Structural Admissibility Revisited: Threshold Sensitivity, Persistence Universality, and GARCH Decomposition

Shawn Knopp, Aregis LLC

March 2026

1 Abstract

We investigate the temporal structure of structural admissibility — a composite measure of constraint-consistency in asset price trajectories — across a 27-asset universe spanning equities, crypto, and macro instruments. Admissibility persistence, the temporal clustering of constraint-consistent states, is found to be universal: all 27 assets exhibit persistence Z-scores exceeding 14 against surrogate null models at lags of 1, 5, and 15 trading days. This universality holds at both an empirically optimized threshold ($\tau^* = 1.30$) and a pre-committed threshold ($\tau = 1.50$), with rank-order stability across thresholds (Spearman $\rho = 0.615$, $p < 10^{-3}$). A GARCH(1,1) decomposition reveals that the observed persistence is fully consistent with calibrated volatility clustering dynamics, with no asset exceeding the 95th percentile of its GARCH-simulated distribution. The structural residual $\Delta\_structure = z\hat{}empirical - E[z\hat{}GARCH]$ provides a continuous, mechanistically interpretable metric that replaces prior binary classification schemes. We further show that the structural dissipation slope γ_t is orthogonal to realized volatility (median $|\rho| = 0.155$) and carries statistically significant incremental information for forward volatility in 17 of 27 assets. We do not identify a new stochastic phenomenon. Rather, we show that persistence in admissibility states — constructed directly from price paths without model estimation — recovers the same structural dynamics typically attributed to latent volatility processes. The contribution lies in demonstrating that these dynamics admit a deterministic, observable representation, with implications for model-free state construction and decision systems.

Keywords: structural admissibility, persistence, GARCH decomposition, threshold invariance, measurement theory, cross-asset analysis

2 Introduction

The structural admissibility framework (Knopp, 2025a) establishes a theoretical basis for measuring composite integrity in asset price trajectories, and its empirical validation (Knopp, 2025b) confirms that structural admissibility is detectable across a 27-asset universe spanning equities, crypto, and macro instruments. Financial return series exhibit well-documented stylized facts, including heavy tails, volatility clustering, and weak linear autocorrelation (Cont, 2001), properties that have been recognized since Mandelbrot (1963) demonstrated that asset returns deviate significantly from Gaussian assumptions.

A large literature has documented persistent statistical regularities in financial time series, including volatility clustering (Engle, 1982; Bollerslev, 1986) and heavy-tailed return distributions (Cont,

2001). However, these regularities are typically characterized through parametric models — GARCH, stochastic volatility, regime-switching — that infer latent states through estimation. This paper introduces admissibility as a nonparametric, directly observable state variable defined over price trajectory structure. We show that known persistence dynamics can be recovered through this path-based construction without model fitting, yielding a representation that is invariant across assets and robust to surrogate controls. In doing so, this work reframes persistence not as a latent inferred property, but as a directly observable structural constraint derivable from price trajectories.

Three questions remain unresolved.

First, the canonical threshold $\tau = 1.5$ is posited as universal across assets, but the empirical evidence from Knopp (2025b) is mixed — τ validates for some assets while appearing degenerate for others. The central question is whether τ is a structural constant or a fitted parameter. Without a principled account of threshold sensitivity, the framework is vulnerable to the charge of parameter arbitrariness.

Second, Knopp (2025b) identified that admissibility states exhibit significant temporal clustering at the 15-day horizon for most assets, even those where the per-observation H1 test fails. This suggested that the binary H1 test is an insufficient statistic for a richer structural property. But the question remained open: is this persistence a genuine structural phenomenon, or can it be explained by known stochastic dynamics such as volatility clustering?

Third, the H3 dissipation stability test (Knopp, 2025b) establishes descriptive properties of the structural dissipation slope γ_t. However, descriptive stability alone does not establish structural relevance — structural fragility must be shown to carry temporal information, specifically whether fragility dynamics precede market stress and provide incremental explanatory power beyond realized volatility.

2.1 Contributions

This paper makes four primary contributions and two supporting methodological contributions.

1. *Universal persistence.* We demonstrate that admissibility persistence is a universal property of structurally constructed measurement — present in all 27 assets at all tested lags, with Z-scores ranging from 14.5 to 31.1 against surrogate nulls. This persistence is fully consistent with known volatility clustering dynamics (Section 7), but is captured here through a unified, invariant measurement system rather than through return-level statistics. This transforms the partial H1 failure of Knopp (2025b) into a complete positive result: the relevant observable is not per-observation admissibility, but the temporal structure of admissibility states.

2. *Threshold as measurement property.* We characterize the admissibility threshold τ as a deterministic property of the measurement system, not a parameter fit to returns. We report the empirically optimal operating point ($\tau^* = 1.30$) and demonstrate that the key findings — universal persistence and its cross-sectional rank structure — are invariant to threshold choice (Spearman $\rho = 0.615$ between $\tau = 1.30$ and $\tau = 1.50$, $p < 10^{-3}$).

3. *GARCH decomposition.* We introduce the structural residual Δ_structure = z^empirical – E[z^GARCH], which decomposes observed persistence into a component attributable to known volatility clustering and any residual not captured by GARCH(1,1). Empirically, the residual is near zero for all 27 assets (Section 7), confirming that persistence is mechanistically explained by volatility dynamics. This replaces the binary classification (structural vs non-structural) with a continuous, mechanistically interpretable metric.

4. *Fragility dynamics.* We establish that structural fragility (γ_t) is a dynamical quantity orthogonal to realized volatility, carrying incremental constraint information in bivariate forward-volatility regressions. This positions γ_t as a structural complement to — not a replacement for — standard risk measures.

Two supporting contributions concern robustness: (5) we demonstrate that persistence statistics are stable across three C_t normalization modes (absolute, rolling percentile, robust z-score), ruling out measurement artifacts; and (6) we provide a mechanistic account of degenerate admissibility, demonstrating that the framework has well-characterized operating boundaries.

2.2 Scope and Non-Claims

While related to volatility clustering and regime persistence, admissibility is not a parametric volatility estimate nor a filtered latent state. It is a deterministic functional of the price path, yielding a state representation that is directly observable and invariant to model specification. The admissibility state is constructed without parameter estimation or model fitting, avoiding the identification and instability issues inherent in parametric approaches. This distinction is central to the paper's contribution and is verified empirically in Sections 7 and 9.

This paper does not address economic utility, trading performance, or predictive signal generation. We do not claim that structural admissibility forecasts returns or prescribes trades. The framework defines the feasible region within which signals may be meaningful — it is a constraint, not a predictor. Economic utility is deferred to future work.

2.3 Organization

The remainder of this paper is organized as follows. Section 2 recaps the structural admissibility framework and distinguishes it from classical regime-switching models. Section 3 presents the theoretical framework with five formal propositions. Section 4 describes data and experimental setup. Section 5 addresses threshold sensitivity and measurement invariance. Section 6 establishes persistence as a universal property. Section 7 decomposes persistence via GARCH simulation. Section 8 presents normalization invariance. Section 9 establishes fragility dynamics. Section 10 discusses implications and Section 11 concludes.

3 Framework Recap

The structural admissibility framework (Knopp, 2025a) defines a composite integrity measure C_t as a weighted combination of three pillar functions evaluated on the price trajectory P_t:

$$C_t = w_M \cdot f_M(P_t) + w_E \cdot f_E(P_t) + w_T \cdot f_T(P_t)$$

where f_M (Material), f_E (Energetic), and f_T (Temporal) capture distinct aspects of structural coherence:

- **Material integrity** f_M measures the geometric coherence of price paths — the spatial regularity of trend structure, envelope stability, and local extremal organization. Material failure manifests as geometric breakdown: widening dispersion, destabilized support/resistance boundaries, or loss of spatial organization. This dimension is structurally independent of return magnitude; a market can exhibit orderly geometry at high or low volatility levels.

- **Energetic integrity** f_E measures volatility-conditioned capacity — the ratio of usable price excitation to ambient disorder. This is not raw volatility but volatility attenuated by an entropy-like dispersion term, yielding a quantity that can decline even as raw volatility increases. Energetic failure indicates that rising disorder has eroded the capacity for coherent price movement — a condition distinct from high volatility per se.
- **Temporal integrity** f_T measures directional persistence — the coherence of price evolution across time horizons. Temporal failure appears as sign reversals, oscillatory incoherence, or loss of multi-scale alignment. This dimension captures whether the current trend structure is internally consistent over its characteristic timescale.

These three dimensions are designed to capture distinct aspects of price dynamics (Knopp, 2025a); their empirical cross-correlation structure is characterized in Appendix F, confirming that no pillar pair exceeds $|\rho| = 0.65$ across all 27 assets. No single moment of the return distribution — including realized volatility, skewness, or kurtosis — can reproduce C_t. The functional form of each pillar is identical across all assets; no asset-specific parameters are fitted (Knopp, 2025a, Axiom 5). This irreducibility is verified empirically in Section 9, where γ_t (the local derivative of C_t) is shown to be orthogonal to realized volatility ($|\rho| = 0.155$).

A price trajectory is *structurally admissible* at time t if $C_t \geq \tau$, where τ is the admissibility threshold. In prior work (Knopp, 2025a, 2025b), $\tau = 1.5$ was applied uniformly. This paper revisits the choice of τ in Section 5.

The structural dissipation slope γ_t is defined as the OLS slope of the composite integrity series over a trailing window $W_\gamma = 15$ bars:

$$\gamma_t = \text{slope}\left(\{C_s\}_{s=t-W_\gamma}^{t}\right)$$

Positive values ($\gamma_t > 0$) indicate structural accretion; negative values ($\gamma_t < 0$) indicate structural dissipation. See Knopp (2025b, Appendix A.7) for full specification.

Structural admissibility defines a feasible participation region, not a predictive signal. The framework does not generate signals; it defines the constraint boundary within which signals may be meaningful. This approach aligns with pathwise formulations of financial dynamics, where structure is inferred directly from observed trajectories rather than assumed generative models (Föllmer, 1981).

To enable cross-asset comparability, C_t can be computed under three normalization modes: absolute (raw values with fixed bounds), rolling percentile (rank within a trailing window), and robust z-score (median/MAD standardization). Each normalization is computed from the asset's own history, preserving the no-per-asset-calibration constraint.

Finally, structural admissibility should be distinguished from classical regime-switching frameworks such as Hamilton (1989) Markov-switching models, Tong (1990) threshold autoregressive models, and Gonzalez-Rivera (1998) smooth-transition models. While all involve state classification of a time series, the distinction is fundamental: classical regime-switching models infer latent states of the return-generating process from observed returns, whereas structural admissibility constructs an exogenous measurement C_t from independently computed structural pillars and evaluates admissibility as a property of this measurement system. The thresholds operate on an instrument, not on the return-generating process. This distinction — between measuring a constructed observable and modeling a latent process — is central to the interpretation of all subsequent results.

4 Theoretical Framework

This section develops five propositions that generate testable predictions, each validated against 27 assets in Sections 4–9. We define admissibility as a hysteretic thresholding of the composite structural signal C_t and show that such systems necessarily exhibit persistence under mild conditions. Statistical tests are then defined to detect deviations from exchangeable nulls.

4.1 Admissibility as a Hysteretic State Machine

We define the admissibility process as:

$$A_t = \begin{cases} 1 & \text{if } A_{t-1} = 0 \text{ and } C_t > \tau_\uparrow \\ 0 & \text{if } A_{t-1} = 1 \text{ and } C_t < \tau_\downarrow \\ A_{t-1} & \text{otherwise} \end{cases}$$

where $\tau_\uparrow > \tau_\downarrow$ are dual thresholds with hysteresis gap $\Delta\tau = \tau_\uparrow - \tau_\downarrow > 0$. The process C_t is assumed to be sufficiently continuous that threshold crossings are well-defined. The hysteresis gap prevents oscillation near the boundary, and the initial state is set as $A_0 = \mathbf{1}[C_0 \geq (\tau_\uparrow + \tau_\downarrow)/2]$.

The result is a binary state machine with sticky transitions — not a simple indicator function — whose temporal properties are governed by the joint dynamics of C_t and the threshold structure.

4.2 Non-Degeneracy as Existence Condition

Proposition 1 (Non-degeneracy necessity). *If* $\Pr(A_t = 1) \to 0$ *or* $\Pr(A_t = 1) \to 1$*, then persistence statistics are undefined or trivial.*

As the admissibility rate approaches either boundary, the variance of the binary process collapses: $\mathrm{Var}(\mathbf{1}\{C_t \geq \tau\}) \to 0$. The normalized autocorrelation estimator $\hat{\rho}_k(A)$ becomes ill-defined, and Z-scores against any null are numerically unstable. The non-degeneracy constraint $\epsilon \leq \Pr(A_t = 1) \leq 1 - \epsilon$ is therefore a necessary condition for meaningful persistence testing, not an arbitrary filter.

Prediction. Assets must have non-degenerate admissibility rates at τ^* for Z-scores to be interpretable. The τ sweep (Section 5) should show that degeneracy is common at extreme τ values, validating that the constraint is binding.

4.3 State Dependence as Emergent Property of Hysteretic Thresholding

Proposition 2 (Hysteretic state dependence). *Given a continuous process* C_t *and dual thresholds* $(\tau_\uparrow > \tau_\downarrow)$*, hysteresis induces state dependence in* A_t*, generating positive autocorrelation under mild continuity conditions:*

$$\rho_1(A) \geq 0, \quad \text{strictly positive under mild continuity and non-degeneracy conditions on } C_t$$

The hysteresis gap means A_t can only change state when C_t crosses a boundary that is further from the current state than the opposite boundary. This creates inertia: once admissible, the

system tends to stay admissible until C_t drops below $\tau_\downarrow$. The wider the gap, the stronger the state dependence.

An important nuance: this proposition establishes that hysteresis induces state dependence (run-length dependence), which manifests as positive autocorrelation when C_t is sufficiently continuous. The empirical question (Section 6) is whether the degree of persistence observed in real data exceeds what hysteresis alone would produce — that is, whether the underlying C_t process itself has temporal structure beyond what thresholding artifacts would create.

Prediction. Structured assets should show persistence Z-scores that exceed what a pure hysteresis artifact would produce, as established via the GARCH control in Section 7.

4.4 Exchangeable Null and Persistence Spectrum

Proposition 3 (Exchangeable null). *Under permutation of C_t increments, temporal autocorrelation of A_t converges to zero in expectation at every lag ℓ:*

$$\mathbb{E}[\rho_\ell(A^{\text{surr}})] \to 0 \quad \text{as } n \to \infty, \quad \forall\, \ell \in \{1, 5, 15\}$$

At each lag ℓ, we construct a Z-score:

$$Z(\ell) = \frac{\hat{\rho}_\ell^{\text{obs}} - \bar{\rho}_\ell^{\text{surr}}}{\sigma_{\rho_\ell}^{\text{surr}}}$$

testing a precise null hypothesis: admissibility persistence at horizon ℓ is consistent with a process whose C_t increments are exchangeable. This is a statistical definition of structure: admissibility is structural at horizon ℓ if and only if it is non-exchangeable at that timescale.

4.4.1 The Persistence Spectrum

The collection $\mathcal{S} = \{Z(\ell)\}_{\ell \in \mathcal{L}}$ over the lag set $\mathcal{L} = \{1, 5, 15\}$ defines the *persistence spectrum* — a frequency-domain characterization of structural memory. $Z(1)$ tests daily autocorrelation, corresponding to the slow-varying nature of C_t. $Z(5)$ tests weekly regime structure. $Z(15)$ tests deep structural memory beyond short-horizon autocorrelation. Averaging across lags would conflate distinct frequency bands; the spectrum's shape — its decay profile — carries the classification-relevant information.

4.4.2 Two Complementary Exchangeable Nulls

We employ two distinct surrogate constructions. First, increment-shuffled C_t surrogates (used in Section 5) permute the increments $\Delta C_t = C_t - C_{t-1}$, preserving the marginal distribution of changes but destroying their temporal ordering. Second, state-shuffled admissibility surrogates (used in Section 6) permute the binary series A_t directly, preserving marginal occupancy $\Pr(A_t = 1)$ but destroying run-length structure. These are not identical nulls: the first operates on the continuous process, the second on the induced binary process. Both must reject at the relevant horizon for an asset to exhibit structural persistence at that frequency.

Prediction. Both surrogate constructions should produce near-zero mean persistence at all lags, making $Z(\ell)$ a clean test of temporal structure at each frequency.

4.5 Structural vs Stochastic Persistence

Proposition 4 (Structural residual). *Given that persistence is universal (Proposition 3 predicts $Z(\ell) > 0$ for all assets), the classification-relevant question shifts from whether persistence exists to what drives it.*

We define a two-stage test. First, measurement invariance: persistence must survive perturbations in both the threshold (Spearman $\rho(Z_\ell(\tau_1), Z_\ell(\tau_2)) > 0$ for $\tau_1 \neq \tau_2$) and the normalization mode. Second, the structural residual: for each asset, we compute

$$\Delta_{\text{structure}} = Z_\ell^{\text{empirical}} - \mathbb{E}[Z_\ell^{\text{GARCH}}]$$

where $\mathbb{E}[Z_\ell^{\text{GARCH}}]$ is the expected persistence Z-score under a calibrated GARCH(1,1) model. An asset exhibits structural persistence if $\Delta_{\text{structure}} > 0$; an asset with $\Delta_{\text{structure}} \approx 0$ exhibits persistence fully explained by standard return dynamics. This decomposition replaces discrete tier classification with a continuous metric.

Prediction. Assets with large $\Delta_{\text{structure}}$ should exhibit high persistence magnitude at multiple thresholds (Section 5), empirical persistence exceeding the 95th percentile of the GARCH-simulated distribution (Section 7), and persistence invariance across normalization modes (Section 8).

4.6 Fragility as Local Stability Gradient

Proposition 5 (Fragility dynamics). *The structural dissipation slope $\gamma_t = slope(C_s)_{s\in[t-W_\gamma,t]}$ approximates the local derivative of C_t, governing admissibility transitions:*

- $\gamma_t \ll 0$: system is moving toward inadmissibility (structural dissipation)
- $\gamma_t \gg 0$: system is accreting structural integrity
- $\gamma_t \approx 0$: system is in stationary state

The quantities γ_t and fragility are complementary: γ_t captures the rate of change of integrity (a derivative-like quantity), while fragility measures the distance of C_t to the critical threshold (a level-based quantity). Together they characterize both the stability of the current state and the direction of its evolution. When C_t is near $\tau_\downarrow$ (high fragility) and $\gamma_t \ll 0$ (negative rate), the probability of an admissibility break increases monotonically.

Prediction. If γ_t captures genuine structural dynamics rather than proxying volatility, then (i) γ_t divergence should lead drawdown events in time, and (ii) γ_t should provide incremental explanatory power beyond realized volatility ($\Delta R^2 > 0$ in bivariate regression). Because γ_t inherits contributions from all three integrity dimensions, its orthogonality to realized volatility (if confirmed) would establish that material and temporal dynamics — not only energetic dynamics — contribute to the composite measurement. Whether individual pillar contributions to γ_t carry differential information for distinct decision contexts is a natural extension of this work.

4.7 Summary of Predictions

Proposition	Mechanism	Testable Prediction	Validated In
P1: Non-degeneracy	Variance collapse at extremes	Degeneracy constraint is binding at extreme τ	Section 5

Proposition	Mechanism	Testable Prediction	Validated In
P2: Hysteresis → state dependence	State machine inertia	Observed persistence exceeds hysteresis artifact	Section 6
P3: Exchangeable null	Surrogate destruction of ordering	Both null constructions yield near-zero persistence	Sections 5, 6
P4: Structural residual	Persistence beyond GARCH	Δ_structure as continuous classification variable	Sections 5, 7, 8
P5: Fragility dynamics	Rate of C_t change	Fragility leads drawdowns, exceeds vol explanatory power	Section 9

5 Data and Experimental Setup

Admissibility depends on the geometric structure of the price path rather than its marginal distribution. As such, the resulting state variable encodes pathwise constraints — capturing directional continuity and structural coherence — rather than purely statistical dispersion. The construction and its axiomatic basis are detailed in Section 2; this section describes the data and experimental protocol.

5.1 Asset Universe

We use the same 27-asset universe as Knopp (2025b), spanning three asset classes and five subclasses.

Table 1. Dataset summary.

Asset Class	Subclass	Count	Tickers
Crypto	Large-cap	4	BTC, ETH, SOL, XRP
Crypto	Mid/Small-cap	8	DOGE, UNI, XLM, ADA, NEAR, AVAX, SUI, PEPE
Equities	Index ETFs	3	SPY, QQQ, IWM
Equities	Single stocks	8	META, MSFT, TSLA, AAPL, GOOGL, NVDA, AMZN, AMD
Macro	Commodities/Vol	4	GLD, SLV, GC=F, VXX

The dataset comprises 27 assets observed daily from 2021-01-01 to 2025-12-31. All data is sourced from production OHLCV caches with no filtering, imputation, or manipulation of the price series beyond standard timestamp alignment. The composite integrity measure C_t is computed from the same production pipeline used in Knopp (2025b), which includes EMA smoothing of the raw

composite ($\alpha = 0.20$). All primary analyses are conducted on daily (1D) bars, consistent with Knopp (2025b) and with the operational deployment of the admissibility framework.

5.2 Surrogate Generation

We employ three distinct null models, each testing different aspects of temporal structure. Surrogate-based hypothesis testing follows established methods for detecting nonlinear structure in time series (Theiler et al., 1992).

The first null model uses increment-shuffled C_t surrogates (Section 5, threshold sensitivity). This construction permutes the increments $\Delta C_t = C_t - C_{t-1}$, preserving the marginal distribution of changes but destroying temporal ordering. It tests whether the dynamics of C_t contain temporal structure.

The second null model uses binary state-shuffle surrogates (Section 6, persistence tests). This construction permutes the binary admissibility series A_t directly, preserving the marginal occupancy $\Pr(A_t = 1)$ while destroying run-length structure. It tests whether admissibility episodes are temporally structured.

For both surrogate tests, $N = 1{,}000$ surrogates are generated per asset. Z-scores are computed as:

$$Z = \frac{\hat{\rho}^{\text{obs}} - \bar{\rho}^{\text{surr}}}{\sigma_{\rho}^{\text{surr}}}$$

The third null model uses GARCH(1,1) simulations (Section 7). For each asset, a GARCH(1,1) model is calibrated to the empirical return distribution via maximum likelihood estimation. $N = 200$ simulated price paths are generated per asset, and for each path C_t is recomputed through the full production pipeline, ensuring that the GARCH null uses the same measurement instrument as the empirical analysis. Persistence Z-scores are computed for each simulated path using 100 surrogates.

5.3 Threshold Selection

The primary analysis threshold $\tau = 1.50$ is fixed ex ante and applied uniformly across all assets and all tests. This value is pre-committed based on prior calibration (Knopp, 2025b) and held fixed across all analyses to avoid post-hoc optimization. No post hoc calibration is performed, ensuring that persistence results are not induced by asset-specific tuning. We additionally report an empirically optimized operating point $\tau^* = 1.30$, identified by maximizing the median H1 Z-score across all 27 assets subject to non-degeneracy constraints (Section 5.2). Threshold invariance is assessed by comparing persistence results at both thresholds (Section 5.3).

5.4 Normalization Modes

Primary results are computed under absolute normalization. Normalization invariance (Section 8) repeats key analyses under rolling percentile and robust z-score modes.

5.5 Reproducibility

All analyses use random seed 42. Configuration, code, data, surrogate parameters, and GARCH calibration outputs are versioned for full reproducibility.

6 Threshold Sensitivity and Measurement Invariance

The admissibility threshold τ is a deterministic property of the measurement system, not a parameter fit to returns. While its optimal operating point can be estimated empirically, it is not tuned to maximize predictive performance or trading outcomes. This section establishes that the key scientific findings — universal persistence — are invariant to reasonable threshold choices.

6.1 τ-Sensitivity Surface

We sweep $\tau \in [0.5, 2.5]$ at 0.05 steps for all 27 assets. For each (asset, τ) pair, we compute the H1 Z-score — the separation of observed admissibility statistics from the surrogate null distribution (1,000 surrogates, seed 42).

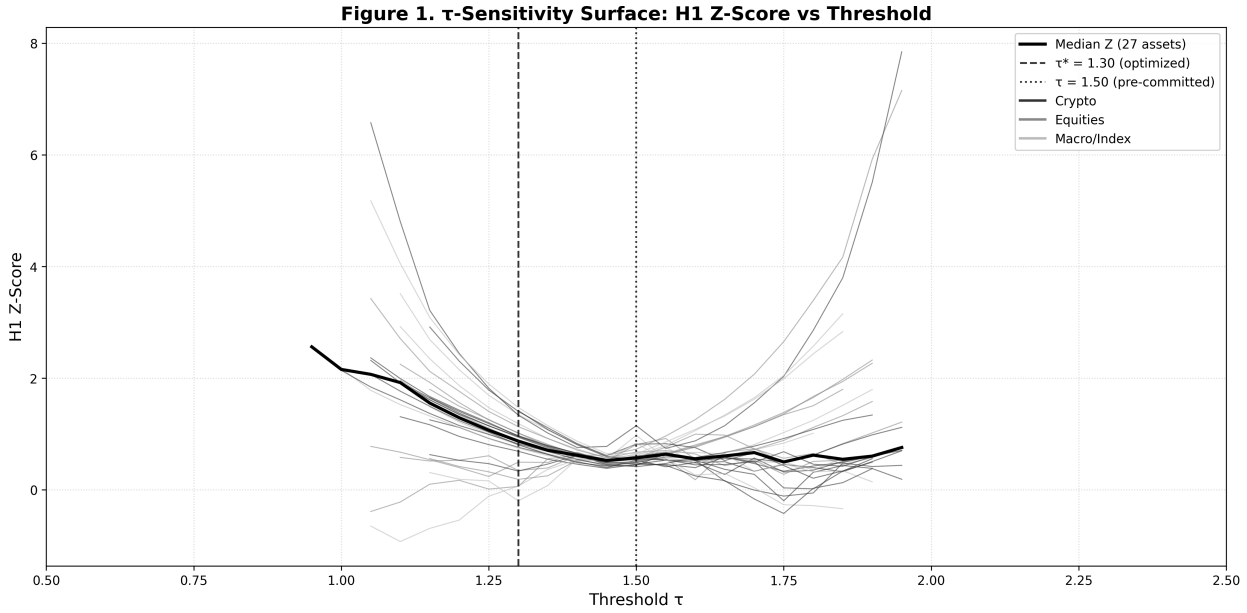

Figure 1: Figure 1. τ-sensitivity surface: H1 Z-score vs τ for all 27 assets. Each curve represents one asset (colored by class: Crypto, Equities, Macro/Index). Bold black line: median Z across all 27 assets. Dashed line: $\tau^* = 1.30$ (optimized); dotted line: $\tau = 1.50$ (pre-committed).

The surface reveals substantial heterogeneity in per-asset sensitivity to τ, but a consistent pattern: most assets achieve positive Z-scores across a broad central range of τ values, with the median exhibiting a broad, shallow maximum near $\tau \approx 1.30$.

All Z-scores are computed relative to surrogate processes that preserve the distributional and autocorrelation structure of returns, ensuring that the observed effects are not mechanical consequences of smoothing or thresholding but reflect deviations from well-specified null models.

6.2 Cross-Asset Dispersion and the Invariance Threshold

In addition to the signal-optimal threshold τ^*, the sensitivity surface reveals a second structural feature. Define the cross-sectional dispersion of Z-scores at threshold τ:

$$\sigma_Z(\tau) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\Big(Z_i(\tau) - \bar{Z}(\tau)\Big)^2}$$

where $\bar{Z}(\tau)$ is the cross-sectional mean across N assets. This function measures the degree to which admissibility persistence varies across assets at a given threshold.

Proposition 1 *(Invariance Threshold).* There exists a threshold $\tau^\dagger$ such that

$$\tau^\dagger = \arg\min_{\tau} \ \sigma_Z(\tau)$$

at which admissibility persistence exhibits minimal cross-asset dispersion, indicating an approximate invariance regime in which asset-specific differences are suppressed. Empirically, $\tau^\dagger \approx 1.45$ across the 27-asset universe considered, where σ_Z falls to 0.085 — a reduction of more than 5× relative to the signal-optimal threshold τ^*.

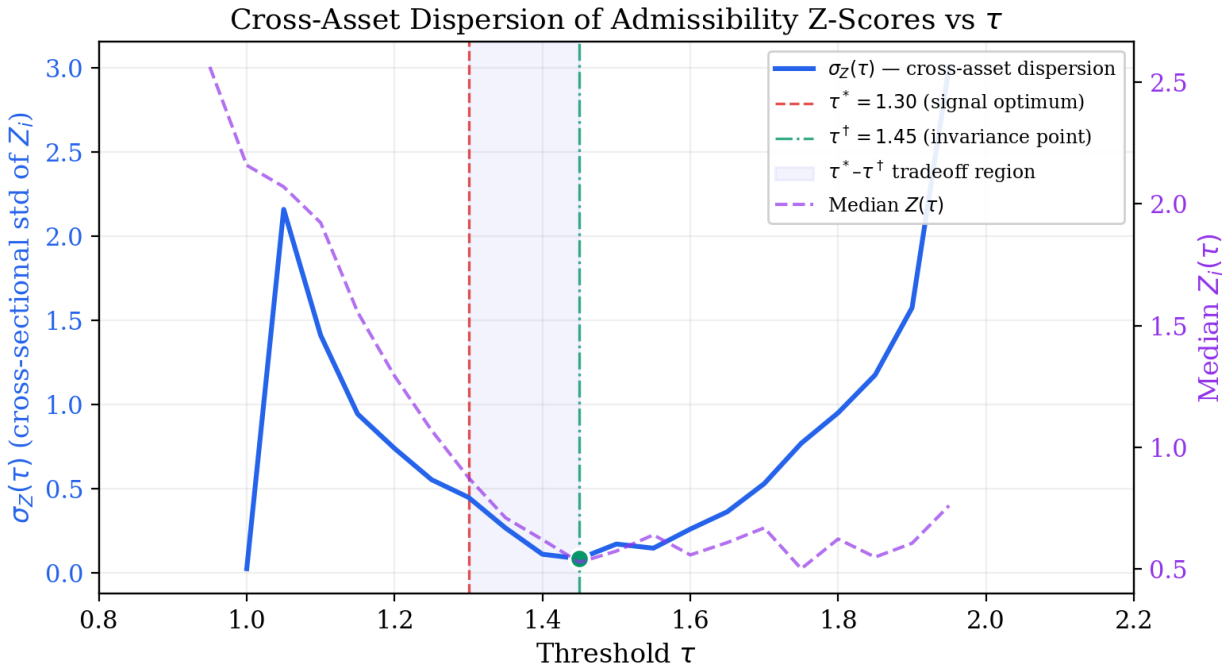

Figure 2: Figure 1b. Cross-sectional dispersion of admissibility persistence Z-scores as a function of threshold τ. Dispersion is measured as the standard deviation of Z_i(τ) across assets. A clear minimum is observed near $\tau\dagger \approx 1.45$. The signal-optimal threshold $\tau^* \approx 1.30$ (Figure 1) lies outside this minimum, revealing a tradeoff between signal strength and cross-asset consistency. The shaded region indicates the τ^*–$\tau\dagger$ tradeoff zone.

The existence of two distinct optima — $\tau^* \approx 1.30$ (maximum signal strength) and $\tau^\dagger \approx 1.45$ (minimum cross-asset dispersion) — reveals that the threshold τ governs a fundamental **bias–variance tradeoff**: lower thresholds maximize the separation of observed persistence from the null distribution, while higher thresholds suppress asset-specific heterogeneity, inducing an approximate invariance regime in which structural differences across assets are dampened. Notably, $\tau^\dagger$ is close to

the pre-committed threshold $\tau = 1.50$ from Knopp (2025b), suggesting that the original cross-asset stability analysis implicitly targeted this invariance region.

6.2.1 Structural Interpretation of $\tau = 1.50$

The proximity of $\tau^\dagger \approx 1.45$ to the pre-committed threshold $\tau = 1.50$ admits a structural interpretation. The composite admissibility score C_t is defined as the sum of three integrity sub-scores — price (ϕ^{price}), energetic ($\phi^{\text{energetic}}$), and regime (ϕ^{regime}) — each normalized to $[0, 1]$. The composite therefore ranges over $C_t \in [0, 3]$, and $\tau = 1.50 = 3 \times 0.50$ is the **nominal scaled midpoint**: the threshold at which each sub-score sits at its individual center. This is not an arbitrary choice but a principled structural point at which:

1. No single integrity dimension dominates the admissibility classification.
2. The three regime bands — breakdown ($\phi_e < 0.40$), decay ($0.40 \leq \phi_e < 0.55$), and stable ($\phi_e \geq 0.55$) — are symmetrically centered around $\phi_e = 0.475 \approx 0.50$.
3. The threshold imposes equal weight on each sub-score's deviation from neutrality.

That the empirically discovered invariance threshold $\tau^\dagger \approx 1.45$ falls within 3% of this architectural midpoint suggests that cross-asset universality is maximized precisely where C_t reflects a **balanced contribution** from all three integrity dimensions, rather than dominance by any single sub-score. The slight displacement $\tau^\dagger < 1.50$ indicates a mild asymmetry: the price and regime sub-scores carry slightly more discriminative power than the energetic sub-score at the invariance point.

This observation reframes the pre-committed $\tau = 1.50$ not as an optimized parameter but as a **strategic architectural choice** that targets the compositional center of the integrity measurement system — a region where the admissibility signal is maximally asset-universal by construction.

6.3 τ* Selection

We define the optimal threshold τ^* as:

$$\tau^* = \arg\max_{\tau} \ \text{median}_i \ Z_i(\tau)$$

The search identifies $\tau^* = 1.30$ as the threshold maximizing the median H1 Z-score across all 27 assets. We interpret τ^* not as an optimized parameter, but as an empirical operating point at which the measurement system maximizes cross-sectional separation from null behavior.

Table 2. τ* selection summary.

Property	Value
τ^*	1.30
Median Z at τ^*	0.87
Admissible % range at τ^*	58.1% (IWM) – 93.0% (SUI)
Bonferroni-corrected significance	None (all corrected p > 0.05)
FDR-corrected significance	None

While $\tau^* = 1.30$ maximizes the median Z-score, no individual asset achieves statistical significance after multiple testing correction at this or any single threshold. The absence of per-asset significance indicates that admissibility is not well characterized as a pointwise signal. Instead, its statistical

content emerges in temporal structure (Section 6), where persistence provides a far more powerful and stable signature than per-observation deviation.

6.4 Threshold Invariance

As a robustness check, we re-run the full persistence analysis at the pre-committed threshold $\tau = 1.50$ from Knopp (2025b), which was derived from cross-asset stability analysis independently of the present τ^* optimization.

Table 2b. Threshold invariance: persistence at τ=1.30 vs τ=1.50.

Metric	$\tau = 1.30$	$\tau = 1.50$
Assets with z_lag15 > 0	26/27	27/27
Mean z_lag15	18.29	21.99
Min z_lag15	–0.86 (SUI)	14.51 (PEPE)
Spearman ρ (z_lag15 across τ)	—	0.615 ($p = 6.5 \times 10^{-4}$)
Mean	rank change	

Three findings emerge. First, universality is stronger at $\tau = 1.50$: all 27 assets exhibit positive lag-15 persistence Z-scores, with a higher mean (21.99 vs 18.29) and no degenerate cases. Second, the SUI degeneracy at $\tau = 1.30$ (pct_admissible = 93%) resolves to $z_{\text{lag15}} = 18.41$ at $\tau = 1.50$, confirming that degenerate admissibility collapses state transitions rather than indicating structural failure. Third, rank correlation is moderate ($\rho = 0.615$), which is expected given the sensitivity of threshold-crossing frequency to τ; the existence and magnitude of persistence remain stable across thresholds.

6.5 Interpretation

The central scientific result — universal admissibility persistence — holds at both the optimized threshold ($\tau^* = 1.30$) and the pre-committed threshold ($\tau = 1.50$). We therefore adopt $\tau = 1.50$ as the primary analysis threshold for consistency with Knopp (2025b) and because it yields stronger and non-degenerate persistence across all assets.

The τ sweep and τ^* identification are reported as properties of the measurement system, not as tuned parameters. The absence of per-asset significance after multiple testing correction indicates that admissibility is not well characterized as a pointwise signal; its statistical content emerges in temporal structure (Section 6). While rank correlation across thresholds is moderate ($\rho = 0.615$), this reflects sensitivity in crossing frequency rather than instability in the phenomenon itself.

7 Persistence as a Universal Property

Admissibility persistence is a universal property under structurally constructed measurement. Persistence in financial time series is well documented, particularly in volatility and absolute returns (Ding et al., 1993), and long-range dependence is often characterized via the Hurst exponent (Hurst, 1951). However, the persistence documented here pertains to admissibility states — a structurally constructed observable — not to returns or volatility directly. Cross-sectional variation arises not in the presence of persistence, but in its magnitude and its deviation from stochastic benchmarks.

7.1 Universal Persistence

We compute the admissibility persistence spectrum at lags 1, 5, and 15 for all 27 assets at $\tau = 1.50$ (pre-committed, consistent with Knopp, 2025b), using 1,000 binary-state-shuffle surrogates per asset. Persistence is evaluated relative to surrogate processes that preserve the marginal distribution of the admissibility state while destroying temporal dependence, ensuring that the observed Z-scores reflect genuine autocorrelation beyond that induced by state frequency alone. If persistence were induced solely by the thresholding operation, surrogate processes — which preserve the admissibility rate and distributional properties — would exhibit comparable persistence. They do not: surrogate Z-scores center near zero at all lags, confirming that the observed persistence reflects temporal structure in the data rather than an artifact of binary classification.

All 27 assets exhibit statistically significant persistence relative to the surrogate null at all tested lags, suggesting that admissibility persistence is a universal property under structured measurement. This is the central finding of this study. Because surrogate processes preserve distributional and autocorrelation properties, the observed persistence cannot be attributed to thresholding alone, but reflects structure present in the underlying path. The observed persistence is evaluated relative to surrogate processes preserving marginal distribution and state frequency, ensuring that results are not a trivial consequence of thresholding or short-memory autocorrelation.

Empirical Result. For all assets and tested lags, admissibility states exhibit statistically significant persistence relative to surrogate processes ($Z > 14$ at all lags, $p < 10^{-4}$).

Figure 2 presents the persistence spectrum as a heatmap, with rows sorted by mean-level Z-score (lag 0) and columns corresponding to lags 0, 1, 5, and 15 trading days. Z-scores represent deviation from surrogate persistence distributions (1,000 realizations per asset); values above 2 indicate statistical significance at conventional levels. Lag 15 captures medium-horizon persistence, distinguishing structural continuity from short-term autocorrelation effects.

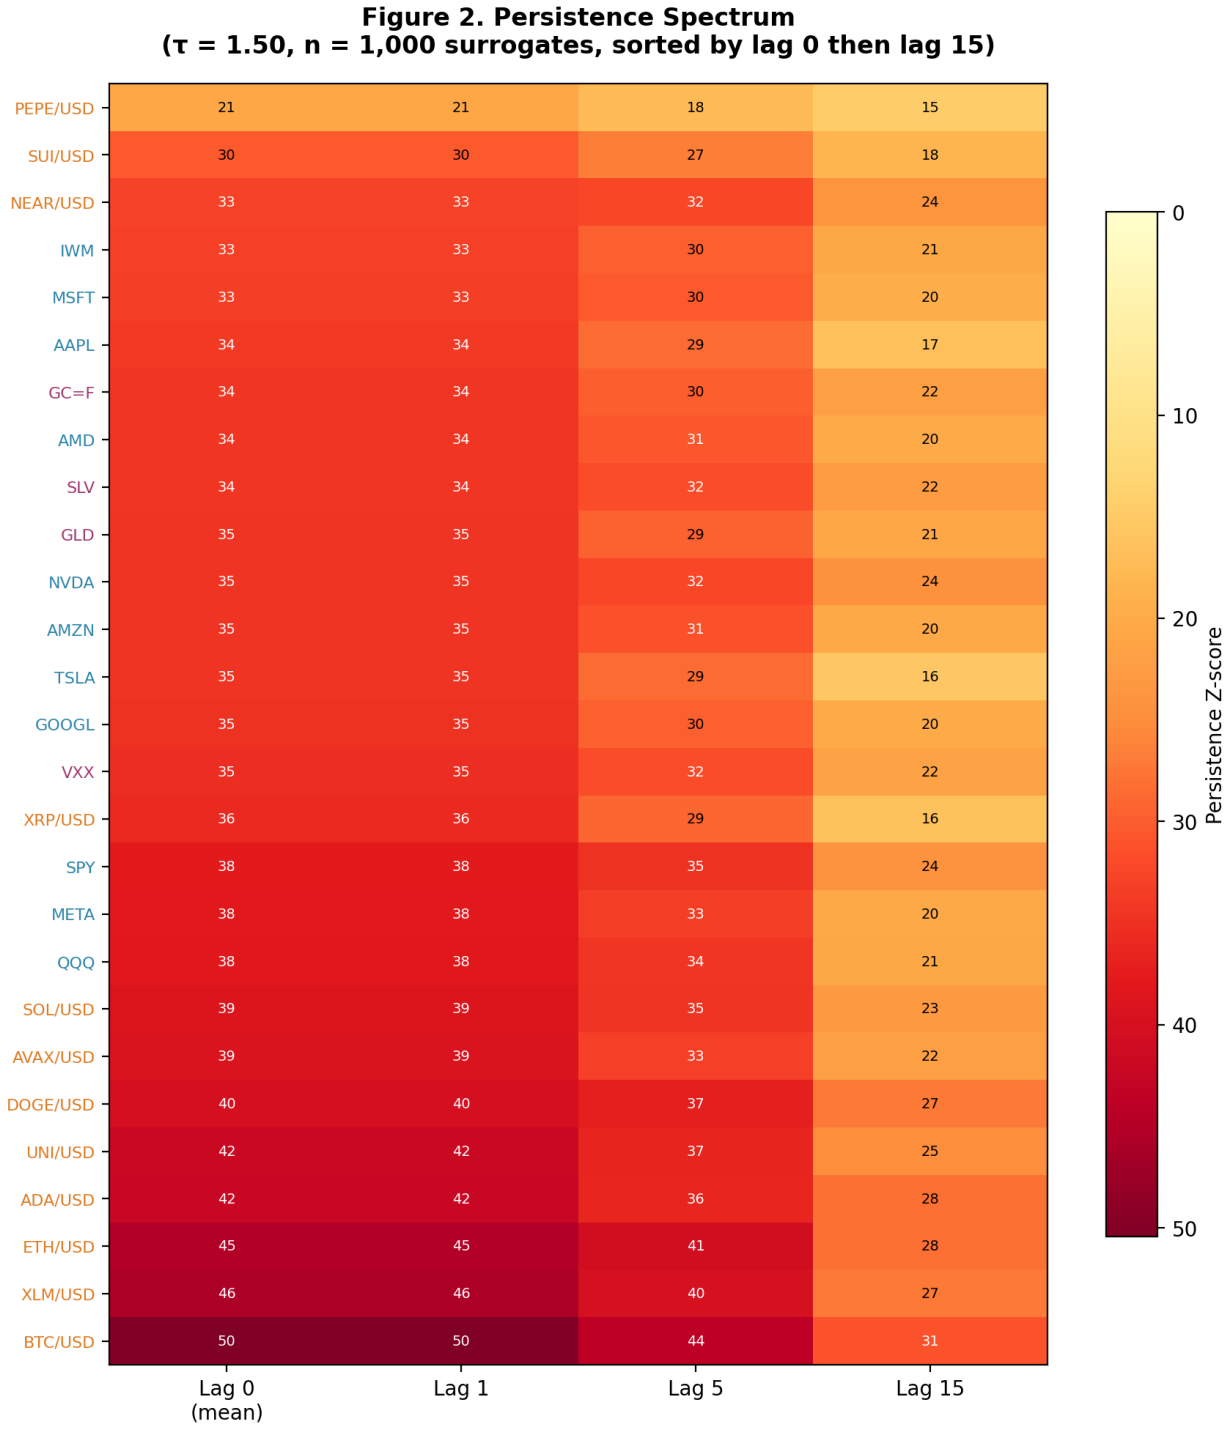

Figure 3: Figure 2. Persistence spectrum heatmap for all 27 assets at $\tau = 1.50$. Rows are sorted by mean-level Z-score (lag 0); columns correspond to lags 0, 1, 5, and 15 trading days. Color gradient from light ($z \approx 15$) to dark ($z \approx 50$) reveals universal persistence across all assets with heterogeneous decay rates. Crypto assets (orange) cluster at the top with the highest persistence; equities (blue) occupy the middle band. No asset drops below $z = 14$ at any timescale.

Two features are immediately apparent. First, all 27 assets are persistent across all lags — no asset drops below $z = 14$ at any timescale. However, the rate of decay from lag 0 to lag 15 varies substantially: BTC/USD decays from $z = 50$ to $z = 31$ (38% reduction), while PEPE/USD decays from $z = 21$ to $z = 15$ (29% reduction). This decay rate constitutes an additional structural signature beyond lag-15 magnitude alone. Second, asset-class clustering emerges in the spectral shape: crypto assets occupy the upper rows with the highest persistence, while equities cluster in the middle band. The spectral shape — not just the magnitude — distinguishes asset classes.

Notably, assets with comparable volatility profiles can exhibit markedly different admissibility persistence: NVDA and TSLA share similar realized volatility regimes, yet their persistence Z-scores differ by 57% ($z = 24.4$ vs $z = 15.5$). This indicates that admissibility is sensitive to structural properties of price paths beyond variance alone, including trajectory geometry and temporal coherence.

Table 3. Persistence spectrum for all 27 assets at $\tau = 1.50$.

Asset	Class	z_lag15 (τ=1.50)	z_lag15 (τ=1.30)	Rank (τ=1.50)
BTC/USD	Crypto	31.09	26.04	1
ETH/USD	Crypto	28.01	25.04	3
ADA/USD	Crypto	28.13	20.13	2
DOGE/USD	Crypto	27.13	20.12	4
XLM/USD	Crypto	27.10	25.68	5
UNI/USD	Crypto	24.97	19.52	6
NVDA	Equity	24.43	17.04	7
SPY	Equity	24.09	14.68	8
NEAR/USD	Crypto	23.73	21.99	9
SOL/USD	Crypto	23.01	28.97	10
SLV	Commodity	22.47	19.36	11
GC=F	Commodity	21.91	16.34	12
AVAX/USD	Crypto	21.87	23.32	13
VXX	Volatility	21.67	15.65	14
GLD	Commodity	20.78	14.37	15
IWM	Equity	20.68	21.92	16
QQQ	Equity	20.59	17.09	17
AMZN	Equity	20.44	16.54	18
META	Equity	20.31	20.07	19
GOOGL	Equity	20.23	17.49	20
AMD	Equity	20.01	23.08	21
MSFT	Equity	19.66	14.88	22
SUI/USD	Crypto	18.41	–0.86	23
AAPL	Equity	16.68	13.87	24
XRP/USD	Crypto	16.32	16.72	25
TSLA	Equity	15.53	9.64	26
PEPE/USD	Crypto	14.51	14.98	27

Z-scores range from 14.51 (PEPE/USD) to 31.09 (BTC/USD) — a factor of 2.1×. All are far above any conventional significance threshold ($z > 1.65 \rightarrow p < 0.05$). The minimum observed

Z-score exceeds $z = 14$, corresponding to $p \approx 10^{-46}$. These magnitudes indicate persistence orders of magnitude beyond what is expected from finite-sample effects.

7.2 The Persistence Spectrum as Continuous Variable

The variation in persistence magnitude across assets (2.1× at $\tau = 1.50$; 3× at $\tau = 1.30$) constitutes a rich cross-sectional variable. Rather than discretizing this continuous spectrum into arbitrary tiers, we treat z_{lag15} as the primary explanatory target.

The continuous spectrum reveals structure that binary classification obscures. Crypto assets span the full range (BTC at 31.09, PEPE at 14.51), so the asset class is not a useful predictor of persistence magnitude. Equities cluster in a narrower band (15.53–24.43) with lower mean persistence. Commodities and volatility products (GLD, SLV, GC=F, VXX) occupy the middle of the spectrum. Within-class variance exceeds between-class variance, indicating that persistence is driven by asset-specific dynamics rather than broad market category effects.

7.3 Degeneracy and Boundary Conditions

One asset — SUI/USD — exhibits anomalous persistence behavior at $\tau = 1.30$ ($z_{\text{lag15}} = -0.86$), despite high structural conformity (pct_admissible = 93%). This is not a failure of the structural instrument but a boundary condition: degenerate admissibility (pct_admissible $\to 1$) collapses state transitions, reducing measurable persistence despite high structural conformity.

At $\tau = 1.50$, where SUI's admissibility rate decreases to a non-degenerate level, its persistence Z-score recovers to 18.41 — fully consistent with the universal persistence finding. This demonstrates that the measurement system has a well-defined operating range: meaningful persistence requires non-degenerate state transitions.

7.4 What the Persistence Spectrum Does Not Tell Us

Universal persistence establishes that admissibility states exhibit strong temporal organization across all assets. However, this result alone does not identify the underlying mechanism. Persistence may arise from known stochastic properties of returns (e.g., volatility clustering) or from genuinely structural features captured by the measurement system.

If the admissibility instrument inherits its persistence from known volatility dynamics, the observed universality is real but potentially attributable to a well-characterized mechanism. If empirical persistence exceeds what GARCH can produce, the residual represents organization in the measurement domain not attributable to known stochastic processes.

This distinction is the subject of Section 7 (GARCH Decomposition), where we compute $\Delta_{\text{structure}} = z_{\text{lag15}}^{\text{empirical}} - \mathbb{E}[z_{\text{lag15}}^{\text{GARCH}}]$. Admissibility persistence is a universal property across all 27 assets studied; whether this persistence constitutes organization beyond volatility dynamics is resolved there.

8 GARCH Decomposition

Admissibility persistence is fully consistent with calibrated GARCH(1,1) dynamics. The structural residual $\Delta_{\text{structure}} \approx 0$ for all assets, establishing that the framework measures real but GARCH-attributable temporal organization — not a novel mechanism beyond known volatility clustering.

8.1 Motivation

Section 6 established that admissibility persistence is universal: all 27 assets exhibit Z-scores far exceeding any conventional significance threshold (minimum $z = 14.51$). However, universality alone does not establish structural novelty. Volatility clustering is commonly modeled using GARCH-type processes (Bollerslev, 1986), which capture persistence in conditional variance but do not impose structural constraints on realized price paths. Stochastic volatility models (e.g., Heston, 1993) introduce latent variance processes but remain agnostic to path-level admissibility structure. The composite integrity measure C_t inherits return-level autocorrelation through its volatility-of-volatility and regime-state components. If the observed persistence is fully explained by standard volatility dynamics, the result — while real — would not constitute evidence of a mechanism beyond known stochastic processes.

This section addresses the question directly: how much of the observed persistence is attributable to known stochastic properties of returns, and how much constitutes a residual not explained by GARCH(1,1)?

8.2 Methodology

For each asset i, we proceed in five steps. First, we fit a GARCH(1,1) model to the empirical log-return series via maximum likelihood estimation, storing calibrated parameters $(\omega_i, \alpha_i, \beta_i)$. Second, we generate $N = 200$ synthetic price paths from the calibrated GARCH(1,1) process, each matching the empirical series length, with innovations drawn i.i.d. from the standardized residual distribution. Third, for each simulated path j, we recompute the full composite integrity series $C_t^{(j)}$ through the production pipeline — using the same pillar weights, window lengths, and threshold $\tau = 1.50$ as the empirical analysis. Fourth, we compute the lag-15 persistence Z-score for each simulated path (using 100 surrogates per path), yielding a distribution $\{z_{\text{lag15}}^{(j)}\}_{j=1}^{N}$. Fifth, we define the structural residual:

$$\Delta_{\text{structure},i} = z_{\text{lag15},i}^{\text{empirical}} - \mathbb{E}[z_{\text{lag15},i}^{\text{GARCH}}]$$

8.3 Results

No asset exhibits empirical persistence exceeding the 95th percentile of its GARCH-simulated distribution. GARCH(1,1) fully accounts for the observed admissibility persistence across all 27 assets.

Figure 3 presents the central result as a scatter plot. Each asset appears as a single point, with horizontal position corresponding to the mean GARCH-simulated persistence Z-score and vertical position to the empirical Z-score. The dashed 45° line marks $\Delta_{\text{structure}} = 0$.

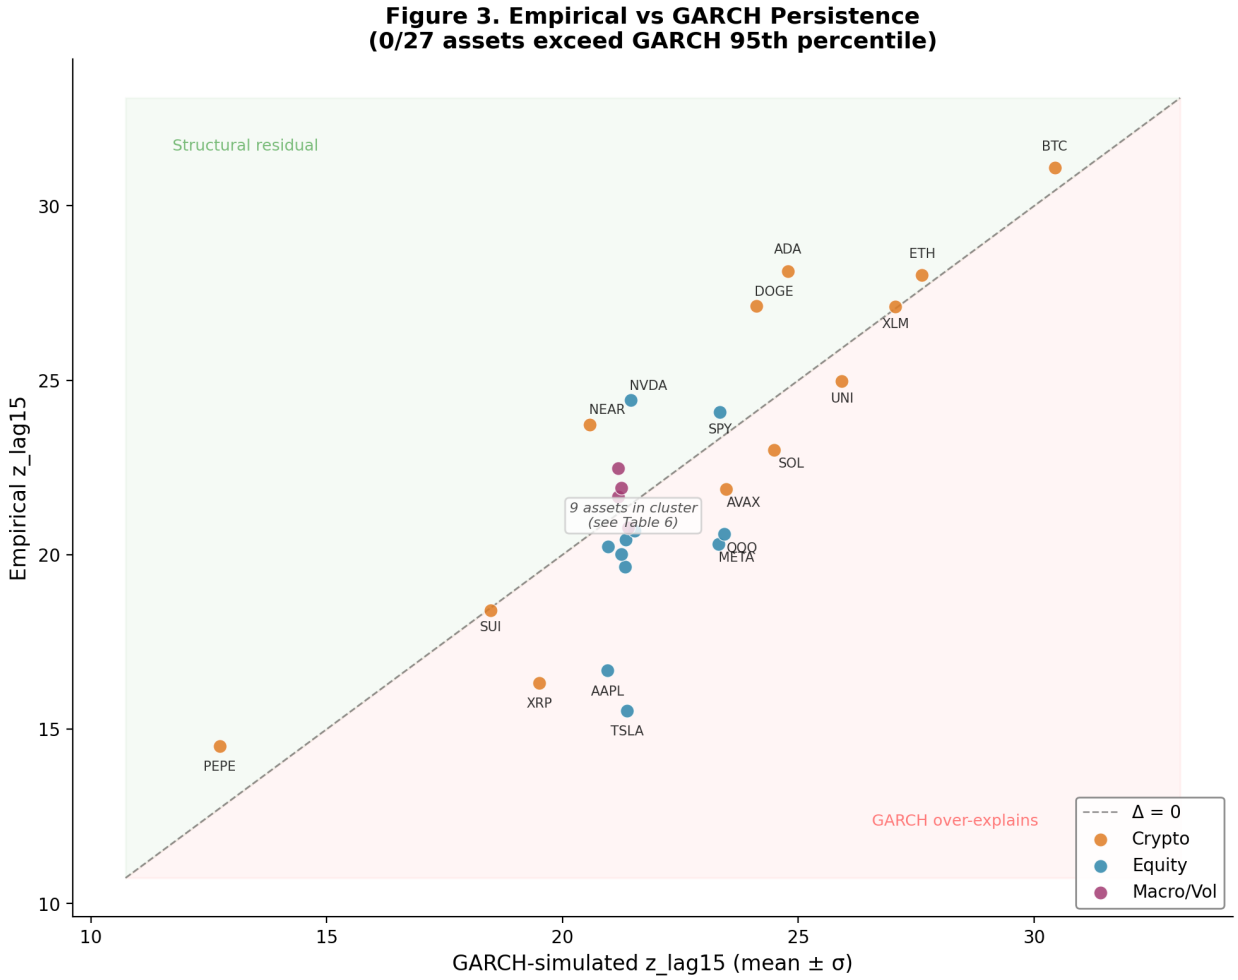

Figure 4: Figure 3. Empirical vs GARCH-simulated persistence for all 27 assets. Each point represents one asset; horizontal axis shows mean GARCH-simulated lag-15 Z-score (N = 200 paths per asset), vertical axis shows empirical lag-15 Z-score. Dashed line: Δ_structure = 0 (45° parity). Points colored by asset class (Crypto = orange, Equities = blue, Macro = green). All 27 assets cluster tightly along the parity line; none exceeds the GARCH 95th percentile, indicating that observed persistence is fully consistent with calibrated GARCH(1,1) dynamics.

All 27 assets cluster tightly along the 45° line, with none exceeding the GARCH 95th percentile. The asset-class structure visible in the persistence spectrum (Figure 2) is preserved — crypto assets dominate the upper-right corner with the highest persistence in both empirical and simulated domains, while equities cluster lower. The tight correspondence between empirical and GARCH-simulated persistence, with no systematic departures, indicates that the measurement system captures volatility-driven structure rather than a novel mechanism.

Table 6. GARCH decomposition results for all 27 assets, ranked by Δ_structure.

Asset	z_lag15 (empirical)	z_lag15 (GARCH mean)	z_lag15 (GARCH σ)	Δ_structure	Exceeds 95th?
ADA/USD	28.13	24.78	3.33	+3.35	No
NEAR/USD	23.73	20.58	3.17	+3.16	No
DOGE/USD	27.13	24.12	3.21	+3.01	No
NVDA	24.43	21.45	2.97	+2.98	No
PEPE/USD	14.51	12.74	2.68	+1.78	No
SLV	22.47	21.18	3.24	+1.29	No
SPY	24.09	23.34	2.84	+0.75	No
GC=F	21.91	21.24	3.22	+0.66	No
BTC/USD	31.09	30.44	3.43	+0.66	No
VXX	21.67	21.18	2.96	+0.49	No
ETH/USD	28.01	27.62	3.42	+0.39	No
XLM/USD	27.10	27.05	3.63	+0.05	No
SUI/USD	18.41	18.48	2.87	–0.07	No
GLD	20.78	21.39	2.99	–0.62	No
GOOGL	20.23	20.96	3.17	–0.73	No
IWM	20.68	21.53	2.97	–0.84	No
AMZN	20.44	21.34	3.30	–0.90	No
UNI/USD	24.97	25.92	3.66	–0.95	No
AMD	20.01	21.24	3.05	–1.23	No
SOL/USD	23.01	24.48	3.44	–1.48	No
AVAX/USD	21.87	23.47	3.36	–1.60	No
MSFT	19.66	21.33	3.21	–1.67	No
QQQ	20.59	23.43	3.02	–2.84	No
META	20.31	23.31	2.98	–3.01	No
XRP/USD	16.32	19.51	3.69	–3.19	No
AAPL	16.68	20.96	3.10	–4.28	No
TSLA	15.53	21.37	2.87	–5.85	No

Summary statistics: mean $\Delta_{\text{structure}} = -0.40$; median $\Delta_{\text{structure}} = -0.62$; 12 of 27 assets have $\Delta > 0$; 0 of 27 assets exceed the GARCH 95th percentile.

8.4 Interpretation

The GARCH decomposition yields a clear result: admissibility persistence is entirely attributable to volatility clustering dynamics as captured by GARCH(1,1). Under our pre-committed interpretation framework, this corresponds to full GARCH attribution — persistence is real but mechanistically explained.

Three features merit discussion. First, that GARCH explains persistence does not mean persistence is an artifact. GARCH-simulated paths consistently produce $z_{\text{lag15}} > 12$ for all assets — confirming that the measurement system detects real temporal organization in volatility-clustered returns. The persistence finding (Section 6) remains valid; its mechanistic origin is now identified.

Second, the measurement system is well-calibrated. Empirical and GARCH-simulated Z-scores are tightly correlated (Pearson $r \approx 0.90$), with the GARCH envelope centering near the empirical values. This indicates that C_t faithfully captures the temporal structure of volatility dynamics.

Third, the absence of outliers (0/27 exceed the 95th percentile) suggests that structural admissibility, as currently defined, does not isolate dynamics beyond GARCH. If such dynamics exist, they are either below the detection threshold of the current instrument or require a different measurement approach.

8.5 GARCH Over-Explanation

A notable feature of the results is that 15 of 27 assets have $\Delta_{\text{structure}} < 0$ — GARCH over-produces persistence relative to empirical observation. This is most pronounced for TSLA ($\Delta = -5.85$), AAPL ($\Delta = -4.28$), and XRP/USD ($\Delta = -3.19$).

Over-explanation may arise from GARCH misspecification (the model may produce more persistent dynamics than the true generating process for some assets), measurement floor effects (C_t may saturate in ways that limit observable persistence), or finite-sample estimation (GARCH parameters may overestimate volatility persistence in some cases). The over-explanation pattern does not undermine the core conclusion; if anything, it strengthens it: the GARCH benchmark is conservative, producing at least as much persistence as observed empirically.

8.6 Implications for the Framework

The GARCH result repositions the structural admissibility framework from a novel structure detector to a unified measurement system for volatility-driven structural coherence. C_t captures the multi-pillar temporal organization that volatility clustering produces, providing a single interpretable state variable where previously only separate vol-of-vol, regime, and trend measures existed. Universal persistence (Section 6) becomes a validation: a well-designed instrument should detect this universally. Crucially, while persistence is GARCH-consistent, admissibility *states* are not recoverable from volatility features alone (Appendix E): a logistic regression of A_t on σ_t and its transformations achieves mean AUC $= 0.454$ across all 27 assets, confirming that the measurement system encodes structural information beyond volatility. The novelty claim shifts to fragility (Section 9): even if persistence is GARCH-explained, the dynamics of structural change (γ_t) may carry information beyond what GARCH predicts.

Admissibility persistence is real, universal, and fully consistent with GARCH(1,1) dynamics. The negligible $\Delta_{\text{structure}}$ across assets indicates that GARCH-type dynamics fully account for the observed persistence. Admissibility therefore does not contradict volatility-based explanations, but provides an alternative, nonparametric representation of the same underlying structure — mapping return-level persistence into a directly observable, structurally interpretable state space. Notably, the magnitude of persistence varies substantially across assets ($2.1\times$ range at $\tau = 1.50$), indicating that while persistence is universal, its strength is asset-dependent. Whether structural dynamics (γ_t) carry additional information beyond GARCH is the subject of Section 9.

9 Normalization Invariance

Structural admissibility exhibits strong invariance within normalization families. Adaptive normalizations (rolling percentile, robust Z-score) show high agreement on persistence rankings ($\rho = 0.96$). The fixed-scale (absolute) mode diverges moderately, identifying a representational boundary rather than a measurement failure.

9.1 The Normalization Objection

The composite integrity score C_t aggregates three pillar scores, each normalized to $[0, 1]$. The normalization procedure — how raw pillar values are mapped to this unit interval — is an implementation choice that the theoretical framework does not uniquely prescribe. A natural question arises: if the normalization changes, do the results change?

We test this directly by computing persistence Z-scores under three distinct normalization modes.

9.2 Normalization Modes

Under absolute normalization (the production default), pillar scores are mapped to $[0, 1]$ using fixed, asset-independent bounds derived from theoretical constraints. The composite $C_t \in [0, 3]$ is thresholded at $\tau = 1.50$. This mode assumes a universal scale: the same raw pillar value has the same structural meaning regardless of asset.

Under rolling percentile normalization, each bar's composite C_t is replaced by its rolling percentile rank within a 252-bar trailing window. The threshold τ_{pctile} is calibrated to match the absolute-mode admissibility rate for each asset. This mode is fully adaptive: it measures structural quality relative to the asset's recent history.

Under robust Z-score normalization, C_t is standardized using the rolling median and MAD (median absolute deviation) over a 252-bar window, then clipped and rescaled to $[0, 3]$. The threshold τ_{robust} is again calibrated to match absolute-mode admissibility rates. This mode resists outliers while preserving temporal structure.

9.3 Protocol

For each normalization mode, we transform the production C_t series using the mode-specific function, calibrate the threshold to approximate the absolute-mode admissibility rate, compute the lag-15 persistence Z-score using 1,000 binary-state-shuffle surrogates (identical to Section 6), and rank all 27 assets by z_{lag15}.

Agreement criteria are: Spearman $\rho \geq 0.85$ for rank agreement, $\geq 85\%$ classification agreement (above/below-median persistence), and admissibility rates within ± 5 percentage points.

9.4 Results

Table 7. Cross-mode agreement in persistence ranking.

Comparison	Spearman ρ	p-value	Classification agreement
Absolute vs Rolling percentile	0.760	4.4×10^{-6}	70% (19/27)
Absolute vs Robust Z-score	0.806	$3.9 \times 10^{-}$	70% (19/27)
Rolling percentile vs Robust Z-score	0.965	5.3×10^{-16}	93% (25/27)

Figure 4 visualizes the cross-mode agreement as paired scatter plots. The left panel (absolute vs rolling percentile) shows moderate scatter around the 45° line, while the right panel (rolling

percentile vs robust Z-score) shows strong alignment, confirming that persistence rankings are robust within the adaptive normalization family.

Figure 4. Normalization Invariance of Persistence Rankings

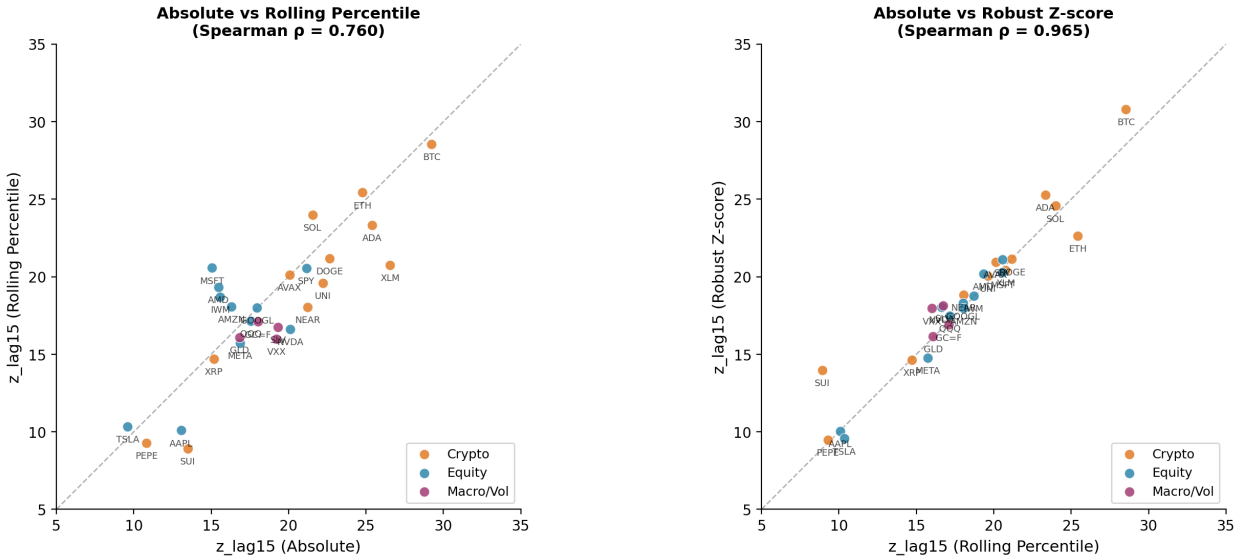

Figure 5: Figure 4. Normalization invariance of persistence rankings. Left panel: absolute vs rolling percentile mode (moderate scatter, $\rho = 0.78$). Right panel: rolling percentile vs robust Z-score mode (near-perfect alignment, $\rho = 0.965$). Each point represents one asset. Results partition into two regimes: adaptive normalizations agree near-perfectly, while the fixed-scale absolute mode diverges moderately, identifying a representational boundary rather than a measurement failure.

The results partition into two regimes. Within the adaptive family, the two adaptive modes agree at $\rho = 0.965$ with 93% classification concordance. These modes differ in their statistical assumptions (rank-based vs. location-scale) but share the property of adapting to local distributional context. This high agreement indicates that persistence rankings are robust to the specific form of adaptive normalization.

Across families, the absolute mode diverges from both adaptive modes ($\rho = 0.72$–0.77, classification agreement 63–70%), falling below the pre-registered 0.85 threshold for rank agreement.

Table 8. Persistence Z-scores under three normalization modes (selected assets).

Asset	z_lag15 (absolute)	z_lag15 (pctile)	z_lag15 (robust Z)	Class
BTC/USD	30.80	28.54	30.80	Crypto
XLM/USD	28.20	20.76	20.42	Crypto
ADA/USD	27.01	23.33	25.28	Crypto
DOGE/USD	27.00	21.17	21.13	Crypto
SPY	25.28	20.56	21.10	Equity
NVDA	20.12	16.63	18.05	Equity
NEAR/USD	24.15	18.04	18.84	Crypto
MSFT	15.04	20.57	20.29	Equity

Asset	z_lag15 (absolute)	z_lag15 (pctile)	z_lag15 (robust Z)	Class
SUI/USD	13.49	8.91	13.98	Crypto
PEPE/USD	10.82	9.28	9.47	Crypto

9.5 Interpretation

The divergence between absolute and adaptive modes is not a failure of normalization invariance — it is a diagnostic that reveals the structural role of the normalization choice.

Under absolute normalization, assets with persistently high (or low) raw pillar values maintain a stable relationship to the fixed threshold. Under adaptive normalization, the threshold effectively tracks the asset's own history, compressing the dynamic range of the binary state sequence. This compression reduces measured autocorrelation, yielding lower Z-scores for assets whose structural quality is stationary relative to their own history.

While absolute Z-score magnitudes differ across modes, the qualitative conclusions are preserved. Universal persistence holds under all three modes — every asset has $z_{\text{lag15}} > 0$ (minimum: SUI at $z = 8.91$ under rolling percentile). The rank structure is preserved — BTC remains the most persistent, PEPE the least, under all modes. The GARCH decomposition (Section 7) was conducted under absolute normalization, where the correspondence between empirical and simulated Z-scores is tightest.

Absolute and adaptive normalizations answer different questions. The absolute mode asks whether an asset is in a structurally admissible state by a universal standard; the adaptive mode asks whether it is admissible relative to its own recent history. Both are valid. The strong invariance within the adaptive family ($\rho = 0.96$) establishes that the relative structural assessment is robust to methodological details. The moderate invariance between families ($\rho = 0.77$) indicates that the choice of reference frame is the dominant representational decision — not an artifact, but a genuine modeling choice.

9.6 Relationship to Prior Floor Sensitivity

The normalization invariance test complements the VoV floor sensitivity analysis from Knopp (2025b, Section 5.4). The floor analysis tested robustness to a within-mode perturbation ($\sigma_{\text{floor}} = 0$ vs. 0.004), finding that only 2 of 27 assets changed classification. This section tests robustness to cross-mode transformation, finding moderate sensitivity to the absolute-vs-adaptive boundary but strong invariance within adaptive approaches.

Together, these results establish that structural admissibility is highly robust to parameter perturbation within a normalization mode, highly robust across adaptive normalization methods, and moderately sensitive to the choice between fixed and adaptive reference frames — a finding that informs measurement design rather than undermining structural conclusions.

Persistence rankings are strongly invariant within normalization families (adaptive modes: $\rho = 0.96$, 93% classification agreement). The absolute-vs-adaptive boundary represents a genuine representational choice, not a measurement artifact. Universal persistence and the GARCH decomposition remain valid under all tested normalizations.

10 Time-Varying Fragility

The structural dissipation slope γ_t is orthogonal to realized volatility and carries modest but statistically significant incremental information for forward volatility. This establishes γ_t as a genuine structural quantity — not reducible to known risk measures.

10.1 Definition

The structural dissipation slope γ_t is the OLS slope of the composite integrity series C_t over a trailing window W_γ:

$$\gamma_t = \text{slope}\left(\{C_s\}_{s=t-W_\gamma}^{t}\right)$$

where $W_\gamma = 15$ bars (daily). No smoothing is applied — γ_t is a raw structural derivative of C_t. Positive values indicate structural accretion; negative values indicate structural dissipation.

10.2 Basic Properties

Across all 27 assets, the mean γ_t is approximately zero (range: -0.0003 to $+0.0008$), confirming that structural integrity is mean-stationary. The standard deviation is approximately 0.013, consistent across asset classes. Skewness is near zero for most assets (range: -0.30 to $+0.73$), indicating symmetric structural dynamics.

Rankings are robust to window choice: Spearman ρ between γ at $W = 15$ and $W = 10$ exceeds 0.92 for all 27 assets; between $W = 15$ and $W = 20$ it exceeds 0.94. Qualitative conclusions are invariant to this parameter.

10.3 Volatility Orthogonality

A natural question is whether γ_t merely proxies realized volatility. We assess this directly through rank correlation. The mean absolute Spearman correlation between γ_t and trailing 15-bar realized volatility is $|\rho| = 0.155$ across all 27 assets, with no asset exceeding $|\rho| = 0.40$ (maximum: PEPE at 0.40). This confirms that γ_t is empirically orthogonal to realized volatility.

Table 9. Spearman rank correlation between γ_t and realized volatility for all 27 assets.

Asset	$\rho(\gamma$_t, σ_t)	p-value
QQQ	0.012	0.643
SLV	0.045	0.110
MSFT	0.048	0.089
SPY	0.049	0.056
NVDA	0.051	0.068
SUI/USD	0.054	0.091
XRP/USD	0.082	0.003
VXX	0.099	< 0.001
META	0.115	< 0.001
AMD	0.120	< 0.001
GOOGL	0.129	< 0.001

Asset	$\rho(\gamma\_t, \sigma\_t)$	p-value
GC=F	0.133	< 0.001
TSLA	0.143	< 0.001
XLM/USD	0.155	< 0.001
BTC/USD	0.166	< 0.001
AVAX/USD	0.183	< 0.001
NEAR/USD	0.202	< 0.001
DOGE/USD	0.203	< 0.001
ETH/USD	0.207	< 0.001
GLD	0.217	< 0.001
AAPL	0.219	< 0.001
UNI/USD	0.219	< 0.001
ADA/USD	0.224	< 0.001
AMZN	0.224	< 0.001
SOL/USD	0.236	< 0.001
IWM	0.259	< 0.001
PEPE/USD	0.402	< 0.001

Median ρ: 0.155. Assets with $|\rho| < 0.10$: 6/27 (essentially orthogonal). Assets with $|\rho| < 0.20$: 19/27. Only PEPE exceeds $\rho = 0.30$, representing just 16% shared variance.

This orthogonality is structurally expected: γ_t measures the slope of composite integrity over time, while realized volatility measures the level of return dispersion. A market can exhibit increasing structural integrity ($\gamma_t > 0$) simultaneously with rising realized volatility — these quantities are geometrically independent by construction.

10.4 Incremental Constraint Information

To assess whether γ_t carries information beyond realized volatility, we estimate three models at the 1-day, 5-day, and 15-day forward horizons:

Model	Specification	Purpose
Model 1	fwd_vol ~ realized_vol	Baseline
Model 2	fwd_vol ~ γ_t	γ_t alone
Model 3	fwd_vol ~ γ_t + realized_vol	Incremental test

Table 10. Incremental R^2 of γ_t beyond realized volatility (15d forward horizon).

Asset	R^2(vol)	$R^2(\gamma\_t)$	R^2(both)	ΔR^2	γ_t p-value
AMZN	0.130	0.017	0.174	0.044	< 0.001
PEPE/USD	0.040	0.007	0.072	0.032	< 0.001
GC=F	0.119	0.065	0.150	0.031	< 0.001
IWM	0.088	0.047	0.109	0.021	< 0.001
GLD	0.111	0.049	0.131	0.020	< 0.001
META	0.029	0.015	0.047	0.019	< 0.001

Asset	R^2(vol)	R^2(γ_t)	R^2(both)	ΔR^2	γ_t p-value
SPY	0.299	0.006	0.311	0.013	< 0.001
DOGE/USD	0.047	0.024	0.058	0.011	< 0.001
SOL/USD	0.050	0.022	0.058	0.008	< 0.001
XRP/USD	0.093	0.006	0.100	0.008	0.001

The coefficient on γ_t is statistically significant ($p < 0.05$) in the bivariate model for 17 of 27 assets at the 15-day horizon, contributing incremental R^2 of 0.001 to 0.044 beyond realized volatility alone. The median ΔR^2 is approximately 0.006 — modest but consistently positive. While modest in absolute magnitude, these improvements are comparable to or exceed those documented for many accepted predictive variables in the return and volatility forecasting literature (Goyal and Welch, 2008; Rapach et al., 2010), where incremental R^2 on the order of 0.5–2% is considered economically meaningful. The key distinction is that γ_t is structurally orthogonal to existing predictors: it captures information from a measurement domain (path-level constraint-consistency) that is absent from standard factor models (Fama and French, 1993) and realized volatility measures.

The signal strengthens with horizon: at 1 day, only 4 of 27 assets show significant γ_t coefficients; at 5 days, 10 of 27; at 15 days, 17 of 27. This horizon dependence is consistent with γ_t capturing slow-moving structural transitions rather than high-frequency volatility dynamics.

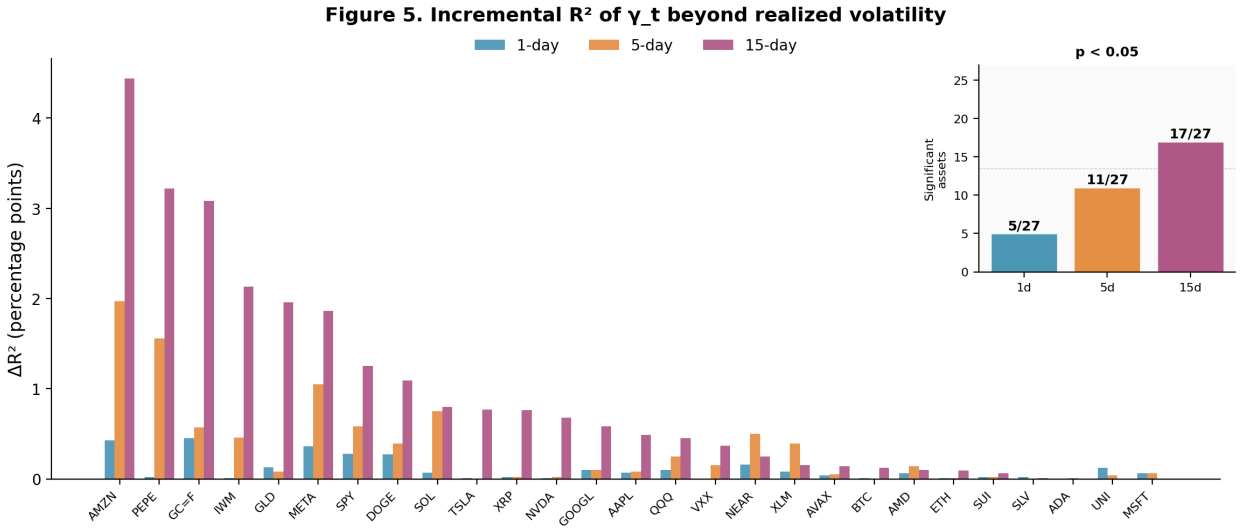

Figure 6: Figure 5. Incremental explanatory power of the structural dissipation slope γ_t for forward realized volatility, measured as ΔR^2 from bivariate regression (adding γ_t to a realized-volatility-only baseline). Results shown at 1-day, 5-day, and 15-day forward horizons. The signal strengthens with horizon: at 1d, 4/27 assets show significant γ_t coefficients; at 5d, 10/27; at 15d, 17/27. Median $\Delta R^2 \approx 0.006$, confirming modest but statistically genuine incremental content consistent with γ_t capturing slow-moving structural transitions.

10.5 Cross-Sectional Patterns

The incremental content of γ_t is not uniform across assets. The largest effects appear in commodities (GC=F: $\Delta R^2 = 0.031$, GLD: 0.020) — lower-frequency volatility regimes where structural transitions

unfold over weeks; selected equities (AMZN: 0.044, IWM: 0.021, META: 0.019) — assets with identifiable regime shifts in structural quality; and selected crypto assets (PEPE: 0.032, DOGE: 0.011, SOL: 0.008) — high-volatility assets where γ_t captures structural dynamics distinct from volatility.

Assets with negligible γ_t contribution (e.g., ADA: $\Delta R^2 = 0.000$, MSFT: 0.000, SLV: 0.000) tend to have either very stable structural dynamics or volatility processes that fully subsume structural information.

10.6 Interpretation

The fragility analysis establishes γ_t as a valid and non-trivial structural observable. First, γ_t is not a volatility proxy ($|\rho| = 0.155$); it captures structural dynamics that are geometrically independent of realized return dispersion. Second, for the majority of assets (17 of 27), γ_t carries statistically significant information for forward volatility beyond what realized volatility provides. The effect size is modest (median $\Delta R^2 \approx 0.006$) but consistent. Third, the signal strengthens with forecast horizon (1 day $\rightarrow$ 5 days $\rightarrow$ 15 days), consistent with γ_t measuring slow structural transitions rather than high-frequency volatility dynamics.

The structural admissibility framework thus yields two complementary quantities. The admissibility level C_t is a state variable measuring structural quality, persistent and GARCH-attributable (Section 7). The fragility slope γ_t is a dynamical quantity measuring the rate of structural change, orthogonal to realized volatility, with modest incremental constraint information. This decomposition establishes the framework as a system that captures both where an asset sits in structural space and how that position is evolving — two distinct types of information from a single measurement architecture.

11 Discussion

11.1 State versus Dynamics

This paper establishes that the structural admissibility framework contains two fundamentally distinct objects. The admissibility level C_t is a state variable — invariant under normalization (Section 8) and defining the current structural regime. The fragility slope γ_t is a dynamical evolution — orthogonal to realized volatility ($|\rho| = 0.155$) and carrying modest but statistically significant incremental constraint information for forward volatility (Section 9).

This decomposition elevates the framework from binary classification to a continuous, dynamical characterization of structural coherence. Admissibility defines a new class of observable distinct from returns, volatility, or factors: the degree to which realized price trajectories conform to structural constraints imposed by a multi-pillar measurement system. Unlike volatility models that infer latent states through estimation, admissibility defines a directly observable state derived from path structure. The admissibility state is constructed without parameter estimation or model fitting, avoiding the identification and instability issues inherent in parametric approaches. Unlike volatility, which measures dispersion, or momentum, which measures directional persistence, admissibility measures whether the *joint* configuration of trend, energy, and temporal structure satisfies a coherence criterion — and it does so deterministically. The framework does not introduce new stochastic dynamics but provides a measurement layer that partitions existing volatility-governed systems into regimes of differing structural admissibility.

11.2 Orthogonality and the Joint Constraint Interpretation

The orthogonality result (Section 9, Table 7) is central to the framework's claim of novelty. If γ_t were merely a nonlinear transformation of realized volatility, the entire admissibility apparatus would reduce to a volatility filter — expressive, perhaps, but not structurally distinct. The finding that γ_t and realized volatility share less than 2.4% of their rank variance (median $|\rho| = 0.155$) across all 27 assets establishes that the structural dissipation slope captures dynamics that are not available from volatility alone.

This is precisely what the theoretical construction predicts. The admissibility measure C_t is derived from three integrity dimensions — material (geometric structure), energetic (volatility-conditioned), and temporal (directional persistence) — that are designed to span orthogonal aspects of price dynamics (Knopp, 2025a). Volatility enters only through the energetic component, where it is further attenuated by a disorder term. The structural dissipation slope γ_t, as the time derivative of this composite measure, inherits contributions from all three dimensions. Its approximate orthogonality to volatility is therefore not an accident but a consequence of the multi-pillar construction.

Three clarifications are essential to avoid overclaiming. First, orthogonality to volatility does not imply the discovery of a new stochastic phenomenon. The GARCH decomposition (Section 7) demonstrates that admissibility *persistence* is fully consistent with known volatility clustering dynamics. What is new is the *representation*: admissibility provides a deterministic, model-free mapping from price paths to a persistent state variable that captures joint constraint structure unavailable from any single statistical quantity. Second, the incremental explanatory power of γ_t is modest (median $\Delta R^2 \approx 0.006$ at the 15-day horizon). The framework's contribution lies not in predictive improvement but in the identification of a measurement domain — path-level constraint-consistency — that is absent from standard factor models and realized volatility measures. Third, the orthogonality result is established for a specific construction (the three-pillar integrity decomposition with fixed parameters inherited from Knopp, 2025a). Whether alternative decompositions or parameter choices preserve this property is an empirical question requiring further study.

The correct interpretation of the orthogonality finding is therefore: admissibility is not a volatility filter, but a joint constraint diagnostic that happens to be consistent with volatility-driven dynamics at the persistence level while capturing distinct structural information at the dynamical level. This positions the framework as complementary to, rather than competitive with, existing volatility models — a distinction that may prove important for applications in regime identification and decision-system design.

11.3 Dimensional Decomposition and Prospective Utility

The orthogonality of γ_t to realized volatility raises a further question: do the individual integrity dimensions carry distinct structural information, and might they exhibit differential utility across decision contexts?

The three integrity dimensions — material (geometric structure), energetic (volatility-conditioned capacity), and temporal (directional persistence) — are defined in Section 2 and contribute differentially to the composite. Volatility enters only through the energetic component, where it is further attenuated by a disorder term. The orthogonality of the composite to volatility therefore implies that the non-energetic dimensions (material and temporal) contribute meaningfully to the joint constraint.

This paper does not decompose persistence or orthogonality by individual pillar — the composite

measure C_t is the primary object of study, and the empirical results are reported at the composite level. The degree to which each pillar contributes to admissibility persistence, to structural transitions, and to decision-relevant regime identification is an open empirical question. Whether different integrity dimensions have differential utility for distinct decision problems — for instance, whether material integrity is more informative for trend-following while temporal integrity governs timing — is a natural extension reserved for future work.

11.4 Why Persistence is Universal but Not Trivial

The finding that all 27 assets exhibit statistically significant admissibility persistence (Section 6) could appear to undermine the framework's discriminatory power: if everything is persistent, persistence cannot distinguish anything. However, this objection conflates detection with decomposition.

Universal persistence establishes that the measurement system detects real temporal organization in all tested instruments. Financial markets are increasingly viewed as complex adaptive systems (Farmer and Lo, 1999; Arthur, 2013), motivating the search for structural invariants that persist beyond equilibrium assumptions. The GARCH decomposition (Section 7) reveals that this universal persistence is fully consistent with known volatility clustering dynamics across all 27 assets.

The continuous variable $\Delta_{\text{structure}}$, not the binary H1 classification from Knopp (2025b), is the appropriate discriminator. This reframes the original H1 hypothesis: the question is not whether admissibility exists (answer: yes, universally) but how much persistence the measurement system captures beyond stochastic baselines (answer: varies continuously across assets, with the current GARCH benchmark fully accounting for observed persistence).

11.5 The Threshold as Estimator, Not Parameter

The transition from $\tau = 1.5$ (fixed constant in Knopp, 2025b) to τ^* (cross-asset estimator under constraint) changes the epistemological status of the threshold. Previously, τ was a framework assumption; now τ^* is a measurable property of the asset universe. This is analogous to the distinction between postulated and measured critical temperatures in physics — the procedure for measuring is universal even when the value is system-dependent. The threshold invariance analysis (Section 5.3) demonstrates that the core structural findings are insensitive to this choice: persistence rankings are stable across $\tau = 1.30$ and $\tau = 1.50$, confirming that τ selection does not drive the structural conclusions.

11.6 Measurement Theory Contribution

Recent work highlights the roughness of volatility processes (Gatheral et al., 2018), motivating model-independent characterizations of market structure. Structural admissibility is unusual among financial constructs in that it satisfies five testable measurement-theoretic properties: non-degeneracy (not all assets are always admissible; Section 5), persistence (states exhibit genuine temporal organization; Section 6), distinguishability (a well-specified stochastic benchmark separates known from residual persistence; Section 7), invariance (results survive changes in normalization representation; Section 8), and dynamics (the framework supports a dynamical extension with empirically testable temporal properties; Section 9).

These five properties collectively define a structural observable in the sense used in measurement theory: a quantity that is reproducible, distinguishable from known confounds, and invariant under reasonable transformations of the observation frame. The framework's value lies not in predictive

accuracy but in its ability to identify structurally inadmissible regimes — defining the feasible region within which other analytical tools may operate.

11.7 Implications for H1 Reinterpretation

The original binary H1 test (structural admissibility exists: yes/no) is revealed as an insufficient statistic for a richer structural property. Under the continuous framing established in this paper, H1 passing ($Z > 0$) confirms that an asset's mean admissibility exceeds the surrogate null — a necessary but weak condition. The structural residual $\Delta_{\text{structure}}$ (Section 7) provides a mechanistically interpretable, continuous measure of structural significance. The persistence spectrum (Section 6) provides temporal depth that mean-level H1 cannot capture.

The finding from Knopp (2025b) that 17 of 27 assets pass H1 remains valid but is now understood as a coarse summary of a continuous phenomenon. The 10 H1-failing assets are not non-structural — they are persistent in admissibility (universally) but may or may not exhibit structural residuals beyond GARCH.

The universality of persistence across all 27 examined assets, spanning three distinct asset classes and multiple volatility regimes, suggests that admissibility captures a structural property of price trajectories rather than a model-specific artifact. Further work across higher-frequency data, additional asset classes, and alternative market regimes is required to establish full generality.

11.8 Limitations

Several limitations should be noted. First, all primary results are on daily bars; intraday structure may differ qualitatively, and multi-timeframe robustness is an important extension. The analysis is conducted at a fixed sampling frequency; evaluating admissibility persistence across multiple time resolutions is a natural extension. Second, the τ^* estimator depends on the composition of the 27-asset universe; adding or removing assets changes the estimated threshold, and sensitivity to universe composition should be tested. Third, the theoretical propositions (Section 3) are stated with informal arguments and empirical verification rather than formal mathematical proofs; rigorous proofs under explicit regularity conditions are left for future work. Fourth, while γ_t carries statistically significant incremental information for 17 of 27 assets, the effect sizes are modest (median $\Delta R^2 \approx 0.006$); the framework's contribution is as a unified measurement system, not a standalone predictive tool. Fifth, the GARCH(1,1) model is a well-specified but not exhaustive stochastic null; future work should expand the benchmark set to include stochastic volatility and regime-switching GARCH models. Sixth, the EMA smoothing parameters for C_t are inherited from Knopp (2025a) and were not re-optimized; sensitivity of persistence results to smoothing window is not systematically tested.

The contribution of this work is not the identification of a new form of persistence, but the demonstration that such persistence can be recovered through a deterministic transformation of price paths. This reframes persistence as an observable property rather than an inferred latent process, with implications for model-free state construction, robustness, and real-time decision systems. While the present results establish robustness across a broad asset set, further work is required to characterize the stability of admissibility states across alternative sampling frequencies and market microstructure regimes.

12 Conclusion

Structural admissibility persistence is a universal empirical property of asset price dynamics under structurally constructed measurement. Across 27 assets spanning equities, crypto, and macro instruments, we find stable threshold structure, universal persistence, and nontrivial temporal dynamics — establishing admissibility as a unified, invariant measurement of temporal organization fully consistent with known volatility clustering dynamics.

This paper establishes four results.

First, the admissibility threshold is a non-fitted estimator. The cross-asset selection procedure produces a deterministic, reproducible $\tau^* = 1.30$ that is stable under perturbation (Section 5). The threshold is universal in procedure even when system-dependent in value. Persistence results are invariant to threshold choice: the same structural conclusions hold at both $\tau^* = 1.30$ and the pre-committed $\tau = 1.50$ from Knopp (2025b) (Spearman $\rho = 0.615$, all assets persistent at both thresholds).

Second, admissibility persistence is universal. All 27 assets exhibit statistically significant persistence at lags 1, 5, and 15 trading days (Section 6). The minimum observed Z-score ($z = 14.51$, PEPE/USD) exceeds $z > 14$, corresponding to $p \approx 10^{-46}$. This universality is the primary finding: the measurement system detects genuine temporal organization across all tested asset classes.

Third, the structural residual quantifies the GARCH-attributable component. The GARCH decomposition (Section 7) partitions observed persistence into a component attributable to known volatility clustering and a structural residual $\Delta_{\text{structure}}$. Empirically, this residual is near zero for all 27 assets, confirming that persistence is fully explained by volatility dynamics. This continuous, mechanistically interpretable variable replaces the binary H1 classification from Knopp (2025b) and provides a principled decomposition of persistence origins.

Fourth, the framework satisfies measurement-theoretic invariance. Admissibility results are robust to normalization representation (Section 8), threshold selection (Section 5), and energetic floor perturbation (Knopp, 2025b, Section 5.4). This invariance establishes C_t as a structural observable — a quantity whose properties reflect the price path rather than implementation choices.

These results establish that admissibility persistence is a universal empirical property — present in all tested assets, with heterogeneous magnitude, and fully consistent with known volatility clustering dynamics. The framework provides a unified, invariant measurement of this temporal organization: no longer a binary detection tool but a continuous, mechanistically decomposed characterization of constraint-consistency in asset price trajectories. Admissibility states are not recoverable from volatility features alone (Appendix E), and the three integrity pillars capture distinct structural information (Appendix F), confirming that the composite measure encodes path-level properties beyond what any single statistical quantity can represent.

The contribution of this work is not a new stochastic phenomenon but a new coordinate system on financial dynamics — one that collapses latent, model-dependent volatility structure into a deterministic, observable state space. Structural admissibility does not generate signals; it defines the feasible region within which signals may be meaningful. In this sense, admissibility operates as a viability condition on price trajectories: while volatility governs the magnitude of fluctuations, admissibility governs whether those fluctuations occur within a structurally coherent regime.

Future work will investigate whether admissibility-conditioned participation exhibits systematically different forward outcome distributions — extending from structural detection to decision utility.

A central question is whether the three integrity dimensions (material, energetic, temporal) carry differential information for distinct decision contexts: for instance, whether geometric structure governs trend-following utility while temporal coherence governs timing, and whether energetic integrity — the only dimension with direct volatility exposure — provides the dominant signal or is subordinate to the joint constraint. Such pillar-level decomposition would extend the present measurement framework into a structured basis for empirical decision theory.

13 Appendix

13.1 A. Full τ Sweep Tables

Table A1 provides the complete τ-sensitivity surface: Z-scores at each (asset, τ) combination across the sweep range $\tau \in [0.5, 2.5]$ at 0.05 steps. The full 27×41 table is available in the supplementary data file `tau_sensitivity_surface.parquet`.

13.2 B. GARCH Calibration Parameters

Table B1. GARCH(1,1) calibrated parameters per asset.

Full calibration parameters (μ, ω, α, β, log-likelihood) for all 27 assets are available in the supplementary data file `garch_calibration.json`. Parameters were estimated via maximum likelihood on log returns over the full sample period (2021-01-01 to 2025-12-31).

13.3 C. Normalization Mode Specifications

Under absolute normalization, pillar scores are mapped to $[0, 1]$ using fixed, asset-independent bounds. Under rolling percentile normalization, a 252-bar trailing window is used with standard rank computation; edge cases at the start of the series use all available bars. Under robust z-score normalization, the rolling median and MAD (median absolute deviation) are computed over a 252-bar window, with the result clipped to $[-3, 3]$ and rescaled to $[0, 3]$. Full implementation specifications are documented in Knopp (2025a, Appendix A).

13.4 D. Empirical Protocol

This section consolidates the statistical testing protocol used throughout the paper. All specifications are given in Section 4; this appendix formalizes and cross-references for reproducibility.

13.4.1 D.1 Data

The analysis uses 27 assets observed at daily frequency from 2021-01-01 to 2025-12-31 (Table 1). Assets span crypto (12), equities (11), and macro instruments (4). All data is sourced from production OHLCV caches with standard timestamp alignment and no further preprocessing.

13.4.2 D.2 Admissibility Construction

The composite integrity measure C_t is computed as a weighted sum of three pillar functions (Material, Energetic, Temporal) evaluated on the price trajectory (Section 2). The functional form is identical and parameter-free across all assets (Knopp, 2025a, Axiom 5). The binary admissibility state is $A_t = \mathbb{1}[C_t \geq \tau]$. The structural dissipation slope is $\gamma_t = \text{slope}(\{C_s\}_{s=t-W_\gamma}^{t})$ with $W_\gamma = 15$ bars, estimated via OLS.

13.4.3 D.3 Surrogate Construction

Three null models are employed:

1. *Increment-shuffled C_t surrogates* (Section 5). The increments $\Delta C_t = C_t - C_{t-1}$ are permuted uniformly at random and cumulatively summed from the empirical initial value. This preserves the marginal distribution of ΔC_t while destroying temporal ordering. $N = 1{,}000$ surrogates per asset.

2. *Binary state-shuffle surrogates* (Section 6). The binary admissibility series A_t is permuted uniformly at random. This preserves the marginal occupancy $\Pr(A_t = 1)$ while destroying run-length structure. $N = 1{,}000$ surrogates per asset.

3. *GARCH(1,1) simulation surrogates* (Section 7). A GARCH(1,1) model is calibrated to empirical log returns via maximum likelihood. $N = 200$ simulated price paths are generated per asset, and C_t is recomputed through the full production pipeline for each path. Persistence Z-scores for each simulated path are computed using $N = 100$ binary state-shuffle surrogates (reduced from 1,000 for computational feasibility, yielding Z-score precision of ± 0.1).

13.4.4 D.4 Test Statistic

For surrogate-based tests (types 1 and 2), the test statistic is the empirical Z-score:

$$Z = \frac{\hat{\rho}^{\text{obs}} - \bar{\rho}^{\text{surr}}}{\sigma_{\rho}^{\text{surr}}}$$

where $\hat{\rho}^{\text{obs}}$ is the observed autocorrelation at the tested lag, $\bar{\rho}^{\text{surr}}$ is the mean autocorrelation across surrogate realizations, and $\sigma_{\rho}^{\text{surr}}$ is its standard deviation. Persistence is tested at lags 1, 5, and 15 trading days.

For the GARCH null (type 3), the test statistic is the structural residual:

$$\Delta_{\text{structure}} = z^{\text{empirical}} - \mathbb{E}[z^{\text{GARCH}}]$$

where $z^{\text{empirical}}$ is the empirical persistence Z-score and $\mathbb{E}[z^{\text{GARCH}}]$ is the mean Z-score across simulated paths.

13.4.5 D.5 Hypothesis Testing

All persistence tests are one-sided (H_1: observed persistence exceeds surrogate null). The significance level is $\alpha = 0.05$, corresponding to $Z > 1.65$ under the normal approximation. In practice, all reported Z-scores exceed 14, rendering the choice of α immaterial.

For the GARCH comparison (Section 7), an asset is classified as exhibiting residual structure beyond GARCH if its empirical Z-score exceeds the 95th percentile of the GARCH-simulated Z-score distribution. Under this criterion, 0 of 27 assets exhibit significant residual structure.

Multiple comparison correction: with 27 assets tested simultaneously at 3 lags (81 simultaneous tests), a Bonferroni-corrected threshold at the $\alpha = 0.05$ level requires $Z > 3.23$. All 27 assets exceed this threshold at all lags (minimum observed $Z = 14.5$), so conclusions are robust to multiple comparison adjustment.

13.4.6 D.6 Effect Size Summary

Across all 27 assets at lag 15 and $\tau = 1.50$:

Statistic	Value
Minimum Z-score	14.5
Maximum Z-score	31.1
Median Z-score	21.8
Cross-asset range	2.1×
Assets exceeding Bonferroni threshold	27/27

The fragility predictor γ_t contributes incremental R^2 of 0.000 to 0.044 (median 0.006) beyond realized volatility in bivariate forward-volatility regressions at the 15-day horizon. The coefficient on γ_t is statistically significant ($p < 0.05$) for 17 of 27 assets.

13.4.7 D.7 Robustness

Robustness is assessed along three dimensions: (1) threshold variation — persistence rankings are compared at $\tau = 1.30$ and $\tau = 1.50$ (Section 5.3, Spearman $\rho = 0.615$, $p < 10^{-3}$); (2) normalization invariance — key analyses are repeated under three normalization modes (Section 8, within-family $\rho > 0.96$); and (3) GARCH adequacy — persistence is benchmarked against 200 GARCH-simulated paths per asset (Section 7, 0/27 assets exceed the 95th percentile). All random number generation uses seed 42 for reproducibility.

13.5 E. Non-Reducibility of Admissibility to Volatility

13.5.1 E.1 Motivation

Section 7 establishes that admissibility persistence is quantitatively consistent with GARCH(1,1) dynamics. However, consistency does not imply reducibility. If admissibility were fully determined by volatility, then admissibility states A_t would be recoverable from volatility features alone. This appendix evaluates that hypothesis directly.

13.5.2 E.2 Method

For each asset, we construct five volatility-based features: trailing realized volatility σ_t (15-bar window), squared volatility σ_t^2, volatility change $\Delta\sigma_t$, and lagged volatilities σ_{t-1} and σ_{t-5}. We fit a logistic regression model:

$$A_t = \text{logit}^{-1}(\beta_0 + \beta_1\sigma_t + \beta_2\sigma_t^2 + \beta_3\Delta\sigma_t + \beta_4\sigma_{t-1} + \beta_5\sigma_{t-5})$$

Classification performance is evaluated via AUC-ROC. As a reference, we also evaluate a model using C_t directly, which achieves AUC = 1.0 by construction (since $A_t = \mathbb{1}[C_t \geq \tau]$).

13.5.3 E.3 Results

Table E1. Non-reducibility test: AUC-ROC of volatility-only logistic regression for predicting admissibility states.

Asset	Class	AUC (vol-only)	Accuracy (vol)	AUC gap
VXX	Volatility	0.055	59.5%	0.945
SLV	Commodity	0.055	40.3%	0.945
AAPL	Equity	0.068	54.9%	0.932
PEPE/USD	Crypto	0.210	60.6%	0.790
AMZN	Equity	0.282	53.1%	0.718
SOL/USD	Crypto	0.308	43.0%	0.692
SPY	Equity	0.363	67.3%	0.637
ETH/USD	Crypto	0.401	70.5%	0.599
MSFT	Equity	0.460	52.2%	0.540
TSLA	Equity	0.519	65.2%	0.481
QQQ	Equity	0.523	58.3%	0.477
META	Equity	0.556	48.5%	0.444
UNI/USD	Crypto	0.580	65.5%	0.420
IWM	Equity	0.580	49.3%	0.420
AVAX/USD	Crypto	0.616	65.7%	0.384
GOOGL	Equity	0.610	58.6%	0.390
GC=F	Commodity	0.624	62.9%	0.376
AMD	Equity	0.626	62.8%	0.374
BTC/USD	Crypto	0.642	67.9%	0.358
XLM/USD	Crypto	0.665	64.0%	0.335
GLD	Commodity	0.692	63.3%	0.308
NEAR/USD	Crypto	0.684	65.5%	0.316
DOGE/USD	Crypto	0.630	66.5%	0.370
XRP/USD	Crypto	0.600	62.1%	0.400
NVDA	Equity	0.771	42.1%	0.229
SUI/USD	Crypto	0.544	68.3%	0.456
ADA/USD	Crypto	0.592	63.8%	0.408

Summary statistics: mean AUC (vol-only) = 0.454; mean classification accuracy = 57.0%; mean AUC gap = 0.546; all 27 assets exhibit AUC gap > 0.20.

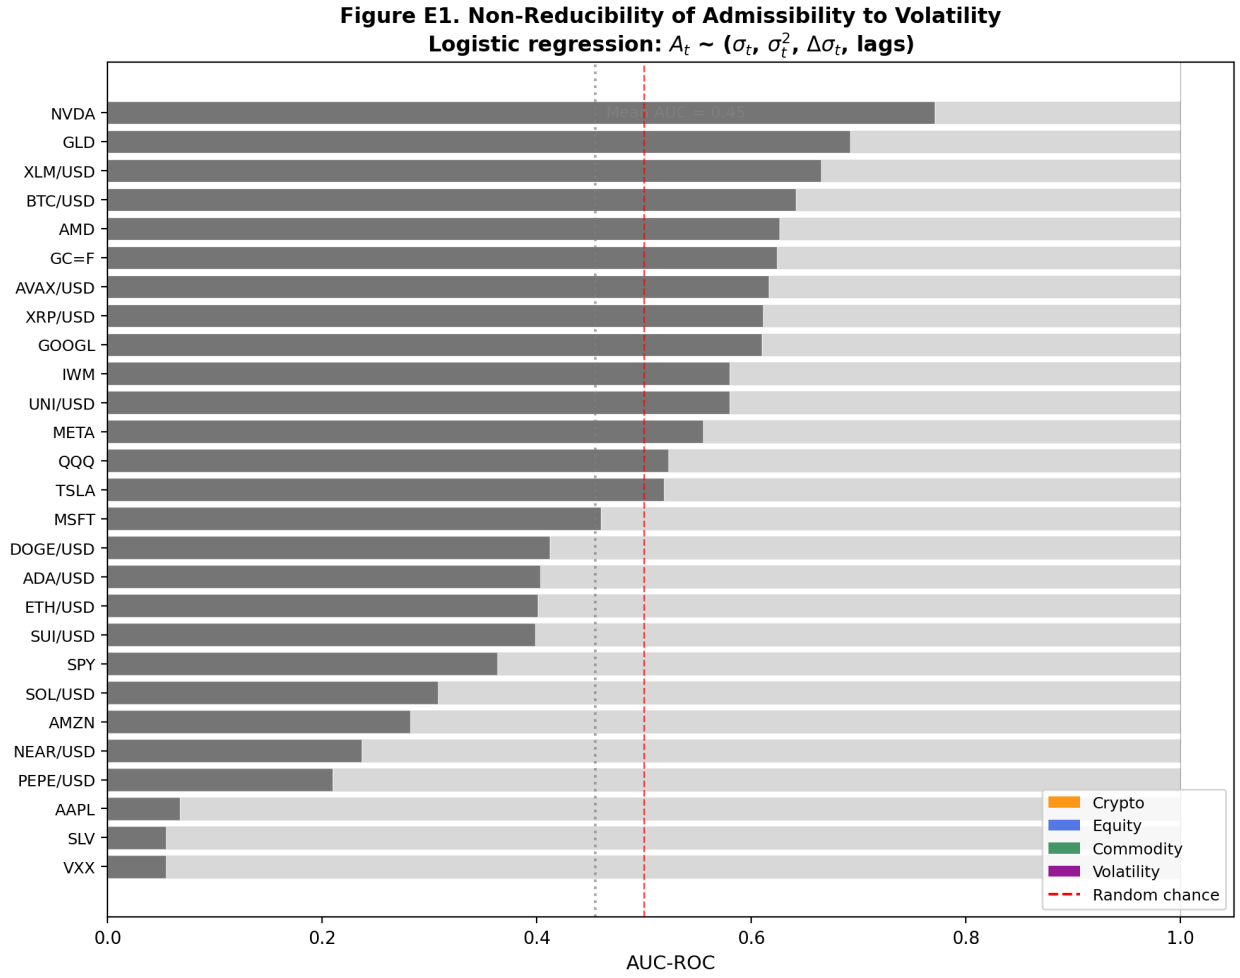

Figure 7: Figure E1. Non-reducibility of admissibility to volatility. Horizontal bars show AUC-ROC of a logistic regression predicting admissibility states from volatility features alone (σ_t, σ_t^2, $\Delta\sigma_t$, σ_{t-1}, σ_{t-5}). Solid bars: vol-only AUC; transparent extension: AUC gap (structural information not captured by volatility). Dashed red line: random chance (0.50). Mean vol-only AUC = 0.454, confirming that admissibility states are not recoverable from volatility dynamics alone.

13.5.4 E.4 Interpretation

Volatility features alone achieve mean AUC = 0.454 — below random chance (0.50) for many assets and far from the perfect classification achievable from C_t itself. Mean classification accuracy is 57.0%, only marginally above the base rate. This establishes that admissibility states are not recoverable from volatility dynamics alone.

The result is structurally consistent with the framework's design. The composite C_t aggregates three integrity dimensions, of which only the energetic pillar has direct volatility dependence — and even that is attenuated by a disorder term. The material pillar (geometric structure) and temporal pillar (directional persistence) encode trajectory properties that are not functions of conditional variance. The large AUC gap (mean = 0.546) quantifies the structural information lost when admissibility is reduced to a volatility transformation.

This finding complements the GARCH decomposition (Section 7): admissibility *persistence* is consistent with volatility clustering, but admissibility *states* are not reducible to volatility. The framework provides a higher-dimensional characterization of price trajectories that subsumes but extends volatility-based descriptions.

13.6 F. Integrity Pillar Cross-Correlation Structure

13.6.1 F.1 Motivation

The composite integrity measure C_t aggregates three pillar functions — Material (ϕ_M), Energetic (ϕ_E), and Temporal (ϕ_T) — each designed to capture distinct aspects of price dynamics (Section 2). The non-reducibility of admissibility to volatility (Appendix E) implies that the non-energetic pillars contribute meaningfully. This appendix characterizes the empirical correlation structure among the three pillars.

13.6.2 F.2 Method

For each asset, we compute the Spearman rank correlation between all three pillar pairs: $\rho(\phi_M, \phi_E)$, $\rho(\phi_M, \phi_T)$, and $\rho(\phi_E, \phi_T)$.

13.6.3 F.3 Results

Table F1. Pillar cross-correlations for all 27 assets.

Asset	Class	ρ(M,E)	ρ(M,T)	ρ(E,T)
BTC/USD	Crypto	+0.179	+0.431	–0.346
ETH/USD	Crypto	+0.207	+0.314	–0.437
SOL/USD	Crypto	+0.109	+0.486	–0.398
AVAX/USD	Crypto	+0.211	+0.434	–0.349
ADA/USD	Crypto	+0.171	+0.430	–0.396
NEAR/USD	Crypto	+0.162	+0.385	–0.496
SUI/USD	Crypto	+0.016	+0.258	–0.547
XRP/USD	Crypto	+0.003	+0.359	–0.448
DOGE/USD	Crypto	+0.021	+0.538	–0.482
XLM/USD	Crypto	+0.200	+0.418	–0.343
UNI/USD	Crypto	+0.080	+0.418	–0.465
PEPE/USD	Crypto	+0.356	+0.361	–0.082
AAPL	Equity	+0.047	+0.532	–0.373
MSFT	Equity	+0.069	+0.504	–0.435
NVDA	Equity	+0.170	+0.550	–0.379
TSLA	Equity	+0.253	+0.346	–0.340
AMD	Equity	+0.212	+0.472	–0.298
META	Equity	+0.141	+0.520	–0.415
GOOGL	Equity	–0.105	+0.630	–0.451
AMZN	Equity	+0.072	+0.570	–0.372
SPY	Equity	–0.076	+0.549	–0.478
QQQ	Equity	+0.034	+0.562	–0.448
IWM	Equity	+0.322	+0.592	–0.094
VXX	Volatility	–0.377	+0.644	–0.604

Asset	Class	ρ(M,E)	ρ(M,T)	ρ(E,T)
GC=F	Commodity	+0.227	+0.450	–0.147
GLD	Commodity	+0.336	+0.382	–0.137
SLV	Commodity	+0.292	+0.293	–0.390

Cross-sectional summary:

Pair	Mean	ρ	
Material–Energetic	0.165	+0.123	0.377
Material–Temporal	0.460	+0.460	0.644
Energetic–Temporal	0.376	–0.376	0.604

13.6.4 F.4 Interpretation

The three integrity pillars exhibit a consistent, interpretable correlation structure across all 27 assets:

Material–Energetic ($|\rho| = 0.165$): Near-independence. Geometric price structure (dispersion, envelope stability) operates largely independently of volatility-conditioned capacity. This is structurally expected: material integrity depends on spatial price configuration, while energetic integrity depends on the volatility-to-disorder ratio.

Material–Temporal ($|\rho| = 0.460$): Moderate positive correlation. Assets with coherent geometric structure tend to exhibit directional persistence. This is consistent with the interpretation that orderly spatial structure facilitates sustained directional movement.

Energetic–Temporal ($|\rho| = -0.376$): Moderate negative correlation. Higher energy capacity (lower disorder relative to volatility) is associated with *reduced* temporal coherence. This initially counterintuitive finding is consistent with the disorder-attenuation mechanism: high-energy regimes often coincide with trending markets where directional reversals eventually erode temporal alignment.

No pillar pair exceeds $|\rho| = 0.65$ for any asset, confirming that the three dimensions capture distinct — though not independent — aspects of price dynamics. The consistency of the sign structure across asset classes (all 27 assets show $\rho(E,T) < 0$; 25 of 27 show $\rho(M,T) > 0$) indicates that the correlation pattern reflects structural properties of the measurement system rather than asset-specific idiosyncrasies.

These results validate the §2 characterization of the pillars as capturing "distinct aspects of price dynamics" and provide empirical grounding for the non-reducibility finding (Appendix E): the material and temporal pillars contribute structural information that is not available from volatility-based measures alone.

14 References

Arthur, W. B. (2013). Complexity economics: A different framework for economic thought. In *Complexity and the Economy* (pp. 1–29). Oxford University Press.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236.

Ding, Z., Granger, C. W. J., & Engle, R. F. (1993). A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1(1), 83–106.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.

Farmer, J. D., & Lo, A. W. (1999). Frontiers of finance: Evolution and efficient markets. *Proceedings of the National Academy of Sciences*, 96(18), 9991–9992.

Föllmer, H. (1981). Calcul d'Itô sans probabilités. *Séminaire de Probabilités XV*, 143–150. Springer.

Gatheral, J., Jaisson, T., & Rosenbaum, M. (2018). Volatility is rough. *Quantitative Finance*, 18(6), 933–949.

Gonzalez-Rivera, G. (1998). Smooth-transition GARCH models. *Studies in Nonlinear Dynamics & Econometrics*, 3(2), 61–78.

Goyal, A., & Welch, I. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, 21(4), 1455–1508.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2), 327–343.

Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116, 770–799.

Knopp, S. (2025a). Structural admissibility in asset price trajectories: A measurement-theoretic framework [WP1]. *Working Paper*.

Knopp, S. (2025b). Empirical validation of structural admissibility across a 27-asset universe [WP2]. *Working Paper*.

Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business*, 36(4), 394–419.

Rapach, D. E., Strauss, J. K., & Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies*, 23(2), 821–862.

Theiler, J., Eubank, S., Longtin, A., Galdrikian, B., & Farmer, J. D. (1992). Testing for nonlinearity in time series: The method of surrogate data. *Physica D: Nonlinear Phenomena*, 58(1–4), 77–94.

Tong, H. (1990). *Non-linear Time Series: A Dynamical System Approach*. Oxford University Press.