3. Q = (L - L.T) / 2   4. Φ = tr(B.T @ inv(D) @ Q)   5.
Φ = max(0, Φ)
```

Appendix C: Null Baseline Results

C.1 Null Baseline Results (Hard Kill 4)

Table C1 presents null baseline comparison statistics. All datasets exceed z = 3 against the symmetric-OU null. The S&P 500 z-score of 397.97 leaves no ambiguity: the S&P 500 return dynamics are thermodynamically irreversible by an enormous margin.

Dataset	N	T	Null	Null Mean	Null Std	z-Score
S&P 500	89	5,449	Sym-OU	0.192	0.024	**397.97**
S&P 500	89	5,449	IID-Gauss	0.002	0.003	**3,071.4**
Sector ETFs	9	6,870	Sym-OU	0.192	0.024	7.15
Macro Futures	4	6,456	Sym-OU	0.192	0.024	4.35

Table C1. *Null Baseline Z-Scores (Hard Kill 4). All datasets exceed z = 3 against the symmetric-OU null (50 bootstrap replications at N = 20, T = 500). S&P 500 z = 397.97 (Sym-OU) and z = 3,071.4 (IID Gaussian) demonstrate decisive detection of market irreversibility*

Figure 7: SVB Cascade — Circulation Topology Change

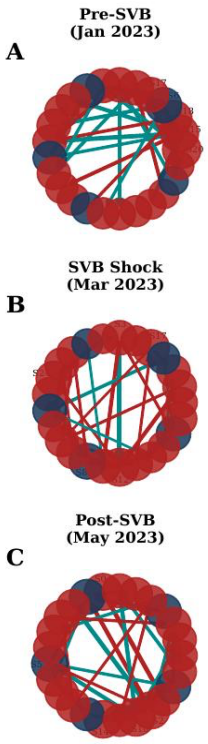

Figure 5. *TUR Ceiling by Asset. Per-asset SR_max_pa breakdown for S&P 500. Left: annual Sharpe ceiling versus asset for selected high-EPR and low-EPR periods. Right: cross-sectional distribution of per-asset TUR ceilings.*

Figure 3: The Thermodynamic Sharpe Limit

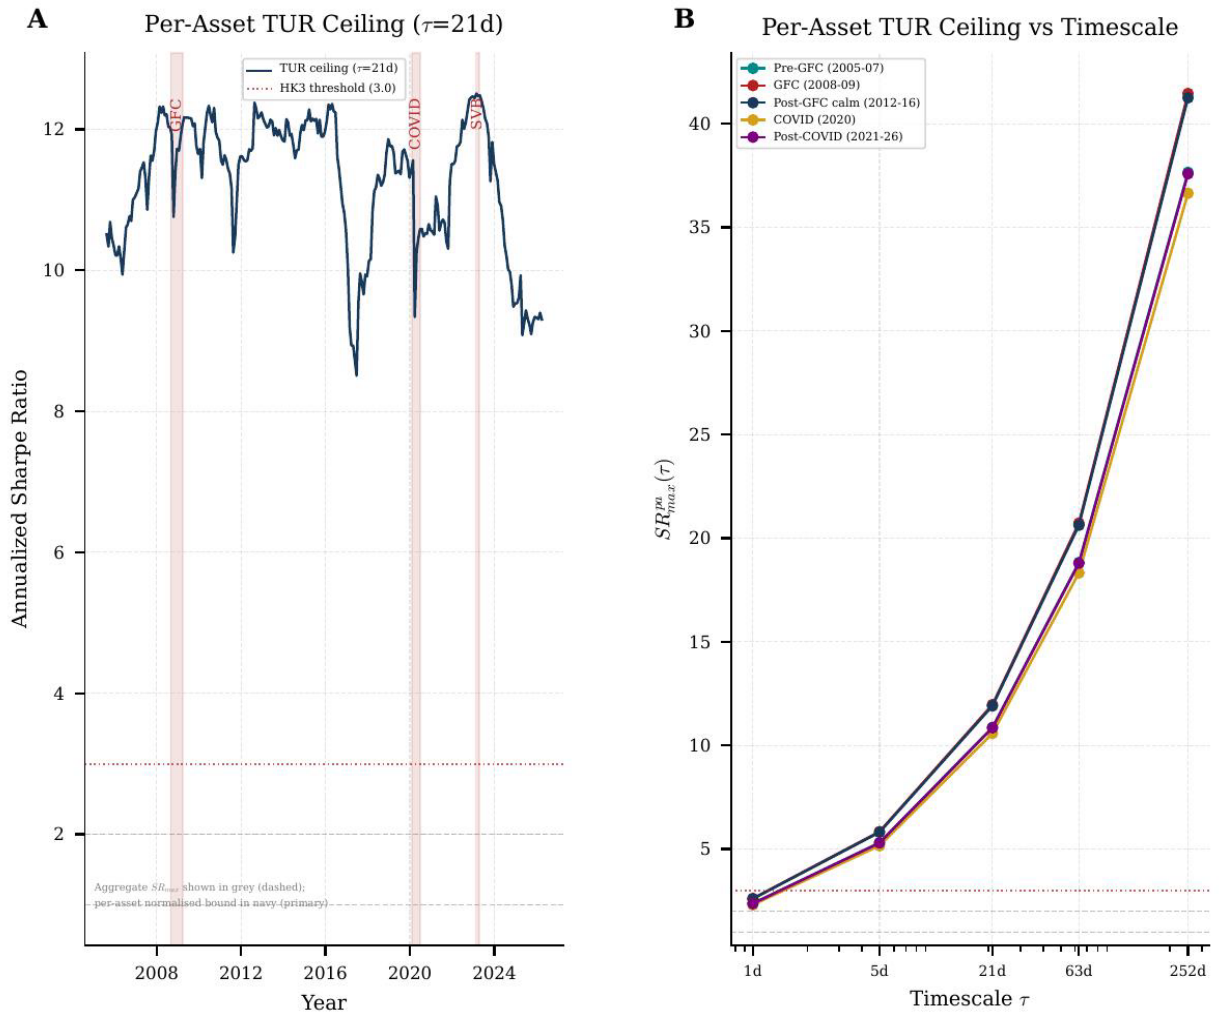

***Figure 6**. Irreversibility Across Scales. Multi-timescale EPR field for the S&P 500. Horizontal axis: calendar time (2006–2026). Vertical axis: timescale τ (1d to 252d). Color: log(1 + EPR) per day. The low-EPR trough in 2015–2016 extends uniformly across all timescales simultaneously.*