

WORKING PAPER

Scale, Shape, Location, and Speed: Decomposing Distributional Volatility Persistence

Fuli Yang

GARCH Institute | Harvard Square, Cambridge, MA 02138, United States

dr@garch.ac | ORCID: 0009-0005-4055-7561 | September 2026

Abstract

Volatility persistence is usually represented by a scalar conditional variance, even though return distributions can remain abnormal through dispersion, center shifts, shape distortions, or continued movement. This paper develops a leakage-safe Wasserstein framework that separates these features into four rolling coordinates: scale, standardized shape, signed location, and distributional speed. Using daily S&P 500, FTSE 100, Nikkei 225, and Hang Seng returns from January 2000 to June 2026, the robust scale coordinate is strongly associated with 22-day realized volatility, whereas standardized shape is approximately orthogonal to the current volatility level. Signed location is negatively associated with realized volatility, and speed occupies an intermediate position because it measures one-step movement rather than state level. Scale is also the most synchronized coordinate across markets; location and speed are moderately synchronized, while shape is largely market-specific. Paired block-bootstrap tests, multiplicity corrections, block-length sensitivity checks, and a variance-matched heavy-tailed placebo support the separation of the coordinates. The results establish measurement nonredundancy rather than forecasting superiority or causal transmission. The framework therefore recasts volatility persistence as a multidimensional financial-risk state and provides testable inputs for subsequent state-model, recovery, and forecasting comparisons.

Keywords: volatility persistence; Wasserstein distance; distributional regimes; realized volatility; market states; optimal transport

JEL classification: C14; C22; C58; G12; G17

1. Introduction

Volatility is persistent, but the object that persists is rarely interrogated. In ARCH, GARCH, stochastic-volatility, and realized-volatility models, persistence is generally attached to a scalar conditional variance or to a function of past realized variation. This scalar treatment is economically useful and empirically powerful. It is nevertheless incomplete whenever two return windows have similar dispersion but different skewness, tail balance, or rates of distributional change.

Consider three markets with the same current 22-day standard deviation. The first may have a nearly symmetric return distribution, the second a pronounced lower tail, and the third a distribution that only arrived at its current configuration after an unusually large one-day displacement. A variance-based state variable assigns these observations nearly the same location in state space. A distributional representation does not. The distinction matters for measurement before it matters for prediction: if the underlying state is multidimensional, a single persistence coefficient combines objects that may have different durations, recovery paths, and degrees of cross-market synchronization.

This paper develops a decomposition of distributional volatility persistence into scale, shape, location, and speed. Scale records current dispersion. Shape compares location- and scale-standardized rolling distributions and therefore isolates asymmetry and tail configuration from volatility level. Location retains the signed shift of the distributional center rather than folding it into scale or shape. Speed measures the one-step Wasserstein motion of the rolling empirical distribution. The decomposition uses only information available at each measurement date, and its reference objects are frozen using pre-2007 data.

The first-stage empirical evidence is deliberately modest but informative. Across four equity markets, scale is strongly associated with current realized volatility, shape is almost uncorrelated with it at the baseline 22-day horizon, signed location is negatively associated with it, and speed is moderately positively related to it. Scale displays substantial window-matched autocorrelation; location persistence is market- and horizon-dependent; and shape persistence is weak at short horizons. Cross-market comovement follows a similarly ordered pattern: realized volatility and scale are highly synchronized, location and speed are moderately synchronized, and shape is mostly market-specific. A controlled placebo shows that the shape statistic detects standardized asymmetry even after sample location and variance are matched.

The contribution is therefore a measurement contribution. The paper does not claim that the proposed coordinates forecast returns, crises, or future volatility; it does not yet claim superiority over an ordered Markov-switching GJR-GARCH model. Those questions require separate out-of-sample and

state-comparison tests. The present evidence establishes that the coordinates are operational, empirically nonredundant, and suitable for those tests.

The paper connects three literatures. First, it complements scalar volatility models and the heterogeneous autoregressive tradition by asking which aspect of a return distribution is persistent. Second, it relates to regime-switching methods but represents state continuously rather than through a single latent probability. Third, it builds on optimal-transport approaches to market regimes while replacing clustering with an interpretable coordinate system. The remainder of the paper develops this positioning, defines the measures, documents the first-stage evidence, and states the boundary of the current claims.

2. Related literature and positioning

2.1 Scalar volatility persistence

The empirical persistence of volatility is foundational to modern financial econometrics. GARCH models encode it through conditional-variance recursion (Bollerslev, 1986), while realized-volatility research measures ex post variation at high frequency and models its long memory and heterogeneous horizons (Andersen et al., 2001, 2003; Barndorff-Nielsen and Shephard, 2002; Corsi, 2009). Model comparison repeatedly shows that parsimonious HAR-type specifications remain difficult benchmarks to beat (Hansen and Lunde, 2005; Patton, 2011). Options-driven rough-volatility augmentation is a recent complementary forecasting direction (Fan, Wang, and Ye, 2026). Recent surveys and foundation-model comparisons reinforce the broader point: flexible models can add value, but gains are uneven and specification discipline remains central (Leushuis and Petkov, 2026; Brini, 2026).

This paper does not replace these models. It asks a prior measurement question: when volatility appears persistent, is persistence concentrated in dispersion, standardized shape, location, or the rate of distributional motion? The distinction is compatible with a HAR or GARCH forecasting layer, but it is not implied by one.

2.2 Regime switching and continuous state measurement

Markov-switching models represent persistence through transition probabilities among latent states (Hamilton, 1989). Regime-switching volatility models extend this logic to conditional distributions and allow the variance response to differ across states (Gray, 1996). The asymmetric GJR-GARCH formulation is a relevant benchmark because it permits negative and positive innovations to have different volatility effects (Glosten, Jagannathan, and Runkle, 1993). These discrete probabilities are natural benchmarks for a continuous distributional state coordinate. Yet a high-state probability ordinarily summarizes the

likelihood of a variance regime, not a decomposition of the observed return distribution. In the full empirical design, the primary benchmark is the filtered high-state probability from an ordered MS-GJR-GARCH-t model. Smoothed probabilities are reserved for robustness because they use future observations.

2.3 Wasserstein geometry and market regimes

Optimal transport supplies a natural geometry for empirical distributions. In one dimension, the first Wasserstein distance is the integral of absolute quantile differences, making it interpretable in return units and straightforward to compute for rolling windows. Wasserstein barycenters and distributional data analysis provide the broader statistical foundation (Agueh and Carlier, 2011; Panaretos and Zemel, 2019). For serially indexed distributions, Wasserstein autoregressive models provide a direct precedent for estimating dependence in density-valued time series (Zhang, Kokoszka, and Petersen, 2022). Horvath, Issa, and Muguruza (2024) use Wasserstein k-means to segment market regimes. The present paper differs in object and estimand: it does not cluster windows into regimes or forecast an unobserved density, but decomposes a continuously observed rolling distribution into economically labeled coordinates and studies their persistence.

This distinction also separates the paper from earlier Distributional Regime Persistence work, where a normalized distance to a frozen prototype was treated as a composite state measure. Here the composite is opened. Scale is defined directly from current dispersion—not as distance to a high-volatility prototype—and shape is constructed only after location and scale normalization. Speed is movement between adjacent observed distributions, not distance from a fixed state.

3. Measurement framework

3.1 Rolling distributions and information set

Let $r_{i,t}$ denote the daily log return for market i . For a rolling window of length L , define the empirical distribution $W_{i,t}$ from observations $\{r_{i,t-L+1}, \dots, r_{i,t}\}$. All transformations, normalizers, and reference distributions used at date t are measurable with respect to information available at or before t . The baseline uses $L = 22$ trading days; $L = 11$ and $L = 66$ are pre-specified sensitivity windows.

$$W_1(F, G) = \int_0^1 |F^{-1}(u) - G^{-1}(u)| du. \quad (1)$$

The empirical implementation evaluates quantile functions on a common probability grid. The frozen reference prototype is constructed from the top 20% of pre-2007 rolling windows ranked by realized volatility. This choice supplies a stable calibration object without using subsequent crisis labels. Importantly, the reference distance is not itself called the scale coordinate.

3.2 Scale

The primary scale coordinate is the logarithm of a robust IQR-based dispersion estimate for the current window. Rolling standard deviation is retained as a benchmark. A robust measure limits the influence of isolated observations while preserving the economic meaning of current dispersion (Rousseeuw and Croux, 1993). Under a Gaussian benchmark, the interquartile range equals $2\Phi^{-1}(0.75)\sigma \approx 1.349\sigma$, which supplies the consistency constant in Equation (2). The absolute gap between current and prototype scale may be reported as a secondary diagnostic, but proximity to a high-volatility prototype can mechanically reverse its relation to current volatility and is therefore unsuitable as the primary scale state variable.

$$Scale_{i,t} = \log[(Q_{0.75}(W_{i,t}) - Q_{0.25}(W_{i,t}))/1.349]. \quad (2)$$

3.3 Shape and tail shape

Let $m_{i,t}$ be the window median and $s_{i,t}$ the robust scale. Standardize every observation as $z_{i,\tau} = (r_{i,\tau} - m_{i,t})/s_{i,t}$ and denote the resulting empirical distribution by $\tilde{W}_{i,t}$. Apply the same location-scale normalization to the frozen reference. The shape coordinate is the first Wasserstein distance between standardized distributions.

$$Shape_{i,t} = W_1(W_{i,t}, P_{i,0}). \quad (3)$$

Because both distributions have standardized center and scale, the remaining distance is attributable to differences in quantile configuration: asymmetry, tail thickness, and other shape features. Lower- and upper-tail versions restrict the quantile integral to pre-specified tail regions. These tail coordinates are secondary until their cutoffs and multiplicity adjustment are frozen.

3.4 Location

Location is the signed difference between the current window median and the frozen reference median, expressed in units of the current robust scale. Positive values indicate an upward shift and negative values a downward shift. Retaining the sign is essential: it prevents directional drift from being misattributed to scale or shape and permits asymmetric association with volatility. A mean-based version is a sensitivity measure, not the primary specification, because it is more exposed to extreme returns.

$$Location_{i,t} = (m_{i,t} - m_{i,0}) / s_{i,t}. \quad (4)$$

3.5 Speed

Distributional speed is the distance between adjacent rolling distributions. It measures motion, not position. A market can display high speed while scale remains moderate if the window is reconfiguring rapidly, and it can display low speed at high scale if a volatile distribution has stabilized.

$$Speed_{i,t} = W_1(W_{i,t}, W_{i,t-1}). \quad (5)$$

A standardized shape-speed statistic repeats this calculation after location-scale normalization, as defined explicitly in Equation (6). Because adjacent rolling windows share $L - 1$ observations, serial dependence is expected by construction; inference therefore uses moving-block resampling and persistence is evaluated at the window-matched lag rather than only at lag one.

$$ShapeSpeed_{i,t} = W_1(\bar{W}_{i,t}, \bar{W}_{i,t-1}). \quad (6)$$

3.6 What the decomposition is—and is not

The four coordinates form an interpretable measurement system, not an exact additive identity for raw Wasserstein distance. In general, raw distance is not equal to scale plus shape plus location, and speed is a temporal derivative rather than a component of cross-sectional position. The empirical tests therefore assess separability, conditional association, persistence heterogeneity, and synchronization—not accounting shares that sum to 100%.

4. Data and leakage-safe empirical design

The sample contains daily percentage log returns for four broad equity indices from 3 January 2000 through 30 June 2026. The first usable return is 4 January 2000. Market calendars are preserved for univariate construction and aligned only for cross-market synchronization. Table 1 reports sample coverage and the baseline window length.

Table 1. Sample coverage and baseline window length

Market	Index	Usable returns	First usable	Last date	Baseline L
United States	S&P 500	6,841	2000-01-04	2026-06-30	22
United Kingdom	FTSE 100	6,852	2000-01-04	2026-06-30	22
Japan	Nikkei 225	6,834	2000-01-04	2026-06-30	22
Hong Kong	Hang Seng	6,530	2000-01-04	2026-06-30	22

Notes: Returns are daily percentage log returns. Component construction uses market-specific trading days. Cross-market analyses use common calendar dates.

The main prototype uses the top quintile of pre-2007 windows by realized volatility. For $L = 22$, this produces 360 prototype windows in each market. The prototype and normalization rules are frozen before evaluating post-calibration observations. No event labels, future crisis dates, or market-specific tuning parameters are used to define the baseline components.

Uncertainty for first-stage correlations and window-matched autocorrelations uses 1,999 circular moving-block bootstrap replications with seed 20260903. The baseline block length is $\max(5, L/2)$. For autocorrelations, blocks resample the paired lag vectors (x_t, x_{t-L}) , preserving the target lag relation. Because adjacent rolling windows share $L - 1$ observations, the baseline is supplemented at $L = 22$ by block lengths 11, 22, and 44 and at $L = 66$ by block lengths 33, 66, and 132. These deterministic sensitivity grids span one-half, one, and two rolling-window lengths and complement data-driven block-length methods (Politis and White, 2004; Patton, Politis, and White, 2009).

4.1 Pre-specified hypotheses

H1—Component separability. Standardized shape retains variation that is not explained by current volatility scale.

H2—Persistence heterogeneity. Scale, shape, location, and speed have different persistence and recovery profiles.

H3—Conditional motion. High speed identifies transition dynamics not reducible to a high current scale level.

H4—Cross-market asymmetry. Synchronization is strongest for scale among the four coordinates; no strict ordering is posited among shape, location, and speed.

H5—Model distinction. The coordinates are not reducible to the filtered high-state probability of an ordered MS-GJR-GARCH-t benchmark.

5. First-stage empirical evidence

5.1 Separation from current volatility

Table 2 reports contemporaneous correlations with 22-day realized volatility for all four coordinates. Scale is strongly related to the current volatility level in every market. Shape behaves differently: point estimates lie between -0.046 and 0.044 and every 95% block interval contains zero. Signed location is negatively associated with realized volatility in every market, with estimates from -0.418 to -0.251 . This pattern indicates that higher-volatility windows tend to have lower medians relative to their frozen high-volatility prototype after normalization by current robust scale. Unlike a classical leverage-effect

specification, which relates past negative returns to subsequent conditional variance and is often studied through asymmetric news-impact responses (Engle and Ng, 1993), the location statistic measures a contemporaneous association between a rolling window's normalized median displacement and realized volatility. It therefore identifies neither a predictive leverage response nor a causal channel. Speed is positively associated with current volatility, with correlations from 0.429 to 0.561, but remains materially less collinear than scale.

Table 2. Baseline component correlations with 22-day realized volatility

Market	Scale-RV22 [95% CI]	Shape-RV22 [95% CI]	Location-RV22 [95% CI]	Speed-RV22 [95% CI]
S&P 500	0.818 [0.800, 0.838]	-0.046 [-0.102, 0.020]	-0.418 [-0.462, -0.374]	0.561 [0.504, 0.608]
FTSE 100	0.854 [0.840, 0.870]	0.026 [-0.037, 0.093]	-0.398 [-0.446, -0.350]	0.511 [0.461, 0.554]
Nikkei 225	0.763 [0.733, 0.792]	0.044 [-0.037, 0.127]	-0.251 [-0.311, -0.194]	0.429 [0.359, 0.489]
Hang Seng	0.817 [0.797, 0.840]	-0.044 [-0.108, 0.027]	-0.350 [-0.405, -0.293]	0.455 [0.409, 0.496]

Notes: Pearson correlations. Circular moving-block bootstrap, 1,999 replications; block length 11 for the baseline 22-day window. Intervals are percentile 95% intervals.

A paired bootstrap directly tests H1 using $\Delta = |\text{corr}(\text{Scale}, \text{RV})| - |\text{corr}(\text{Shape}, \text{RV})|$. The estimated gaps are 0.772 for the S&P 500 (95% CI [0.721, 0.819]), 0.829 for the FTSE 100 ([0.755, 0.862]), 0.719 for the Nikkei 225 ([0.616, 0.781]), and 0.773 for the Hang Seng ([0.712, 0.820]). Every one-sided bootstrap p-value equals 0.0005, the smallest attainable value under the add-one adjustment with 1,999 draws. Thus H1 is supported by a direct contrast, not only by separate confidence intervals. This does not yet establish a distinct economic mechanism. The speed result is likewise interpreted as concurrent transition intensity, not an early-warning or forecasting statistic.

5.2 Persistence across windows

Table 3 summarizes window sensitivity and persistence at the corresponding lag L. Scale remains strongly associated with realized volatility across all twelve market-window combinations. Its lag-L autocorrelation ranges from 0.380 to 0.685. Signed location is negatively associated with realized volatility throughout and becomes more persistent as the horizon lengthens, although short-horizon persistence is market-specific. Shape persistence is much weaker at 11 and 22 days and becomes more visible only in selected 66-day specifications. Speed displays moderate lag-L persistence, approximately 0.37–0.46, while its correlation with current volatility declines as the window length rises.

Table 3. Window sensitivity and window-matched persistence

Panel A. Correlations with window-matched realized volatility

Market	L	Scale-RV [95% CI]	Shape-RV [95% CI]	Location-RV [95% CI]	Speed-RV [95% CI]
S&P 500	11	0.796 [0.782, 0.813]	-0.039 [-0.074, -0.006]	-0.328 [-0.361, -0.295]	0.556 [0.509, 0.601]

DISTRIBUTIONAL VOLATILITY COORDINATES

Market	L	Scale-RV [95% CI]	Shape-RV [95% CI]	Location-RV [95% CI]	Speed-RV [95% CI]
S&P 500	22	0.818 [0.800, 0.838]	-0.046 [-0.102, 0.020]	-0.418 [-0.462, -0.374]	0.561 [0.504, 0.608]
S&P 500	66	0.844 [0.811, 0.877]	0.019 [-0.166, 0.211]	-0.486 [-0.550, -0.426]	0.468 [0.405, 0.520]
FTSE 100	11	0.813 [0.799, 0.827]	-0.027 [-0.061, 0.012]	-0.217 [-0.253, -0.177]	0.552 [0.510, 0.590]
FTSE 100	22	0.854 [0.840, 0.870]	0.026 [-0.037, 0.093]	-0.398 [-0.446, -0.350]	0.511 [0.461, 0.554]
FTSE 100	66	0.889 [0.868, 0.911]	0.186 [0.063, 0.311]	-0.595 [-0.652, -0.536]	0.429 [0.376, 0.472]
Nikkei 225	11	0.735 [0.715, 0.756]	-0.024 [-0.062, 0.017]	-0.152 [-0.195, -0.114]	0.443 [0.396, 0.488]
Nikkei 225	22	0.763 [0.733, 0.792]	0.044 [-0.037, 0.127]	-0.251 [-0.311, -0.194]	0.429 [0.359, 0.489]
Nikkei 225	66	0.798 [0.758, 0.835]	0.137 [0.006, 0.272]	-0.306 [-0.411, -0.203]	0.322 [0.250, 0.389]
Hang Seng	11	0.756 [0.734, 0.780]	-0.016 [-0.059, 0.033]	-0.251 [-0.298, -0.204]	0.462 [0.421, 0.502]
Hang Seng	22	0.817 [0.797, 0.840]	-0.044 [-0.108, 0.027]	-0.350 [-0.405, -0.293]	0.455 [0.409, 0.496]
Hang Seng	66	0.876 [0.853, 0.899]	0.088 [-0.008, 0.186]	-0.518 [-0.590, -0.447]	0.392 [0.333, 0.443]

Panel B. Window-matched autocorrelations

Market	L	ACF _t (scale) [95% CI]	ACF _t (shape) [95% CI]	ACF _t (location) [95% CI]	ACF _t (speed) [95% CI]
S&P 500	11	0.578 [0.542, 0.611]	-0.007 [-0.036, 0.026]	0.090 [0.051, 0.131]	0.431 [0.390, 0.469]
S&P 500	22	0.613 [0.562, 0.657]	0.043 [-0.022, 0.114]	0.154 [0.092, 0.219]	0.436 [0.378, 0.493]
S&P 500	66	0.586 [0.487, 0.677]	0.168 [0.074, 0.278]	0.285 [0.180, 0.387]	0.420 [0.318, 0.515]
FTSE 100	11	0.492 [0.453, 0.530]	-0.031 [-0.062, -0.001]	-0.024 [-0.066, 0.019]	0.455 [0.411, 0.498]
FTSE 100	22	0.579 [0.527, 0.622]	0.003 [-0.042, 0.053]	0.049 [-0.017, 0.112]	0.430 [0.378, 0.476]
FTSE 100	66	0.647 [0.578, 0.706]	0.090 [0.002, 0.192]	0.333 [0.237, 0.420]	0.421 [0.343, 0.497]
Nikkei 225	11	0.380 [0.340, 0.423]	0.017 [-0.020, 0.053]	0.018 [-0.020, 0.058]	0.387 [0.346, 0.427]
Nikkei 225	22	0.450 [0.394, 0.498]	-0.022 [-0.067, 0.024]	0.022 [-0.040, 0.081]	0.386 [0.330, 0.442]
Nikkei 225	66	0.530 [0.438, 0.616]	-0.057 [-0.127, 0.023]	0.167 [0.071, 0.262]	0.393 [0.316, 0.477]
Hang Seng	11	0.528 [0.492, 0.561]	0.008 [-0.029, 0.045]	0.093 [0.049, 0.136]	0.399 [0.357, 0.438]
Hang Seng	22	0.631 [0.586, 0.670]	0.038 [-0.011, 0.088]	0.218 [0.158, 0.275]	0.406 [0.357, 0.452]
Hang Seng	66	0.685 [0.610, 0.748]	-0.078 [-0.167, 0.013]	0.339 [0.255, 0.422]	0.373 [0.302, 0.445]

Notes: Each cell reports the point estimate and percentile 95% circular-block interval beneath it. ACF_t uses the same lag as the rolling-window length and resamples paired lag vectors. Correlations are contemporaneous with window-matched realized volatility. Bootstrap replications: 1,999.

These statistics provide preliminary support for H2 in the sense of heterogeneous empirical memory. Scale and speed survive the window-matched lag. Location persistence rises at longer horizons but is weak in some 11- and 22-day panels, while short-horizon standardized shape is generally weaker still. The result is not a contradiction: shape may be episodic rather than smoothly autoregressive, and a 66-day window may average transient configurations into a slower-moving object. Event-duration and local-projection evidence are needed before converting this pattern into a structural interpretation.

The revised multiplicity protocol treats coordinate–RV associations in Panel A and window-matched persistence statistics in Panel B as two conceptually distinct families because they address different hypotheses. Each family contains 48 statistics after completing the location evidence. For associations, Bonferroni uses $\alpha = 0.00104$ and retains 36 results; Benjamini–Hochberg control at 5% retains 39. Every scale, location, and speed correlation survives both corrections. Three shape correlations survive FDR—the S&P 500 at $L = 11$, FTSE 100 at $L = 66$, and Nikkei 225 at $L = 66$ —and none survives Bonferroni. For persistence, Bonferroni retains 32 results and FDR retains 33: all 12 scale and all 12 speed autocorrelations survive both, seven location autocorrelations survive Bonferroni and eight survive FDR, and only the S&P 500 $L = 66$ shape autocorrelation survives either correction.

Block-length sensitivity is also extended to the longest window. At $L = 22$, using blocks of 11, 22, or 44 leaves every inferential classification unchanged. At $L = 66$, blocks of 33, 66, or 132 preserve the positive scale and speed associations and persistence, the negative location associations, and positive location persistence in all markets. Shape remains the least stable coordinate: the Nikkei 225 $L = 66$ shape–RV interval excludes zero with block 33 but includes zero with blocks 66 and 132. This exception is reported rather than smoothed over and reinforces the cautious interpretation of long-window shape evidence.

5.3 Variance-matched shape placebo

A shape statistic should distinguish distributions whose first two moments are equalized but whose standardized quantile configurations differ. The revised placebo uses 5,000 paired replications with $n = 22$ and random seed 97531. In each replication, the reference and symmetric control are independent Student-t draws with five degrees of freedom. The asymmetric alternative starts from the same t_5 family, multiplying negative observations by 1.8 and nonnegative observations by 0.55. Every sample is then recentered to sample mean zero and rescaled to sample standard deviation one before the robust median/IQR shape normalization is applied.

The mean distance from the reference is 0.543 for the asymmetric alternative (95% Monte Carlo CI [0.536, 0.550]) and 0.389 for the symmetric control ([0.384, 0.394]). The paired mean discrimination

effect is 0.154 ([0.146, 0.161]). Reporting uncertainty for the paired mean replaces the earlier arbitrary 0.10 exceedance threshold and distinguishes systematic shape sensitivity from the positive finite-sample distance between any two empirical distributions. This validates identification capability under a controlled alternative; it does not prove that empirical market shape is economically meaningful. A fuller placebo should also preserve time-series dependence, empirical tail frequencies, and window overlap.

5.4 Cross-market synchronization

The four panels are aligned on common calendar dates and pairwise Pearson correlations are averaged across the six unordered market pairs. Table 4 shows a clear ordering. Realized volatility is most synchronized, followed by scale, signed location, speed, and shape. Shape speed is similarly weak. To remove dependence on arbitrary pair naming, the one-day statistic averages both lead directions within each pair before averaging across pairs. No directional transmission claim is made.

Table 4. Mean cross-market synchronization

Component	Same day	Symmetric ± 1 -day mean	Interpretive status
RV22	0.725	0.723	Common volatility level
Scale	0.568	0.567	Strong shared component
Shape	0.049	0.047	Mostly market-specific
Location	0.350	0.342	Moderate shared direction
Speed	0.290	0.243	Intermediate synchronization
Shape speed	0.032	0.021	Weak synchronization

Notes: Means across six pairwise Pearson correlations on common dates. For each unordered pair, the ± 1 -day statistic averages both lead directions and is invariant to left/right naming. No causal or directional interpretation is made.

The ordering supports the primary contrast in H4: scale is the most synchronized of the four coordinates. H4 deliberately imposes no strict ex ante ranking among the remaining coordinates. In the observed sample, signed location is moderately synchronized, consistent with some shared directional movement after scale normalization; speed lies below location but above shape, consistent with partially shared transition intensity and partially local adjustment. Standardized distributional shape is not a global volatility factor reproduced across markets. Formal claims require date-block bootstrap intervals and a pre-specified multiple-testing correction.

6. Interpretation and model boundary

The evidence suggests a four-coordinate empirical organization. Scale is the slow, globally synchronized coordinate closest to conventional volatility. Signed location captures directional center shifts and is negatively associated with volatility, with persistence that strengthens at longer horizons.

Shape is comparatively orthogonal to the current level and weakly synchronized across markets. Speed records distributional motion: it co-moves with volatility but is not merely volatility, and it has more cross-market structure than shape.

This organization changes the language of persistence. A high-volatility episode may persist because dispersion remains elevated even after the distribution stops moving. A transition may display high speed before settling into either a high- or low-scale state. A local downside-tail distortion may appear in shape without a comparable global scale shock. These are measurable distinctions. They are not yet claims about economic causes.

6.1 What remains for the Markov-switching comparison

The frozen design specifies an ordered MS-GJR-GARCH-t benchmark and treats the filtered probability of the high-volatility state as the primary comparison. That daily filtered path is not available in the current evidence package, so no correlation, nested regression, classification, or state-recovery result against the benchmark is reported here. Substituting a different state model after inspecting the component results would weaken the design. The comparison will therefore enter only when the pre-specified daily state path and its estimation diagnostics are reproduced.

The benchmark stage should test whether each coordinate retains conditional variation after controlling for filtered state probability, whether speed peaks during state transitions rather than stable high-state periods, and whether shape distinguishes windows with similar state probabilities. Smoothed probabilities should appear only in robustness tables.

The present-stage contribution does not depend on winning that benchmark comparison. It supplies observable, unit-specific coordinates with transparent information timing; formally documents scale-shape separation; and reveals different persistence and synchronization profiles that a single latent high-state probability does not identify by construction. The future Markov-switching exercise is therefore an incremental-redundancy test—whether the new coordinates add information conditional on a mainstream state model—not the source of their definitions or their first-stage empirical validity.

6.2 Forecasting, early warning, and state discrimination

Association with current realized volatility is not forecasting. Separation of historical event windows is not early warning. Classification of a contemporaneous latent state is not return predictability. The current draft makes only measurement and descriptive state claims. Any forecasting extension must use expanding or rolling estimation, strictly out-of-sample losses, a competitive HAR or Log-HAR benchmark, and formal forecast-comparison tests. This standard is especially important given recent evidence that

even large pretrained time-series models achieve uneven and often narrow gains over disciplined econometric baselines (Brini, 2026).

7. Robustness and inference roadmap

The revised first-stage analysis now attaches block intervals to all four coordinates, reports Bonferroni and false-discovery-rate sensitivity for two separately defined 48-test families, and checks both the baseline and longest windows against longer blocks. Four additions remain for a final empirical paper. First, date-block bootstrap intervals and multiplicity control should be attached to the full cross-market matrix. Second, event-level persistence should be estimated with pre-specified start rules, recovery thresholds, and local-projection decay rather than inferred only from autocorrelation. Third, shape identification should be challenged with dependence-preserving and empirical-tail-frequency-matched placebos. Fourth, the filtered MS-GJR-GARCH-t state path must be reproduced from a frozen estimation script and compared with the coordinates using nested models and state-transition diagnostics.

Optional forecasting tests should remain secondary. They may ask whether lagged coordinates improve forecasts of future realized volatility conditional on HAR terms, but a null result would not invalidate the measurement contribution. Conversely, an in-sample association should never be presented as a forecasting gain.

8. Limitations

The baseline rolling distributions contain only 22 observations. This supports timely state measurement but limits quantile resolution, especially in the tails. Window overlap induces mechanical dependence, which is partly addressed by block resampling and window-matched lags but not eliminated. The pre-2007 high-volatility prototype is deliberately frozen; it may become less representative as market microstructure and policy regimes change. Robust scale normalization separates first-order dispersion from shape, but no finite-sample standardization can make the coordinates perfectly orthogonal.

The evidence is drawn from four equity indices. Cross-asset generalization to rates, foreign exchange, commodities, or cryptocurrencies is not established. Cross-market synchronization may also reflect common trading hours, asynchronous closes, exchange-rate exposure, and shared news. Finally, the present draft lacks the daily ordered MS-GJR-GARCH-t benchmark, event-recovery analysis, and out-of-sample tests. Its conclusions are correspondingly limited to first-stage measurement evidence.

9. Conclusion

Volatility persistence is not necessarily one-dimensional. A rolling return distribution can remain unusual because its dispersion stays high, because its standardized shape remains distorted, because its center shifts, or because it continues to move rapidly through distribution space. The proposed scale–shape–location–speed framework makes these possibilities observable without using future information.

Across four major equity markets, scale behaves like a persistent and globally synchronized volatility coordinate. Signed location is negatively associated with current realized volatility and displays moderate cross-market synchronization, while its persistence strengthens with window length. Shape is nearly orthogonal to current realized volatility at the baseline horizon and weakly synchronized across markets. Speed is a distinct transition measure with intermediate association and synchronization. Formal paired tests and multiplicity corrections preserve these contrasts. A paired heavy-tailed placebo confirms that standardized shape can identify asymmetry after first and second moments are equalized.

The conclusion is precise: distributional volatility persistence can be decomposed into empirically distinguishable coordinates. Whether those coordinates improve state models, reveal different recovery mechanisms, or add out-of-sample forecasting value remains to be tested under the frozen benchmark design.

Declarations

Funding

No external funding was received for this study.

Declarations of Interest

The author reports no conflicts of interest and is solely responsible for the content and writing of the paper.

Data and code availability

The analysis uses daily index return panels and reproducible Python scripts for component construction, bootstrap diagnostics, placebo simulation, and cross-market synchronization. Code, derived component panels, and source-specific access details are available from the author, subject to the redistribution terms of the underlying market data.

Use of generative artificial intelligence

OpenAI ChatGPT was used as an editing aid for language, document structure, and code review. The author specified the research design, verified the calculations and references, reviewed all substantive claims, and remains responsible for the integrity and accuracy of the work.

References

- Agueh, M., & Carlier, G. (2011). Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2), 904–924. <https://doi.org/10.1137/100805741>
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2001). The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, 96(453), 42–55. <https://doi.org/10.1198/016214501750332965>
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2), 579–625. <https://doi.org/10.1111/1468-0262.00418>
- Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B*, 64(2), 253–280. <https://doi.org/10.1111/1467-9868.00336>
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327. [https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1)
- Brini, A. (2026). Forecasting realized volatility with time series foundation models: A comparison with econometric benchmarks. *arXiv:2607.05291*.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174–196. <https://doi.org/10.1093/jffinec/nbp001>
- Engle, R. F., & Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The Journal of Finance*, 48(5), 1749–1778. <https://doi.org/10.1111/j.1540-6261.1993.tb05127.x>
- Fan, Z., Wang, M. (Melody), & Ye, Y. (2026). On options-driven realized volatility forecasting: Information gains via rough volatility model. *arXiv:2604.02743*.
- Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5), 1779–1801. <https://doi.org/10.1111/j.1540-6261.1993.tb05128.x>
- Gray, S. F. (1996). Modeling the conditional distribution of interest rates as a regime-switching process. *Journal of Financial Economics*, 42(1), 27–62. [https://doi.org/10.1016/0304-405X\(96\)00875-6](https://doi.org/10.1016/0304-405X(96)00875-6)
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384. <https://doi.org/10.2307/1912559>
- Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, 20(7), 873–889. <https://doi.org/10.1002/jae.800>
- Horvath, B., Issa, Z., & Muguruza, A. (2024). Clustering market regimes using the Wasserstein distance. *Journal of Computational Finance*, 28(1), 1–39. <https://doi.org/10.21314/JCF.2024.005>
- Leushuis, R. M., & Petkov, N. (2026). Advances in forecasting realized volatility: A review of methodologies. *Financial Innovation*, 12, Article 14. <https://doi.org/10.1186/s40854-025-00809-5>
- Panaretos, V. M., & Zemel, Y. (2019). Statistical aspects of Wasserstein distances. *Annual Review of Statistics and Its Application*, 6, 405–431. <https://doi.org/10.1146/annurev-statistics-030718-104938>
- Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256. <https://doi.org/10.1016/j.jeconom.2010.03.034>
- Patton, A. J., Politis, D. N., & White, H. (2009). Correction to “Automatic block-length selection for the dependent bootstrap.” *Econometric Reviews*, 28(4), 372–375. <https://doi.org/10.1080/07474930802459016>
- Politis, D. N., & White, H. (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 23(1), 53–70. <https://doi.org/10.1081/ETC-120028836>
- Rousseeuw, P. J., & Croux, C. (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424), 1273–1283. <https://doi.org/10.1080/01621459.1993.10476408>
- Zhang, C., Kokoszka, P., & Petersen, A. (2022). Wasserstein autoregressive models for density time series. *Journal of Time Series Analysis*, 43(1), 30–52. <https://doi.org/10.1111/jtsa.12590>