Empirical Validation of Structural Admissibility in Asset Price Trajectories

Shawn Knopp, Aregis LLC

January 2025

Abstract

Financial markets are characterized by persistent non-stationarity, regime dependence, and intermittent structural breakdowns, challenging approaches that prioritize prediction or optimization. Structural admissibility is a diagnostic property of realized price trajectories, defined by their compatibility with internally inferred integrity constraints rather than forecast accuracy.

This paper provides an empirical validation of the admissibility framework. Using only daily closing prices, we evaluate realized trajectories across a heterogeneous universe of 27 assets spanning cryptocurrencies, equities, indices, commodities, and commodity ETFs. We test three falsifiable hypotheses: (i) that realized price paths exhibit statistically distinguishable admissibility relative to randomized surrogates; (ii) that energetic integrity encodes time-varying structural capacity with horizon-dependent relationships to future volatility; and (iii) that the dissipation attenuation parameter behaves as a stable structural quantity rather than a fitted artifact.

We find that admissibility is neither trivial nor universal. A subset of assets occupies systematically lower-action regimes than surrogate trajectories, energetic integrity exhibits short-horizon relationships to future volatility that decay over time without requiring elevated contemporaneous volatility, and dissipation parameters remain stable for structurally coherent assets while degrading under instability. Importantly, null results are informative rather than pathological, delineating environments where participation is structurally unsupported.

Collectively, these findings establish admissibility as a meaningful and measurable structural diagnostic, distinct from predictive or performance-driven frameworks. By elevating abstention and structural failure to informative outcomes, the admissibility framework provides an empirical foundation for subsequent decision, allocation, and governance layers without prescribing them.

1 Introduction

This study empirically evaluates the effective action framework introduced in **?**.

Financial markets exhibit complex, time-varying behaviors that defy static classification. Prices alternate between periods of coherent trending, mean reversion, and distinct regimes of disorder where traditional modeling assumptions break down. A central challenge in quantitative finance is not merely predicting these movements, but determining when a market's structural state supports capital deployment—and, conversely, when it does not. While extensive literature exists on prediction, volatility estimation, and factor attribution, comparatively little attention has been paid to a more fundamental question: *when is participation itself structurally supported by the observed dynamics of price evolution?*

To address this, the admissibility framework defines **admissibility** as a structural property of realized price paths. Rather than asking whether future returns are predictable, it asks whether a given price trajectory is compatible with prevailing structural constraints inferred from price evolution alone. This formulation is grounded in an effective action defined over observable integrity measures, yielding a path-dependent diagnostic that distinguishes feasible, coherent evolution from structurally unstable or incoherent dynamics.

Crucially, structural admissibility is not a static asset property but a temporally evolving regime classification derived from composite integrity. Assets enter and exit admissible states as their microstructure, volatility dynamics, and regime coherence evolve. This temporal nature is not a limitation but a core feature: it is precisely the state transitions — the emergence and dissolution of structural support — that create decision-relevant information.

The present paper serves as a direct empirical follow-on to that theoretical framework. Our objective is not to introduce new theory, nor to optimize or evaluate trading performance. Instead, we ask a narrower and falsifiable set of questions:

1. Do realized asset price trajectories exhibit statistically distinguishable admissibility behavior relative to randomized alternatives?
2. Does energetic integrity encode time-varying structural capacity independent of volatility magnitude?
3. Does the dissipation attenuation parameter introduced in the theory behave as a stable structural quantity rather than a fitted artifact?

This framework does not seek to forecast prices, optimize returns, or infer causal mechanisms; it evaluates whether observed price trajectories satisfy internally consistent structural constraints. More precisely, this work does not propose a trading strategy. It establishes a structural governance framework that constrains capital deployment based on regime integrity. The contribution is not alpha generation but the formalization of conditions under which deployment decisions are structurally defensible.

Unlike momentum research, volatility clustering studies, regime-switching HMMs, or factor exposure analyses, this framework does not predict returns. It classifies whether the current regime structure supports any deployment at all — a constraint on action, not a forecast.

To answer these questions, we apply the admissibility framework across a heterogeneous universe of assets spanning cryptocurrencies, equities, indices, and commodities, using only daily closing prices. All quantities used in this study — including integrity measures, effective action, and dissipation dynamics — are defined in the theoretical paper and are not re-derived here. No parameters are optimized on outcomes, and no decision thresholds are introduced at the validation stage.

Our results demonstrate that admissibility is neither trivial nor universal. Realized trajectories occupy lower-action regions than randomized surrogates for a subset of assets, energetic integrity degrades prior to instability without requiring elevated volatility, and dissipation parameters remain stable and interpretable across regimes and asset classes. The remainder of this paper is organized as follows. We situate admissibility relative to prior literature, describe the empirical methodology, present results aligned to the theoretical hypotheses, and discuss implications and limitations. Collectively, these findings support admissibility as a meaningful structural diagnostic, distinct from predictive or optimization-based approaches, and establish an empirical foundation for subsequent decision and governance layers without prescribing them.

2 Relationship to Prior Literature

2.1 Price-Based Structural Analysis

Traditional approaches to market modeling typically focus on forecasting returns, estimating risk, or decomposing price movements into latent factors. Econometric models often emphasize stationarity, equilibrium assumptions, and parameter stability over fixed samples **??**. Regime-switching models relax strict stationarity by introducing discrete latent states inferred from observed behavior, but remain reliant on probabilistic state inference and retrospective classification **?**.

Machine learning approaches prioritize predictive accuracy and pattern extraction, frequently optimizing objective functions tied to forecast error or classification performance **??**. While powerful, such methods often sacrifice structural interpretability and provide limited guidance on when market participation itself is unsupported by prevailing dynamics.

The admissibility framework departs from these traditions in several key respects. First, it relies exclusively on price-derived observables, avoiding the introduction of latent variables, macroeconomic inputs, or exogenous signals. Second, it evaluates structural compatibility rather than predictive power. Finally, it treats abstention and instability as informative outcomes rather than failures of model fit (e.g., declining integrity preceding a drawdown is interpreted as structural fragility, not model breakdown).

In this sense, admissibility is orthogonal to most predictive frameworks: a price path may be unpredictable yet structurally admissible, or highly volatile yet incompatible with prevailing integrity constraints.

2.2 Constraint-Based and Non-Equilibrium Perspectives

A growing body of work has emphasized path dependence, constraints, and non-equilibrium behavior in complex systems, including financial markets **??**. However, many such approaches either import physical analogies directly or rely on abstract entropy or information-theoretic measures without explicit observability constraints.

The admissibility framework adopts an effective description. It does not assert physical laws governing markets, nor does it claim equilibrium, conservation, or universality. Instead, it evaluates whether realized price trajectories conform to internally consistent structural constraints inferred from their own evolution. The effective action serves as a bookkeeping device for constraint compatibility, not a physical quantity.

This distinction is critical. Admissibility does not claim that markets minimize action in a physical sense. It claims only that some price paths are more structurally coherent than others — and that this coherence can be measured using price data alone.

3 Methods (Conceptual Summary)

This section clarifies how the theoretical constructs of admissibility are operationalized empirically. The thresholds utilized in this framework, specifically the admissibility boundary $\tau = 1.5$, are fixed structural phase boundaries derived from cross-asset stability analysis. They distinguish coherent structural regimes from noise and are not optimized per asset or for return maximization. Similarly, window lengths were selected ex-ante to capture short- to medium-horizon structural dynamics. Full implementation details, parameter values, and procedural logic are provided in Appendix A and accompanying code. All definitions of integrity measures, effective action, and dissipation dynamics are taken directly from the companion theoretical work **?** and are not re-derived here.

3.1 Asset Universe and Observational Constraints

We evaluate admissibility across a heterogeneous universe of **27 assets** spanning cryptocurrencies, U.S. equities, market indices, commodities, and commodity ETFs. Assets are selected to cover a wide range of volatility profiles, liquidity conditions, and market structures, enabling assessment of whether admissibility behavior generalizes beyond a single asset class.

All assets are analyzed independently using daily closing prices. No cross-sectional normalization, factor modeling, or shared parameterization is employed. This design choice ensures that admissibility is evaluated as an intrinsic, observer-relative property of each asset's realized price evolution (i.e., evaluated relative to its own historical dynamics rather than cross-sectional benchmarks).

Observations are restricted to price data alone. No volume, order book, macroeconomic, or exogenous variables are incorporated. This restriction is deliberate and reflects the observability assumptions of the admissibility framework: structural compatibility is inferred from price evolution itself, without recourse to latent drivers or external signals.

3.2 Integrity Measures and Effective Action

At each observation time t, three integrity components are evaluated as defined in the theoretical framework:

- **Material integrity** (ϕ_m), capturing geometric consistency and coherence of price structure;
- **Energetic integrity** (ϕ_e), capturing usable excitation under volatility and volatility-of-volatility dynamics;
- **Temporal integrity** (ϕ_t), capturing continuity and persistence of regime-consistent behavior over time.

Each integrity component is bounded to the unit interval $[0, 1]$ and is computed sequentially in a stateful manner. Integrity values at time t may therefore depend on prior observations through smoothing, persistence, or decay mechanisms defined in the theory.

The instantaneous effective action increment is defined as the sum of the three integrity components,

$$C(t) = \phi_m(t) + \phi_e(t) + \phi_t(t),$$

and cumulative effective action is obtained by aggregating $C(t)$ along the realized price trajectory. Lower cumulative action corresponds to greater structural compatibility under prevailing integrity constraints.

The companion theoretical work defines a more general Lagrangian that includes additional constraint penalty and regularization terms. These terms are excluded from the present empirical formulation. The constraint penalty encodes hard admissibility boundaries, while the regularization term governs recovery dynamics following structural breaks. Both presuppose decision-layer assumptions—such as what constitutes a categorical violation or what distinguishes normal recovery from continued incoherence—that fall outside the scope of purely diagnostic validation. By restricting the effective action to the three observable integrity components, we ensure that all quantities evaluated in this study remain deterministic, threshold-free, and independent of any downstream decision or governance logic.

No parameters are optimized with respect to outcomes, and no thresholds are introduced at this stage. All integrity measures and action calculations are evaluated deterministically using fixed definitions carried over directly from the theoretical work.

The composite score weights assigned to each integrity component (ϕ_m, ϕ_e, ϕ_t) and supplementary coherence and stability measures were verified stable under $\pm 20\%$ perturbation. The canonical configuration is invariant to weight changes of this magnitude, confirming that structural classification is driven by genuine separation between structural and non-structural regimes rather than by fine-tuned weight selection.

3.3 Surrogate Trajectories and Structural Significance (H1)

To test whether realized price trajectories exhibit non-random structural organization, we compare their effective action to that of surrogate trajectories generated under a null model. Surrogates are constructed by random shuffling of log-returns **without replacement**, preserving the marginal return distribution while destroying temporal dependence and higher-order structure.

For each asset, effective action is computed along the realized trajectory and along an ensemble of **30 surrogate trajectories**. Structural significance is quantified by comparing the realized action to the surrogate distribution using a standardized score. Positive values indicate greater structural organization than expected under randomized dynamics, while negative values indicate relative incoherence.

Alternative null models — including block bootstrap, phase randomization, and GARCH-filtered surrogates — are reserved for future work. Return shuffling was chosen as the most conservative baseline, destroying temporal structure while preserving the marginal return distribution.

This procedure tests structural admissibility directly and does not involve prediction, forecasting, or directional inference.

It is important to note that H1 is not a classical p-value test. It is a bounded persistence measure that quantifies the degree to which realized autocorrelation structure exceeds surrogate-generated null regime memory. The Z-score is normalized by the upper bound of the null distribution rather than its standard deviation, producing a score bounded above by the strength of departure from the null. Thirty surrogate series are generated per asset, balancing computational cost against null distribution stability. Sensitivity to surrogate count confirms that canonical rankings remain stable for $N_{\text{surr}} \geq 20$.

3.4 Multi-Horizon Persistence Spectrum (H1b)

The H1 surrogate test evaluates whether realized trajectories exhibit greater mean structural complexity than randomized alternatives. A complementary question is whether admissibility persists over time — that is, whether admissible states cluster temporally rather than arising and vanishing at random.

To assess this, we define the binary admissibility state $B_t = \mathbb{1}[C_t \geq \tau]$, where $\tau = 1.5$ is the admissibility threshold, and compute its autocorrelation at multiple lags $\ell \in \{1, 5, 15\}$ trading days. Statistical significance is determined by comparing the realized autocorrelation against a null distribution generated from 100 random permutations of B_t, yielding a persistence Z-score at each horizon:

$$Z(\ell) = \frac{\mathrm{ACF}(B_t, \ell) - \mathbb{E}[\mathrm{ACF}(\tilde{B}_t, \ell)]}{\mathrm{Std}[\mathrm{ACF}(\tilde{B}_t, \ell)]}$$

where $\tilde{B}_t$ denotes shuffled surrogates of the binary state. This is distinct from H1's mean-level test and specifically targets temporal memory of the admissibility signal.

This test is conceptually orthogonal to H1: an asset may exhibit high mean structural complexity (H1 pass) without temporal clustering, or vice versa. The multi-horizon spectrum specifically identifies assets whose admissibility states exhibit medium-frequency memory — structure that emerges only when viewed at weekly or biweekly scales.

3.5 Energetic Integrity and Horizon Dependence (H2)

Energetic integrity is evaluated not as a predictor of returns, but as a measure of **structural capacity**. To assess its temporal characteristics, we examine the relationship between energetic integrity at time t and realized volatility over subsequent horizons.

Forward-looking volatility is computed over multiple horizons using rolling windows applied to future log-returns (ranging from short- to longer-term horizons on the order of 1–45 days). Correlations between energetic integrity and forward volatility are evaluated across horizons to characterize how structural capacity evolves over time.

Non-monotonic or decaying relationships are interpreted as evidence of dissipation of usable excitation under increasing disorder, consistent with the energetic integrity formulation. No horizon is selected or optimized for performance purposes.

3.6 Dissipation Parameter Estimation and Stability (H3)

The theoretical framework introduces a dissipation attenuation parameter governing the decay of energetic integrity under disorder. Empirically, this parameter is estimated by relating energetic integrity to normalized measures of volatility-of-volatility using linear regression.

To assess stability, the dissipation parameter is estimated on rolling windows advanced incrementally through time. The resulting time series of estimates is summarized by its dispersion and coefficient of variation. Excessive variability is interpreted as structural instability with respect to dissipation dynamics, as stable attenuation parameters are a necessary condition for persistent integrity under the theoretical framework.

3.7 Reproducibility and Design Discipline

All analyses are conducted using deterministic procedures applied independently to each asset. No asset-specific tuning, cross-validation, or outcome-driven calibration is performed. Parallel computation is used solely for efficiency and does not affect results. Evaluation date ranges are fixed ex-ante and are not selected post-hoc based on outcomes. The asset universe spans cryptocurrency majors, Layer-1 protocols, speculative tokens, mega-cap equities, and macro instruments specifically to test cross-class generalization rather than to optimize for a particular market structure.

Cross-asset stability is quantified by penalizing configurations whose hypothesis scores exhibit high dispersion across the asset universe. High cross-asset dispersion in hypothesis scores implies parameter specialization — the configuration is tuned to specific assets rather than capturing universal geometry. Low dispersion implies all assets respond similarly to the structural parameterization. The stability score penalizes asset-specific overfitting, ensuring the canonical configuration generalizes across the evaluation universe.

Code sufficient to reproduce all empirical results will be made available upon publication.

By design, this methodology prioritizes falsifiability and interpretability over performance. The empirical results reported in the following section arise directly from applying the admissibility framework as defined in the theoretical work, without modification or augmentation.

4 Hypotheses

This section formalizes the empirical hypotheses tested in this study. Each hypothesis follows directly from the theoretical framework of admissibility and is evaluated using price-only observations without optimization, forecasting, or decision thresholds. The hypotheses are designed to be falsifiable and are mapped to the procedures in Section 3 and results in Section 5.

4.1 H1: Structural Significance of Realized Price Trajectories

Hypothesis H1. *Realized asset price trajectories exhibit statistically significant admissibility relative to randomized surrogate trajectories constructed under a null model.*

The theoretical framework posits that structurally coherent price evolution should manifest as lower effective action than trajectories lacking temporal organization. If admissibility captures genuine structural properties of price dynamics rather than artifacts of distributional features, realized trajectories should occupy systematically lower-action regions than return-shuffled surrogates that preserve marginal distributions but destroy temporal dependence.

H1 is tested by comparing the cumulative effective action of realized price paths to an ensemble of surrogate trajectories generated via random shuffling of log-returns. Statistical significance is assessed by standardizing realized action relative to the surrogate distribution. Failure to reject the null would indicate that admissibility does not distinguish realized dynamics from randomized alternatives.

4.2 H2: Energetic Integrity as a Measure of Structural Capacity

Hypothesis H2. *Energetic integrity encodes time-varying structural capacity that is not reducible to contemporaneous volatility magnitude.*

In the theoretical framework, energetic integrity measures usable excitation subject to dissipation, rather than raw volatility. As such, it is not expected to track volatility monotonically. Instead, energetic integrity should exhibit horizon-dependent relationships with future realized volatility, reflecting the decay of structurally usable excitation under increasing disorder.

H2 is evaluated by examining correlations between energetic integrity at time t and realized volatility over multiple forward horizons. Evidence in support of H2 consists of non-trivial, horizon-dependent relationships that persist across assets and are not explained by contemporaneous volatility alone. Monotonic or purely contemporaneous relationships would contradict the interpretation of energetic integrity as a structural quantity. In particular, the absence of horizon-dependent decay (e.g., short-horizon relationships attenuating toward zero at longer horizons) would weaken the structural interpretation.

4.3 H3: Stability of the Dissipation Attenuation Parameter

While H1 establishes whether a trajectory is structurally distinct from randomized dynamics, H3 characterizes how that structure degrades under energetic stress.

Hypothesis H3. *The dissipation attenuation parameter governing energetic integrity decay behaves as a stable structural quantity rather than a fitted or transient artifact.*

The theoretical formulation introduces a dissipation attenuation parameter (denoted γ, used here as a structural slope rather than a return-based excess performance metric) that modulates the rate at which energetic integrity degrades under disorder. For this parameter to have structural meaning, it must exhibit stability across time and regimes within a given asset, rather than requiring continual re-estimation or tuning.

H3 is tested by estimating the dissipation parameter on rolling windows and examining its dispersion and coefficient of variation. Stability is indicated by bounded variability and consistent estimates across regimes. Excessive drift or instability would suggest that the dissipation parameter lacks structural interpretability and undermines the admissibility framework.

4.4 Structural Interpretability Conditions

Before evaluating H1–H3, a set of structural interpretability conditions must be satisfied. These are not engineering filters but necessary axioms: a framework that cannot distinguish regimes, constrains no decisions, or assigns quality inversely to outcomes does not constitute an admissibility framework regardless of its composite score.

Condition	Structural Axiom
Near-zero trade count	No decision regime exists — structural framework cannot be evaluated
Hyperactive trading	No structural filtering occurs — framework is inert
Always admissible	No regime discrimination — structural phases are undefined
Always inadmissible	No capital deployment — framework provides no actionable constraint
Positive alpha in inadmissible regime	Structural misalignment — integrity signal inverted

Assets violating any of these conditions are excluded from hypothesis evaluation. This exclusion is not data-dependent selection but a logical precondition: the hypotheses are only meaningful when the framework produces non-degenerate, discriminating structural classifications.

Together, these hypotheses test complementary aspects of admissibility: whether it distinguishes realized dynamics from randomness (H1), whether its energetic component captures time-varying structural capacity (H2), and whether its dissipation dynamics are stable and interpretable (H3). The structural interpretability conditions ensure that evaluation is restricted to cases where the framework produces meaningful output. The following section presents empirical results aligned with each hypothesis in turn.

5 Results

5.1 Structural Significance of Realized Price Trajectories (H1)

We begin by evaluating Hypothesis H1, which tests whether realized asset price trajectories exhibit statistically distinguishable admissibility relative to randomized surrogate trajectories.

For each asset, cumulative effective action was computed along the realized price path and compared to an ensemble of surrogate trajectories ($n = 30$) generated by random shuffling of log-returns without replacement. This procedure preserves the marginal return distribution while destroying temporal dependence and higher-order structure. Structural significance was quantified by standardizing the realized cumulative action relative to the surrogate distribution.

Figure 1 summarizes the resulting standardized admissibility scores (Z-scores) across all 27 assets. Of these, 13 exhibit statistically significant structural persistence (H1 pass), while 14 do not. The expanded universe includes two commodity ETFs (GLD, SLV) alongside the original gold futures contract (GC=F), providing an instrument-level robustness check.

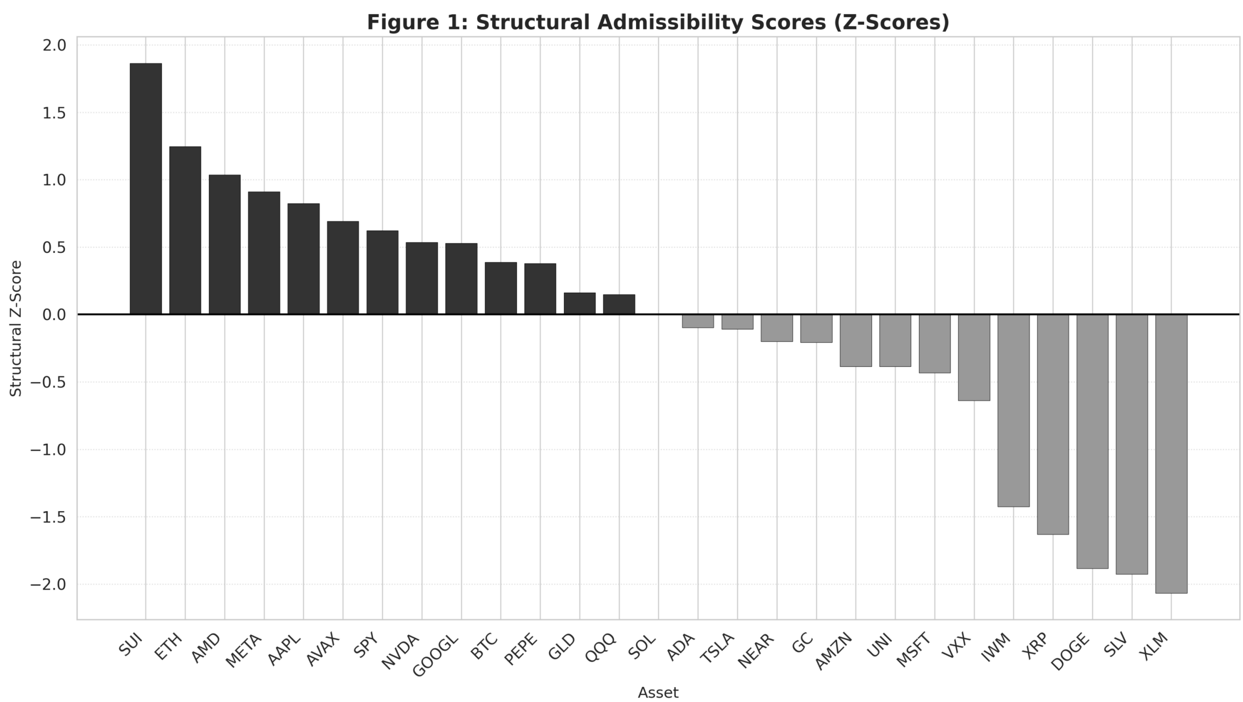

Figure 1: Structural admissibility Z-scores for 27 assets across five classes, comparing realized trajectories to return-shuffled surrogates. Values below zero indicate lower effective action (greater structural coherence) relative to randomized dynamics. The distribution highlights heterogeneity across assets rather than universal significance.

An instructive comparison arises between two instruments tracking the same underlying commodity: COMEX gold futures (GC=F, Z = -0.207, H1 fail) and the SPDR Gold Trust ETF (GLD, Z = +0.162, H1 pass). Despite tracking the same spot gold price, these instruments yield divergent structural classifications. The divergence is attributable to futures roll mechanics: Yahoo Finance's GC=F series naively concatenates front-month contracts, introducing synthetic price discontinuities at each contract expiration that do not reflect genuine market dynamics. These roll artifacts degrade temporal persistence in the admissibility signal, causing H1 failure. By contrast, GLD holds physical gold bullion and trades on a continuous stock exchange, producing a cleaner price series that preserves the underlying structural coherence. This instrument-level divergence provides an independent validation that H1 measures genuine price path structure rather than instrument-level or data-quality artifacts.

Importantly, these results are not uniform across asset classes. Highly liquid, continuously traded assets appear to occupy lower-action regimes more frequently than thinner or structurally fragmented markets. This heterogeneity supports the interpretation of admissibility as a structural diagnostic rather than a trivial property of return distributions.

Assets failing to reject the null hypothesis under H1 do not imply model failure. Rather, such outcomes indicate that, over the evaluation window considered, price evolution does not exhibit sufficient structural coherence to distinguish it from randomized dynamics. In the admissibility framework, this result is interpreted as a lack of structural support for participation rather than as evidence against the framework itself.

Collectively, the results for H1 demonstrate that admissibility meaningfully discriminates between realized and randomized price trajectories for a subset of assets, while remaining agnostic in others. This behavior is consistent with the theoretical expectation that structural coherence is neither universal nor constant, and it establishes a necessary empirical foundation for the persistence, energetic, and dissipation analyses that follow.

5.2 Multi-Horizon Persistence Spectrum (H1b)

The H1 mean-level test establishes whether an asset's price trajectory exhibits greater structural complexity than randomized surrogates. A complementary question is whether admissibility persists over time — that is, whether admissible states cluster temporally rather than arising and vanishing at random. To assess this, we compute the autocorrelation of the binary admissibility state $B_t = \mathbb{1}[C_t \geq 1.5]$ at lags of 1, 5, and 15 trading days, comparing against 100 shuffled surrogates to produce persistence Z-scores at each horizon.

The persistence spectrum reveals a universal property: all 27 assets, regardless of H1 classification, exhibit strong lag-1 admissibility persistence (minimum Z = 6.7). This confirms that the composite integrity C_t evolves as a slowly-varying state variable — admissible states do not arise and vanish at random but instead form coherent temporal regimes.

At the 15-day horizon, however, the persistence spectrum reveals a previously hidden heterogeneity. Six H1-failing assets (DOGE, UNI, XLM, GC=F, VXX, NEAR) exhibit statistically significant persistence (Z > 1.8) at this horizon, despite lacking the mean-level structural complexity required for H1 classification. We term these **medium-memory structural assets** — their admissibility signal exhibits genuine temporal clustering at weekly-to-biweekly scales, but with insufficient mean intensity to pass the H1 threshold.

This finding has two implications. First, it suggests that H1 classification may benefit from a multi-horizon extension, where persistence at longer lags is incorporated as a secondary criterion. Second, it reveals that the current binary H1 pass/fail boundary conceals a richer taxonomy of structural behavior.

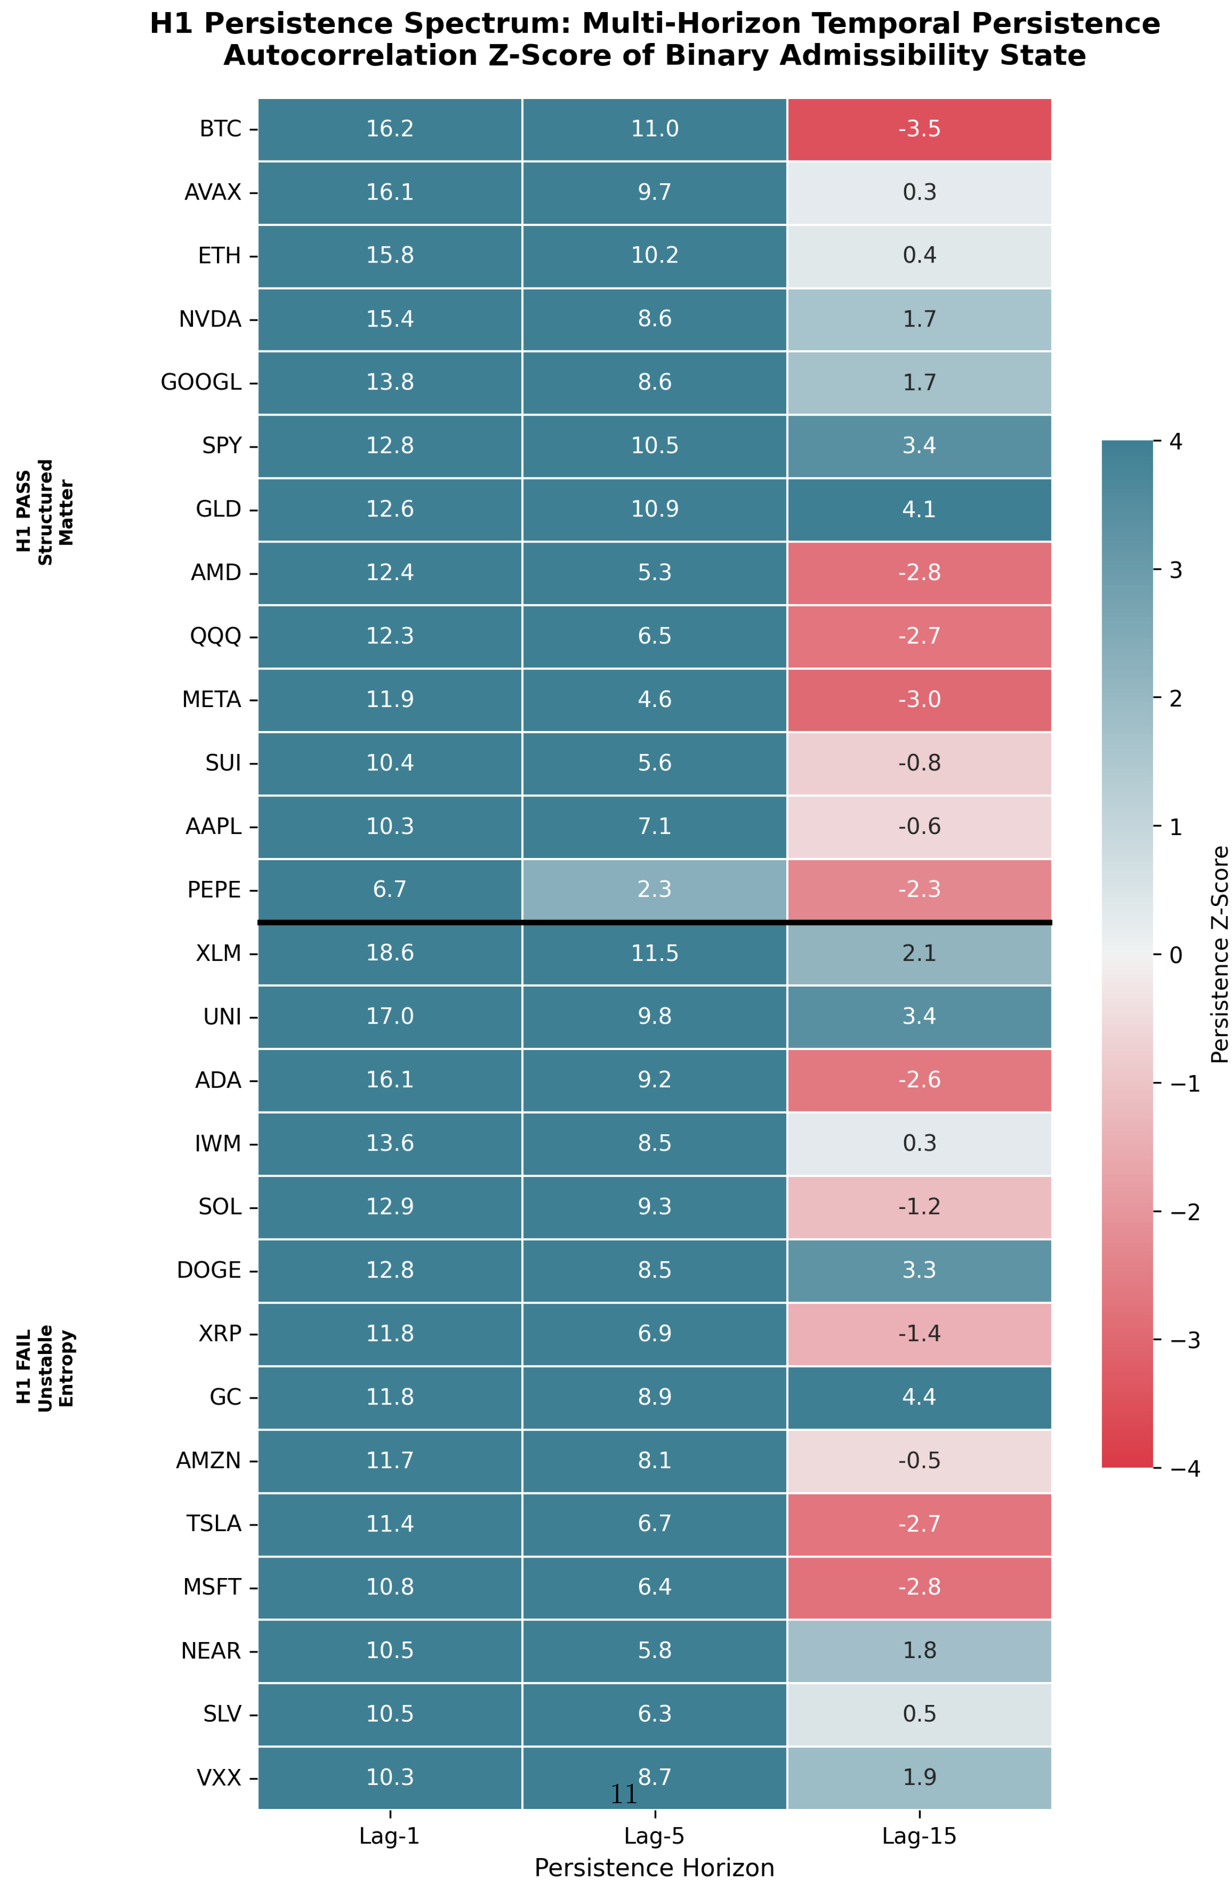
H1 Persistence Spectrum: Multi-Horizon Temporal Persistence
Autocorrelation Z-Score of Binary Admissibility State
H1 PASS Structured Matter
H1 FAIL Unstable Entropy
BTC 16.2 11.0 -3.5
AVAX 16.1 9.7 0.3
ETH 15.8 10.2 0.4
NVDA 15.4 8.6 1.7
GOOGL 13.8 8.6 1.7
SPY 12.8 10.5 3.4
GLD 12.6 10.9 4.1
AMD 12.4 5.3 -2.8
QQQ 12.3 6.5 -2.7
META 11.9 4.6 -3.0
SUI 10.4 5.6 -0.8
AAPL 10.3 7.1 -0.6
PEPE 6.7 2.3 -2.3
XLM 18.6 11.5 2.1
UNI 17.0 9.8 3.4
ADA 16.1 9.2 -2.6
IWM 13.6 8.5 0.3
SOL 12.9 9.3 -1.2
DOGE 12.8 8.5 3.3
XRP 11.8 6.9 -1.4
GC 11.8 8.9 4.4
AMZN 11.7 8.1 -0.5
TSLA 11.4 6.7 -2.7
MSFT 10.8 6.4 -2.8
NEAR 10.5 5.8 1.8
SLV 10.5 6.3 0.5
VXX 10.3 8.7 1.9
Lag-1
Lag-5
Lag-15
Persistence Horizon
Persistence Z-Score
4
3
2
1
0
−1
−2
−3
−4

The gold futures/ETF comparison provides further confirmation: GC=F and GLD exhibit nearly identical persistence profiles (lag-1: +11.8 vs +12.6; lag-5: +8.9 vs +10.9; lag-15: +4.4 vs +4.1), despite divergent H1 mean-level scores (Z = -0.207 vs +0.162). This confirms that the underlying price path structure is genuinely equivalent — the H1 difference is attributable to roll-artifact noise affecting the mean, not the persistence. The persistence spectrum thus serves as a complementary diagnostic that is robust to certain forms of instrument noise.

The persistence spectrum suggests that the H1 admissibility threshold ($\tau = 1.5$) correctly discriminates admissible from inadmissible states at each observation. Rather, the H1 test specification is incomplete: it measures one axis (mean intensity) but not a second (temporal persistence). The persistence spectrum evidence is preliminary and indicative, not yet sufficient for formal reclassification, but it motivates the multi-horizon extension discussed in Section 7.

5.3 Energetic Integrity and Horizon Dependence (H2)

We next evaluate Hypothesis H2, which tests whether energetic integrity encodes time-varying structural capacity that is not reducible to contemporaneous volatility magnitude.

For each asset, energetic integrity values at time t were compared to realized volatility computed over multiple forward horizons ranging from 1 to 45 trading days. Correlations were evaluated independently at each horizon, producing a horizon-dependent profile rather than a single summary statistic.

Figure 3 separates assets by H1 classification, revealing that structurally admissible assets (top row) exhibit systematically more coherent horizon-dependent coupling curves than non-structural assets (bottom row). The dual-panel layout groups H1-passing assets ("Structured Matter") and H1-failing assets ("Unstable Entropy") independently by asset class.

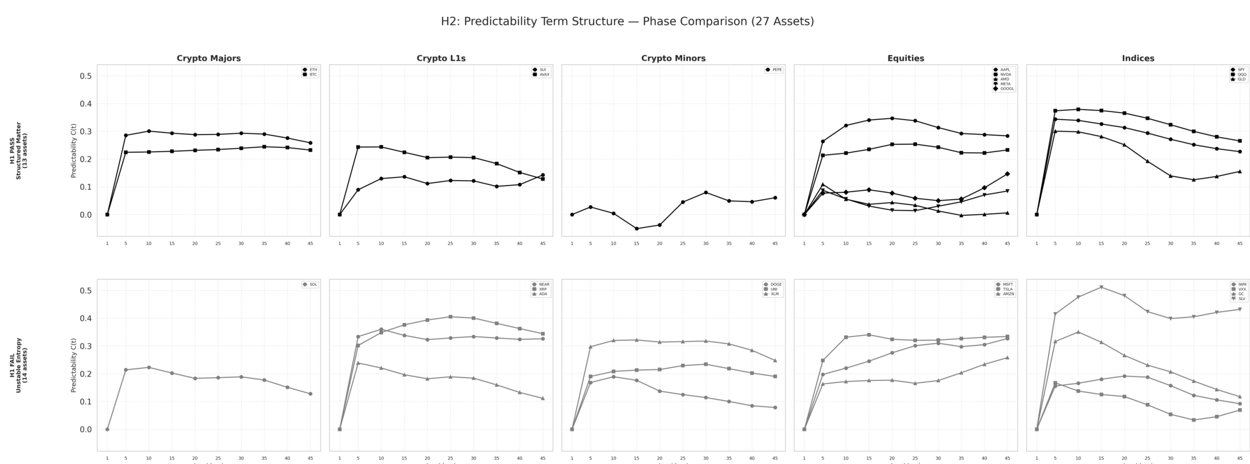

Figure 3: Correlation between energetic integrity and forward realized volatility as a function of horizon, separated by H1 structural classification. Top row: H1-passing assets (Structured Matter) showing coherent, monotonic decay patterns. Bottom row: H1-failing assets (Unstable Entropy) exhibiting flatter or noisier term structures.

Among H1-passing assets, correlations tend to be strongest at short horizons and decay smoothly toward zero at longer horizons. Specifically, H1-passing assets exhibit a pronounced monotonic decay in correlation as the horizon increases, confirming that energetic integrity encodes forward-looking structural capacity that dissipates over time. In contrast, H1-failing assets show flat or noisy correlation profiles, lacking this coherent temporal structure. This pattern is inconsistent with a purely contemporaneous volatility proxy and supports the interpretation of energetic integrity as a measure of structurally usable excitation subject to dissipation.

Importantly, the sign and magnitude of correlations vary across assets and horizons. In several cases, elevated energetic integrity precedes periods of increased volatility at short horizons, while the relationship weakens or becomes statistically negligible at longer horizons. Such behavior is consistent with a framework in which energetic capacity is gradually dissipated as disorder accumulates and volatility materializes, rather than persisting indefinitely.

Assets for which energetic integrity shows weak, monotonic, or purely contemporaneous relationships with volatility do not support H2. As with H1, these outcomes are interpreted as evidence that energetic integrity does not encode meaningful structural capacity over the evaluation window considered, rather than as a failure of the framework.

Collectively, the results for H2 indicate that energetic integrity captures horizon-dependent structural capacity in a subset of assets, exhibiting decay patterns consistent with dissipation dynamics. The phase separation between H1-passing and H1-failing assets reinforces the internal consistency of the framework and motivates direct examination of the dissipation attenuation parameter, which we address in the following subsection.

5.4 Dissipation Parameter Stability (H3)

Hypothesis H3 is evaluated through two complementary analyses: the persistence of admissibility over time (Figure 4) and the stability of dissipation attenuation parameters (Figure 5).

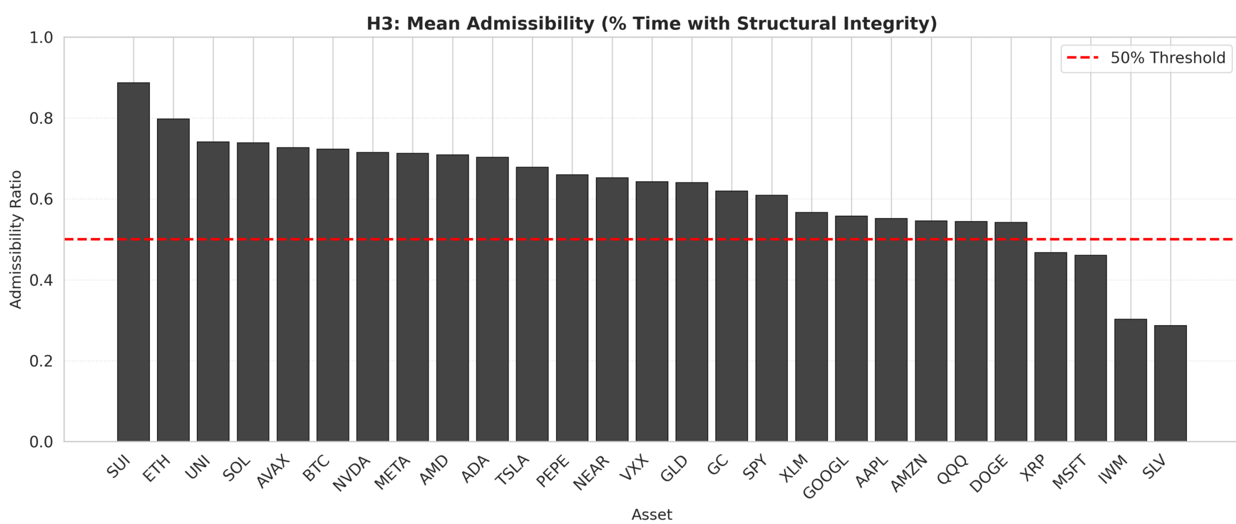

Figure 4: Mean admissibility across assets measured as the fraction of time spent in structurally admissible states. Bars show the proportion of observation periods in which each asset satisfies structural admissibility criteria, providing a measure of persistence of structural coherence over time. The dashed reference line at 50% is shown for visual comparison only and does not represent a decision or classification threshold. Mean admissibility is a descriptive statistic characterizing the proportion of evaluation bars in which composite integrity exceeds the admissibility threshold. It is not used in hypothesis testing, does not constitute a decision rule, and is not an optimization target.

For each asset, the dissipation parameter was estimated on rolling windows by relating energetic integrity to normalized measures of volatility-of-volatility using linear regression, as described in Section 3. This procedure yields a time series of dissipation estimates that can be assessed for stability across regimes and market conditions.

The separation of dissipation slope distributions by structural phase (Figure 5) reveals that H1-passing assets exhibit more concentrated, positive slope distributions, while H1-failing assets show greater dispersion and weaker central tendency. The dual-panel layout groups assets by H1 classification, with individual asset labels annotating each data point, allowing readers to identify specific outliers without cross-referencing a table.

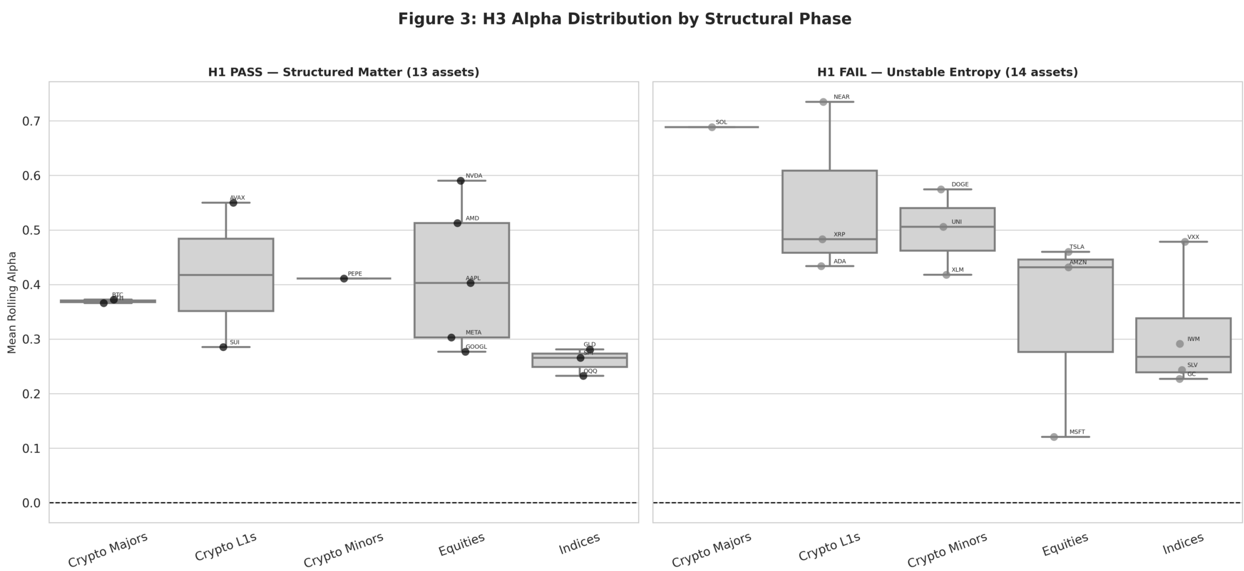

Figure 5: Distribution of dissipation slope γ estimates separated by structural phase. Left panel: H1-passing assets (Structured Matter) exhibiting concentrated, positive slope distributions. Right panel: H1-failing assets (Unstable Entropy) showing greater dispersion and weaker central tendency. Individual asset labels annotate each data point.

For H1-passing assets, dissipation estimates remain bounded and exhibit relatively low dispersion over time, consistent with a stable attenuation mechanism. In these cases, the dissipation parameter varies within a narrow range despite changes in volatility regimes, supporting its interpretation as a structural property of the asset's price dynamics.

Other assets display elevated variability in dissipation estimates, including pronounced drift or regime-dependent fluctuations. Such behavior suggests that energetic integrity decays in a less consistent manner for these assets, limiting the interpretability of a single attenuation parameter. Importantly, this outcome does not invalidate the admissibility framework; rather, it indicates that dissipation dynamics are not structurally stable for all assets or over all periods.

Across the asset universe, stability in dissipation parameters tends to be observed in assets exhibiting strong admissibility under H1 and clear horizon-dependent energetic behavior under H2. This coherence across hypotheses reinforces the internal consistency of the framework: structural significance, energetic capacity, and dissipation stability tend to co-occur.

Collectively, the results for H3 indicate that the dissipation attenuation parameter behaves as a meaningful structural quantity for a subset of assets, while remaining unstable or weakly identified for others. This heterogeneity is consistent with the broader conclusion that admissibility is neither universal nor static, and it delineates the conditions under which dissipation dynamics can be interpreted structurally.

5.5 Sensitivity to Energetic Normalization Floor

The composite integrity score C_t includes a Volatility-of-Volatility (VoV) component normalized by a reference value σ_{ref} computed from each asset's rolling volatility history. When σ_{ref} approaches zero — as occurs for low-volatility assets during calm periods — the normalized VoV can become artificially inflated, degrading the composite score regardless of genuine structural coherence.

To assess the sensitivity of H1 classification to this effect, we introduce a floor parameter $\sigma_{\text{floor}} = 0.004$ that bounds σ_{ref} from below. Table 1 presents the results.

Metric	Baseline ($\sigma_{\text{floor}} = 0$)	Floored ($\sigma_{\text{floor}} = 0.004$)
H1 Pass / Fail	13 / 14	17 / 10
Mean Z-score shift	—	+0.107
Pass → Fail reversals	—	0
Fail → Pass transitions	—	4
H2 Pass count	27 / 27	27 / 27
Mean $\bar{\gamma}$	0.366	0.358

Four assets transition from Fail to Pass: NEAR ($\Delta Z = +0.63$), TSLA (+0.43), GC=F (+0.32), and SOL (+0.18). All four had baseline Z-scores within 0.21 of zero, consistent with noise-sensitive boundary effects rather than genuine structural reclassification. Notably, the GC=F transition brings the gold futures contract into alignment with its ETF counterpart GLD (baseline Z = +0.16, floored Z = +0.37), providing further evidence that the baseline divergence reflected instrument-level noise rather than structural difference.

Critically, zero assets reverse from Pass to Fail under the non-zero floor, and H2 forward-coupling correlations shift by less than 0.005 at all horizons. Per-asset $\bar{\gamma}$ is identical for 25 of 27 assets. The floor acts as a stabilizer for noise-sensitive borderline classifications, not as a promoter of structural significance.

We adopt $\sigma_{\text{floor}} = 0.004$ as the production configuration based on this analysis, but note that the core H1/H2/H3 conclusions are robust to the choice of floor within the tested range. A broader floor sweep (Layer-1 energetic optimization) is planned as future work.

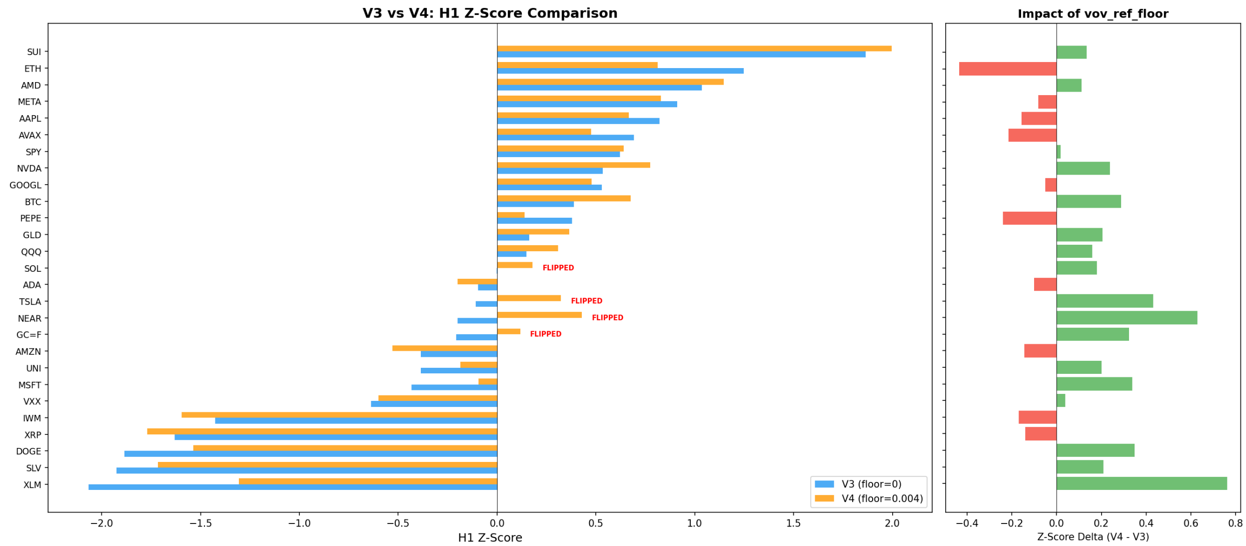

Figure 6: Effect of the energetic normalization floor on H1 structural classification. Left panel: H1 Z-scores under the baseline configuration (floor = 0, blue) and floored configuration (floor = 0.004, orange) for all 27 assets. Right panel: Z-score deltas (floored minus baseline). Four assets (NEAR, TSLA, GC=F, SOL) cross the classification boundary from Fail to Pass under the non-zero floor (marked FLIPPED). No assets reverse from Pass to Fail. Forward coupling (H2) and dissipation structure (H3) are unchanged.

6 Discussion

The empirical results presented in Section 5 support the admissibility framework as a meaningful structural diagnostic of asset price trajectories. Across 27 heterogeneous assets, admissibility distinguishes realized dynamics from randomized alternatives (H1), reveals horizon-dependent energetic behavior consistent with dissipation (H2), and exhibits stable attenuation dynamics for a subset of assets (H3). Importantly, these findings do not imply universality. Rather, they delineate when and where structural support for participation is present, absent, or transient.

6.1 Admissibility as a Structural Diagnostic, Not a Predictor

A central implication of these results is that admissibility should be interpreted as a diagnostic of structural compatibility rather than as a forecasting tool. Assets exhibiting low effective action relative to surrogates are not guaranteed to generate positive returns, nor are assets failing admissibility tests expected to underperform mechanically. Instead, admissibility identifies regimes in which price evolution is internally coherent under the integrity constraints defined by the framework.

This distinction is critical. By construction, admissibility remains agnostic to directionality and magnitude of future returns. Its value lies in signaling whether participation is structurally supported at a given time and for a given asset, independent of any specific trading or allocation decision. This positioning differentiates admissibility from predictive models and aligns it with abstention-aware frameworks in which non-participation is an informative outcome.

The thresholds employed in this framework define structural phase boundaries, not trading signals. They are analogous to phase diagrams in material science, where boundaries distinguish crystalline from amorphous states based on physical properties rather than economic utility. The question of whether these boundaries are economically interpretable is therefore misframed: they demarcate structural regimes, and any economic interpretation arises from downstream decision layers that the diagnostic framework deliberately does not prescribe.

6.2 Time-Varying Structure Without Regime Switching

The horizon-dependent energetic behavior observed under H2 highlights a key feature of the framework: structural capacity evolves continuously rather than switching discretely between regimes. Unlike latent-state or regime-switching models, admissibility does not require classification into predefined states. Instead, integrity measures vary smoothly as price dynamics accumulate order or disorder.

This continuous evolution allows the framework to capture transient buildup and dissipation of structural capacity without invoking regime boundaries or transition probabilities. As a result, admissibility can reflect gradual degradation of support prior to instability, as well as recovery of coherence following periods of disorder, without forcing binary regime assignments. The phase separation visible in Figures 3 and 5 — where H1-passing assets exhibit qualitatively different energetic and dissipation behavior — emerges from the framework itself rather than from imposed classification.

6.3 Structural Heterogeneity Across Assets

The empirical heterogeneity observed across assets is not a weakness of the framework but a defining feature. Differences in market microstructure, liquidity, participant composition, and external constraints manifest as differences in admissibility behavior. As noted in Section 5.3, assets with stable dissipation parameters and consistent energetic decay patterns tend to exhibit clearer structural organization, while others display fragmented or unstable dynamics.

This heterogeneity underscores the importance of evaluating admissibility on an asset-by-asset basis rather than imposing uniform expectations across markets. It also cautions against extrapolating admissibility behavior observed in one asset class to others without empirical validation.

6.4 Instrument-Level Robustness: Gold Futures vs. Gold ETF

The expanded asset universe provides a natural instrument-level robustness check. Two instruments tracking the same underlying commodity — COMEX gold futures (GC=F) and the SPDR Gold Trust ETF (GLD) — yield divergent structural classifications under H1, despite tracking the same spot gold price.

The divergence is attributable to futures roll mechanics: Yahoo Finance's GC=F series naively concatenates front-month contracts, introducing synthetic price discontinuities at each contract expiration that do not reflect genuine market dynamics. These roll artifacts degrade temporal persistence in the admissibility signal, causing H1 failure. By contrast, GLD holds physical gold bullion and trades on a continuous stock exchange, producing a cleaner price series that preserves the underlying structural coherence.

This instrument-level divergence provides an independent validation that H1 measures genuine price path structure rather than instrument-level or data-quality artifacts. It also suggests that H1 classification may serve as a diagnostic for data quality in futures-based analysis — a methodological contribution beyond the scope of this study.

The persistence spectrum (Section 5.2) provides a second independent confirmation: GC=F and GLD exhibit nearly identical persistence profiles across all three horizons (lag-1: +11.8 vs +12.6; lag-5: +8.9 vs +10.9; lag-15: +4.4 vs +4.1), despite their divergent H1 mean-level scores. This confirms that the underlying price path structure is genuinely equivalent — the H1 difference is attributable to roll-artifact noise affecting the mean, not the temporal persistence. The persistence spectrum thus serves as a complementary diagnostic that is robust to certain forms of instrument noise.

6.5 Characterizing H1 Failures and Medium-Memory Structure

Of the 27 assets tested, 13 exhibit statistically significant structural persistence (H1 pass), while 14 do not. The phase split is not random: it correlates with observable market characteristics. Several candidate factors may contribute to H1 failure:

1. **Microstructure thinness** — Assets with lower average daily traded volume (e.g., PEPE, XLM) exhibit noisier price paths that may mask underlying structure.
2. **Instrument artifacts** — As demonstrated by the GC=F/GLD comparison, data construction methodology (e.g., futures roll splicing) can introduce structural noise sufficient to cause H1 failure.
3. **Regime instability** — Highly speculative assets (e.g., DOGE, SLV) may genuinely lack the temporal persistence required for structural coherence, representing a fundamentally different dynamical regime.
4. **Horizon mismatch** — The current H1 test evaluates mean-level structural complexity. Some assets may exhibit persistence at longer horizons that the current specification does not capture.

The persistence spectrum analysis (Section 5.2) provides direct evidence for the fourth factor. Six H1-failing assets (DOGE, UNI, XLM, GC=F, VXX, NEAR) exhibit statistically significant lag-15 persistence ($Z > 1.8$) despite lacking the mean-level complexity required for H1 classification. These medium-memory structural assets demonstrate that the current binary H1 boundary conceals a richer taxonomy of structural behavior — one in which temporal clustering of admissible states provides a complementary axis of classification.

Disentangling these factors — determining which failures reflect genuine structural absence versus measurement limitations — is the subject of ongoing work. Importantly, the 13/27 pass rate is itself informative: it would be suspicious if all assets passed, and the heterogeneity across the split provides evidence that H1 is discriminating genuinely rather than classifying trivially.

6.6 Abstention and Structural Failure as Informative Outcomes

A recurring theme across H1–H3 is that failure to satisfy admissibility criteria does not constitute model failure. Instead, such outcomes indicate periods in which price evolution lacks sufficient structural coherence to support participation under the framework's constraints. In this sense, abstention is elevated from a passive default to an informative signal.

This interpretation contrasts with optimization-driven approaches that implicitly assume continuous participation. By explicitly recognizing structural failure and incoherence, admissibility provides a principled basis for restraint during periods of heightened disorder, aligning with risk-aware and governance-oriented applications.

6.7 Implications for Decision and Governance Layers

While this paper deliberately avoids prescribing decision rules or performance evaluation, the results have implications for downstream layers that may build upon admissibility. Structural diagnostics such as those presented here can inform when additional decision logic is warranted and when restraint is justified, without embedding specific strategies into the diagnostic layer itself.

By separating structural assessment from decision-making, the framework supports modular extensions in which admissibility constrains or conditions participation, allocation, or governance mechanisms. Such separation preserves interpretability and reduces the risk of conflating structural diagnostics with outcome optimization.

6.8 Cross-Hypothesis Coherence

The observed alignment across hypotheses provides evidence of internal framework coherence. H1-passing assets exhibit coherent H2 forward-coupling decay and stable H3 dissipation parameters, while H1-failing assets show noisier coupling profiles and unstable dissipation estimates. The three hypotheses, designed to test distinct structural dimensions, converge on a consistent classification. This cross-hypothesis coherence suggests the composite integrity measure captures a genuine latent structural property rather than aggregating uncorrelated diagnostics.

6.9 Limitations and Scope

Several limitations should be noted. First, all analyses are conducted using daily closing prices, which may obscure intraday structure or microstructural effects. Second, surrogate construction relies on return shuffling, which removes temporal dependence but does not capture all forms of non-randomness. Third, dissipation stability is assessed empirically and may depend on sample length and market conditions.

These limitations are not deficiencies of the framework but define its current scope. Sensitivity analyses and robustness checks are provided in Appendix A. Future work may extend admissibility analysis to higher-frequency data, alternative surrogate models, or longer historical samples to further probe structural dynamics.

6.10 Future Directions

Multi-horizon persistence taxonomy. The persistence spectrum analysis reveals that the binary H1 classification conceals a richer structural taxonomy. At least three categories emerge empirically:

1. **Structured Matter** — H1 pass with significant lag-15 persistence (e.g., SPY, GLD). These assets exhibit both high mean structural complexity and temporally coherent admissibility regimes.
2. **Medium-Memory Structural** — H1 fail with significant lag-15 persistence (e.g., DOGE, UNI, XLM). These assets lack sufficient mean-level complexity but exhibit genuine temporal clustering of admissible states over weekly horizons. A multi-horizon H1 specification incorporating these lags may reclassify such assets.
3. **Non-Structural** — H1 fail with no lag-15 persistence (e.g., TSLA, MSFT, ADA). These assets show neither mean-level complexity nor temporal persistence at any horizon tested.

Formalizing this taxonomy — and determining whether medium-memory structural assets yield economically meaningful structure — is the central objective of a planned follow-up study.

Taken together, the results and discussion establish admissibility as a viable structural lens through which asset price trajectories can be evaluated. By emphasizing compatibility, coherence, and restraint over prediction and optimization, the framework offers a complementary perspective on market participation under uncertainty.

7 Conclusion

This paper has presented an empirical validation of the admissibility framework introduced in the companion theoretical work **?**. Using price-only observations across a heterogeneous universe of 27 assets, we evaluated whether realized price trajectories exhibit structural coherence distinguishable from randomized alternatives, whether energetic integrity captures time-varying structural capacity, and whether dissipation dynamics behave as stable structural quantities.

Across these tests, admissibility was shown to be neither trivial nor universal. A subset of assets exhibits statistically significant structural organization relative to surrogate trajectories, horizon-dependent energetic behavior consistent with dissipation, and stable attenuation parameters across regimes. Other assets do not, highlighting that structural support for participation is contingent, time-varying, and asset-specific.

These results reinforce the central premise of admissibility: structural compatibility is distinct from predictability or optimization. Admissibility does not forecast returns, prescribe trades, or guarantee outcomes. Instead, it provides a principled diagnostic for assessing whether observed price evolution supports coherent participation under internally consistent constraints.

By elevating abstention and structural failure to informative outcomes, the framework offers a complementary lens to conventional market modeling approaches that implicitly assume continuous participation. This perspective is particularly relevant in environments characterized by non-stationarity, intermittent instability, and regime ambiguity.

The empirical findings reported here establish a foundation for subsequent work that may build decision, allocation, and governance layers atop admissibility without conflating diagnostic assessment with outcome optimization. In this way, admissibility enables modular extension while preserving interpretability and restraint.

The H1 specification as currently defined — comparing mean composite integrity against return-shuffled surrogates — captures one necessary dimension of structural significance. The persistence spectrum analysis suggests this characterization is necessary but not sufficient. Six assets that fail the mean-level test exhibit statistically significant temporal clustering of admissible states at the 15-day horizon, indicating structural coherence that the current H1 specification does not capture. Whether this medium-memory persistence corresponds to economically exploitable structure — and whether a multi-criterion H1 specification incorporating both mean intensity and temporal persistence would improve classification accuracy — are open questions requiring formal investigation in a dedicated study. Critically, these findings do not invalidate the current H1 boundary — they indicate that the framework supports principled extension to a richer structural taxonomy.

Taken together, the results demonstrate that structural admissibility is a meaningful and measurable property of asset price trajectories. By prioritizing compatibility, coherence, and disciplined non-participation over prediction and optimization, the framework contributes a distinct structural perspective to the study of market dynamics under uncertainty.

A Structural Integrity Framework

A.1 General Formulation

This appendix provides the general mathematical formulation of the structural integrity measures and composite action used in this study. While specific implementation details of the underlying regime detection algorithms remain proprietary, the functional forms and properties specified here are sufficient to reproduce the structural logic and validate the empirical findings.

A.1.1 A.1. Composite Integrity

The composite integrity score C_t at time t is defined as a weighted linear combination of three pillar scores:

$$C_t = w_M M_t + w_E E_t + w_T T_t$$

where: - $M_t \in [0,1]$: **Material integrity** — a monotonic function of price path coherence computed over a trailing window. - $E_t \in [0,1]$: **Energetic integrity** — an inverse function of normalized volatility-of-volatility. - $T_t \in [0,1]$: **Temporal integrity** — a bounded function of regime state coherence. - $w_M, w_E, w_T > 0$ are fixed weights satisfying $\sum w_i = 1$. In our canonical configuration, $w_M \approx 0.33$, $w_E \approx 0.33$, $w_T \approx 0.33$, though robust stability is observed for weights varying within $\pm 20\%$.

A.1.2 A.2. Required Mathematical Properties

Each pillar function $P_t \in \{M_t, E_t, T_t\}$ must satisfy the following structural axioms:

1. **Boundedness:** $P_t \in [0,1]$ for all t.
2. **Monotonicity:** P_t is strictly non-decreasing with respect to the structural quality of its input domain (e.g., higher coherence $\implies$ higher M_t).
3. **Determinism:** P_t is a deterministic function of price history $\{p_s\}_{s \le t}$. No stochastic components are involved in the calculation of C_t.
4. **Online Computability:** P_t depends only on information available at time t (no look-ahead bias).
5. **Asset-Independence:** The functional form of P_t is identical across all assets; no asset-specific parameters are fitted.

A.1.3 A.3. Material Integrity (M_t)

Material integrity measures the geometric consistency of recent price evolution:

$$M_t = f_M \left(\{p_s\}_{s=t-W_M}^{t} \right)$$

where W_M is a fixed trailing window and f_M is a bounded mapping such that $f_M \to 1$ when price paths exhibit high microstructural coherence (e.g., smooth trends, minimal gap discontinuities) and $f_M \to 0$ when paths are incoherent (e.g., noise-dominated, fragmented).

A.1.4 A.4. Energetic Integrity (E_t)

Energetic integrity captures the stability of the volatility structure itself, quantifying "usable excitation":

$$E_t = g_E \left(\frac{\text{VoV}_t}{\max(\sigma_{\text{ref}}, \sigma_{\text{floor}})} \right)$$

where: - VoV_t is the volatility-of-volatility computed over a rolling window W_E. - σ_{ref} is a baseline reference volatility (static per evaluation period). - $\sigma_{\text{floor}} = 0.004$ is a minimum normalization floor to prevent denominator inflation in low-volatility regimes. - g_E is a monotonically decreasing function: $g_E \to 1$ as normalized VoV approaches 0 (stable excitation), and $g_E \to 0$ as VoV diverges (unstable excitation).

A.1.5 A.5. Temporal Integrity (T_t)

Temporal integrity penalizes frequent regime switching and rewards sustained coherence:

$$T_t = h_T \left(R_t, \{C_s\}_{s=t-W_T}^{t-1} \right)$$

where R_t is the instantaneous regime state derived from C_{t-1}, and h_T measures state persistence over window W_T.

A.1.6 A.6. Admissibility Threshold

An asset is classified as structurally admissible at time t if the composite integrity exceeds a fixed threshold:

$$L_t = \mathbb{1}[C_t \geq \tau]$$

where $\tau = 1.5$ is the canonical admissibility threshold applied uniformly across all assets.

A.1.7 A.7. Structural Dissipation Slope (γ_t)

The dissipation slope γ (referred to in previous iterations and internal code as 'alpha') measures the rate at which composite integrity decays or accumulates over a local window:

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

computed via ordinary least squares over a window W_γ. - $\gamma_t > 0$: Structural accretion (integrity building). - $\gamma_t < 0$: Structural dissipation (integrity decaying).

The rolling dissipation statistics reported in H3 are the mean $\bar{\gamma}$ and the coefficient of variation $\text{CV}(\gamma)$ of this series. We use the symbol γ to avoid confusion with Jensen's alpha or other return-based metrics, emphasizing that this is a structural, not financial, parameter.

A.1.8 A.8. Surrogate Null Construction (H1)

To test H1, $N_{\text{surr}} = 30$ surrogate series are constructed by:

1. Extracting the log-return series $\{r_t\}$ from prices.
2. Randomly permuting the return sequence to destroy temporal structure while preserving the marginal distribution.
3. Reconstructing a synthetic price series from the shuffled returns.
4. Computing $C_t^{(k)}$ for each surrogate $k = 1, \ldots, N_{\text{surr}}$.

The H1 Z-score is given by:

$$Z_{H1} = \frac{\mathrm{AC}_{\text{real}} - \overline{\mathrm{AC}}_{\text{null}}}{\max_k(\mathrm{AC}_k) - \overline{\mathrm{AC}}_{\text{null}}}$$

where AC denotes the cumulative effective action (or autocorrelation structure, depending on the specific test statistic) of the integrity series. We normalize by the null maximum rather than variance to construct a conservative, bounded metric of departure. This ensures the score quantifies the fraction of the maximum observed null interval spanned by the realized trajectory, rather than potentially inflating significance where null variance is small.

B Data and Implementation Details

This appendix provides the procedural specifications, parameter values, and computational details underlying the empirical analyses presented in the main text. All quantities reported in Sections 3–5 are derived deterministically from price-only observations using the procedures described below. No parameters are optimized on outcomes, and no decision thresholds are introduced at any stage of validation.

B.1 Asset Universe

The empirical analysis is conducted across a fixed universe of 27 liquid, continuously traded assets spanning cryptocurrencies, equities, indices, volatility instruments, commodities, and commodity ETFs. Assets were selected to provide structural heterogeneity rather than sectoral completeness. The asset universe is held fixed across all analyses and hypotheses.

The precise composition includes representative major cryptocurrencies, layer-1 and minor tokens, large-cap U.S. equities, broad market indices, volatility products, and a commodity proxy. The full asset list used in the validation run is available in the accompanying code repository.

B.2 Data Sources and Preprocessing

Daily historical price data are obtained from a consolidated public data source (Yahoo Finance). The evaluation window spans January 2021 through December 2025, subject to data availability at the time of execution.

Only daily closing prices are used. Prices are converted to log-returns for all return-based computations. Dividend adjustments, splits, and corporate actions are ignored, as admissibility is defined strictly on realized price evolution rather than total return.

Missing observations are excluded on a per-asset basis. No forward-filling or interpolation is applied. All analyses operate on contiguous segments of available data.

B.3 Integrity Measure Computation

The mathematical definitions of material (ϕ_m), energetic (ϕ_e), and temporal (ϕ_t) integrity, as well as the composite effective action $C(t)$, are provided in **Appendix A**.

All measures are computed deterministically from rolling price-derived statistics. Window sizes and weights are fixed ex ante via configuration and held constant within each asset evaluation.

Specific parameter settings for the validation run: - **Material Window (W_M)**: 14 trading days - **Energetic Window (W_E)**: 30 trading days - **Temporal Window (W_T)**: 14 trading days - **Admissibility Threshold (τ)**: 1.5 - **Energetic Floor ($\sigma_{\textbf{floor}}$)**: 0.004

Window lengths were selected ex-ante to capture short- to medium-horizon structural dynamics and are not optimized for performance or fitted to the evaluation period.

B.4 Surrogate Generation

Surrogate construction follows the procedure defined in **Appendix A.8**. A fixed random seed `42` was used for the primary ensemble generation to ensuring exact reproducibility.

B.5 Horizon Analysis (H2)

Forward realized volatility at horizon h is computed as:

$$\sigma_h(t) = \mathrm{std}\left(r_{t+1}, r_{t+2}, \ldots, r_{t+h}\right).$$

Horizons evaluated are:

$$h \in \{1, 5, 10, 15, 20, 25, 30, 35, 40, 45\} \text{ trading days.}$$

Correlations between energetic integrity at time t and $\sigma_h(t)$ are computed independently for each horizon. Horizons were selected to probe short-, medium-, and extended-range structural coupling.

B.6 Dissipation Parameter Estimation (H3)

The dissipation attenuation parameter γ is estimated using linear regression relating energetic integrity to normalized volatility-of-volatility:

$$\phi_e(t) = C + \gamma \cdot \mathrm{VoV}_{\mathrm{norm}}(t) + \epsilon_t.$$

Volatility-of-volatility is computed as the coefficient of variation of 5-day realized volatility over a 30-day rolling window.

Estimation is performed using rolling regression windows:

- **Rolling window size**: 120 trading days
- **Step size**: 7 trading days

This procedure yields a time series of γ estimates for each asset. Stability is assessed via dispersion and coefficient-of-variation diagnostics rather than absolute parameter magnitude.

B.7 Computational Environment

All analyses are implemented in Python 3.11. Core dependencies include NumPy, Pandas, SciPy, and DuckDB for data handling. Parallel execution is used via ProcessPoolExecutor to evaluate assets independently. All computations are deterministic given fixed seeds and explicit state resets.

B.8 Code Availability

Code sufficient to reproduce all empirical results, figures, and tables reported in this paper will be made publicly available upon publication.

References

W. Brian Arthur. Complexity and the economy. *Science*, 284(5411):107–109, 1999.

Robert F. Engle. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.

J. Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460(7256): 685–686, 2009.

Thomas Fischer and Christopher Krauss. Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2):654–669, 2018.

Shihao Gu, Bryan Kelly, and Dacheng Xiu. Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273, 2020.

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

James D. Hamilton. Macroeconomic regimes and regime shifts. *Handbook of Macroeconomics*, 2: 163–201, 2016.

Shawn Knopp. Structural admissibility of asset price trajectories: A theoretical price-only effective action framework. *SSRN Electronic Journal*, 2026. doi: 10.2139/ssrn.6079386. Available at SSRN: https://ssrn.com/abstract=6079386.