

A ROUGH THEORY OF MARKETS LATENT CONTAGION, MARKET PHASES, AND MARKET-VARIANCE RISK

J. VIDAL LLAURADÓ 

ABSTRACT. Microscopic volatility can remain continuous while market-level observations become rough, jump-like at coarse resolution, and only partially hedgeable. This paper studies that separation in an N -asset Gaussian Volterra market with directed latent-contagion kernels and aggregate market-volatility functional M . Market-scale admissibility is selected by the projected experiment, not imposed as a rough subclass. An active off-diagonal channel is admissible exactly when it remains source-screened visible, carries nonzero market loading, and admits a genuinely shrinking projected local normalization; under the high-frequency market array this is equivalent to $H_{ij} < 1/2$ together with source-screened visibility.

Within the admissible class, smooth directional contagion is locally degenerate at high frequency. Rough channels are the stable observable high-frequency contagion phase. If the market-dominant admissible channels aggregate without leading-order cancellation, the aggregate market-volatility exponent is the roughest active source-screened exponent and is strictly below one half. The same dominant edges define a screened synchronization matrix on the $N \times N$ network. Crossing its Perron threshold aligns latent stress along a common mode and turns local rough propagation into macroscopic amplification. Under threshold-triggered activation, synchronized bursts whose widths contract faster than observation resolution converge for coarse observers to a stochastic jump limit, although the primitive market paths remain continuous.

The synchronized mode also carries the pricing and hedging obstruction. For integrated market variance $Q_M = \int_0^T M(t)^2 dt$, compression to any visible pricing summary leaves a lower-bounded latent convexity gap, and visible hedges for variance-linked claims satisfy an explicit lower bound on residual risk whenever the synchronized latent mode retains conditional variance. The reduction from Volterra primitives to these claims is derived from the canonical Perron projection of the latent envelope and an aggregate sufficient-state sigma-field; no pricing-side factor split is postulated. Leading-order short-maturity skew statistics of the type isolated in [1] are visible $\mathbb{Q}$ -summaries, but they remain incomplete relative to latent market-variance convexity. The resulting phase classification connects asymmetric latent contagion, roughness, Perron synchronization, coarse-observation jump limits, and public compression of market-variance risk.

1. INTRODUCTION

Rough-volatility models give a tractable language for fractional memory and short-maturity scaling [2, 3]. They usually begin with a rough volatility law and then study pricing, simulation, filtering, or calibration [4, 5, 6]. That route does not decide when a market must become rough because its visible high-frequency experiment admits directed latent-contagion channels. The present construction starts from the directed market experiment. Roughness is the consequence of source-screened local admissibility, not the primitive regularity assumption.

The observable market object is the aggregate market-volatility functional M . The latent object is an $N \times N$ Gaussian Volterra system whose off-diagonal kernels transmit volatility from source drivers to target assets. A channel is economically active only when its market loading is nonzero. It is statistically admissible only when it remains visible after source screening and admits a shrinking projected local normalization in the minimax experiment of [7]. Under the high-frequency market array, these requirements select rough off-diagonal channels with

Date: April 2026.

2020 Mathematics Subject Classification. 60G22, 62M15, 91G80, 62F12.

Key words and phrases. rough volatility, latent contagion, Gaussian Volterra markets, market phases, Perron synchronization, market-variance risk, visible pricing and hedging.

DOI: 10.2139/ssrn.6616121.

$H_{ij} < 1/2$. Smooth channels do not disappear from the primitive model; they fail to generate a nondegenerate shrinking projected experiment.

This separates roughness from specification. Section 3 characterizes the projected-minimax admissible class before any aggregate roughness theorem is used. Source-screened visibility, nonzero market loading, and $H_{ij} < 1/2$ are necessary and sufficient for an active off-diagonal channel to produce a finite nonzero projected local experiment. The local information geometry therefore selects the market-scale class rather than imposing roughness as a convenient subclass. The projected-score evidence in [7] supplies the corresponding finite-panel diagnostic for realized-volatility data.

Three literatures determine the boundary of the construction. Rough-volatility models supply the Volterra scaling language, but in standard pricing applications the rough law is specified rather than recovered from a market experiment [2, 4, 8]. Network contagion and connectedness methods summarize spillovers, variance decompositions, and crisis comovement, but they usually operate on visible dependence objects rather than on a source-screened local likelihood geometry [9, 10, 11]. The latent-contagion observability strand separates source effects from target-only confounding and identifies the projected local information geometry [1, 12, 7]. The present paper places these three roles in one N -asset Volterra market: the Volterra kernels define the primitive channels, the projected experiment selects the admissible active edges, and the market aggregation turns those edges into roughness, synchronization, and public compression.

The comparison with the classic “no contagion, only interdependence” argument of [11] is at the level of the observed object. Forbes and Rigobon ask whether crisis comovement in returns exceeds visible interdependence after correcting for heteroskedasticity. Here the object is a directional latent-volatility channel that survives source screening at the local scale. A synchronized latent network can amplify market-wide stress while remaining close to interdependence in reduced-form correlation measures. Visible crisis correlation alone does not certify contagion. The theorem instead isolates latent directed channels whose aggregate effects are roughness, Perron synchronization, and incomplete visible pricing and hedging summaries.

The comparison with incomplete-information pricing is different. The hidden object is the synchronized volatility mode generated by the directed contagion network, rather than a private signal or an exogenous equilibrium state. Prices are treated as public compressions, not as noisy measurements of an external factor. Visible pricing summaries are incomplete because public compression removes part of a latent rough mode that remains only partially screenable under $\mathbb{Q}$. The compression statements below are instantiated for CoVaR-type tail summaries [13], SRISK-style shortfall functionals [14], and Diebold–Yilmaz connectedness measures [9, 10].

The empirical layer is deliberately narrow. Appendix B reports an Oxford-Man Perron illustration with no unused-sample trading or theorem-validation status. The theoretical input comes from the projected observability results of [7]. Admissibility fixes the active local experiment. Roughness is the resulting high-frequency phase. The screened synchronization matrix then separates rough local propagation from Perron synchronization. Above the Perron threshold, the dominant rough edges align latent stress along a common mode. Under threshold-triggered activation, that mode can become jump-like for coarse observers. The same mode carries the market-variance convexity and hedge shortfall left by visible compression.

The pricing statement stays at the compression level. The threshold analysis in [1] isolates visible $\mathbb{Q}$ -side statistics, such as leading-order short-maturity skew, that can fail to identify latent contagion even when the physical law remains distinguishable. The present paper extends that asymmetry to market scale. Visible pricing summaries may be well defined and fully observable while latent market-variance convexity and hedging exposure remain outside the visible sigma-field.

The results follow the objects they identify. Section 3 characterizes the active off-diagonal channels that are projected-minimax admissible. Within that class, Theorem 4.2 proves that market-dominant admissible asymmetry forces $H_M < 1/2$. Theorem 5.6 separates local roughness from systemic synchronization through the Perron threshold of the screened synchronization matrix. Theorem 6.10 and Corollary 6.11 characterize the jump-like limit produced by synchronized rough amplification under activated concentration. Theorem 7.3 derives the aggregate

sufficient-state sigma-field from primitive coupling, canonical Perron projection, and off-Perron screening. Proposition 7.6 extracts the centered synchronized factor content of integrated market variance. Theorem 7.11 then proves the visible-price and visible-hedge compression bounds. Theorem 6.15 collects these statements into a market-scale classification of contagion phases.

The phase classification is deliberately narrow. A *smooth-degenerate* regime has only smooth active off-diagonal channels and no nondegenerate shrinking projected local contagion experiment. A *rough-observable* regime has admissible dominant asymmetry and rough aggregate volatility but no synchronized amplification. A *rough-synchronized* regime crosses the screened Perron threshold and aligns stress along a common latent mode. An *activated synchronized* regime adds threshold-triggered concentration of that mode, so continuous latent synchronization becomes jump-like at coarse observation scales. These phases classify directional contagion, not all market dislocations. They do not rule out rough idiosyncratic diagonal channels. What the smooth-degenerate phase excludes is a nondegenerate off-diagonal projected contagion experiment.

The remaining sections follow this construction. Section 2 introduces the N -asset directed Volterra market. Section 3 derives local observability and admissibility at market scale, and Section 4 proves the aggregate roughness law. Section 5 studies the screened synchronization matrix and its Perron threshold. Section 6 derives the endogenous jump appearance at coarse observation scales. Section 7 develops the pricing, convexity, and hedging compression results, and Section 8 concludes.

2. AN N -ASSET DIRECTED VOLTERRA MARKET

The primitive object is a directed Volterra volatility network. Its off-diagonal kernels are the candidate contagion channels; market aggregation will later decide which of them survive the projected local experiment.

Let $N \geq 2$ be fixed and let $W = (W^1, \dots, W^N)$ be an N -dimensional Brownian motion with independent coordinates on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$. Asset i carries latent volatility channel

$$V_i(t) = \mu_i(t) + \sum_{j=1}^N \int_0^t K_{ij}(t, s) dW_s^j, \quad 1 \leq i \leq N, \quad 0 \leq t \leq T.$$

The kernel K_{ij} encodes the influence of latent driver j on asset i . The diagonal terms K_{ii} are idiosyncratic; the off-diagonal terms K_{ij} , $i \neq j$, are directed latent contagion channels.

Assumption 2.1 (Directed Volterra regularity). For every i, j there exist exponent $H_{ij} \in (0, 1)$, amplitude $\beta_{ij} \in \mathbb{R}$, and a continuous function $a_{ij}: [0, T]^2 \rightarrow \mathbb{R}$ such that for $0 \leq s < t \leq T$,

$$K_{ij}(t, s) = \beta_{ij} a_{ij}(t, s) (t - s)^{H_{ij}-1/2}.$$

Moreover $a_{ij}(t, t) > 0$ whenever $\beta_{ij} \neq 0$, and the drift $t \mapsto \mu_i(t)$ is deterministic and continuous.

Definition 2.2 (Aggregate market volatility). Fix deterministic weights $w = (w_1, \dots, w_N) \in (0, \infty)^N$. The aggregate market-volatility process is

$$M(t) = \sum_{i=1}^N w_i V_i(t).$$

Its effective short-scale exponent H_M is defined by the requirement that for every interior time $t \in (0, T)$,

$$\text{Var}(M(t+h) - M(t)) = c_M(t) h^{2H_M} + o(h^{2H_M}) \quad \text{as } h \downarrow 0,$$

for some deterministic $c_M(t) > 0$.

3. LOCAL OBSERVABILITY AND ADMISSIBILITY AT MARKET SCALE

Admissibility is selected before aggregate roughness is proved. The projected experiment fixes the local scale of an active edge; the shrinking-scale condition then separates rough source-screened channels from smooth or only $O(1)$ -visible channels.

Assumption 3.1 (High-frequency market resolution). There exists a triangular array of local market blocks indexed by n , with block width $h_n \downarrow 0$ and block count $N_n \uparrow \infty$, such that

$$N_n h_n \rightarrow \tau \in (0, \infty).$$

For each active edge, $I_{ij,n}^{\text{eff}}$ denotes the effective information count in the projected local experiment. In an independent-block idealization $I_{ij,n}^{\text{eff}} = N_n$. In an overlapping-block implementation it is the corresponding HAC-effective count, denoted $N_{ij,n}^{\text{HAC}}$. The market-admissibility statements below use $I_{ij,n}^{\text{eff}} h_n$ in place of a raw grid count and assume, for active visible edges, that it stays bounded away from zero and infinity; write

$$I_{ij,n}^{\text{eff}} h_n \rightarrow \tau_{ij} \in (0, \infty)$$

when a limit is needed.

Proposition 3.2 (Projected local scale from the market experiment). *Fix an edge (i, j) with nonzero market loading, $w_i \beta_{ij} \neq 0$. Suppose that along the high-frequency array of Assumption 3.1, the source-screened projected experiment for that edge admits blockwise score contributions $S_{ij,\ell,n}$, $1 \leq \ell \leq N_n$, such that under the null*

$$\mathbb{E}[S_{ij,\ell,n}] = 0, \quad \text{Var}(S_{ij,\ell,n}) = \kappa_{ij} h_n^{2H_{ij}} + o(h_n^{2H_{ij}})$$

for some $\kappa_{ij} > 0$, and

$$\text{Var}\left(\sum_{\ell=1}^{N_n} S_{ij,\ell,n}\right) = I_{ij,n}^{\text{eff}} \kappa_{ij} h_n^{2H_{ij}} + o(I_{ij,n}^{\text{eff}} h_n^{2H_{ij}}).$$

Writing $S_{ij,n} := \sum_{\ell=1}^{N_n} S_{ij,\ell,n}$, assume moreover that for every shrinking perturbation sequence $\delta_n \rightarrow 0$, the projected log-likelihood ratio admits the local quadratic expansion

$$\log \frac{d\mathbb{P}_{ij,n}^{(\delta_n)}}{d\mathbb{P}_{ij,n}^{(0)}} = \delta_n S_{ij,n} - \frac{1}{2} \delta_n^2 \mathcal{I}_{ij,n} + o_{\mathbb{P}}(1),$$

with

$$\mathcal{I}_{ij,n} = I_{ij,n}^{\text{eff}} \kappa_{ij} h_n^{2H_{ij}} + o(I_{ij,n}^{\text{eff}} h_n^{2H_{ij}}).$$

Then a local perturbation of amplitude δ_n remains in the finite local quadratic regime if and only if

$$\delta_n \sqrt{I_{ij,n}^{\text{eff}}} h_n^{H_{ij}} = O(1).$$

It is nondegenerate on a local sequence when this product is bounded away from zero and infinity. Consequently the unique local normalization, up to constant multiples, that produces a finite nonzero local experiment is

$$\theta_{ij,n} = (I_{ij,n}^{\text{eff}})^{-1/2} h_n^{-H_{ij}}.$$

Proof. Under the stated variance expansion, the aggregated projected score $S_{ij,n}$ has standard deviation

$$\sqrt{\text{Var}(S_{ij,n})} = \sqrt{\kappa_{ij}} \sqrt{I_{ij,n}^{\text{eff}}} h_n^{H_{ij}} (1 + o(1)).$$

Likewise,

$$\delta_n^2 \mathcal{I}_{ij,n} = \kappa_{ij} (\delta_n \sqrt{I_{ij,n}^{\text{eff}}} h_n^{H_{ij}})^2 (1 + o(1)).$$

Set

$$\ell_n := \delta_n \sqrt{I_{ij,n}^{\text{eff}}} h_n^{H_{ij}}.$$

The linear score term has stochastic size $O_{\mathbb{P}}(|\ell_n|)$, and the quadratic information term has deterministic size $O(\ell_n^2)$. If $\ell_n \rightarrow 0$, both leading terms vanish and the perturbation is asymptotically

negligible. If $|\ell_n| \rightarrow \infty$, the quadratic term leaves the finite local regime. Finite nonzero local alternatives are obtained by keeping ℓ_n bounded and non-vanishing along the local sequence under consideration. This fixes the local scale, up to constant multiples, as

$$\theta_{ij,n} = (I_{ij,n}^{\text{eff}})^{-1/2} h_n^{-H_{ij}}.$$

□

Remark 3.3 (Imported input and market-level step). The imported object from [7] is the source-screened projected experiment: visible edges have strictly positive local information and satisfy the local quadratic expansion used above. Proposition 3.2 performs the market-level step. Once the projected experiment is placed on a high-frequency market array, the local normalization is not a chosen scale; it is the unique finite nonzero scale of the local log-likelihood.

Proposition 3.4 (Pairwise projected import is dimension-honest). *Fix an ordered pair (i, j) . The admissibility analysis in the present section uses from [7] only the pairwise source-screened projected experiment attached to that pair: source-screened visibility, a strictly positive projected information lower bound when the channel is visible, and the corresponding local quadratic expansion on the high-frequency array. Conditional on that pairwise input, Proposition 3.2, Lemma 3.5, Definition 3.8, and Proposition 3.9 are dimension-agnostic. In particular, the admissibility characterization does not rely on any bivariate ambient market restriction.*

Proof. Every step in the admissibility block fixes one ordered pair (i, j) and then uses only the projected score variance, the projected information lower bound, and the projected local normalization for that pair. The ambient dimension N enters only through the choice of the pair and the market-loading coefficient $w_i \beta_{ij}$. No argument reduces the full market to a two-asset system. □

Lemma 3.5 (Observability failure for smooth channels). *Fix an off-diagonal edge (i, j) with $i \neq j$ and nonzero market loading $w_i \beta_{ij} \neq 0$, and assume the blockwise score and local quadratic expansion hypotheses of Proposition 3.2. If $H_{ij} \geq \frac{1}{2}$, then for every shrinking perturbation sequence $\delta_n \rightarrow 0$,*

$$\delta_n S_{ij,n} \xrightarrow{\mathbb{P}} 0$$

and

$$\delta_n^2 \mathcal{I}_{ij,n} \rightarrow 0.$$

Consequently the projected local log-likelihood shift vanishes along every shrinking alternative, so the edge does not generate a nondegenerate projected local experiment at high frequency.

Proof. Proposition 3.2 gives

$$\sqrt{\text{Var}(S_{ij,n})} = \sqrt{\kappa_{ij}} \sqrt{I_{ij,n}^{\text{eff}} h_n^{H_{ij}}} (1 + o(1)) = \sqrt{\kappa_{ij}} (I_{ij,n}^{\text{eff}} h_n)^{1/2} h_n^{H_{ij}-1/2} (1 + o(1)).$$

Because $I_{ij,n}^{\text{eff}} h_n \rightarrow \tau_{ij} \in (0, \infty)$, the factor $(I_{ij,n}^{\text{eff}} h_n)^{1/2}$ stays bounded away from zero and infinity. If $H_{ij} > \frac{1}{2}$, then $h_n^{H_{ij}-1/2} \rightarrow 0$, so

$$\sqrt{\text{Var}(S_{ij,n})} \rightarrow 0.$$

If $H_{ij} = \frac{1}{2}$, then $\sqrt{\text{Var}(S_{ij,n})} \rightarrow \sqrt{\kappa_{ij} \tau} \in (0, \infty)$. Therefore in either case

$$\delta_n \sqrt{\text{Var}(S_{ij,n})} \rightarrow 0$$

for every sequence $\delta_n \rightarrow 0$. Chebyshev's inequality then yields

$$\delta_n S_{ij,n} \xrightarrow{\mathbb{P}} 0.$$

Moreover,

$$\mathcal{I}_{ij,n} = \kappa_{ij} (I_{ij,n}^{\text{eff}} h_n)^{2H_{ij}-1} (1 + o(1)).$$

If $H_{ij} > \frac{1}{2}$, then $\mathcal{I}_{ij,n} \rightarrow 0$. If $H_{ij} = \frac{1}{2}$, then $\mathcal{I}_{ij,n} \rightarrow \kappa_{ij} \tau_{ij} \in (0, \infty)$. Therefore in either case

$$\delta_n^2 \mathcal{I}_{ij,n} \rightarrow 0.$$

The local quadratic expansion in Proposition 3.2 now yields

$$\log \frac{d\mathbb{P}_{ij,n}^{(\delta_n)}}{d\mathbb{P}_{ij,n}^{(0)}} = o_{\mathbb{P}}(1),$$

so both the linear score term and the quadratic information term vanish along every shrinking alternative. This proves the degeneracy claim. $\square$

Definition 3.6 (Active and dominant edges). The active edge set is

$$\mathcal{E}_{\text{act}} = \{(i, j) : w_i \beta_{ij} \neq 0\}.$$

When $\mathcal{E}_{\text{act}} \neq \emptyset$, its dominant exponent and dominant edge set are

$$H_* = \min_{(i,j) \in \mathcal{E}_{\text{act}}} H_{ij}, \quad \mathcal{E}_* = \{(i, j) \in \mathcal{E}_{\text{act}} : H_{ij} = H_*\}.$$

All subsequent uses of H_* occur under assumptions that make the active set nonempty.

Assumption 3.7 (Noncancelling aggregation). The active edge set is nonempty. For every interior time $t \in (0, T)$,

$$\sum_{j=1}^N \left(\sum_{i: (i,j) \in \mathcal{E}_*} w_i \beta_{ij} a_{ij}(t, t) \right)^2 > 0.$$

Equivalently, the dominant short-scale contributions do not cancel at leading order in the variance of $M(t+h) - M(t)$.

Definition 3.8 (Projected-minimax admissibility). An active off-diagonal edge $(i, j) \in \mathcal{E}_{\text{act}}$ with $i \neq j$ is called *projected-minimax admissible* if it satisfies both conditions.

- (i) the corresponding source-to-target channel remains source-screened visible in the projected information geometry of [7], so its projected information lower bound is strictly positive at the local scale;
- (ii) along the high-frequency array of Assumption 3.1, the projected normalization

$$\theta_{ij,n} = (I_{ij,n}^{\text{eff}})^{-1/2} h_n^{-H_{ij}}$$

identified in Proposition 3.2 is a genuinely local scale, in the sense that $\theta_{ij,n} \rightarrow 0$.

The set of admissible edges is denoted $\mathcal{E}_{\text{adm}}$.

3.1. Admissible market classes. Proposition 3.9 characterizes the admissible class inside the projected experiment. Its proof uses only Definition 3.8, Proposition 3.2, and Lemma 3.5. Later roughness, synchronization, and pricing results depend on this class; they do not enter its definition.

Proposition 3.9 (Characterization of projected-minimax admissibility). *Assume Assumption 3.1 and fix an active off-diagonal edge $(i, j) \in \mathcal{E}_{\text{act}}$ with $i \neq j$. Suppose the projected experiment attached to that edge inherits a strictly positive projected information lower bound whenever the channel is source-screened visible, as in [7], and satisfies the blockwise score and local quadratic expansion hypotheses of Proposition 3.2. Then the following are equivalent:*

- (i) (i, j) is projected-minimax admissible.
- (ii) The channel remains source-screened visible in the projected experiment and satisfies

$$H_{ij} < \frac{1}{2}.$$

Proof. If (i, j) is projected-minimax admissible, source-screened visibility is part (i) of Definition 3.8. Proposition 3.2 identifies $\theta_{ij,n} = (I_{ij,n}^{\text{eff}})^{-1/2} h_n^{-H_{ij}}$ as the only candidate finite nonzero local normalization. Under Assumption 3.1,

$$\theta_{ij,n} = (I_{ij,n}^{\text{eff}} h_n)^{-1/2} h_n^{1/2 - H_{ij}}.$$

If $H_{ij} = 1/2$, this scale converges to $\tau_{ij}^{-1/2}$; if $H_{ij} > 1/2$, it diverges. It is therefore not a shrinking local scale. Lemma 3.5 gives the equivalent likelihood statement: every shrinking

perturbation is asymptotically invisible in the projected experiment. Definition 3.8(ii) is violated, and hence $H_{ij} < 1/2$.

Conversely, suppose (i, j) is source-screened visible and satisfies $H_{ij} < 1/2$. Proposition 3.2 identifies the unique finite nonzero local normalization as

$$\theta_{ij,n} = (I_{ij,n}^{\text{eff}})^{-1/2} h_n^{-H_{ij}} = (I_{ij,n}^{\text{eff}} h_n)^{-1/2} h_n^{1/2-H_{ij}}.$$

Since $I_{ij,n}^{\text{eff}} h_n \rightarrow \tau_{ij} \in (0, \infty)$ and $H_{ij} < 1/2$, one has $\theta_{ij,n} \rightarrow 0$. The positive projected information lower bound makes this shrinking scale nondegenerate in the projected experiment. Together with source-screened visibility, this gives Definition 3.8. $\square$

Corollary 3.10 (Non-vacuity of the admissible class). *Assume Assumption 3.1. If there exists an off-diagonal edge $(i, j) \in \mathcal{E}_{\text{act}}$ such that $i \neq j$, the channel is source-screened visible in the projected experiment with strictly positive projected information, the corresponding projected experiment satisfies the blockwise score and local quadratic expansion hypotheses of Proposition 3.2, and*

$$H_{ij} < \frac{1}{2},$$

then

$$\mathcal{E}_{\text{adm}} \neq \emptyset.$$

In particular, the admissibility class is nonempty whenever the market contains at least one economically active rough source-screened contagion channel. If, in addition, the dominant market-scale edges lie in that same class, Theorem 4.2 applies nontrivially.

Proof. This is the implication (ii) $\Rightarrow$ (i) in Proposition 3.9, applied to the given off-diagonal edge. $\square$

Remark 3.11 (Which channels are excluded). The admissibility class excludes two economically distinct cases.

- (i) Smooth off-diagonal contagion channels with $H_{ij} \geq \frac{1}{2}$, because they generate degenerate shrinking projected experiments.
- (ii) Off-diagonal channels that are visible only at an $O(1)$ scale, because they do not admit a genuinely local projected normalization.

A channel with zero market loading does not belong to the active market-edge set in the first place, so it is screened out before admissibility is assessed.

Remark 3.12 (Economic content of the admissible class). Proposition 3.9 and Corollary 3.10 make the admissible class non-vacuous. The class is not an exogenous rough subclass. It consists of source-screened channels that are statistically visible, economically active, and genuinely local at high frequency. The projected-score evidence in [7] documents rough visible edges of this form in realized-volatility panels, so the admissible class has empirical content as well as formal content.

Remark 3.13 (What the admissibility condition means). Definition 3.8 imports the source-screened observability and minimax results of [1, 12, 7] into the market model and selects the directed channels that remain source-screened visible and admit a genuinely local projected experiment. The shrinking-scale condition is the link to the minimax theory. An edge visible only at an $O(1)$ scale is not a latent high-frequency channel in the projected LAN sense.

Lemma 3.14 (Local projected admissibility forces roughness). *Assume Assumption 3.1. If an active off-diagonal edge (i, j) is projected-minimax admissible, then*

$$H_{ij} < \frac{1}{2}.$$

Proof. Projected-minimax admissibility requires the normalization

$$\theta_{ij,n} = (I_{ij,n}^{\text{eff}})^{-1/2} h_n^{-H_{ij}} = (I_{ij,n}^{\text{eff}} h_n)^{-1/2} h_n^{1/2-H_{ij}}$$

to be genuinely local, hence $\theta_{ij,n} \rightarrow 0$. Since $I_{ij,n}^{\text{eff}} h_n \rightarrow \tau_{ij} \in (0, \infty)$, this is possible only when $1/2 - H_{ij} > 0$. Therefore $H_{ij} < 1/2$. $\square$

Lemma 3.15 (Driverwise dominant short-scale expansion). *Under Assumption 2.1, assume $\mathcal{E}_{\text{act}} \neq \emptyset$, fix an interior time $t \in (0, T)$ and a driver $j \in \{1, \dots, N\}$. Define*

$$G_{j,h,t}(s) := \sum_{i=1}^N w_i [K_{ij}(t+h, s) - K_{ij}(t, s)].$$

Then there exists a constant $\kappa(H_) > 0$, depending only on H_* , such that*

$$\int_0^{t+h} G_{j,h,t}(s)^2 ds = \kappa(H_*) \left(\sum_{i: (i,j) \in \mathcal{E}_*} w_i \beta_{ij} a_{ij}(t, t) \right)^2 h^{2H_*} + o(h^{2H_*})$$

as $h \downarrow 0$.

Proof. For $H \in (0, 1)$, write

$$\Gamma_{H,t,h}(s) := (t+h-s)^{H-1/2} \mathbf{1}_{\{s < t+h\}} - (t-s)^{H-1/2} \mathbf{1}_{\{s < t\}}.$$

Assumption 2.1 and continuity of a_{ij} imply the local expansion

$$K_{ij}(t+h, s) - K_{ij}(t, s) = \beta_{ij} a_{ij}(t, t) \Gamma_{H_{ij},t,h}(s) + r_{ij,h,t}(s),$$

where $\|r_{ij,h,t}\|_{L^2([0,t+h])} = o(h^{H_{ij}})$. The Riemann–Liouville increment kernel satisfies the short-scale expansion from Lemma A.1. Edges with $w_i \beta_{ij} = 0$ contribute no leading aggregate term, regardless of their exponent. Active edges outside $\mathcal{E}_*$ have exponent strictly larger than H_* , hence contribute only $o(h^{H_*})$ in L^2 . The edges in $\mathcal{E}_*$ share the same leading kernel $\Gamma_{H_*,t,h}$. It follows, using the finiteness of the network, that

$$G_{j,h,t}(s) = \left(\sum_{i: (i,j) \in \mathcal{E}_*} w_i \beta_{ij} a_{ij}(t, t) \right) \Gamma_{H_*,t,h}(s) + o_{L^2}(h^{H_*}).$$

Denote the coefficient in parentheses by $A_j(t)$. Since $\|\Gamma_{H_*,t,h}\|_{L^2} = O(h^{H_*})$, Cauchy–Schwarz gives

$$\|G_{j,h,t}\|_2^2 = A_j(t)^2 \|\Gamma_{H_*,t,h}\|_2^2 + o(h^{2H_*}).$$

Lemma A.1 then gives the claimed squared-norm expansion. $\square$

Proposition 3.16 (Leading-order aggregate variance expansion). *Under Assumptions 2.1 and 3.7, for every interior time $t \in (0, T)$,*

$$\text{Var}(M(t+h) - M(t)) = c_M(t) h^{2H_*} + o(h^{2H_*}) \quad \text{as } h \downarrow 0,$$

where the leading coefficient is

$$c_M(t) = \kappa(H_*) \sum_{j=1}^N \left(\sum_{i: (i,j) \in \mathcal{E}_*} w_i \beta_{ij} a_{ij}(t, t) \right)^2 > 0.$$

Proof. Write $\Delta_h M(t) = M(t+h) - M(t)$. Since each μ_i is deterministic, it contributes no variance. By linearity and Itô isometry,

$$\text{Var}(\Delta_h M(t)) = \sum_{j=1}^N \int_0^{t+h} \left(\sum_{i=1}^N w_i [K_{ij}(t+h, s) - K_{ij}(t, s)] \right)^2 ds.$$

Applying Lemma 3.15 to each driver j yields

$$\text{Var}(\Delta_h M(t)) = \kappa(H_*) \sum_{j=1}^N \left(\sum_{i: (i,j) \in \mathcal{E}_*} w_i \beta_{ij} a_{ij}(t, t) \right)^2 h^{2H_*} + o(h^{2H_*}).$$

Assumption 3.7 makes the coefficient of h^{2H_*} strictly positive, which gives the claimed expansion. $\square$

Corollary 3.17 (Aggregate exponent is the minimum active exponent). *Under Assumptions 2.1 and 3.7, the aggregate market volatility satisfies*

$$H_M = H_*.$$

Proof. Proposition 3.16 identifies the leading variance order as h^{2H_*} with strictly positive coefficient. The definition of H_M therefore gives $H_M = H_*$. $\square$

4. ROUGHNESS NECESSITY FOR MARKET VOLATILITY

The aggregate short-scale exponent is the minimum active edge exponent. The admissibility restriction now changes that algebraic fact into a roughness theorem. If the market-dominant directed edges are source-screened visible, economically active, and genuinely local at high frequency, the aggregate exponent cannot remain smooth.

Assumption 4.1 (Admissible dominant asymmetry). The market has at least one projected-minimax admissible off-diagonal edge, that is,

$$\mathcal{E}_{\text{adm}} \neq \emptyset.$$

Moreover every market-dominant active edge is an admissible off-diagonal edge.

$$\mathcal{E}_* \subseteq \mathcal{E}_{\text{adm}}.$$

In particular, the dominant short-scale market asymmetry is carried by admissible contagion channels rather than by diagonal or inadmissible edges.

Theorem 4.2 (Market roughness necessity). *Assume Assumptions 2.1, 3.1, 3.7, and 4.1. Assume moreover that every edge in $\mathcal{E}_*$ satisfies the local quadratic projected-experiment hypotheses of Proposition 3.2. Write*

$$H_{\text{adm}} = \min_{(i,j) \in \mathcal{E}_{\text{adm}}} H_{ij}.$$

Then the aggregate market volatility is necessarily rough.

$$H_M = H_{\text{adm}} < \frac{1}{2}.$$

In particular, any market satisfying the observability/minimax admissibility condition cannot have a smooth aggregate volatility law once its market-dominant contagion channels are economically active, source-screened, and shrinking-scale.

Proof. By Assumption 4.1, every market-dominant edge belongs to $\mathcal{E}_{\text{adm}}$. Lemma 3.14 therefore proves that every edge in $\mathcal{E}_*$ satisfies $H_{ij} < 1/2$. Since H_* is the minimum exponent over $\mathcal{E}_*$, we have

$$H_* < \frac{1}{2}.$$

By Definition 3.8, admissible edges are active off-diagonal edges, so $\mathcal{E}_{\text{adm}} \subseteq \mathcal{E}_{\text{act}}$. Hence $H_{\text{adm}} \geq H_*$. The inclusion $\mathcal{E}_* \subseteq \mathcal{E}_{\text{adm}}$ from Assumption 4.1 gives the reverse inequality, and therefore $H_{\text{adm}} = H_*$. Corollary 3.17 then yields

$$H_M = H_* = H_{\text{adm}} < \frac{1}{2},$$

which proves the claim. $\square$

Corollary 4.3 (Recovery of the bivariate theory). *In the case $N = 2$, Theorem 4.2 reduces to the low-dimensional face of the bivariate latent-contagion results. A nondegenerate admissible channel $1 \rightarrow 2$ or $2 \rightarrow 1$ forces the aggregate exponent to be the exponent of that active source-screened channel, so the market-level roughness law collapses to the bivariate latent-contagion geometry studied in [1, 12, 7].*

Remark 4.4 (Necessity is relative to the admissibility class). Theorem 4.2 applies to markets whose active directional channels are source-screened observable and informative for minimax discrimination. Propositions 3.9 and 3.10 characterize that class rather than stipulate it. Within the class, the admissibility criterion forces roughness instead of treating it as an empirical coincidence or free regularity choice.

5. SYSTEMIC SYNCHRONIZATION IN THE SCREENED NETWORK

Roughness identifies the dominant market-scale edges. Synchronization asks whether those edges remain local spillovers or align into a common Perron mode. The screened synchronization matrix is the object that separates these two regimes.

Fix an interior time $t_0 \in (0, T)$. The synchronization analysis is local around t_0 , just as the roughness theorem is local in the short-scale limit.

Definition 5.1 (Dominant signed edge matrix and screened synchronization matrix). Let

$$\mathcal{E}_M = \{(i, j) \in \mathcal{E}_{\text{adm}} : H_{ij} = H_M\}$$

be the set of admissible market-scale edges. For each $(i, j) \in \mathcal{E}_M$, let $\pi_{ij} \in (0, \infty)$ denote its projected visibility weight, inherited from the source-screened information geometry and normalized so that the matrix below is a dimensionless one-step amplification operator for the local stress envelope. The signed directed edge matrix is

$$B_{ij} = \beta_{ij} a_{ij}(t_0, t_0) \mathbf{1}_{\{(i,j) \in \mathcal{E}_M\}}, \quad 1 \leq i, j \leq N.$$

The screened synchronization matrix is the nonnegative matrix $A = (A_{ij})_{1 \leq i, j \leq N}$ with

$$A_{ij} = |B_{ij}| \pi_{ij}.$$

The matrix B keeps signs and governs signed aggregate contributions. The matrix A forgets signs and governs Perron synchronization.

Definition 5.2 (Screened contagion graph). The *screened contagion graph* is the directed graph G_M on vertices $\{1, \dots, N\}$ with an arrow $j \rightarrow i$ whenever $A_{ij} > 0$. It is called *strongly connected* if every ordered pair of vertices is linked by a directed path, and *aperiodic* if the greatest common divisor of its directed cycle lengths is one.

Definition 5.3 (Dominant rough stress envelope). An H_M -scale stress envelope at time t_0 is a vector $x = (x_1, \dots, x_N) \in [0, \infty)^N$ such that the dominant short-scale contribution to asset i takes the form

$$x_i(t - s)^{H_M - 1/2}$$

up to $o((t - s)^{H_M - 1/2})$ as $s \uparrow t_0$. The corresponding market-scale reduced dynamics are governed by repeated action of the screened synchronization matrix A . The network is said to be in the *synchronized regime* when $\text{spr}(A) > 1$.

Lemma 5.4 (Market-to-envelope reduction). *Let $x_m \in [0, \infty)^N$ be the coefficient vector of the dominant H_M -scale contribution generated by the market after m rounds of admissible contagion around t_0 . Then one further round of dominant contagion produces*

$$x_{m+1} = Ax_m + o(1)$$

at the level of leading-order coefficients.

Proof. At time t_0 , the leading contribution of a dominant admissible edge $(i, j) \in \mathcal{E}_M$ is

$$\beta_{ij} a_{ij}(t_0, t_0) (t - s)^{H_M - 1/2}.$$

Multiplying by the projected visibility weight π_{ij} and the incoming envelope coefficient $x_{m,j}$, then summing over dominant incoming edges, gives the next leading contribution for asset i ,

$$\sum_{j=1}^N |B_{ij}| \pi_{ij} x_{m,j}.$$

All non-dominant edges have exponent strictly larger than H_M , hence contribute only lower order terms in the short-scale expansion. Therefore the dominant market coefficients are reduced, at leading order, by the matrix A , which yields $x_{m+1} = Ax_m + o(1)$. $\square$

Proposition 5.5 (Finite-horizon reduction to the screened recursion). *Fix $M \in \mathbb{N}$. For each short-scale parameter $h \downarrow 0$, let $(x_0^{(h)}, \dots, x_M^{(h)}) \subset [0, \infty)^N$ satisfy*

$$x_{m+1}^{(h)} = Ax_m^{(h)} + r_m^{(h)}, \quad 0 \leq m < M,$$

with

$$x_0^{(h)} \rightarrow x_0 \quad \text{and} \quad \max_{0 \leq m < M} \|r_m^{(h)}\| \rightarrow 0$$

for some $x_0 \in [0, \infty)^N$. Then

$$\max_{0 \leq m \leq M} \|x_m^{(h)} - A^m x_0\| \rightarrow 0.$$

In particular, on every fixed contagion depth, the actual market envelope and the exact screened recursion have the same high-frequency limit.

Proof. Write

$$e_m^{(h)} := x_m^{(h)} - A^m x_0.$$

Then

$$e_{m+1}^{(h)} = Ae_m^{(h)} + r_m^{(h)}, \quad 0 \leq m < M,$$

with $e_0^{(h)} = x_0^{(h)} - x_0 \rightarrow 0$. Iterating the recursion yields

$$e_m^{(h)} = A^m e_0^{(h)} + \sum_{\ell=0}^{m-1} A^{m-1-\ell} r_\ell^{(h)}.$$

Fix any matrix norm. Since M is fixed,

$$C_M := \max_{0 \leq m \leq M} \|A^m\| < \infty.$$

Therefore for $0 \leq m \leq M$,

$$\|e_m^{(h)}\| \leq C_M \|e_0^{(h)}\| + MC_M \max_{0 \leq \ell < M} \|r_\ell^{(h)}\|.$$

Taking the maximum over $m \leq M$ gives

$$\max_{0 \leq m \leq M} \|e_m^{(h)}\| \leq C_M \|e_0^{(h)}\| + MC_M \max_{0 \leq \ell < M} \|r_\ell^{(h)}\|.$$

The assumed convergence of $\|e_0^{(h)}\|$ and $\max_{0 \leq \ell < M} \|r_\ell^{(h)}\|$ then proves the claim. $\square$

Theorem 5.6 (Synchronization threshold for the screened market network). *Assume the screened contagion graph G_M is strongly connected. Let $x_{m+1} = Ax_m$, $m \geq 0$, be the market-scale reduced dynamics associated with an initial dominant envelope $x_0 \in [0, \infty)^N$.*

- (i) *If $\text{spr}(A) < 1$, then $x_m \rightarrow 0$ exponentially fast for every initial condition $x_0 \in [0, \infty)^N$ interpreted as the initial market-scale reduced envelope.*
- (ii) *If G_M is moreover aperiodic, $\text{spr}(A) > 1$, and $x_0 \neq 0$, then there exist positive left and right Perron–Frobenius eigenvectors u, v , normalized by $u^\top v = 1$, such that*

$$\text{spr}(A)^{-m} x_m \longrightarrow (u^\top x_0) v \quad \text{as } m \rightarrow \infty.$$

Consequently all coordinates align with the same latent mode and the aggregate stress $w^\top x_m$ grows at rate $\text{spr}(A)^m$.

Proof. Lemma 5.4 and Proposition 5.5 show that, on any fixed contagion depth, the actual market envelope is asymptotically governed by the exact screened recursion $x_{m+1} = Ax_m$. The theorem therefore analyzes that exact asymptotic reduced system. Strong connectivity of G_M is equivalent to irreducibility of the nonnegative matrix A . The subcritical part is a spectral-radius statement for the screened linear recursion.

The subcritical bound follows from the spectral-radius formula. If $\text{spr}(A) < 1$, then there exists a norm $\|\cdot\|_*$ on $\mathbb{R}^N$ such that $\|A\|_* < 1$, hence $\|x_m\|_* \leq \|A\|_*^m \|x_0\|_* \rightarrow 0$.

For part (ii), aperiodicity of G_M is equivalent to primitivity of A . Perron–Frobenius therefore yields strictly positive eigenvectors u, v with

$$Av = \text{spr}(A)v, \quad u^\top A = \text{spr}(A)u^\top, \quad u^\top v = 1,$$

and the rank-one asymptotic

$$\text{spr}(A)^{-m} A^m \rightarrow vu^\top.$$

Applying this to $x_m = A^m x_0$ yields

$$\text{spr}(A)^{-m} x_m = \text{spr}(A)^{-m} A^m x_0 \rightarrow vu^\top x_0 = (u^\top x_0)v.$$

Since u is strictly positive and $x_0 \in [0, \infty)^N$ is nonzero, one has $u^\top x_0 > 0$. Since v has strictly positive coordinates, every coordinate aligns componentwise with the same mode in the primitive case. Multiplying by the positive weight vector w gives

$$w^\top x_m \sim (u^\top x_0)(w^\top v) \text{spr}(A)^m,$$

which is the announced macroscopic stress concentration. $\square$

Remark 5.7 (Network interpretation of irreducibility and primitivity). Strong connectivity means that every asset can be reached from every other asset by a chain of market-scale admissible contagion channels. Aperiodicity rules out deterministic cycling across disjoint classes and makes the Perron mode the unique asymptotic direction of amplified stress. These are economically interpretable network restrictions, not auxiliary linear-algebra normalizations.

Remark 5.8 (Synchronization and crisis appearance). Theorem 5.6 concerns latent alignment, not literal price crashes. Once the dominant rough edges become supercritical, the market no longer behaves as a collection of local spillovers; its leading envelope aligns with one amplified mode. Visible correlation can still look like interdependence in the sense of [11]. The theorem identifies the latent screened propagation structure that reduced-form correlation does not separate. Section 6 derives the coarse-observation jump limit from that synchronized mode.

5.1. Primitive Perron projection of aggregate envelopes. The synchronization theorem identifies the asymptotic direction of amplified stress. The Perron projection extracts the scalar coordinate carried by that direction. This coordinate is the primitive synchronized amplitude used later in the jump, pricing, and hedging results.

Proposition 5.9 (Canonical Perron decomposition of a latent envelope). *Assume the hypotheses of Theorem 5.6(ii), and normalize the Perron eigenvectors so that*

$$u^\top v = 1.$$

Let $x: [0, T] \rightarrow \mathbb{R}^N$ be any adapted latent envelope with square-integrable coordinates. Define

$$\varpi(t) := u^\top x(t), \quad r(t) := x(t) - \varpi(t)v.$$

Then

$$x(t) = \varpi(t)v + r(t), \quad u^\top r(t) = 0, \quad 0 \leq t \leq T.$$

Moreover, this decomposition is unique among decompositions of the form

$$x(t) = \tilde{\varpi}(t)v + \tilde{r}(t), \quad u^\top \tilde{r}(t) = 0,$$

and the path-generated sigma-field

$$\sigma(\varpi) := \sigma(\varpi(s) : 0 \leq s \leq T)$$

is automatically contained in

$$\sigma(x) := \sigma(x(s) : 0 \leq s \leq T).$$

Proof. The identity

$$x(t) = \varpi(t)v + r(t)$$

is just the definition of $r(t)$. Since $u^\top v = 1$,

$$u^\top r(t) = u^\top x(t) - \varpi(t)u^\top v = u^\top x(t) - \varpi(t) = 0.$$

If $x(t) = \tilde{\varpi}(t)v + \tilde{r}(t)$ with $u^\top \tilde{r}(t) = 0$, then applying $u^\top$ gives

$$\tilde{\varpi}(t) = u^\top x(t) = \varpi(t),$$

and hence $\tilde{r}(t) = r(t)$. Finally, $\varpi(t) = u^\top x(t)$ is a measurable functional of $x(t)$ at each time, so $\sigma(\varpi) \subseteq \sigma(x)$. $\square$

Theorem 5.10 (Spectral-gap control of the off-Perron aggregate remainder). *Assume the hypotheses of Theorem 5.6(ii), and let u, v be normalized so that $u^\top v = 1$. Let $x_{m+1} = Ax_m$, $m \geq 0$, be the exact screened recursion associated with an initial envelope $x_0 \in [0, \infty)^N$. Define*

$$\varpi_m := u^\top x_m, \quad r_m := x_m - \varpi_m v.$$

Let

$$\rho_2 := \max\{|\zeta| : \zeta \in \text{spec}(A), \zeta \neq \text{spr}(A)\}.$$

Then $\rho_2 < \text{spr}(A)$. For every ρ with $\rho_2 < \rho < \text{spr}(A)$, there exists a constant $C_\rho < \infty$, depending only on the matrix A , ρ , and the chosen norm, such that for every $m \geq 0$,

$$\begin{aligned} \varpi_m &= \text{spr}(A)^m (u^\top x_0), \\ x_m &= \text{spr}(A)^m (u^\top x_0) v + r_m, \\ \|r_m\| &\leq C_\rho \rho^m \|x_0\|. \end{aligned}$$

Consequently,

$$w^\top x_m = (w^\top v) \text{spr}(A)^m (u^\top x_0) + O(\rho^m)$$

as $m \rightarrow \infty$, for every such ρ . In particular, the Perron amplitude is the aggregate leading state of the exact screened recursion, and the off-Perron part is spectrally lower order.

Proof. Because $u^\top A = \text{spr}(A)u^\top$,

$$\varpi_m = u^\top x_m = u^\top A^m x_0 = \text{spr}(A)^m (u^\top x_0).$$

By Proposition 5.9,

$$x_m = \varpi_m v + r_m = \text{spr}(A)^m (u^\top x_0) v + r_m,$$

and $u^\top r_m = 0$. Let

$$P := vu^\top, \quad Q := I - P.$$

Then $P^2 = P$, $Q^2 = Q$, $PQ = QP = 0$, $AP = PA = \text{spr}(A)P$, and

$$r_m = Qx_m = QA^m x_0 = (AQ)^m x_0.$$

Since A is primitive, Perron–Frobenius theory gives a simple dominant eigenvalue $\text{spr}(A)$. The operator AQ acts on the off-Perron subspace, and its spectrum is $\text{spec}(A) \setminus \{\text{spr}(A)\}$ together with possible zero eigenvalues. Hence

$$\text{spr}(AQ) = \rho_2 < \text{spr}(A).$$

Fix ρ with $\rho_2 < \rho < \text{spr}(A)$. The finite-dimensional spectral-radius formula yields a constant $C_\rho < \infty$ such that

$$\|(AQ)^m\| \leq C_\rho \rho^m, \quad m \geq 0.$$

Therefore

$$\|r_m\| = \|(AQ)^m x_0\| \leq C_\rho \rho^m \|x_0\|.$$

Multiplying the decomposition of x_m by $w^\top$ gives

$$w^\top x_m = (w^\top v) \text{spr}(A)^m (u^\top x_0) + w^\top r_m,$$

and the bound on $w^\top r_m$ follows from the estimate on r_m , with the same admissible choice of ρ . $\square$

6. ENDOGENOUS JUMP APPEARANCE FOR COARSE OBSERVERS

The primitive latent market remains continuous. Jump appearance requires two additional operations: Perron-mode synchronization must concentrate stress into thin bursts, and the observer must sample at a coarser resolution than the burst width.

Assumption 6.1 (Synchronized mode reduction). Let $v \in (0, \infty)^N$ be the right Perron vector from the synchronized regime in Theorem 5.6, and let $u \in (0, \infty)^N$ be the corresponding left Perron vector normalized so that

$$u^\top v = 1, \quad w^\top v > 0.$$

Let x_n be the continuous dominant rough stress envelope and define its canonical Perron coordinates by

$$\varpi_n(t) := u^\top x_n(t), \quad r_n(t) := x_n(t) - \varpi_n(t)v, \quad 0 \leq t \leq T.$$

Assume that

$$\int_0^T \varpi_n(t)^2 dt < \infty \quad \text{almost surely,}$$

and that the off-Perron aggregate remainder is negligible at observation scale in the sense that

$$\sup_{0 \leq t \leq T} \left| \int_0^t w^\top r_n(s) ds \right| \xrightarrow{\mathbb{P}} 0.$$

Assumption 6.2 (Burst-concentration regime). Let $\Delta_n \downarrow 0$. Let $\tau_1, \dots, \tau_K$ be $(\mathcal{F}_t)$ -stopping times such that

$$0 < \tau_1 < \dots < \tau_K < T \quad \text{almost surely,}$$

and let $J_1, \dots, J_K$ be strictly positive square-integrable random variables. Let $\psi_1, \dots, \psi_K: \Omega \times \mathbb{R} \rightarrow [0, \infty)$ be $\mathcal{F}_T$ -measurable random profiles such that, almost surely, each path $u \mapsto \psi_k(\omega, u)$ is continuous, compactly supported, and satisfies $\int_{\mathbb{R}} \psi_k(\omega, s) ds = 1$. Let $\delta_{k,n} > 0$ be random scales such that

$$\max_{1 \leq k \leq K} \frac{\delta_{k,n}}{\Delta_n} \rightarrow 0 \quad \text{almost surely,}$$

and let ϖ_n be the scalar synchronized mode amplitude from Assumption 6.1. Suppose that

$$\varpi_n(t) = \sum_{k=1}^K \frac{J_k}{(w^\top v)\delta_{k,n}} \psi_k\left(\frac{t - \tau_k}{\delta_{k,n}}\right) + q_n(t), \quad 0 \leq t \leq T,$$

where

$$\int_0^T |q_n(t)| dt \xrightarrow{\mathbb{P}} 0.$$

Definition 6.3 (Activated synchronized regime). A rough-synchronized market is said to be in an *activated synchronized regime* if the Perron-mode amplitude ϖ_n is driven by a common scalar trigger process whose crossings of finitely many activation levels are magnified through a steep nonlinear response at a width δ_n that shrinks faster than the observation mesh Δ_n . Assumption 6.4 records the reduced-form version of this regime used in the jump theorem.

Assumption 6.4 (Reduced-form threshold-triggered activation). Let $\psi: \mathbb{R} \rightarrow [0, \infty)$ be continuous, compactly supported, and satisfy $\int_{\mathbb{R}} \psi(s) ds = 1$. Let $\delta_n \downarrow 0$. Suppose there exists an adapted process Z with almost surely C^1 sample paths, deterministic levels $c_1, \dots, c_K \in \mathbb{R}$, and strictly positive square-integrable random variables $J_1, \dots, J_K$ such that:

- (i) $Z(\tau_k) = c_k$ almost surely for stopping times $0 < \tau_1 < \dots < \tau_K < T$;
- (ii) the crossings are transversal, in the sense that

$$\gamma_k := \dot{Z}(\tau_k) > 0$$

almost surely for every k ;

(iii) there exist pairwise disjoint random neighborhoods I_k of τ_k on which

$$\inf_{t \in I_k} \dot{Z}(t) \geq \frac{\gamma_k}{2} > 0$$

almost surely, and outside $\bigcup_k I_k$ one has $|Z(t) - c_k| > \varepsilon$ for every k , for some $\varepsilon > 0$ independent of n ;

(iv) the synchronized mode amplitude is generated by the steep activation derivative

$$\varpi_n(t) = \sum_{k=1}^K \frac{J_k}{w^\top v} \frac{\dot{Z}(t)}{\delta_n} \psi\left(\frac{Z(t) - c_k}{\delta_n}\right), \quad 0 \leq t \leq T.$$

Proposition 6.5 (Threshold activation implies burst concentration). *Under Assumption 6.4, define*

$$\delta_{k,n} := \frac{\delta_n}{\gamma_k}, \quad \psi_k(u) := \gamma_k \psi(\gamma_k u), \quad 1 \leq k \leq K.$$

Then $\int_{\mathbb{R}} \psi_k(s) ds = 1$, $\max_k \delta_{k,n}/\Delta_n \rightarrow 0$ whenever $\delta_n/\Delta_n \rightarrow 0$, both almost surely, and there exists a remainder q_n such that

$$\int_0^T |q_n(t)| dt \xrightarrow{\mathbb{P}} 0$$

and

$$\varpi_n(t) = \sum_{k=1}^K \frac{J_k}{(w^\top v) \delta_{k,n}} \psi_k\left(\frac{t - \tau_k}{\delta_{k,n}}\right) + q_n(t).$$

Consequently, whenever $\delta_n/\Delta_n \rightarrow 0$, Assumption 6.4 implies Assumption 6.2.

Proof. Work on the probability-one event on which the paths of Z are C^1 , the crossings are transversal, the neighborhoods I_k are disjoint, and the separation condition holds. Let the support of ψ be contained in $[-L, L]$. For the k -th summand in Assumption 6.4, the support is contained in

$$A_{k,n} := \{t \in I_k : |Z(t) - c_k| \leq L\delta_n\}$$

for all large n ; outside $\bigcup_k I_k$ the separation condition makes the summand vanish. The lower derivative bound on I_k gives

$$A_{k,n} \subseteq \left\{t : |t - \tau_k| \leq \frac{2L}{\gamma_k} \delta_n\right\}.$$

By C^1 -regularity,

$$\omega_{k,n} := \sup_{|t - \tau_k| \leq (2L/\gamma_k)\delta_n} |\dot{Z}(t) - \gamma_k| \rightarrow 0.$$

On $A_{k,n}$, write

$$Z(t) - c_k = \gamma_k(t - \tau_k) + \rho_{k,n}(t).$$

Then

$$\sup_{t \in A_{k,n}} \frac{|\rho_{k,n}(t)|}{\delta_n} \rightarrow 0, \quad \sup_{t \in A_{k,n}} |\dot{Z}(t) - \gamma_k| \rightarrow 0.$$

Uniform continuity of ψ on its compact support therefore yields

$$\sup_{t \in A_{k,n}} \left| \psi\left(\frac{Z(t) - c_k}{\delta_n}\right) - \psi\left(\gamma_k \frac{t - \tau_k}{\delta_n}\right) \right| \rightarrow 0.$$

Consequently, uniformly on $A_{k,n}$,

$$\frac{\dot{Z}(t)}{\delta_n} \psi\left(\frac{Z(t) - c_k}{\delta_n}\right) = \frac{\gamma_k}{\delta_n} \psi\left(\gamma_k \frac{t - \tau_k}{\delta_n}\right) + e_{k,n}(t) = \frac{1}{\delta_{k,n}} \psi_k\left(\frac{t - \tau_k}{\delta_{k,n}}\right) + e_{k,n}(t),$$

where $e_{k,n}$ is set to zero outside $A_{k,n}$ and satisfies

$$\int_0^T |e_{k,n}(t)| dt \rightarrow 0.$$

Indeed, the support length is $O(\delta_n)$, while the preceding uniform error is $o(1)/\delta_n$. Summing over the finite set of crossings and setting

$$q_n(t) := \sum_{k=1}^K \frac{J_k}{w^\top v} e_{k,n}(t)$$

proves the claimed representation and gives $\int_0^T |q_n(t)| dt \rightarrow 0$ on the same event; in particular the convergence holds in probability. The normalization of ψ_k follows from

$$\int_{\mathbb{R}} \gamma_k \psi(\gamma_k s) ds = \int_{\mathbb{R}} \psi(r) dr = 1,$$

and $\delta_{k,n}/\Delta_n = (\delta_n/\Delta_n)/\gamma_k \rightarrow 0$ whenever $\delta_n/\Delta_n \rightarrow 0$, because $\gamma_k > 0$ almost surely. $\square$

Definition 6.6 (Coarse observer). Let B be a background process with almost surely continuous sample paths and define the coarse observer by

$$Y_n(t) = B(\lfloor t/\Delta_n \rfloor \Delta_n) + \sum_{m=1}^{\lfloor t/\Delta_n \rfloor} \int_{(m-1)\Delta_n}^{m\Delta_n} w^\top x_n(s) ds.$$

Proposition 6.7 (Burst concentration reduction). *Under Assumptions 6.1 and 6.2, and for the coarse observer in Definition 6.6,*

$$Y_n = B_n + S_n + R_n + Q_n,$$

where

$$\begin{aligned} B_n(t) &= B(\lfloor t/\Delta_n \rfloor \Delta_n), \\ R_n(t) &= \sum_{m=1}^{\lfloor t/\Delta_n \rfloor} \int_{(m-1)\Delta_n}^{m\Delta_n} w^\top r_n(s) ds, \end{aligned}$$

and

$$S_n(t) = \sum_{k=1}^K J_k \left[F_k \left(\frac{\lfloor t/\Delta_n \rfloor \Delta_n - \tau_k}{\delta_{k,n}} \right) - F_k \left(\frac{-\tau_k}{\delta_{k,n}} \right) \right], \quad F_k(x) := \int_{-\infty}^x \psi_k(u) du,$$

and the q_n -remainder is

$$Q_n(t) := \sum_{m=1}^{\lfloor t/\Delta_n \rfloor} \int_{(m-1)\Delta_n}^{m\Delta_n} (w^\top v) q_n(s) ds.$$

Moreover,

$$\sup_{0 \leq t \leq T} |R_n(t)| \xrightarrow{\mathbb{P}} 0,$$

and

$$\sup_{0 \leq t \leq T} |Q_n(t)| \xrightarrow{\mathbb{P}} 0.$$

Proof. Since $x_n(t) = \varpi_n(t)v + r_n(t)$, the definition of Y_n gives

$$Y_n(t) = B_n(t) + \sum_{m=1}^{\lfloor t/\Delta_n \rfloor} \int_{(m-1)\Delta_n}^{m\Delta_n} (w^\top v) \varpi_n(s) ds + R_n(t).$$

Substituting the burst representation of ϖ_n yields

$$Y_n(t) = B_n(t) + S_n(t) + R_n(t) + Q_n(t),$$

where

$$Q_n(t) := \sum_{m=1}^{\lfloor t/\Delta_n \rfloor} \int_{(m-1)\Delta_n}^{m\Delta_n} (w^\top v) q_n(s) ds.$$

Integrating the localized profiles gives the stated form of S_n , including the lower-limit correction $F_k(-\tau_k/\delta_{k,n})$. Because $\tau_k > 0$ almost surely and each ψ_k has compact support, that correction

is eventually zero, but the displayed formula is exact. The bound on R_n follows directly from Assumption 6.1, and

$$\sup_{0 \leq t \leq T} |Q_n(t)| \leq (w^\top v) \int_0^T |q_n(s)| ds \xrightarrow{\mathbb{P}} 0$$

by Assumption 6.2. $\square$

Remark 6.8 (Hierarchy of the jump limit). The jump block has two distinct restrictions. Assumption 6.1 is the baseline Perron projection suggested by Theorem 5.6 and made canonical by Proposition 5.9. Latent stress is decomposed into the Perron amplitude $\varpi_n = u^\top x_n$ and an off-Perron remainder whose aggregate effect is negligible after integration. Assumption 6.2 is the sharper regime restriction that turns that synchronized mode into narrow bursts. Assumption 6.4 and Proposition 6.5 give one endogenous route through common threshold-triggered activation of the synchronized scalar mode. Theorem 6.10 then derives coarse jumps from the simultaneous presence of mode alignment, burst concentration, and coarse observation.

Lemma 6.9 (Burst profiles collapse to jumps under coarse observation). *Under Assumption 6.2, the process*

$$S_n(t) = \sum_{k=1}^K J_k \left[F_k \left(\frac{\lfloor t/\Delta_n \rfloor \Delta_n - \tau_k}{\delta_{k,n}} \right) - F_k \left(\frac{-\tau_k}{\delta_{k,n}} \right) \right]$$

converges almost surely in $D([0, T])$ under the Skorokhod J_1 topology to

$$S(t) = \sum_{k=1}^K J_k \mathbf{1}_{\{\tau_k \leq t\}}.$$

Proof. Fix an ω in the almost-sure event on which the stopping times are ordered, the profiles are continuous probability densities with compact support, and $\max_k \delta_{k,n}(\omega)/\Delta_n \rightarrow 0$. In particular,

$$0 < \tau_1(\omega) < \dots < \tau_K(\omega) < T.$$

For each k , define

$$S_{k,n}(t, \omega) := J_k(\omega) \left[F_k \left(\frac{\lfloor t/\Delta_n \rfloor \Delta_n - \tau_k(\omega)}{\delta_{k,n}} \right) - F_k \left(\frac{-\tau_k(\omega)}{\delta_{k,n}} \right) \right].$$

Because each F_k is nondecreasing, $S_{k,n}(\cdot, \omega)$ is nondecreasing and càdlàg. Since $\tau_k(\omega) > 0$, $\delta_{k,n}(\omega) \rightarrow 0$, and ψ_k has compact support, the correction term $F_k(-\tau_k(\omega)/\delta_{k,n})$ vanishes for all large n . If $t < \tau_k(\omega)$, then eventually $\lfloor t/\Delta_n \rfloor \Delta_n < \tau_k(\omega)$, so the argument of F_k tends to $-\infty$ and $S_{k,n}(t, \omega) \rightarrow 0$. If $t > \tau_k(\omega)$, then eventually $\lfloor t/\Delta_n \rfloor \Delta_n > \tau_k(\omega)$, so the argument tends to $+\infty$ and $S_{k,n}(t, \omega) \rightarrow J_k(\omega)$. Thus

$$S_{k,n}(\cdot, \omega) \rightarrow J_k(\omega) \mathbf{1}_{\{\tau_k(\omega) \leq \cdot\}}$$

at every continuity point of the limit. By the standard J_1 convergence criterion for monotone càdlàg functions, this implies

$$S_{k,n}(\cdot, \omega) \rightarrow J_k(\omega) \mathbf{1}_{\{\tau_k(\omega) \leq \cdot\}}$$

in $D([0, T])$ under J_1 .

The jump times $\tau_1(\omega), \dots, \tau_K(\omega)$ are distinct, so the addition map is continuous at the limit vector

$$(J_1(\omega) \mathbf{1}_{\{\tau_1(\omega) \leq \cdot\}}, \dots, J_K(\omega) \mathbf{1}_{\{\tau_K(\omega) \leq \cdot\}})$$

under the product J_1 topology. Summing the K convergences therefore gives

$$S_n(\cdot, \omega) \rightarrow \sum_{k=1}^K J_k(\omega) \mathbf{1}_{\{\tau_k(\omega) \leq \cdot\}}$$

in $D([0, T])$ under J_1 . Since the exceptional set does not depend on n , the convergence is almost sure. $\square$

Theorem 6.10 (Stochastic endogenous jump appearance from synchronized modes). *Under Assumptions 6.1 and 6.2, the coarse observer from Definition 6.6 converges in probability in $D([0, T])$ under the Skorokhod J_1 topology to*

$$Y(t) = B(t) + \sum_{k: \tau_k \leq t} J_k.$$

In particular, a continuous latent synchronized market generates a random jump process for a coarse observer once synchronized bursts contract faster than the observation resolution.

Proof. By Proposition 6.7,

$$Y_n = B_n + S_n + R_n + Q_n.$$

Because B has almost surely continuous sample paths on the compact interval $[0, T]$,

$$\sup_{0 \leq t \leq T} |B_n(t) - B(t)| \rightarrow 0 \quad \text{almost surely.}$$

Lemma 6.9 gives

$$S_n \rightarrow \sum_{k=1}^K J_k \mathbf{1}_{\{\tau_k \leq \cdot\}} \quad \text{almost surely in } D([0, T]).$$

Proposition 6.7 also gives $\sup_t |R_n(t)| + \sup_t |Q_n(t)| \xrightarrow{\mathbb{P}} 0$, hence $R_n + Q_n \xrightarrow{\mathbb{P}} 0$ in the J_1 topology. Combining this with the almost sure uniform convergence $B_n \rightarrow B$ and the J_1 -limit of S_n gives

$$Y_n = B_n + S_n + R_n + Q_n \xrightarrow{\mathbb{P}} B + \sum_{k=1}^K J_k \mathbf{1}_{\{\tau_k \leq \cdot\}}$$

in $D([0, T])$ under the J_1 topology. $\square$

Corollary 6.11 (Jump appearance under synchronized threshold activation). *Under Assumptions 6.1 and 6.4, and provided $\delta_n/\Delta_n \rightarrow 0$, the conclusion of Theorem 6.10 holds.*

Proof. Proposition 6.5 implies Assumption 6.2. The claim therefore follows directly from Theorem 6.10. $\square$

Remark 6.12 (Endogeneity of the jump limit). Endogeneity has four steps. Theorem 5.6 uses $\text{spr}(A) > 1$ to produce Perron-mode synchronization. Definition 6.3 and Assumption 6.4 add a common nonlinear trigger for that synchronized mode. Proposition 6.5 converts the trigger into burst concentration. Theorem 6.10 then converts concentration plus coarse observation into a jump limit. Spectral supercriticality alone does not produce jumps. The jump limit is endogenous only in this activated synchronized sense.

Corollary 6.13 (Jump-like aggregate episodes under synchronized activation). *Under the synchronized regime from Theorem 5.6 and the concentration assumptions of Section 6, every jump time τ_k in the coarse observer is generated by concentration of the Perron mode of the screened synchronization matrix A . Consequently, within this activated synchronized regime, the jump-like episodes recorded at market resolution arise endogenously from synchronization of the latent network rather than from primitive discontinuities of the microscopic market model.*

Proof. Theorem 5.6 identifies the synchronized regime with alignment along the Perron mode v , while Theorem 6.10 identifies the coarse jump mass at time τ_k is produced by concentration of the corresponding scalar amplitude ϖ_n . The observed jump times are therefore the times at which the synchronized latent mode concentrates at a scale finer than observation resolution. $\square$

Definition 6.14 (Contagion phases at market scale). Under the high-frequency market array, the market's directional contagion regime is said to be in one of the following phases.

- (i) the *smooth-degenerate phase* if every active off-diagonal edge satisfies $H_{ij} \geq \frac{1}{2}$;
- (ii) the *rough-observable phase* if Assumption 4.1 holds and $\text{spr}(A) \leq 1$;
- (iii) the *rough-synchronized phase* if Assumption 4.1 holds and $\text{spr}(A) > 1$;
- (iv) the *activated synchronized phase* if the market is rough-synchronized and, in addition, Assumptions 6.1 and 6.4 hold with $\delta_n/\Delta_n \rightarrow 0$.

Theorem 6.15 (Phase structure of admissible contagion regimes). *Assume Assumption 2.1 and the high-frequency array of Assumption 3.1. Assume moreover that every active off-diagonal edge satisfies the local quadratic projected-experiment hypotheses of Proposition 3.2.*

- (i) *If every active off-diagonal edge satisfies $H_{ij} \geq \frac{1}{2}$, then the market lies in the smooth-degenerate phase with respect to its directional contagion channels; no active off-diagonal channel generates a nondegenerate shrinking projected experiment.*
- (ii) *If Assumptions 3.7 and 4.1 hold and $\text{spr}(A) \leq 1$, then the market lies in the rough-observable phase. In particular,*

$$H_M < \frac{1}{2},$$

but no synchronized systemic amplification is implied.

- (iii) *If Assumptions 3.7 and 4.1 hold, the screened contagion graph G_M is strongly connected and aperiodic, and $\text{spr}(A) > 1$, then the market lies in the rough-synchronized phase; for every nonzero dominant rough envelope in the reduced dynamics, latent stress aligns along the Perron mode of the screened synchronization matrix and aggregate stress becomes macroscopically amplifying.*
- (iv) *If the assumptions of part (iii) hold and, in addition, Assumptions 6.1 and 6.4 hold with $\delta_n/\Delta_n \rightarrow 0$, then the market lies in the activated synchronized phase. The coarse observer from Definition 6.6 converges in probability in $D([0, T])$ under J_1 to*

$$Y(t) = B(t) + \sum_{k: T_k \leq t} J_k.$$

Proof. If every active off-diagonal edge satisfies $H_{ij} \geq \frac{1}{2}$, then Lemma 3.5 applies to each such edge and proves that every shrinking perturbation is asymptotically invisible in the projected experiment. Part (i) follows.

Under Assumptions 3.7 and 4.1, Theorem 4.2 yields $H_M < 1/2$. If moreover $\text{spr}(A) \leq 1$, the market is rough but not in the supercritical synchronized regime of Theorem 5.6, which is the rough-observable phase in part (ii).

For part (iii), Theorem 5.6 implies that when G_M is strongly connected, aperiodic, and $\text{spr}(A) > 1$, every nonzero dominant rough envelope in the reduced dynamics aligns along the Perron mode and aggregate stress grows at rate $\text{spr}(A)^m$. This is the rough-synchronized phase.

Finally, under the assumptions of part (iv), Corollary 6.11 applies and gives the stated stochastic jump limit for the coarse observer. This is the activated synchronized phase. $\square$

Remark 6.16 (What the smooth-degenerate phase excludes). Part (i) of Theorem 6.15 is a statement about directional contagion, not about every possible source of roughness in the aggregate law. The theorem excludes a nondegenerate off-diagonal projected contagion experiment when all active contagion channels are smooth. It does not rule out rough diagonal or idiosyncratic channels from contributing to the aggregate exponent H_M .

7. PRICING SUMMARIES, CONVEXITY, AND HEDGING UNDER LATENT SYNCHRONIZATION

The synchronized phase enters pricing and hedging through public compression. Visible prices and visible hedges do not recover the full latent market state. They project it onto $\mathcal{G}_T$ or onto a visible payoff span. The governing underlying is integrated market variance

$$Q_M := \int_0^T M(t)^2 dt,$$

the payoff underlying a market variance swap. Convex functions of Q_M describe options on market variance. The pricing and hedging statements are controlled by the Perron-projected scalar amplitude of the synchronized market envelope.

The pricing reduction has two steps. Aggregate market volatility first couples to the full latent stress envelope, not directly to a pricing factor. In the synchronized regime, Perron projection then decomposes that envelope into its leading amplitude and an aggregate-centered off-Perron

remainder. These reductions identify $\mathcal{H}_\varpi := \mathcal{G}_T \vee \sigma(\varpi)$ as the sufficient aggregate sigma-field for the pricing objects treated below.

Assumption 7.1 (Primitive aggregate visible-plus-envelope coupling). There exist a visible process M^{vis} , an adapted latent envelope $x = (x_1, \dots, x_N)$, and a residual process η such that

$$M(t) = M^{\text{vis}}(t) + w^\top x(t) + \eta(t), \quad 0 \leq t \leq T,$$

where $M^{\text{vis}}(t)$ is $\mathcal{G}_t$ -measurable for every t . Let

$$\mathcal{H}_x := \mathcal{G}_T \vee \sigma(x(s) : 0 \leq s \leq T).$$

Assume the path-generated sigma-field in the display above is well defined and that

$$\mathbb{E}[\eta(t) \mid \mathcal{H}_x] = 0 \quad \text{for almost every } t \in [0, T].$$

Assume moreover that

$$\int_0^T (M^{\text{vis}}(t))^2 dt, \quad \int_0^T (w^\top x(t))^2 dt, \quad \int_0^T \eta(t)^2 dt \in L^2(\mathcal{F}_T).$$

Assumption 7.2 (Aggregate off-Perron screening). Assume the synchronized regime of Theorem 5.6(ii), and let u, v be the left and right Perron vectors normalized so that

$$u^\top v = 1, \quad w^\top v > 0.$$

Let x be the latent envelope from Assumption 7.1, and define

$$\varpi(t) := u^\top x(t), \quad r(t) := x(t) - \varpi(t)v, \quad 0 \leq t \leq T.$$

Assume that the path-generated sigma-field

$$\sigma(\varpi) := \sigma(\varpi(s) : 0 \leq s \leq T)$$

is well defined and that

$$\mathbb{E}[w^\top r(t) \mid \mathcal{G}_T \vee \sigma(\varpi)] = 0 \quad \text{for almost every } t \in [0, T].$$

Assume also that

$$\int_0^T \varpi(t)^2 dt, \quad \int_0^T (w^\top r(t))^2 dt \in L^2(\mathcal{F}_T).$$

Theorem 7.3 (Perron-projected sufficient-state reduction of aggregate market volatility). Assume Assumptions 7.1 and 7.2. Set

$$\mathcal{H}_\varpi := \mathcal{G}_T \vee \sigma(\varpi), \quad \eta_\varpi(t) := w^\top r(t) + \eta(t).$$

Then

$$M(t) = M^{\text{vis}}(t) + (w^\top v)\varpi(t) + \eta_\varpi(t), \quad 0 \leq t \leq T,$$

and

$$\mathbb{E}[\eta_\varpi(t) \mid \mathcal{H}_\varpi] = 0 \quad \text{for almost every } t \in [0, T].$$

Equivalently,

$$\mathbb{E}[M(t) \mid \mathcal{H}_\varpi] = M^{\text{vis}}(t) + (w^\top v)\varpi(t) \quad \text{for almost every } t \in [0, T].$$

Moreover,

$$\int_0^T \eta_\varpi(t)^2 dt \in L^2(\mathcal{F}_T), \quad Q_M = \int_0^T M(t)^2 dt \in L^2(\mathcal{F}_T).$$

Thus $\mathcal{H}_\varpi$ is a derived sufficient aggregate sigma-field for the pricing objects built from M .

Proof. By Proposition 5.9, the latent envelope from Assumption 7.1 admits the canonical decomposition

$$x(t) = \varpi(t)v + r(t), \quad u^\top r(t) = 0.$$

Combining this with Assumption 7.1 gives

$$M(t) = M^{\text{vis}}(t) + w^\top (\varpi(t)v + r(t)) + \eta(t) = M^{\text{vis}}(t) + (w^\top v)\varpi(t) + \eta_\varpi(t).$$

Because $\sigma(\varpi) \subseteq \sigma(x(s) : 0 \leq s \leq T)$, one has

$$\mathcal{H}_\varpi = \mathcal{G}_T \vee \sigma(\varpi) \subseteq \mathcal{H}_x.$$

Hence the tower property and Assumption 7.1 imply

$$\mathbb{E}[\eta(t) \mid \mathcal{H}_\varpi] = \mathbb{E}[\mathbb{E}[\eta(t) \mid \mathcal{H}_x] \mid \mathcal{H}_\varpi] = 0.$$

Assumption 7.2 gives

$$\mathbb{E}[w^\top r(t) \mid \mathcal{H}_\varpi] = 0.$$

Therefore

$$\mathbb{E}[\eta_\varpi(t) \mid \mathcal{H}_\varpi] = \mathbb{E}[w^\top r(t) \mid \mathcal{H}_\varpi] + \mathbb{E}[\eta(t) \mid \mathcal{H}_\varpi] = 0,$$

which proves the sufficient-state identity for $M(t)$. The conditional expectation display for $M(t)$ follows immediately.

Finally, from $(a + b)^2 \leq 2a^2 + 2b^2$,

$$\int_0^T \eta_\varpi(t)^2 dt \leq 2 \int_0^T (w^\top r(t))^2 dt + 2 \int_0^T \eta(t)^2 dt,$$

so $\int_0^T \eta_\varpi(t)^2 dt \in L^2(\mathcal{F}_T)$. Also

$$Q_M \leq 3 \int_0^T (M^{\text{vis}}(t))^2 dt + 3(w^\top v)^2 \int_0^T \varpi(t)^2 dt + 3 \int_0^T \eta_\varpi(t)^2 dt.$$

The right-hand side belongs to $L^2(\mathcal{F}_T)$. Hence $Q_M \in L^2(\mathcal{F}_T)$. $\square$

Corollary 7.4 (Derived visible aggregate-volatility reduction). *Under Theorem 7.3, set*

$$\zeta(t) := (w^\top v)\varpi(t),$$

and define

$$A_M := \int_0^T (M^{\text{vis}}(t))^2 dt,$$

$$\Xi_M := 2 \int_0^T M^{\text{vis}}(t)\zeta(t) dt + \int_0^T \zeta(t)^2 dt,$$

and

$$R_M := 2 \int_0^T M^{\text{vis}}(t)\eta_\varpi(t) dt + 2 \int_0^T \zeta(t)\eta_\varpi(t) dt + \int_0^T \eta_\varpi(t)^2 dt.$$

Then

$$Q_M = A_M + \Xi_M + R_M,$$

the random variables A_M , Ξ_M , and R_M belong to $L^2(\mathcal{F}_T)$, and

$$\mathbb{E}[\eta_\varpi(t) \mid \mathcal{H}_\varpi] = 0 \quad \text{for almost every } t \in [0, T].$$

This is the visible aggregate-volatility reduction used below. It is derived from primitive aggregate coupling, canonical Perron projection, and off-Perron aggregate screening. Appendix B reports an Oxford-Man finite-panel illustration of the same Perron-projected sufficient-state reduction using a rolling source-screened market operator built from the projected-score construction of [7].

Proof. The identity for Q_M is the square expansion of $M(t) = M^{\text{vis}}(t) + \zeta(t) + \eta_\varpi(t)$. The square terms belong to L^2 by Theorem 7.3. For the product terms, Cauchy–Schwarz gives, for example,

$$\left| \int_0^T M^{\text{vis}}(t)\eta_\varpi(t) dt \right|^2 \leq \left(\int_0^T (M^{\text{vis}}(t))^2 dt \right) \left(\int_0^T \eta_\varpi(t)^2 dt \right),$$

and the expectation of the right-hand side is finite by another Cauchy–Schwarz inequality in $L^2(\Omega)$. The same argument applies to $\int_0^T \zeta(t)\eta_\varpi(t) dt$ and $\int_0^T M^{\text{vis}}(t)\zeta(t) dt$. Thus A_M , Ξ_M , and R_M are in $L^2(\mathcal{F}_T)$. The conditional centering of η_ϖ is the theorem’s conclusion. $\square$

Definition 7.5 (Synchronized latent factor and variance-linked claims). Let $\mathcal{G}_T \subset \mathcal{F}_T$ be the visible market filtration at maturity. A square-integrable random variable $\Lambda \in L^2(\mathcal{F}_T)$ is called a *synchronized latent factor* if there exists a canonically Perron-projected amplitude process ϖ from the sufficient-state reduction above such that Λ is measurable with respect to $\mathcal{H}_\varpi = \mathcal{G}_T \vee \sigma(\varpi)$ and satisfies

$$\mathbb{E}[\Lambda \mid \mathcal{G}_T] = 0.$$

More generally, a *variance-linked claim* is any square-integrable claim of the form $X = \Psi(Q_M)$ for some measurable Ψ , and it admits a synchronized decomposition if

$$X = X^{\text{vis}} + \alpha_X \Lambda + \varepsilon_X,$$

with $X^{\text{vis}} \in L^2(\mathcal{G}_T)$ and

$$\mathbb{E}[\varepsilon_X \mid \mathcal{G}_T, \Lambda] = 0.$$

For any sub- σ -field $\mathcal{H} \subseteq \mathcal{F}_T$, write $C_{\mathcal{H}}(Y) := \mathbb{E}[Y \mid \mathcal{H}]$ for the L^2 -projection of Y . The factor extraction is a chain of projections. Theorem 7.3 derives $\mathcal{H}_{\varpi}$ from primitive aggregate coupling and Perron projection, and Corollary 7.4 turns that reduction into the square-expansion inputs for Q_M . The market-variance underlying then inherits the centered synchronized factor $C_{\mathcal{H}_{\varpi}}(Q_M) - C_{\mathcal{G}_T}(Q_M)$. Each variance-linked claim $X = \Psi(Q_M)$ then inherits its own synchronized decomposition by projection onto the visible sigma-field enlarged by the market factor Λ_M . The compression theorem acts on these decompositions.

Proposition 7.6 (Canonical synchronized factor extraction for integrated market variance). *Assume Corollary 7.4, and let ζ , A_M , Ξ_M , and R_M be the random objects introduced there. Let $\mathcal{H}_{\varpi} := \mathcal{G}_T \vee \sigma(\varpi)$, and define*

$$\begin{aligned} L_{Q_M} &:= \mathbb{E}[Q_M \mid \mathcal{H}_{\varpi}], \\ Q_M^{\text{vis}} &:= \mathbb{E}[Q_M \mid \mathcal{G}_T], \\ \Delta_{Q_M} &:= L_{Q_M} - Q_M^{\text{vis}}, \\ \tilde{\Lambda}_M &:= \Delta_{Q_M}, \end{aligned}$$

and

$$\varepsilon_M := Q_M - L_{Q_M} = R_M - \mathbb{E}[R_M \mid \mathcal{H}_{\varpi}].$$

If $\tilde{\Lambda}_M = 0$ almost surely, set $\alpha_M := 0$ and $\Lambda_M := 0$. Otherwise set

$$\alpha_M := \|\tilde{\Lambda}_M\|_{L^2}, \quad \Lambda_M := \frac{\tilde{\Lambda}_M}{\alpha_M}.$$

Then

$$Q_M = Q_M^{\text{vis}} + \alpha_M \Lambda_M + \varepsilon_M,$$

where $Q_M^{\text{vis}} \in L^2(\mathcal{G}_T)$, $\Lambda_M \in L^2(\mathcal{F}_T)$ is a synchronized latent factor, and

$$L_{Q_M} = Q_M^{\text{vis}} + \alpha_M \Lambda_M,$$

$$Q_M^{\text{vis}} = \mathbb{E}[Q_M \mid \mathcal{G}_T], \quad \Delta_{Q_M} = C_{\mathcal{H}_{\varpi}}(Q_M) - C_{\mathcal{G}_T}(Q_M), \quad \mathbb{E}[\varepsilon_M \mid \mathcal{G}_T, \Lambda_M] = 0.$$

Moreover, Δ_{Q_M} is the unique element of $L^2(\mathcal{H}_{\varpi})$ satisfying

$$\mathbb{E}[\Delta_{Q_M} \mid \mathcal{G}_T] = 0$$

and

$$\mathbb{E}[Q_M - Q_M^{\text{vis}} - \Delta_{Q_M} \mid \mathcal{H}_{\varpi}] = 0.$$

Equivalently, the centered projection

$$C_{\mathcal{H}_{\varpi}}(Q_M) - C_{\mathcal{G}_T}(Q_M)$$

is the canonical synchronized factor content of Q_M extracted from the sufficient-state reduction.

Proof. Expanding the square in $Q_M = \int_0^T (M^{\text{vis}} + \zeta + \eta_{\varpi})^2 dt$ gives

$$Q_M = A_M + \Xi_M + R_M.$$

By construction, $A_M \in L^2(\mathcal{G}_T)$. Since Ξ_M and R_M are square-integrable, all conditional expectations below are well defined.

Set $\mathcal{H}_{\varpi} := \mathcal{G}_T \vee \sigma(\varpi)$. Then

$$L_{Q_M} = \mathbb{E}[Q_M \mid \mathcal{H}_{\varpi}] = A_M + \Xi_M + \mathbb{E}[R_M \mid \mathcal{H}_{\varpi}]$$

because A_M and Ξ_M are $\mathcal{H}_{\varpi}$ -measurable. Taking one more conditional expectation gives

$$Q_M^{\text{vis}} = \mathbb{E}[Q_M \mid \mathcal{G}_T] = A_M + \mathbb{E}[\Xi_M + \mathbb{E}[R_M \mid \mathcal{H}_{\varpi}] \mid \mathcal{G}_T].$$

Therefore

$$\Delta_{Q_M} = L_{Q_M} - Q_M^{\text{vis}} = \Xi_M + \mathbb{E}[R_M \mid \mathcal{H}_\varpi] - \mathbb{E}[\Xi_M + \mathbb{E}[R_M \mid \mathcal{H}_\varpi] \mid \mathcal{G}_T],$$

so $\tilde{\Lambda}_M = \Delta_{Q_M}$. Since

$$\varepsilon_M = Q_M - L_{Q_M} = R_M - \mathbb{E}[R_M \mid \mathcal{H}_\varpi],$$

one has

$$Q_M = Q_M^{\text{vis}} + \tilde{\Lambda}_M + \varepsilon_M,$$

which gives the claimed decomposition after normalizing $\tilde{\Lambda}_M$.

It remains to verify the stated properties. The definition of Q_M^{vis} makes it $\mathcal{G}_T$ -measurable. Likewise $\tilde{\Lambda}_M$ is $\mathcal{H}_\varpi$ -measurable, and

$$\mathbb{E}[\tilde{\Lambda}_M \mid \mathcal{G}_T] = 0$$

by construction, so Λ_M is a synchronized latent factor whenever $\alpha_M > 0$. If $\alpha_M = 0$, the statement is immediate.

Finally,

$$\mathbb{E}[\varepsilon_M \mid \mathcal{H}_\varpi] = \mathbb{E}[R_M - \mathbb{E}[R_M \mid \mathcal{H}_\varpi] \mid \mathcal{H}_\varpi] = 0.$$

Since Λ_M is $\mathcal{H}_\varpi$ -measurable, $\sigma(\mathcal{G}_T, \Lambda_M) \subseteq \mathcal{H}_\varpi$, and therefore the tower property gives

$$\mathbb{E}[\varepsilon_M \mid \mathcal{G}_T, \Lambda_M] = \mathbb{E}[\mathbb{E}[\varepsilon_M \mid \mathcal{H}_\varpi] \mid \mathcal{G}_T, \Lambda_M] = 0.$$

Applying the tower property once more gives

$$\mathbb{E}[Q_M \mid \mathcal{G}_T] = Q_M^{\text{vis}} + \mathbb{E}[\alpha_M \Lambda_M + \varepsilon_M \mid \mathcal{G}_T] = Q_M^{\text{vis}},$$

which proves the stated identity for the visible component.

Finally, if $\Delta \in L^2(\mathcal{H}_\varpi)$ also satisfies

$$\mathbb{E}[\Delta \mid \mathcal{G}_T] = 0$$

and

$$\mathbb{E}[Q_M - Q_M^{\text{vis}} - \Delta \mid \mathcal{H}_\varpi] = 0,$$

then taking $\mathbb{E}[\cdot \mid \mathcal{H}_\varpi]$ yields

$$L_{Q_M} - Q_M^{\text{vis}} - \Delta = 0,$$

so $\Delta = \Delta_{Q_M}$ almost surely. Uniqueness and the canonical centered projection claim follow. $\square$

Remark 7.7 (What is canonical and what is still claim-specific). The canonical objects derived from the primitive market structure are the Perron eigenvectors u, v , the Perron-projected amplitude $\varpi = u^\top x$, the sufficient sigma-field $\mathcal{H}_\varpi = \mathcal{G}_T \vee \sigma(\varpi)$, and, for integrated market variance, the centered increment

$$\Delta_{Q_M} = C_{\mathcal{H}_\varpi}(Q_M) - C_{\mathcal{G}_T}(Q_M).$$

The sufficient-state reduction for Q_M therefore extracts Λ_M as a one-dimensional factor rather than choosing it ad hoc. The later variance-linked loadings and normalized directions remain claim-specific because they are inherited by projection from that aggregate object. The construction derives canonical aggregate factor content, not a universal claim-independent factor theorem.

Remark 7.8 (Scope of the sufficient-state reduction). The sufficient-state reduction first couples aggregate market volatility to the full latent envelope x , then extracts the synchronized amplitude by the Perron projection $\varpi = u^\top x$, and only then imposes a pricing-side screening condition on the aggregate off-Perron remainder. The resulting sufficient state is therefore a derived reduction rather than a direct visible-plus-factor postulate. The pricing-side centering that remains is

$$\mathbb{E}[w^\top r(t) \mid \mathcal{G}_T \vee \sigma(\varpi)] = 0$$

for the aggregate remainder, together with

$$\mathbb{E}[\eta(t) \mid \mathcal{H}_x] = 0$$

for the primitive residual. These two statements are enough for the market-variance decomposition in Proposition 7.6 and for the generic compression theorem, Theorem 7.11. The stronger

conditional independence restriction between ε_M and Λ_M enters only later, in the loading-transfer and convex-claim results such as Proposition 7.13, Corollary 7.15, and Corollary 7.23. It is satisfied, for example, when the synchronized amplitude ϖ is generated by the common trigger driving the Perron mode while the residual process η is built from innovations orthogonal to that trigger after conditioning on the visible filtration. In that form, the condition is a measurement-model separation between a common latent factor and idiosyncratic residual noise. It is used only when the generic compression bound is turned into sharper comparative-statics statements for convex claims.

Proposition 7.9 (Synchronized decomposition of variance-linked claims). *Assume Corollary 7.4, let $X = \Psi(Q_M) \in L^2(\mathcal{F}_T)$ be a variance-linked claim, and let*

$$Q_M = Q_M^{\text{vis}} + \alpha_M \Lambda_M + \varepsilon_M$$

be the decomposition from Proposition 7.6. Define

$$\begin{aligned} \mathcal{H}_M &:= \mathcal{G}_T \vee \sigma(\Lambda_M), \\ X^{\text{vis}} &:= \mathbb{E}[X \mid \mathcal{G}_T], \\ \tilde{\Lambda}_X &:= \mathbb{E}[X \mid \mathcal{H}_M] - \mathbb{E}[X \mid \mathcal{G}_T], \end{aligned}$$

and

$$\varepsilon_X := X - \mathbb{E}[X \mid \mathcal{H}_M].$$

If $\tilde{\Lambda}_X = 0$ almost surely, set $\alpha_X := 0$ and $\Lambda_X := 0$. Otherwise set

$$\alpha_X := \|\tilde{\Lambda}_X\|_{L^2}, \quad \Lambda_X := \frac{\tilde{\Lambda}_X}{\alpha_X}.$$

Then

$$X = X^{\text{vis}} + \alpha_X \Lambda_X + \varepsilon_X,$$

where $X^{\text{vis}} \in L^2(\mathcal{G}_T)$, $\Lambda_X \in L^2(\mathcal{F}_T)$ is a synchronized latent factor, and

$$X^{\text{vis}} = \mathbb{E}[X \mid \mathcal{G}_T], \quad \mathbb{E}[\varepsilon_X \mid \mathcal{G}_T, \Lambda_X] = 0.$$

Proof. The tower property gives

$$X = \mathbb{E}[X \mid \mathcal{G}_T] + (\mathbb{E}[X \mid \mathcal{H}_M] - \mathbb{E}[X \mid \mathcal{G}_T]) + (X - \mathbb{E}[X \mid \mathcal{H}_M]),$$

which gives the claimed decomposition after normalizing $\tilde{\Lambda}_X$.

By construction, $X^{\text{vis}} \in L^2(\mathcal{G}_T)$, $\tilde{\Lambda}_X$ is $\mathcal{H}_M$ -measurable, and

$$\mathbb{E}[\tilde{\Lambda}_X \mid \mathcal{G}_T] = 0.$$

Since Λ_M is measurable with respect to $\mathcal{G}_T \vee \sigma(\varpi)$, $\mathcal{H}_M \subseteq \mathcal{G}_T \vee \sigma(\varpi)$, so Λ_X is a synchronized latent factor whenever $\alpha_X > 0$. If $\alpha_X = 0$, the statement is immediate.

Finally,

$$\mathbb{E}[\varepsilon_X \mid \mathcal{H}_M] = \mathbb{E}[X - \mathbb{E}[X \mid \mathcal{H}_M] \mid \mathcal{H}_M] = 0.$$

Since Λ_X is $\mathcal{H}_M$ -measurable, $\sigma(\mathcal{G}_T, \Lambda_X) \subseteq \mathcal{H}_M$, and therefore the tower property gives

$$\mathbb{E}[\varepsilon_X \mid \mathcal{G}_T, \Lambda_X] = \mathbb{E}[\mathbb{E}[\varepsilon_X \mid \mathcal{H}_M] \mid \mathcal{G}_T, \Lambda_X] = 0.$$

□

Definition 7.10 (Visible compression operator). For any sub- σ -field $\mathcal{H} \subseteq \mathcal{F}_T$ and any square-integrable random variable $X \in L^2(\mathcal{F}_T)$, define the conditional compression of X onto $\mathcal{H}$ by

$$C_{\mathcal{H}}(X) := \mathbb{E}[X \mid \mathcal{H}].$$

It is a visible compression when $\mathcal{H} \subseteq \mathcal{G}_T$. Visible pricing summaries correspond to choosing $\mathcal{H} = \sigma(\mathfrak{S})$ for some $\mathfrak{S} \in L^2(\mathcal{G}_T)$, while visible hedge classes are closed linear subspaces of $L^2(\mathcal{G}_T)$.

Theorem 7.11 (Visible markets are compressions of latent market-variance risk). *Let $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ be convex and twice continuously differentiable with $\inf_x \Phi''(x) \geq \underline{\kappa} > 0$. In part (i), the displayed expectations involving Φ are assumed finite for the visible compression under consideration.*

(i) Suppose Corollary 7.4 holds, and let

$$Q_M = Q_M^{\text{vis}} + \alpha_M \Lambda_M + \varepsilon_M$$

be the decomposition from Proposition 7.6. Then for every sub- σ -field $\mathcal{H} \subseteq \mathcal{G}_T$,

$$\mathbb{E}[\Phi(Q_M)] - \mathbb{E}[\Phi(C_{\mathcal{H}}(Q_M))] \geq \frac{\kappa}{2} \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \mathcal{H})].$$

In particular, compression to any visible pricing summary leaves a lower-bounded latent convexity gap.

(ii) Suppose Corollary 7.4 holds, let $\mathcal{S} \subseteq L^2(\mathcal{G}_T)$ be any closed linear subspace representing visible hedge payoffs, let $X = \Psi(Q_M) \in L^2(\mathcal{F}_T)$ be a variance-linked claim, and let the decomposition

$$X = X^{\text{vis}} + \alpha_X \Lambda_X + \varepsilon_X$$

be given by Proposition 7.9. Then

$$\inf_{H \in \mathcal{S}} \mathbb{E}[(X - H)^2] \geq \alpha_X^2 \mathbb{E}[\text{Var}(\Lambda_X \mid \mathcal{G}_T)].$$

Hence every visible hedge class leaves residual risk whenever the synchronized latent mode retains conditional variance and loads nontrivially on the claim.

Proof. Fix $\mathcal{H} \subseteq \mathcal{G}_T$. Since $\mathbb{E}[\Lambda_M \mid \mathcal{G}_T] = 0$ and $\mathbb{E}[\varepsilon_M \mid \mathcal{G}_T, \Lambda_M] = 0$, the tower property gives

$$C_{\mathcal{H}}(Q_M) = \mathbb{E}[Q_M \mid \mathcal{H}] = \mathbb{E}[Q_M^{\text{vis}} \mid \mathcal{H}].$$

Set $Y := Q_M$. Since Φ is κ -strongly convex, the function $x \mapsto \Phi(x) - \frac{\kappa}{2}x^2$ is convex. Conditional Jensen gives

$$\mathbb{E}[\Phi(Y) \mid \mathcal{H}] - \Phi(\mathbb{E}[Y \mid \mathcal{H}]) \geq \frac{\kappa}{2} (\mathbb{E}[Y^2 \mid \mathcal{H}] - \mathbb{E}[Y \mid \mathcal{H}]^2).$$

After taking expectations and using $\mathbb{E}[Y \mid \mathcal{H}] = C_{\mathcal{H}}(Q_M)$,

$$\mathbb{E}[\Phi(Q_M)] - \mathbb{E}[\Phi(C_{\mathcal{H}}(Q_M))] \geq \frac{\kappa}{2} \mathbb{E}[\text{Var}(Q_M \mid \mathcal{H})].$$

Set

$$D_{\mathcal{H}} := Q_M^{\text{vis}} - \mathbb{E}[Q_M^{\text{vis}} \mid \mathcal{H}].$$

Then

$$Q_M - C_{\mathcal{H}}(Q_M) = D_{\mathcal{H}} + \alpha_M \Lambda_M + \varepsilon_M.$$

The variable $D_{\mathcal{H}}$ has conditional mean zero by construction. Since $\mathbb{E}[\Lambda_M \mid \mathcal{G}_T] = 0$ and $\mathcal{H} \subseteq \mathcal{G}_T$, also $\mathbb{E}[\Lambda_M \mid \mathcal{H}] = 0$. Since $\mathbb{E}[\varepsilon_M \mid \mathcal{G}_T, \Lambda_M] = 0$, one has $\mathbb{E}[\varepsilon_M \mid \mathcal{G}_T] = 0$ and hence $\mathbb{E}[\varepsilon_M \mid \mathcal{H}] = 0$. The cross terms vanish conditionally. Indeed,

$$\mathbb{E}[D_{\mathcal{H}} \Lambda_M \mid \mathcal{H}] = \mathbb{E}[D_{\mathcal{H}} \mathbb{E}[\Lambda_M \mid \mathcal{G}_T] \mid \mathcal{H}] = 0,$$

while $\mathbb{E}[\varepsilon_M \mid \mathcal{G}_T] = 0$ gives

$$\mathbb{E}[D_{\mathcal{H}} \varepsilon_M \mid \mathcal{H}] = \mathbb{E}[D_{\mathcal{H}} \mathbb{E}[\varepsilon_M \mid \mathcal{G}_T] \mid \mathcal{H}] = 0.$$

Finally,

$$\mathbb{E}[\Lambda_M \varepsilon_M \mid \mathcal{H}] = \mathbb{E}[\Lambda_M \mathbb{E}[\varepsilon_M \mid \mathcal{G}_T, \Lambda_M] \mid \mathcal{H}] = 0.$$

Consequently,

$$\text{Var}(Q_M \mid \mathcal{H}) = \mathbb{E}[D_{\mathcal{H}}^2 \mid \mathcal{H}] + \alpha_M^2 \text{Var}(\Lambda_M \mid \mathcal{H}) + \mathbb{E}[\varepsilon_M^2 \mid \mathcal{H}] \geq \alpha_M^2 \text{Var}(\Lambda_M \mid \mathcal{H}),$$

which proves part (i).

For part (ii), because $\mathcal{S} \subseteq L^2(\mathcal{G}_T)$,

$$\inf_{H \in \mathcal{S}} \mathbb{E}[(X - H)^2] \geq \inf_{H \in L^2(\mathcal{G}_T)} \mathbb{E}[(X - H)^2] = \mathbb{E}[\text{Var}(X \mid \mathcal{G}_T)].$$

Since $\mathbb{E}[\Lambda_X \mid \mathcal{G}_T] = 0$ and $\mathbb{E}[\varepsilon_X \mid \mathcal{G}_T, \Lambda_X] = 0$,

$$\mathbb{E}[\Lambda_X \varepsilon_X \mid \mathcal{G}_T] = \mathbb{E}[\Lambda_X \mathbb{E}[\varepsilon_X \mid \mathcal{G}_T, \Lambda_X] \mid \mathcal{G}_T] = 0,$$

and therefore

$$\text{Var}(X \mid \mathcal{G}_T) = \alpha_X^2 \text{Var}(\Lambda_X \mid \mathcal{G}_T) + \mathbb{E}[\text{Var}(\varepsilon_X \mid \mathcal{G}_T, \Lambda_X) \mid \mathcal{G}_T] \geq \alpha_X^2 \text{Var}(\Lambda_X \mid \mathcal{G}_T).$$

Taking expectations gives the bound. $\square$

Remark 7.12 (Where the economic content of compression lives). Theorem 7.11 contains a Jensen/projection step and an economic step. The Jensen inequality is generic. Strictness is carried by the synchronized latent factor itself. The compression theorem is economically restrictive only when Λ_M or Λ_X retains conditional variance after visible conditioning. The economic content is the non-visibility of the synchronized latent mode.

Proposition 7.13 (Claim loadings inherit market-variance loading). *Assume Corollary 7.4, let $X = \Psi(Q_M) \in L^2(\mathcal{F}_T)$ be a variance-linked claim, and let*

$$Q_M = Q_M^{\text{vis}} + \alpha_M \Lambda_M + \varepsilon_M, \quad X = X^{\text{vis}} + \alpha_X \Lambda_X + \varepsilon_X$$

be the decompositions from Propositions 7.6 and 7.9. Set

$$\mathcal{H}_M := \mathcal{G}_T \vee \sigma(\Lambda_M), \quad U_M := \mathbb{E}[Q_M \mid \mathcal{H}_M] = Q_M^{\text{vis}} + \alpha_M \Lambda_M.$$

Then there exists a $\mathcal{G}_T$ -measurable version of the conditional response map $\psi(\omega, \cdot)$ such that

$$\mathbb{E}[X \mid \mathcal{H}_M] = \psi(\omega, U_M)$$

almost surely. Suppose that for almost every ω , the map $u \mapsto \psi(\omega, u)$ is continuously differentiable on the essential range of U_M , and that there exist constants $0 \leq m \leq L < \infty$ such that

$$m \leq \partial_u \psi(\omega, u) \leq L$$

for almost every ω and every such u . Then

$$m^2 \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \mathcal{G}_T)] \leq \alpha_X^2 \leq L^2 \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \mathcal{G}_T)].$$

In particular, if $m > 0$, $\alpha_M \neq 0$, and Λ_M is conditionally nondegenerate given $\mathcal{G}_T$, then $\alpha_X > 0$.

Proof. Proposition 7.6 gives $\mathbb{E}[\varepsilon_M \mid \mathcal{H}_M] = 0$, so $U_M = \mathbb{E}[Q_M \mid \mathcal{H}_M]$. By Proposition 7.9,

$$\alpha_X^2 = \mathbb{E}[(\mathbb{E}[X \mid \mathcal{H}_M] - \mathbb{E}[X \mid \mathcal{G}_T])^2] = \mathbb{E}[\text{Var}(\mathbb{E}[X \mid \mathcal{H}_M] \mid \mathcal{G}_T)].$$

If $\alpha_M = 0$, then $U_M = Q_M^{\text{vis}}$ is $\mathcal{G}_T$ -measurable, so $\mathcal{H}_M = \mathcal{G}_T$ and $\alpha_X = 0$; the claimed bounds are immediate. Hence assume $\alpha_M > 0$. Because Q_M^{vis} is $\mathcal{G}_T$ -measurable and

$$\Lambda_M = \frac{U_M - Q_M^{\text{vis}}}{\alpha_M},$$

one has $\mathcal{H}_M = \mathcal{G}_T \vee \sigma(U_M)$. The Doob–Dynkin lemma, conditionally on $\mathcal{G}_T$, gives a version of $\mathbb{E}[X \mid \mathcal{H}_M]$ of the form $\psi(\omega, U_M)$, and

$$\alpha_X^2 = \mathbb{E}[\text{Var}(\psi(\omega, U_M) \mid \mathcal{G}_T)].$$

Fix a regular conditional distribution of U_M given $\mathcal{G}_T$, and let U'_M be a conditionally independent copy of U_M given $\mathcal{G}_T$. Then

$$2 \text{Var}(\psi(\omega, U_M) \mid \mathcal{G}_T) = \mathbb{E}[(\psi(\omega, U_M) - \psi(\omega, U'_M))^2 \mid \mathcal{G}_T].$$

By the mean value theorem and the derivative bounds,

$$m^2 (U_M - U'_M)^2 \leq (\psi(\omega, U_M) - \psi(\omega, U'_M))^2 \leq L^2 (U_M - U'_M)^2,$$

almost surely conditionally on $\mathcal{G}_T$. Therefore

$$m^2 \text{Var}(U_M \mid \mathcal{G}_T) \leq \text{Var}(\psi(\omega, U_M) \mid \mathcal{G}_T) \leq L^2 \text{Var}(U_M \mid \mathcal{G}_T).$$

Since $U_M = Q_M^{\text{vis}} + \alpha_M \Lambda_M$ and Q_M^{vis} is $\mathcal{G}_T$ -measurable,

$$\text{Var}(U_M \mid \mathcal{G}_T) = \alpha_M^2 \text{Var}(\Lambda_M \mid \mathcal{G}_T).$$

Taking expectations yields the claimed bounds. $\square$

Corollary 7.14 (Affine market-variance claims inherit exact loading). *Assume Corollary 7.4 and let $X = a + bQ_M$ with $a, b \in \mathbb{R}$. The decomposition can be chosen as*

$$X^{\text{vis}} = a + bQ_M^{\text{vis}}, \quad \Lambda_X = \begin{cases} \text{sgn}(b)\Lambda_M, & b \neq 0, \\ 0, & b = 0, \end{cases} \quad \varepsilon_X = b\varepsilon_M,$$

and

$$\alpha_X = |b|\alpha_M.$$

In particular,

$$\inf_{H \in \mathcal{S}} \mathbb{E}[(X - H)^2] \geq b^2 \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \mathcal{G}_T)]$$

for every closed visible hedge class $\mathcal{S} \subseteq L^2(\mathcal{G}_T)$.

Proof. The claimed decomposition follows by multiplying the decomposition of Proposition 7.6 by b and adding a . The identity $\alpha_X = |b|\alpha_M$ is immediate, and the residual-risk bound then follows from Theorem 7.11(ii). $\square$

Corollary 7.15 (Increasing convex claims inherit synchronized loading). *Assume Corollary 7.4, let $X = \Phi(Q_M) \in L^2(\mathcal{F}_T)$, and let*

$$Q_M = Q_M^{\text{vis}} + \alpha_M \Lambda_M + \varepsilon_M, \quad X = X^{\text{vis}} + \alpha_X \Lambda_X + \varepsilon_X$$

be the decompositions from Propositions 7.6 and 7.9. Suppose ε_M is conditionally independent of Λ_M given $\mathcal{G}_T$, and that $\Phi \in C^1(\mathbb{R})$ is convex with

$$0 \leq \Phi'(q) \leq L$$

for every q in the essential range of Q_M . Then

$$\alpha_X^2 \leq L^2 \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \mathcal{G}_T)].$$

If, in addition, there exists $m > 0$ such that

$$m \leq \Phi'(q)$$

throughout that same essential range, then

$$\alpha_X^2 \geq m^2 \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \mathcal{G}_T)].$$

In particular, if $\alpha_M \neq 0$ and Λ_M is conditionally nondegenerate given $\mathcal{G}_T$, then every such strictly increasing convex claim has $\alpha_X > 0$.

Proof. Set $\mathcal{H}_M := \mathcal{G}_T \vee \sigma(\Lambda_M)$ and

$$U_M := \mathbb{E}[Q_M \mid \mathcal{H}_M] = Q_M^{\text{vis}} + \alpha_M \Lambda_M.$$

Conditional independence gives

$$\mathbb{E}[X \mid \mathcal{H}_M] = \mathbb{E}[\Phi(U_M + \varepsilon_M) \mid \mathcal{H}_M] = \psi(\omega, U_M),$$

where

$$\psi(\omega, u) := \mathbb{E}[\Phi(u + \varepsilon_M) \mid \mathcal{G}_T].$$

Because Φ' is bounded between 0 and L on the essential range of Q_M , bounded convergence permits differentiation under the conditional expectation. The map $u \mapsto \psi(\omega, u)$ is continuously differentiable on the essential range of U_M and satisfies

$$0 \leq \partial_u \psi(\omega, u) \leq L$$

for almost every ω . If $\Phi' \geq m$, then likewise

$$m \leq \partial_u \psi(\omega, u) \leq L.$$

Applying Proposition 7.13 yields the stated upper and lower bounds. $\square$

Corollary 7.16 (Observer-facing consequences by contagion phase). *Assume Corollary 7.4. Within the phase structure of Theorem 6.15,*

- (i) *in the smooth-degenerate phase, no informative projected local contagion object survives at shrinking scales;*

- (ii) in the rough-observable phase, aggregate roughness is present, but the result does not imply macroscopic Perron-mode amplification;
- (iii) in the rough-synchronized phase, the market-variance decomposition of Proposition 7.6 and the inherited claim decomposition of Proposition 7.9 produce positive lower bounds of the form in Theorem 7.11 for any variance-linked claim with nonzero latent loading and a conditionally nondegenerate synchronized latent factor;
- (iv) in the activated synchronized phase, those same compression bounds coexist with the jump-like coarse behavior of Theorem 6.15(iv).

Proof. Part (i) is Theorem 6.15(i). Part (ii) is Theorem 6.15(ii). For part (iii), Theorem 6.15(iii) identifies the rough-synchronized regime, and under Corollary 7.4, Theorem 7.11 gives the stated strictly positive lower bounds whenever the synchronized factor loads nontrivially and retains conditional variance. Part (iv) combines that same compression theorem, still under Corollary 7.4, with Theorem 6.15(iv). $\square$

Proposition 7.17 (Latent convexity gap for market-variance claims). *Assume Corollary 7.4. Let $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ be convex and twice continuously differentiable, with $\inf_x \Phi''(x) \geq \underline{\kappa} > 0$, and assume that the displayed expectations below are finite. Then*

$$\mathbb{E}[\Phi(Q_M)] - \mathbb{E}[\Phi(Q_M^{\text{vis}})] \geq \frac{\underline{\kappa}}{2} \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \mathcal{G}_T)].$$

If $\alpha_M \neq 0$ and Λ_M is conditionally nondegenerate, the inequality is strict.

Proof. Apply Theorem 7.11(i) with $\mathcal{H} = \mathcal{G}_T$. Then

$$C_{\mathcal{G}_T}(Q_M) = \mathbb{E}[Q_M \mid \mathcal{G}_T] = Q_M^{\text{vis}},$$

which gives the claim. $\square$

Definition 7.18 (Visible pricing summary). A *visible pricing summary* is any square-integrable random variable $\mathfrak{S} \in L^2(\mathcal{G}_T)$. In particular, the leading-order short-maturity skew statistics identified on the pricing side in [1] are visible pricing summaries. In the latent contagion regime they are measurable with respect to the visible pricing filtration and do not recover the synchronized latent mode.

The compression bound is a projection statement. The following corollaries place that statement on standard systemic-risk objects. CoVaR-type tail summaries and Diebold–Yilmaz connectedness measures enter as visible summaries. SRISK-style capital-shortfall functionals enter as increasing convex claims on the market-variance state.

Corollary 7.19 (Visible skew is only a partial $\mathbb{Q}$ -summary). *Assume Corollary 7.4, let $\Phi: \mathbb{R} \rightarrow \mathbb{R}$ be convex and twice continuously differentiable with $\inf_x \Phi''(x) \geq \underline{\kappa} > 0$, and let $\mathfrak{S} \in L^2(\mathcal{G}_T)$ be any visible pricing summary. Assume that the displayed expectations below are finite. Then*

$$\mathbb{E}[\Phi(Q_M)] - \mathbb{E}[\Phi(\mathbb{E}[Q_M \mid \sigma(\mathfrak{S})])] \geq \frac{\underline{\kappa}}{2} \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \sigma(\mathfrak{S}))].$$

If $\alpha_M \neq 0$ and Λ_M is conditionally nondegenerate given $\sigma(\mathfrak{S})$, the inequality is strict.

Proof. Apply Theorem 7.11(i) with $\mathcal{H} = \sigma(\mathfrak{S})$. $\square$

Corollary 7.20 (CoVaR-type visible summaries remain incomplete). *Let $\mathfrak{S}_{\text{CoVaR}} \in L^2(\mathcal{G}_T)$ be any visible CoVaR-type summary in the sense of [13]. Under the assumptions of Corollary 7.19,*

$$\mathbb{E}[\Phi(Q_M)] - \mathbb{E}[\Phi(\mathbb{E}[Q_M \mid \sigma(\mathfrak{S}_{\text{CoVaR}})])] \geq \frac{\underline{\kappa}}{2} \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \sigma(\mathfrak{S}_{\text{CoVaR}}))].$$

If $\alpha_M \neq 0$ and Λ_M is conditionally nondegenerate given $\sigma(\mathfrak{S}_{\text{CoVaR}})$, the inequality is strict. Thus conditioning on a visible CoVaR-type tail-risk statistic does not recover latent market-variance convexity.

Proof. Apply Corollary 7.19 with $\mathfrak{S} = \mathfrak{S}_{\text{CoVaR}}$. $\square$

Corollary 7.21 (Connectedness summaries remain partial). *Let $\mathfrak{S}_{\text{DY}} \in L^2(\mathcal{G}_T)$ be any visible connectedness or variance-decomposition summary of the type studied by [9, 10]. Under the assumptions of Corollary 7.19,*

$$\mathbb{E}[\Phi(Q_M)] - \mathbb{E}[\Phi(\mathbb{E}[Q_M \mid \sigma(\mathfrak{S}_{\text{DY}})])] \geq \frac{\kappa}{2} \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \sigma(\mathfrak{S}_{\text{DY}}))].$$

If $\alpha_M \neq 0$ and Λ_M is conditionally nondegenerate given $\sigma(\mathfrak{S}_{\text{DY}})$, the inequality is strict. Thus a visible connectedness statistic can summarize spillovers without recovering the synchronized latent mode.

Proof. Apply Corollary 7.19 with $\mathfrak{S} = \mathfrak{S}_{\text{DY}}$. $\square$

Corollary 7.22 (Visible skew can be fully visible and still incomplete). *Let $\mathfrak{S}_{\text{skew}} \in L^2(\mathcal{G}_T)$ be any leading-order short-maturity skew summary of the type identified in [1]. Under the assumptions of Corollary 7.19,*

$$\mathbb{E}[\Phi(Q_M)] - \mathbb{E}[\Phi(\mathbb{E}[Q_M \mid \sigma(\mathfrak{S}_{\text{skew}})])] \geq \frac{\kappa}{2} \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \sigma(\mathfrak{S}_{\text{skew}}))].$$

If $\alpha_M \neq 0$ and Λ_M is conditionally nondegenerate given $\sigma(\mathfrak{S}_{\text{skew}})$, the inequality is strict. Hence a leading-order skew statistic can be entirely visible under $\mathbb{Q}$ while latent market-variance convexity remains unrecovered.

Proof. Apply Corollary 7.19 with $\mathfrak{S} = \mathfrak{S}_{\text{skew}}$, using the fact that $\mathfrak{S}_{\text{skew}}$ is by construction a visible pricing summary. $\square$

Corollary 7.23 (SRISK-style shortfall functionals remain compressed and under-hedged). *Let $\Phi_{\text{SR}} \in C^2(\mathbb{R})$ be increasing and convex, and suppose there exist constants $\kappa_{\text{SR}}, m_{\text{SR}}, L_{\text{SR}} > 0$ such that*

$$\kappa_{\text{SR}} \leq \Phi_{\text{SR}}''(q), \quad m_{\text{SR}} \leq \Phi_{\text{SR}}'(q) \leq L_{\text{SR}}$$

for every q in the essential range of Q_M . Set

$$X_{\text{SR}} := \Phi_{\text{SR}}(Q_M),$$

and interpret X_{SR} as an SRISK-style capital-shortfall functional in the sense of [14]. Let $\mathfrak{S}_{\text{SR}} \in L^2(\mathcal{G}_T)$ be any visible shortfall summary and let $\mathcal{S} \subseteq L^2(\mathcal{G}_T)$ be any closed visible hedge class. Under Corollary 7.4 and conditional independence of ε_M and Λ_M given $\mathcal{G}_T$,

$$\mathbb{E}[\Phi_{\text{SR}}(Q_M)] - \mathbb{E}[\Phi_{\text{SR}}(\mathbb{E}[Q_M \mid \sigma(\mathfrak{S}_{\text{SR}})])] \geq \frac{\kappa_{\text{SR}}}{2} \alpha_M^2 \mathbb{E}[\text{Var}(\Lambda_M \mid \sigma(\mathfrak{S}_{\text{SR}}))],$$

and

$$\inf_{H \in \mathcal{S}} \mathbb{E}[(X_{\text{SR}} - H)^2] \geq \alpha_{X_{\text{SR}}}^2 \mathbb{E}[\text{Var}(\Lambda_{X_{\text{SR}}} \mid \mathcal{G}_T)].$$

If $\alpha_M \neq 0$ and Λ_M is conditionally nondegenerate given $\mathcal{G}_T$, then $\alpha_{X_{\text{SR}}} > 0$, so both the compression gap and the hedge shortfall are strictly positive whenever visible conditioning fails to identify the synchronized latent mode.

Proof. Apply Theorem 7.11(i) with $\Phi = \Phi_{\text{SR}}$ and $\mathcal{H} = \sigma(\mathfrak{S}_{\text{SR}})$. Apply Theorem 7.11(ii) to the variance-linked claim $X_{\text{SR}} = \Phi_{\text{SR}}(Q_M)$. The lower bound $\alpha_{X_{\text{SR}}} > 0$ follows from Corollary 7.15. $\square$

Remark 7.24 (Visible skew and market-variance convexity). Proposition 7.17 formalizes the failure of visible skew to capture full market-variance convexity. A variance swap corresponds to the linear payoff $\Phi(x) = x$, so the Jensen gap is zero. The missing convexity appears only for nonlinear convex claims written on Q_M . Any pricing summary restricted to Q_M^{vis} misses the conditional variance carried by the synchronized latent factor Λ_M . The gap is largest when the Perron mode concentrates near the jump-like regime.

Corollary 7.25 (Visible hedge shortfall in complete visible markets). *Assume the visible trading market is complete relative to $\mathcal{G}_T$, so that the attainable visible payoffs coincide with $L^2(\mathcal{G}_T)$. Let $\Pi(X)$ be the L^2 -projection of X onto the closed span of visible self-financing gains. Then*

$$\mathbb{E}[(X - \Pi(X))^2] \geq \alpha_X^2 \mathbb{E}[\text{Var}(\Lambda_X \mid \mathcal{G}_T)].$$

Proof. Apply Theorem 7.11(ii) with $\mathcal{S} = L^2(\mathcal{G}_T)$. Under completeness, $\Pi(X) = \mathbb{E}[X \mid \mathcal{G}_T]$, so the lower bound is attained at the best visible hedge. $\square$

8. CONCLUSION

The market experiment studied here is directed and local. The conclusion concerns the class in which roughness is unavoidable, not arbitrary stochastic-volatility models: active off-diagonal Volterra channels that remain source-screened visible, carry market loading, and admit a genuinely shrinking projected normalization. Smooth contagion is then locally degenerate at high frequency. Once the market keeps active admissible asymmetric channels without leading-order aggregate cancellation, the market-volatility law inherits the roughest market-dominant exponent.

The same admissible edges determine the systemic phase. The screened synchronization operator has a Perron threshold that separates rough-local propagation from rough-synchronized amplification. Above the threshold, the leading latent envelope aligns with a common mode. Under threshold-triggered activation, that continuous synchronized mode becomes jump-like for a coarse observer. The theorem isolates one endogenous route from continuous microscopic Volterra dynamics to jump-like aggregate episodes. It does not exclude exogenous crises, primitive discontinuities, or dislocations outside the activated synchronized channel.

Pricing and hedging stop at public compression. For integrated market variance $Q_M = \int_0^T M(t)^2 dt$, visible pricing summaries recover a compression of market-variance convexity, not the full latent convexity carried by the synchronized mode. Visible hedge classes leave a lower-bounded residual exposure whenever that mode retains conditional variance. Leading-order skew-type summaries remain valid $\mathbb{Q}$ -objects, but they are incomplete relative to latent market-variance convexity and hedge risk. The same projection boundary applies to CoVaR-type summaries, SRISK-style shortfall functionals, and visible connectedness measures. The corresponding risk-neutral problem is to price the same Perron state on matched option panels, where the $N \times N$ structure is read through $\mathbb{P}/\mathbb{Q}$ compression rather than realized-volatility geometry alone.

Appendix B is an empirical illustration, not a theorem. On the Oxford-Man panel, the Perron-projected aggregate component built from the source-screened operator tracks the equal-weight aggregate stress proxy closely, while the aggregate off-Perron remainder is small across rolling windows. The appendix does not prove the phase results or full latent recovery. It audits the finite-panel Perron projection that the theory makes central.

APPENDIX A. A SUPPORTING SHORT-SCALE VOLTERRA LEMMA

Lemma A.1 (Riemann–Liouville increment scaling). *Fix $H \in (0, 1)$ and $t \in (0, T)$, and define*

$$\Gamma_{H,t,h}(s) := (t + h - s)^{H-1/2} \mathbf{1}_{\{s < t+h\}} - (t - s)^{H-1/2} \mathbf{1}_{\{s < t\}}.$$

Then

$$\|\Gamma_{H,t,h}\|_{L^2([0,t+h])}^2 = \kappa(H)h^{2H} + o(h^{2H})$$

as $h \downarrow 0$, where

$$\kappa(H) := \int_0^\infty ((1+r)^{H-1/2} - r^{H-1/2})^2 dr + \frac{1}{2H} \in (0, \infty).$$

Proof. Split the square norm at time t :

$$\|\Gamma_{H,t,h}\|_{L^2([0,t+h])}^2 = \int_0^t ((t+h-s)^{H-1/2} - (t-s)^{H-1/2})^2 ds + \int_t^{t+h} (t+h-s)^{2H-1} ds.$$

In the first term, substitute $r = (t-s)/h$; in the second, substitute $r = (t+h-s)/h$. This gives

$$\|\Gamma_{H,t,h}\|_{L^2([0,t+h])}^2 = h^{2H} \int_0^{t/h} ((1+r)^{H-1/2} - r^{H-1/2})^2 dr + \frac{h^{2H}}{2H}.$$

Near zero the integrand is bounded by an integrable multiple of $1 + r^{2H-1}$, and at infinity it behaves like r^{2H-3} . It is therefore integrable for every $H \in (0, 1)$. Monotone convergence on the expanding interval $[0, t/h]$, equivalently dominated convergence after truncation, yields

$$\int_0^{t/h} ((1+r)^{H-1/2} - r^{H-1/2})^2 dr \longrightarrow \int_0^\infty ((1+r)^{H-1/2} - r^{H-1/2})^2 dr,$$

which proves the expansion. $\square$

APPENDIX B. OXFORD-MAN PERRON ILLUSTRATION

This appendix reports an empirical illustration of the Perron-projected sufficient-state reduction developed in Section 7. Its role is narrow. On the balanced eight-index Oxford-Man panel of [15], the market-scale screened operator has a stable dominant Perron direction, the corresponding Perron amplitude tracks an aggregate volatility-stress proxy closely, and the aggregate off-Perron remainder is quantitatively small. The appendix supplies neither a proof of the theorems nor a full empirical section.

Setup. The illustration uses the balanced eight-index panel

$$(.SPX, .IXIC, .FTSE, .GDAXI, .FCHI, .STOXX50E, .N225, .HSI)$$

from 2000-02-02 through 2018-06-27 with the realized-kernel proxy `rk_parzen`. For each rolling window of length 504 trading days and step 21, define the directed empirical screened operator

$$\hat{A}_{ij,t} := \sqrt{\max\{\widehat{\text{projinfo}}_{ij,t}, 0\}}, \quad i \neq j,$$

from the pairwise projected-information outputs of the projected-score procedure in [7], and set $\hat{A}_{ii,t} = 0$. The left and right Perron vectors $(\hat{u}_t, \hat{v}_t)$ are normalized by $\hat{u}_t^\top \hat{v}_t = 1$. The rolling-centered log realized-volatility vector $\hat{x}_d^{(t)}$ is used as an observable envelope proxy. The empirical Perron amplitude and off-Perron remainder are

$$\hat{\omega}_d^{(t)} := \hat{u}_t^\top \hat{x}_d^{(t)}, \quad \hat{r}_d^{(t)} := \hat{x}_d^{(t)} - \hat{\omega}_d^{(t)} \hat{v}_t.$$

With equal deterministic weights $w_i = 1/8$, the aggregate stress proxy is

$$\hat{M}_d^{(t)} := w^\top \hat{x}_d^{(t)},$$

with Perron-implied aggregate component $\hat{M}_{P,d}^{(t)} := (w^\top \hat{v}_t) \hat{\omega}_d^{(t)}$ and off-Perron aggregate remainder $\hat{M}_{R,d}^{(t)} := w^\top \hat{r}_d^{(t)}$.

The empirical compression is summarized by the full-sample spectral radius $\hat{\rho}_1$, the concentration score $1 - \hat{\rho}_2/\hat{\rho}_1$, the rolling explained share

$$R_t^2 := 1 - \frac{\sum_d (\hat{M}_d^{(t)} - \hat{M}_{P,d}^{(t)})^2}{\sum_d (\hat{M}_d^{(t)})^2},$$

and the rolling off-Perron share

$$S_t^{\text{off}} := \frac{\sum_d (\hat{M}_{R,d}^{(t)})^2}{\sum_d (\hat{M}_d^{(t)})^2}.$$

The rolling analysis contains 171 windows.

Figure 1 reports the full-sample screened operator. The dominant eigenvalue is separated from the subleading radius, and the Perron weights concentrate on the same cross-market core that dominates the market-level roughness analysis.

Table 1 reports a full-sample spectral radius of 9.2354 and a concentration score of 0.8438. For the aggregate reduction, the Perron-implied component captures nearly all of the equal-weight aggregate stress proxy in this construction. The full-sample explained share is 0.9999, while the off-Perron share is 0.0001. This is consistent with the aggregate-first sufficiency claim made in Section 7; it is not a claim that the Perron component exhausts all latent market structure.

Figure 2 reports rolling stability of the Perron construction. Across the 171 rolling windows, the spectral radius stays between 6.3277 and 11.0111, the concentration score remains between




FIGURE 1. Full-sample empirical Perron operator on the balanced eight-index Oxford-Man panel. The operator entries are $\sqrt{\widehat{\text{projinfo}}_{ij}}$ on directed off-diagonal edges, with zero diagonal.

Statistic	Value
Full-sample spectral radius $\hat{\rho}_1$	9.2354
Subleading spectral radius $\hat{\rho}_2$	1.4425
Concentration score $1 - \hat{\rho}_2/\hat{\rho}_1$	0.8438
Aggregate explained share R^2	0.9999
Off-Perron share S^{off}	0.0001
Top right-load components $\hat{v}$	Nasdaq (0.39), DAX (0.37), S&P 500 (0.36)
Top left-load components $\hat{u}$	CAC 40 (0.37), DAX (0.37), Euro Stoxx 50 (0.36)

TABLE 1. Full-sample Perron summary for the Oxford-Man illustration.

Statistic	Crisis	Non-crisis
Mean spectral radius $\hat{\rho}_1$	7.8085	7.5213
Mean concentration $1 - \hat{\rho}_2/\hat{\rho}_1$	0.8354	0.8308
Mean aggregate explained share R_t^2	0.9994	0.9995
Mean off-Perron share S_t^{off}	0.0006	0.0005

TABLE 2. Rolling crisis versus non-crisis Perron summaries.

0.7996 and 0.8495, the aggregate explained share never falls below 0.9977, and the off-Perron share never rises above 0.0023. The aggregate explained share remains high in both crisis and non-crisis windows.




FIGURE 2. Rolling Perron summaries on 504-day windows stepped every 21 trading days. The panels show the spectral radius $\hat{\rho}_{1,t}$, the concentration score $1 - \hat{\rho}_{2,t}/\hat{\rho}_{1,t}$, and the aggregate explained share R_t^2 together with the off-Perron share S_t^{off} .

Table 2 reports the same quantities at the regime level. Crisis windows have a somewhat larger mean spectral radius, but both crisis and non-crisis periods retain strong Perron concentration and a very small aggregate off-Perron share. The empirical statement is quantitative. On this realized-volatility panel, the Perron projection captures the dominant aggregate direction while leaving only a small aggregate remainder.

REFERENCES

- [1] J. Vidal Llauredó (2026). *Latent Volatility Contagion in Rough Volatility Models: Spectral Thresholds and Detectability*. SSRN. <https://doi.org/10.2139/ssrn.6520718>.
- [2] J. Gatheral, T. Jaisson, and M. Rosenbaum, *Volatility is rough*, *Quantitative Finance* **18** (2018), no. 6, 933–949, doi:10.1080/14697688.2017.1393551.

- [3] G. Livieri, S. Mouti, A. Pallavicini, and M. Rosenbaum, *Rough volatility: Evidence from option prices*, ISE Transactions **50** (2018), no. 9, 767–776, doi:10.1080/24725854.2018.1444297.
- [4] C. Bayer, P. Friz, and J. Gatheral, *Pricing under rough volatility*, Quantitative Finance **16** (2016), no. 6, 887–904, doi:10.1080/14697688.2015.1099717.
- [5] O. El Euch and M. Rosenbaum, *The characteristic function of rough Heston models*, Mathematical Finance **29** (2019), no. 1, 3–38, doi:10.1111/mafi.12173.
- [6] E. Abi Jaber, M. Larsson, and S. Pulido, *Affine Volterra processes*, Annals of Applied Probability **29** (2019), no. 5, 3155–3200, doi:10.1214/19-AAP1477.
- [7] J. Vidal Llauradó (2026). *Detecting Latent Volatility Contagion: A Projected Score Estimator for Rough Volatility Models*. SSRN. <https://doi.org/10.2139/ssrn.6616118>.
- [8] R. Cont and P. Das, *Rough volatility: Fact or artefact?*, Sankhyā B **86** (2024), no. 1, 191–223, doi:10.1007/s13571-024-00322-2.
- [9] F. X. Diebold and K. Yilmaz, *Measuring financial asset return and volatility spillovers, with application to global equity markets*, The Economic Journal **119** (2009), no. 534, 158–171, doi:10.1111/j.1468-0297.2008.02208.x.
- [10] F. X. Diebold and K. Yilmaz, *On the network topology of variance decompositions: Measuring the connectedness of financial firms*, Journal of Econometrics **182** (2014), no. 1, 119–134, doi:10.1016/j.jeconom.2014.04.012.
- [11] K. J. Forbes and R. Rigobon, *No contagion, only interdependence: Measuring stock market comovements*, The Journal of Finance **57** (2002), no. 5, 2223–2261, doi:10.1111/0022-1082.00494.
- [12] J. Vidal Llauradó (2026). *Dynamic Observability of Latent Contagion: Sequential Revelation and Observable Screening*. SSRN. <https://doi.org/10.2139/ssrn.6539199>.
- [13] T. Adrian and M. K. Brunnermeier, *CoVaR*, American Economic Review **106** (2016), no. 7, 1705–1741, doi:10.1257/aer.20120555.
- [14] C. Brownlees and R. F. Engle, *SRISK: A conditional capital shortfall measure of systemic risk*, The Review of Financial Studies **30** (2017), no. 1, 48–79, doi:10.1093/rfs/hhw060.
- [15] G. Heber, A. Lunde, N. Shephard, and K. Sheppard, *Oxford-Man Institute’s realized library*, version 0.1, Oxford-Man Institute, University of Oxford, 2009, <https://realized.oxford-man.ox.ac.uk/>.

Email address: vidal.llaurado@gmail.com